

Performance evaluation of beam breakup instability analysis of energy recovery linac using physics-inspired neural networks

M. K. Joshi^{1,2,*}, S. P. Sethi,³ S. Setiniyaz⁴, N. Nuchsirikulaphong^{1,2} and R. Apsimon^{1,2,†}

¹*Engineering Department, Lancaster University, Lancaster, LA1 4YW, United Kingdom*

²*Cockcroft Institute, Daresbury Laboratory, Warrington, WA4 4AD, United Kingdom*

³*Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal, Karnataka, 576104, India*

⁴*Center for Advanced Studies of Accelerators, Jefferson Lab, Newport News, USA*



(Received 24 March 2026; accepted 1 September 2026; published 17 September 2026)

Beam breakup (BBU) instability limits the performance of energy recovery linacs (ERLs). This study presents a new approach to estimating the BBU threshold current using a physics-inspired neural network. The application of neural networks to the highly nonlinear beam dynamics of ERLs opens a new paradigm for the development of efficient future particle accelerators. It is found that figure-of-merit-based data can predict the threshold current with higher accuracy than directly feeding the proposed network with raw parameter-based data. This study investigates how embedding such a physically meaningful feature influences predictive performance and generalization. In this work, the effect of the variation in higher-order modes (HOM) parameters due to the geometrical tolerances in the rf cavity fabrication on the BBU threshold current is analyzed with the proposed multilayered bidirectional long short-term memory (BiLSTM) neural network.

DOI: [10.1103/9kpd-tlbg](https://doi.org/10.1103/9kpd-tlbg)

I. INTRODUCTION

Energy recovery linacs (ERLs) represent a transformative technology for next-generation particle accelerators, enabling high beam currents and powers with exceptional energy efficiency [1,2]. In an ERL, electron bunches are accelerated through superconducting radio-frequency (rf) cavities and subsequently decelerated through the same cavities, transferring their energy back to the rf field to accelerate fresh bunches [3]. This energy recovery mechanism makes ERLs attractive for applications ranging from future electron-hadron colliders, such as the large hadron-electron collider and the future circular collider electron-hadron [4–6], to synchrotron radiation sources and free-electron lasers (FELs). The European Strategy for Particle Physics (ESPP), 2020, identified research and development in superconducting ERLs as a priority area for future accelerator-based research facilities [7].

A primary challenge limiting the performance of high-current multipass ERLs is beam breakup (BBU) instability, a regenerative instability which arises from the coupling

between the electron beam and higher-order modes (HOMs) in the accelerating cavities [8,9]. The BBU threshold current is the maximum sustainable beam current that arises before instability onset. It depends on numerous interrelated parameters, including higher-order modes (HOM) properties, beam optics, filling patterns, and machine configuration [10–13]. Recent studies have shown that filling patterns and bunch timing can affect the BBU threshold current by factors of five or more [13–15], highlighting the complex high-dimensional parameter space governing BBU behavior.

Traditional analytical and numerical simulation methods for predicting BBU threshold currents, while physically rigorous, can be computationally intensive and may struggle to capture complex nonlinear interactions across the full parameter space. Artificial intelligence (AI) and machine learning (ML) techniques offer promising complementary approaches for accelerator physics applications [16,17], with demonstrated success in areas, such as beam dynamics optimization, fault classification, and real-time control systems. These methods excel at identifying patterns and correlations in high-dimensional data that conventional approaches may struggle to capture.

Deep neural networks and Bayesian optimization have been used to develop surrogate models that replicate complex nonlinear beam dynamics physics much faster than full simulations [16,18,19]. Neural network-based surrogate models have demonstrated the ability to reproduce

*Contact author: m.k.joshi@lancaster.ac.uk

†Contact author: r.apsimon@lancaster.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

expensive beam physics simulations with orders-of-magnitude reductions in computation time, enabling fast multiobjective optimization and real-time parameter exploration. Neural network-inspired virtual diagnostics approaches have also been proposed to efficiently predict quantities, such as phase-space distributions. This improves real-time system visibility and decision support without interfering with beam operations [18,20].

The performance improvement, especially in ERLs, using machine learning methods has not been widely explored in the literature. However, the broader trends in AI/ML for accelerator physics show clear pathways for ERL applications. In particular, models based on physics-inspired neural networks have the potential to achieve reasonably accurate predictions in the beam dynamics analysis of particle accelerators [21,22].

This work presents a ML framework for predicting BBU threshold currents in multipass ERLs. We use the Powerful ERL for Experiments (PERLE) as a case study, which is a novel multipass ERL test facility designed to demonstrate high-current (20 mA injection), continuous-wave operation with 802 MHz superconducting cavities at 10 MW beam power [3,4]. Although the proposed analysis in this work focuses on PERLE's specific architecture, the ML methodologies developed are broadly applicable to any multipass ERL design. In this study, the BBU threshold current is predicted using a physics-inspired neural network. The proposed method offers a computationally efficient tool for ERL design optimization and operational planning applicable to future facilities.

The objective of this work is not to replace conventional BBU simulations. Instead, the proposed bidirectional long short-term memory (BiLSTM) model serves as a computational surrogate for applications that require repeated evaluations of the BBU threshold current, such as large-scale tolerance studies, uncertainty quantification, machine optimization, and the investigation of various filling patterns. Once trained, the proposed neural network predicts the threshold current almost instantaneously, making it suitable for rapid exploration of high-dimensional parameter spaces and optimization studies while maintaining consistency with the underlying BBU simulations. Future work will focus on validating the proposed surrogate model using physically correlated perturbation scenarios.

This article is structured in six sections. The BBU instabilities and threshold current in the ERLs are discussed in Sec. II. The dependencies of the key HOM parameters on the threshold current are also presented in this section. The BBU threshold current depends on the resonant frequency (f_0), loaded quality factor (Q_L), and geometric shunt impedance (R/Q) of the HOMs. These crucial HOM parameters are used to train and predict the BBU threshold current using the proposed neural network. Section III deals with the generation of data for the desired neural network. In this section, the data acquisition process

for generating samples for efficient neural network training is discussed. The methods used for data preprocessing before feeding the data to the neural network are also presented. Section IV provides the architecture of the proposed neural network. The results of the proposed neural network are discussed in Sec. V. Section VI concludes this article.

II. BEAM BREAKUP INSTABILITY

BBU instability is a regenerative instability that fundamentally limits the achievable beam current in multipass energy recovery linacs. The physical mechanism proceeds as follows: when an electron bunch traverses an rf cavity with a transverse displacement from the cavity axis, it excites higher-order dipole modes (HOMs) in the cavity through wake fields. In this work, only dipole HOMs are considered as they cause transverse BBU, which is critical to ERL machines. Monopole HOMs primarily affect longitudinal beam dynamics and are therefore not included in the present analysis. Throughout this article, the term HOM refers exclusively to dipole HOMs unless otherwise stated. In a multipass ERL, these excited HOMs persist due to their finite quality factors and deliver transverse momentum kicks to subsequent bunches passing through the same cavity. When these kicked bunches recirculate and return to the cavity with increased transverse displacement and appropriate phase relative to the HOM, they deposit additional energy into the mode, creating a positive feedback mechanism. If the beam current exceeds a critical threshold value, this regenerative process leads to exponential growth of both the HOM amplitude and beam oscillations, ultimately resulting in beam loss [8,10].

The BBU threshold current depends on several key parameters, such as HOM properties, transfer matrix elements, recirculation time, and beam energy. The HOM properties are characterized by the resonant frequency of higher-order mode (f_H) (or angular frequency $\omega_H = 2\pi f_H$), the geometric shunt impedance ($\frac{R}{Q}_H$), which quantifies the coupling strength between the beam and the mode, and the loaded quality factor Q_L , which determines how long the mode persists in the cavity. The beam optics plays a crucial role through the transfer matrix elements R_{12} and R_{34} , which relate the angular kick received by a particle to its transverse offset after one complete recirculation through the machine. The recirculation time (or revolution period) T_{rev} determines the phase relationship between successive bunch passages through the cavity. Finally, the beam energy E_{beam} affects the threshold current, as higher-energy beams are more resistant to transverse deflections.

A. Current convention

It is important to clarify the convention used for reporting threshold currents in multipass ERLs. In previous work by

Setiniyaz *et al.* [15], the threshold current refers to the beam current passing through the cavities during operation. For a multipass ERL, this cavity current is larger than the injected current by a factor equal to the number of passes. In contrast, the present work adopts the injected-current convention, in which the threshold current refers to the beam current injected into the machine. For PERLE, which is a six-pass ERL (three accelerating passes and three decelerating passes), the relationship between these two conventions is

$$I_{\text{th, cavity}} = 6 \times I_{\text{th, injection}}. \quad (1)$$

This distinction is crucial when comparing results across different studies and must be carefully considered when applying threshold current predictions to operational scenarios.

B. BBU simulation algorithm

The BBU simulation tool used in this work is based on the algorithm developed in [13,14] and extended to handle multicavity, multimode configurations with filling pattern and bunch timing dependence [15]. The algorithm builds upon the ERLBBU framework originally developed by Pozdeyev *et al.* [8,10] and Tennant [23], with key enhancements to capture the effects of different filling patterns and recombination schemes on BBU dynamics [12–15].

The simulation tracks the evolution of both the HOM amplitude in each cavity and the transverse displacement of each bunch as it passes through the linac multiple times. For a given machine configuration, HOM spectrum, filling pattern, and bunch timing, the algorithm iteratively computes: (i) The transverse kick imparted to each bunch by the existing HOM fields in each cavity. (ii) The evolution of each bunch's transverse position through the recirculation arcs, governed by the beam optics transfer matrices. (iii) The energy deposited into each HOM by each passing bunch, which depends on the bunch's transverse offset and the HOM properties. (iv) The natural decay of HOM amplitudes between bunch passages, determined by the loaded quality factors.

The threshold current is determined by incrementally increasing the beam current until the coupled bunch-HOM system exhibits exponential growth, indicating the onset of instability. The simulation accounts for the complex phase relationships that arise from different filling patterns, where bunches on different passes may arrive at cavities at different rf phases, leading to varying recirculation times and modified HOM-beam coupling dynamics [12,14].

C. Threshold current frequency dependence

The BBU threshold current equation for a simple two-pass ERL offers important physical insight into the instability mechanism [23]:

$$I_{\text{th}} = \frac{-2E_{\text{beam}}}{e \frac{\omega_{\text{H}}}{c} \left(\frac{R}{Q}\right)_{\text{H}} Q_{\text{L}} R_{12} \sin(\omega_{\text{H}} T_{\text{rev}})}, \quad (2)$$

where E_{beam} is the beam energy in electron volts, e is the elementary charge, c is the speed of light in vacuum, ω_{H} is the HOM angular frequency, $\left(\frac{R}{Q}\right)_{\text{H}}$ is the geometric shunt impedance of the HOM, Q_{L} is the loaded quality factor of the HOM, R_{12} is the (1, 2) element of the one-turn transfer matrix relating an angular kick at the cavity to the horizontal offset upon return (replaced by R_{34} for the vertical plane), and T_{rev} is the revolution period. Q_{L}

Although Eq. (2) is strictly valid only for a simple two-turn ERL and does not capture the full complexity of multiturn machines, it nonetheless reveals several key features of the threshold behavior. Most notably, the $\sin(\omega_{\text{H}} T_{\text{rev}})$ factor in the denominator introduces a periodic dependence of I_{th} on the HOM frequency: as this term oscillates between -1 and $+1$, the threshold current can vary by orders of magnitude—and even become negative, which is conventionally interpreted as the beam being unconditionally stable with respect to that HOM at that frequency. This highly nonlinear frequency sensitivity stands in sharp contrast to the straightforward inverse dependence on parameters such as $\left(\frac{R}{Q}\right)_{\text{H}}$ and Q_{L} , and it is precisely this behavior that makes accurate BBU threshold prediction with a neural network particularly challenging.

The dominance of a given HOM can be characterized by the figure-of-merit (ζ) [24], defined as

$$\zeta = \frac{\left(\frac{R}{Q}\right)_{\text{H}} \sqrt{Q_{\text{L}}}}{f_{\text{H}}}. \quad (3)$$

This quantity provides a physically meaningful measure of the relative impact of a HOM on beam stability. In particular, it scales inversely with the BBU threshold current (I_{th}), so modes with larger ζ values tend to drive instability thresholds and therefore are more likely to dominate the BBU behavior. As the effects of only HOMs are studied, for brevity, HOM parameters f_{H} and $\left(\frac{R}{Q}\right)_{\text{H}}$ are denoted as f_0 and R/Q in the remaining sections of this article.

III. DATA GENERATION

The objective of the data generation strategy is to capture the range of variations expected from manufacturing tolerances in a multicavity ERL, rather than modeling a single physical perturbation mechanism. Samples for training the neural network (NN) are generated from PERLE HOM data [25,26]. In this data set, resonant frequency of the HOMs (f_0), loaded quality factor (Q_{L}), and shunt impedance (R/Q) of the first 30 HOMs are considered. The flow diagram of the data generation is depicted in Fig. 1. The original HOM data set is used to generate N samples ($N = 3000$) for the neural network.




FIG. 1. Data (sample) generation flow.

Further, these input samples ($HOM_1, HOM_2, \dots, HOM_N$) are fed into BBU simulator [13,14]. The BBU simulator generates a threshold current (I_{th}^i) corresponding to each HOM input sample (HOM_i). The remaining ERL parameters and simulation constraints are kept the same for the BBU analysis. The details of the HOM sample generator are provided in the next sub-section.

A. Data acquisition

HOM input samples are generated based on the following constraints on the parameters ($f_0, Q_L, R/Q$). For each HOM, the figure-of-merit (ζ_i , where $i = 1, 2, \dots, 30$) is determined from Eq. (3). The larger ζ_i governs more BBU drive. The BBU threshold current depends on whether there is a single dominant HOM or multiple competing HOMs. Mode dominance is controlled to realize the realistic scenario in the generated HOM samples, and for this, a parameter fractional dominance (D_{frac}) is defined, as expressed in Eq. (4),

$$D_{frac} = \frac{\zeta_1 - \zeta_2}{\zeta_1}, \quad (4)$$

where ζ_1 and ζ_2 are figures-of-merit of first and second dominant HOM, respectively. $D_{frac} = 0$ denotes that there is no single dominant HOM, and two or more HOMs are equally likely to be responsible for BBU. Similarly, $D_{frac} = 1$ indicates that one mode is extremely dominant compared to the rest of the HOMs. In this study, mild dominance ($D_{frac} = 0.05$ or 5%) is chosen to account for rf cavity fabrication tolerances. In the HOM sample generator, rf cavity parameters (HOM frequency (f_0), loaded quality factor (Q_L), and shunt impedance (R/Q)) vary in such a way that the first dominant mode's ζ is at least 5% higher than the next competing HOM. The upper limit on ζ_i is also set to 7000 to avoid the overshoots.

In this study, fabrication tolerances are modeled by introducing controlled variations in the HOM parameters of otherwise identical cavities. The variation ranges of all these HOM parameters are derived from the PERLE cavity HOM dataset [25,26]. The resonant frequencies (f_0) of HOMs vary relatively from their original values. The HOM frequency spread (σ_f) with respect to the original HOM frequency is chosen as 1×10^{-4} . Small and pair-correlated perturbations are introduced to reflect realistic

machining-induced frequency shifts while preserving the natural mode pairing between consecutive HOMs (HOM_1 and HOM_2 , HOM_3 and HOM_4 , $\dots$). HOMs occur in pairs considering the horizontal and vertical polarization, so the resonant frequencies of the consecutive HOMs are nearly the same. The loaded quality factors (Q_L) are varied log-uniformly over the range from 150 to 2.9×10^5 in accordance with the reference HOM data. This variation captures the wide dynamic span and multiplicative nature of damping variations arising from coupler and surface effects in the rf cavity. The transverse shunt impedance values are sampled uniformly within the physically admissible interval ($0.1 \, \Omega$ to $55 \, \Omega$) defined by the minimum and maximum R/Q values of the reference PERLE HOM dataset. Consequently, every generated HOM sample remains confined to the parameter space represented by the original cavity rather than exploring arbitrary values outside the physically relevant domain. This represents bounded and approximately symmetric deviations due to geometric fabrication errors. Together, these variations provide a physically consistent ensemble of rf cavities for studying the impact of manufacturing tolerances on beam breakup behavior and for training data-driven prediction models.

B. Data preprocessing

The input parameters are preprocessed before being fed to the NN for efficient learning. The relative variation from the mean of the respective parameter is considered instead of the raw input samples ($f_0, Q_L, R/Q$). For each generated HOM sample, the mean value of a given HOM parameter is first computed over the 30 dipole HOMs. The deviation of each HOM from this sample-specific mean is then used as the NN input, which is expressed in Eq. (5),

$$\Delta x_i^j = x_i^j - \bar{x}^j, \quad \forall i = 1, 2, \dots, 30$$

$$\bar{x}^j = \frac{1}{30} \sum_{i=1}^{30} x_i^j, \quad (5)$$

where x denotes any HOM parameter ($f_0, Q_L, R/Q$), i is the HOM mode index, and j is the sample index. Thus, $\bar{x}^j$ is the mean value of the selected HOM parameter over all 30 HOMs within the j th generated sample, and Δx_i^j represents the deviation of the i th HOM from this sample-specific mean. Based on this, the relative variations of the parameters ($\Delta f_0, \Delta Q_L, \Delta R/Q$) are determined. This preprocessing removes the common offset of a given parameter across each generated HOM sample while preserving the relative distribution, which governs mode competition and, consequently, the BBU threshold current.

To introduce structured physical information into the learning process of the NN, an additional input feature (ζ) is also incorporated. Hence, the additional input parameter ($\Delta \zeta$), which is a physics-motivated parameter, determines

the relative variation of the figure-of-merit (ζ) for the NN feed.

These four parameters ($\Delta f_0, \Delta Q_L, \Delta R/Q, \Delta \zeta$) are the inputs, and the corresponding I_{th} is the output of the proposed neural network, which is discussed in the next section.

Out of 3000 samples, 70% are randomly selected for training, and the remaining 30% for NN testing. To improve generalization, min-max normalization is applied to the training data before feeding it to the neural network, as presented in Eq. (6),

$$X_{\text{norm}} = \frac{X - x_{\min}}{x_{\max} - x_{\min}}, \quad (6)$$

where X is any input parameter ($\Delta f_0, \Delta Q_L, \Delta R/Q, \Delta \zeta$) for a sample. $x_{\max}$ and $x_{\min}$ are the maximum and minimum values of the given input parameter, respectively. In the study, min-max normalization technique is applied solely to rescale the input features to a common range to improve numerical stability and facilitate network training. The network is trained exclusively on the normalized feature vectors.

IV. PHYSICS-INSPIRED NN-BASED BBU THRESHOLD CURRENT ESTIMATION

To evaluate the beam current performance under varying input parameters, a physics-inspired neural network is implemented. Since the acquired data exhibit sequential characteristics and standard machine learning approaches for sequence data modeling, namely recurrent neural network (RNN)-based models and their variants, including long short-term memory (LSTM) and BiLSTM, are considered in the proposed model (Fig. 2). Although RNNs are widely used for sequence modeling, they often suffer from the vanishing gradient problem and rely solely on past information for weight updates, limiting their effectiveness for the present application. While LSTM mitigates the vanishing gradient issue, it processes information only in the forward direction. Therefore, a BiLSTM, which leverages both forward and backward contextual dependencies, is adopted to enhance predictive performance. Initially, the acquired raw data are directly fed to the neural networks, but the prediction accuracy was not desirable; therefore, the raw data are preprocessed using a physics-inspired technique based on the figure-of-merit (ζ), as defined in Eq. (3). This approach does not impose physical equations into the loss function. Instead, it enhances the input feature space with a derived quantity that reflects known physics. Therefore, the model is described as a physics-inspired neural network.

A. Model setup

In this study, a multilayer BiLSTM neural network architecture is considered. The proposed model estimates




FIG. 2. Architecture of the neural network (NN).

the BBU threshold current from sequence data. BiLSTM neural network is an optimized variant of the conventional LSTM architecture (Fig. 2). Unlike the standard LSTM model, which generates predictions solely from past sequence data, a BiLSTM network integrates information from both past and future sequence data. This is achieved by combining forward and backward LSTM layers, enabling the model to process sequence information in both temporal directions and thereby improving its robustness and predictive capability. Therefore, in this work, a two-layer BiLSTM network (BiLSTM 1 and BiLSTM 2) is statistically considered to learn bidirectional sequential features generated from LSTM models. Also, it effectively captures long-term dependencies in the data, with the final threshold current generated by a fully connected layer. A composite loss function [expressed in Eq. (7)] layer is used after a fully connected layer to estimate the correlation between the input and output sequence data.

$$\text{Loss function} = \alpha \sigma_{\text{MSE}} + (1 - \alpha)(1 - R), \quad (7)$$

where σ_{MSE} is the mean square error, R is the Pearson correlation coefficient between the predicted and actual output, and α is the weight component for σ_{MSE} and R . The introduction of the correlation coefficient in the formulation of the loss function is crucial. Since the physical objective of the proposed NN involves capturing nonlinear scaling behavior rather than merely minimizing absolute deviation from the actual data, maximizing correlation

between prediction and actual data is considered physically meaningful.

The implemented BiLSTM architecture is shown in Fig. 2. The mathematical operation of forward and backward hidden states of the BiLSTM network is represented in Eqs. (8a) and (8b), respectively. Also, Eq. (8c) presents the input data sequence in the BiLSTM network, and Eq. (8d) presents the output of the neural network. The BiLSTM network processes the input sequence in both forward and backward temporal directions. Let LSTM_f denote the hidden representation generated by the forward LSTM layer, which processes the sequence from the first to the last time step, and LSTM_b denote the hidden representation generated by the backward LSTM layer, which processes the sequence in the reverse direction. The outputs of these two directions are concatenated to obtain the final BiLSTM feature representation, as expressed in Eq. (8),

$$h_t^f = \text{LSTM}_f(x_t, h_{t-1}^f), \quad (8a)$$

$$h_t^b = \text{LSTM}_b(x_t, h_{t-1}^b), \quad (8b)$$

$$x_t = [h_t^f; h_t^b], \quad (8c)$$

$$y_t = Wh_t + b, \quad (8d)$$

where x_t and x_{t+1} are the input data sequence for time steps t and $t+1$, respectively, which are fed to the LSTM networks. h_t^f and h_t^b are the forward and backward hidden states estimated for the t time step for the $t-1$ previous time step from the LSTM models. LSTM_f and LSTM_b are forward and backward LSTM layers, respectively. W is the weight matrix, b is the bias, and y_t and y_{t+1} are the output of the network in the time steps t and $t+1$, respectively.

In this study, raw data analysis, preprocessing, and model implementation are performed in a MATLAB environment and MATLAB's Deep Learning Toolbox [27].

B. Hyperparameter tuning

Hyperparameter tuning plays a crucial role in achieving optimal convergence with minimal mean squared error (σ_{MSE}) and improved correlation coefficients (R) for the BiLSTM network. The proposed network hyperparameters include the number of BiLSTM units and layers, dropout rates for dropout layers, number of epochs, mini-batch size, learning rate, and L2-regularization. Additionally, the tuning parameter α serves as a weighting factor in the loss function.

To effectively capture complex temporal dependencies while controlling computational complexity and mitigating overfitting, the number of BiLSTM units is varied across different experimental cases. In contrast, the number of layers is kept constant for all input parameter variations. To enhance generalization performance,

dropout regularization is applied, with the dropout rate determining the degree of regularization strength. The mini-batch size is also adjusted to improve learning stability and generalization capability. Other hyperparameters are also tuned statistically through extensive experimentation involving multiple combinations and variations. To ensure stable and efficient convergence, the adaptive moment estimation (ADAM) optimizer is implemented. The optimal set of hyperparameters supplied to the ADAM optimizer is obtained through Bayesian optimization, where the predefined search ranges of each hyperparameter are assigned. The optimal values of the hyperparameters are mentioned in Table I.

V. RESULTS & DISCUSSION

The NN discussed above is trained on three cases: (A) when a single input HOM parameter varies, (B) when two input HOM parameters vary, and (C) when all three input HOM parameters vary. All three cases and their respective sub-cases are discussed. The case-wise performance along with the optimized hyperparameters of the proposed NN is presented in Table I.

A. Single variable HOM parameter

In this case, any one of the three HOM parameters (i.e., f_0 , Q_L , R/Q) is varied while keeping the rest of the two parameters unchanged during the sample generation process. In the first experiment, R/Q is varied as discussed in Sec. III, while f_0 and Q_L are held constant as in the original HOM data of the PERLE. The proposed NN is fed with the $\Delta R/Q$ as input and the corresponding I_{th} as output. Here, f_0 and Q_L do not vary sample-wise and are therefore not considered for NN input. The predicted and actual outputs (I_{th}) are plotted in Fig. 3(a), which shows an excellent correlation with negligible error. The correlation coefficient of 0.9439 is achieved on the test data, with a mean squared error (σ_{MSE}) of 2.9×10^{-4} . The probability density function (PDF) of the predicted and actual data is shown in Fig. 3(a). A good overlap between the probability density functions (PDFs) of predicted and actual data, as shown in Fig. 3(c), indicates that the predicted output spans the entire range of the actual output. In the next phase of the experiment, the performance of the NN is improved when the NN is fed with a physics-inspired parameter ($\Delta\zeta$).

To introduce structured physical information into the learning process, the parameter $\Delta\zeta$ is incorporated as an input feature. The performance of the proposed NN with $\Delta\zeta$ feed is shown in Figs. 3(b) and 3(d). The NN performance improves as the network learns meaningful relations between input parameters and the output. The improved test correlation coefficient of 0.9923 is achieved with σ_{MSE} as 1.4×10^{-4} . Although the exact relationship between the entire set of input parameters and the output

TABLE I. Performance evaluation of the proposed physics-inspired BiLSTM model.

Cases	HOM var.	NN I/P feed	Hyperparameters									Performance		
			bilstm1 Units	bilstm2 Units	drop1	drop2	α	epochs	mini Batch	learn Rate	l2reg	σ_{MSE}	Train Corr.	Test Corr.
A	R/Q	$\Delta R/Q$	175	070	0.4449	0.5766	0.5605	360	032	3.3×10^{-3}	4.8×10^{-7}	2.9×10^{-4}	0.9730	0.9439
	R/Q	$\Delta \zeta$	130	033	0.1346	0.1006	0.3130	286	123	4.6×10^{-3}	1.0×10^{-7}	1.4×10^{-4}	0.9923	0.9698
	Q_L	ΔQ_L	113	329	0.3300	0.5958	0.6689	376	016	7.6×10^{-3}	5.3×10^{-7}	5.8×10^{-3}	0.9926	0.9404
	Q_L	$\Delta \zeta$	211	021	0.2743	0.1187	0.3057	448	036	6.3×10^{-4}	1.2×10^{-7}	2.8×10^{-3}	0.9946	0.9701
	f_0	Δf_0	115	093	0.1921	0.1690	0.7796	394	070	3.4×10^{-3}	2.4×10^{-7}	6.7×10^{-5}	0.9981	0.9971
	f_0	$\Delta \zeta$	145	232	0.1183	0.2827	0.6938	459	077	2.8×10^{-3}	1.1×10^{-7}	2.6×10^{-5}	0.9988	0.9986
B	f_0, Q_L	$\Delta f_0, \Delta Q_L$	108	040	0.1876	0.3820	0.8015	193	122	2.5×10^{-3}	1.7×10^{-6}	2.9×10^{-2}	0.6768	0.4830
	f_0, Q_L	$\Delta \zeta$	217	078	0.4494	0.1715	0.7280	355	036	1.0×10^{-3}	7.1×10^{-6}	8.7×10^{-3}	0.9787	0.8772
	$f_0, R/Q$	$\Delta f_0, \Delta R/Q$	280	061	0.4326	0.3075	0.8841	152	100	3.0×10^{-3}	2.5×10^{-7}	2.9×10^{-3}	0.3998	0.2983
	$f_0, R/Q$	$\Delta \zeta$	143	041	0.3525	0.3505	0.3553	183	075	1.2×10^{-3}	2.1×10^{-5}	1.4×10^{-3}	0.8400	0.7649
	$Q_L, R/Q$	$\Delta Q_L, \Delta R/Q$	130	042	0.4130	0.3577	0.4262	341	120	4.9×10^{-4}	2.8×10^{-6}	2.7×10^{-3}	0.6069	0.5393
	$Q_L, R/Q$	$\Delta \zeta$	174	126	0.5655	0.5413	0.4874	183	028	4.5×10^{-4}	1.5×10^{-5}	2.4×10^{-3}	0.7658	0.6314
C	$f_0, Q_L, R/Q$	$\Delta f_0, \Delta Q_L, \Delta R/Q$	308	075	0.2307	0.3568	0.4466	209	059	1.8×10^{-3}	3.6×10^{-7}	2.7×10^{-3}	0.5781	0.5016
	$f_0, Q_L, R/Q$	$\Delta \zeta$	290	067	0.5177	0.8502	0.1074	528	064	1.8×10^{-4}	1.4×10^{-5}	1.4×10^{-2}	0.6228	0.6175
	$f_0, Q_L, R/Q$	$\Delta \zeta, \Delta f_0, \Delta Q_L, \Delta R/Q$	301	166	0.4330	0.8980	0.4665	467	053	7.5×10^{-3}	2.3×10^{-6}	1.0×10^{-3}	0.9853	0.8518
	$f_0, Q_L, R/Q$ ($\sigma_f = 1 \times 10^{-3}$)	$\Delta \zeta, \Delta f_0, \Delta Q_L, \Delta R/Q$	1106	88	0.6528	0.4376	0.8954	483	119	1.0×10^{-4}	1.8×10^{-5}	5.0×10^{-3}	0.5629	0.5512
	$f_0, Q_L, R/Q$ ($\sigma_f = 1 \times 10^{-2}$)	$\Delta \zeta, \Delta f_0, \Delta Q_L, \Delta R/Q$	241	152	0.5380	0.7091	0.4784	279	021	1.9×10^{-3}	2.6×10^{-4}	6.9×10^{-3}	0.5901	0.5067
	$f_0, Q_L, R/Q$ ($\sigma_f = 1 \times 10^{-2}$)	$\Delta \zeta, \Delta f_0, \Delta Q_L, \Delta R/Q$	241	152	0.5380	0.7091	0.4784	279	021	1.9×10^{-3}	2.6×10^{-4}	6.9×10^{-3}	0.5901	0.5067

Note: Fraction dominance (D_{frac}) of 5% is considered in all the cases.

remains strongly nonlinear, ζ encapsulates physical insight regarding the modal excitation strength.

Similarly, in the next experiment, the parameter Q_L is varied while the remaining parameters are held constant.

The test correlation coefficient of 0.9404 is obtained, and the σ_{MSE} is obtained as 5.8×10^{-3} when the proposed NN is fed with $\Delta R/Q$ as input data. The performance of predicted and actual I_{th} is shown in Figs. 4(a) and 4(c).




FIG. 3. R/Q variation. Predicted and actual I_{th} with NN input as (a) $\Delta R/Q$, (b) $\Delta \zeta$. Probability density functions (PDFs) of the predicted and actual data with NN input as (c) $\Delta R/Q$, (d) $\Delta \zeta$.




FIG. 4. Q_L variation. Predicted and actual I_{th} with NN input as (a) ΔQ_L , (b) $\Delta \zeta$. PDFs of the predicted and actual data with NN input as (c) ΔQ_L , (d) $\Delta \zeta$.




FIG. 5. f_0 variation. Predicted and actual I_{th} with NN input as (a) Δf_0 , (b) $\Delta \zeta$. PDFs of the predicted and actual data with NN input as (c) Δf_0 , (d) $\Delta \zeta$.

If the NN is fed with $\Delta \zeta$, performance improves, as shown in Figs. 4(b) and 4(d). The improved correlation coefficient of 0.9701 and the σ_{MSE} of 2.8×10^{-3} is achieved with the physics-inspired $\Delta \zeta$ feed to the NN.

In continuation of this, the third parameter f_0 is varied, and the output of the NN is analyzed when input to the NN is given as Δf_0 and $\Delta \zeta$, respectively. The HOM frequency spread (σ_f) is taken as 1×10^{-4} while generating HOM samples for this case. The performance of NN is shown in Fig. 5. The test correlation coefficient and σ_{MSE} of 0.9971 and 6.7×10^{-5} are obtained when the input feed is given as Δf_0 . Similarly, the test correlation coefficient and σ_{MSE} of 0.9986 and 2.6×10^{-5} are obtained when the input feed is given as $\Delta \zeta$.

B. Two-variable HOM parameters

In the second case, while generating HOM samples, any two HOM parameters are varied while keeping the third parameter fixed. So, three combinations are possible. In the first experiment, f_0 and Q_L are varied. The values of R/Q in all N samples remain the same as provided in the original HOM data. The NN's performance is shown in Figs. 6(a), 6(c) for the input feed of Δf_0 and ΔQ_L . Figures 6(b) and 6(d) show the performance of $\Delta \zeta$ input feed. Test correlation coefficient of 0.4830 and σ_{MSE} of 2.9×10^{-2} is obtained when only HOM parameters (Δf_0 and ΔQ_L) are fed to NN. However, when $\Delta \zeta$ is fed to NN, the performance of prediction is improved, and test correlation coefficient and σ_{MSE} are obtained as 0.8772 and 8.7×10^{-3} , respectively.




FIG. 6. f_0 and Q_L variation. Predicted and actual I_{th} with NN input as (a) Δf_0 and ΔQ_L , (b) $\Delta \zeta$. PDFs of the predicted and actual data with NN input as (c) Δf_0 and ΔQ_L , (d) $\Delta \zeta$.




FIG. 7. f_0 and R/Q variation analysis. Predicted and actual I_{th} with NN input as (a) Δf_0 and $\Delta R/Q$, (b) $\Delta \zeta$. PDFs of the predicted and actual data with NN input as (c) Δf_0 and $\Delta R/Q$, (d) $\Delta \zeta$.

Physics-inspired NN performed better even in the next experiment when f_0 and R/Q are varied. Performance is shown in Fig. 7. The test correlation coefficient and σ_{MSE} are achieved as 0.2983 and 2.9×10^{-3} , respectively, when




FIG. 8. Q_L and R/Q variation analysis. Predicted and actual I_{th} with NN input as (a) ΔQ_L and $\Delta R/Q$, (b) $\Delta \zeta$. PDFs of the predicted and actual data with NN input as (c) ΔQ_L and $\Delta R/Q$, (d) $\Delta \zeta$.

Δf_0 and $\Delta R/Q$ are fed as input to NN. The performance of the proposed NN is proved when $\Delta \zeta$ is fed to the NN, the test correlation coefficient and σ_{MSE} are obtained as 0.7649 and 1.4×10^{-3} , respectively. Similarly, in the next experiment, Q_L and R/Q are varied while keeping f_0 fixed. The test correlation coefficient and σ_{MSE} are obtained as 0.5393 and 2.7×10^{-3} , respectively, with inputs to the NN provided as ΔQ_L and $\Delta R/Q$. However, physics-inspired $\Delta \zeta$ improves the test correlation coefficient to 0.6314 and the σ_{MSE} to 2.4×10^{-3} . The network's performance is shown in Fig. 8.

The network's performance in the two-variable case is lower than in the single-variable case due to increased randomness. Previously, only a single parameter was varied; in this case, two are, so increasing the randomness reduces the accuracy of the threshold current prediction.

C. All variable HOM parameters

All HOM parameters are varied in this case. In the first experiment, Δf_0 , ΔQ_L , $\Delta R/Q$ are fed to the proposed NN, and the test correlation coefficient and σ_{MSE} are observed as 0.5016 and 2.7×10^{-3} , respectively. When the network is fed with $\Delta \zeta$ as input, the test correlation coefficient and σ_{MSE} are observed as 0.6175 and 1.4×10^{-2} , respectively. The network performance is shown in Fig. 9. Although physics-inspired $\Delta \zeta$ as an input feed to the NN improves performance, alone, the $\Delta \zeta$ input feed is not sufficient for efficient learning in the network. As all three HOM parameters vary, it could be possible that the same value of




FIG. 9. f_0 , Q_L , and R/Q variation analysis. Predicted and actual I_{th} with NN input as (a) Δf_0 , ΔQ_L , and $\Delta R/Q$, (b) $\Delta \zeta$. PDFs of the predicted and actual data with NN input as (c) Δf_0 , ΔQ_L , and $\Delta R/Q$, (d) $\Delta \zeta$.

ζ is obtained from multiple combinations of HOM parameters. So, in addition to $\Delta \zeta$, all three HOM parameters (Δf_0 , ΔQ_L , $\Delta R/Q$) are also incorporated to improve network learning. This significantly enhances the prediction of the threshold current. The test correlation coefficient and σ_{MSE} are improved to 0.8518 and 1.0×10^{-3} , respectively. The predicted and actual threshold currents variation are shown in Fig. 10(a). The trend of the predicted output follows the straight line. The probability density function (PDF) of the relative prediction error ($100 \times \frac{I_{th}^{pred} - I_{th}^{act}}{I_{th}^{act}} \%$) is shown in Fig. 10(b). To avoid artificially large relative errors associated with very small actual threshold currents, values below 0.03 A are regularized in Fig. 10(b). In most cases, samples exhibit a relative prediction error within a reasonable range. Figure 10(c) shows the distribution of the predicted current along with the actual current. It shows that the predicted current follows the entire range of the actual current. Also, the predicted current follows the current of the sorted test samples, as shown in Fig. 10(d). These results indicate that the proposed network captures the dominant nonlinear structure governing the BBU threshold current rather than merely interpolating local patterns.

The BBU threshold current and the input HOMs have a highly nonlinear relationship. Moreover, the threshold current is quite sensitive to the HOM frequency. So far, in all cases, the HOM frequency spread (σ_f) is chosen as 1×10^{-4} , which corresponds to the general fabrication tolerances in the rf cavity. This model is also tested with




FIG. 10. f_0 , Q_L , and R/Q variation with $\Delta\zeta$, Δf_0 , ΔQ_L , and $\Delta(R/Q)$ as inputs to NN. (a) Predicted and actual I_{th} , (b) PDF of the relative prediction error (%), (c) PDFs of the predicted and actual data, and (d) comparison of the predicted and actual values for the sorted test samples.

the extreme cases of fabrication effects, where σ_f is set to 1×10^{-3} and 1×10^{-2} . The performance of the network with $\sigma_f = 1 \times 10^{-3}$ variation is shown in Figs. 11(a) and 11(c), and considering with $\sigma_f = 1 \times 10^{-2}$ variation is shown in Figs. 11(b) and 11(d). The test correlation coefficient and σ_{MSE} in these experiments are observed as 0.5512 and 5.0×10^{-3} (for $\sigma_f = 1 \times 10^{-3}$), 0.5067 and 6.9×10^{-3} (for $\sigma_f = 1 \times 10^{-2}$), respectively. As the randomness in the HOM frequency increases, the prediction also decreases. However, a higher σ_f is the extreme testing of the proposed network.

In this study, two modeling strategies are explored. In the first approach, the neural network is trained using only raw parameters. Although the proposed NN model is able to learn general trends with raw input parameters, performance metrics indicate moderate predictive capability. The training and testing correlation coefficients also remained limited. The raw input parameters do not explicitly encode the modal-dominance behavior that governs BBU growth. In the second approach, the input space is augmented with a derived physics-based parameter (ζ). The parameter ζ represents a figure-of-merit associated with the dominant mode excitation in the BBU dynamics. When $\Delta\zeta$ is included in the input space of the proposed NN, the network demonstrates a noticeable improvement in predictive stability and correlation metrics. This confirms that embedding a physically interpretable structure in the input




FIG. 11. f_0 ($\sigma_f = 1 \times 10^{-3}$ and 1×10^{-2}), Q_L , and R/Q variation analyses. Predicted and actual I_{th} with NN input as $\Delta\zeta$, Δf_0 , ΔQ_L , and $\Delta R/Q$ (a) for $\sigma_f = 1 \times 10^{-3}$, (b) for $\sigma_f = 1 \times 10^{-2}$. PDF of the predicted and actual data (c) for $\sigma_f = 1 \times 10^{-3}$, (d) for $\sigma_f = 1 \times 10^{-2}$.

space of the neural network improves the nonlinear relationships in the network.

VI. CONCLUSION

This article proposes a novel method for evaluating the BBU threshold current using a physics-inspired neural network. To improve the prediction accuracy of BBU threshold current, the raw HOM parameters (f_0 , Q_L , R/Q) are interpolated with physics-derived parameter (figure-of-merit: ζ) before being evaluated through a proposed BiLSTM neural network. This study demonstrates that embedding physically meaningful derived parameters in the input space of a neural network significantly enhances predictive performance in highly nonlinear accelerator stability problems. In this study, the prediction accuracy of the threshold current is also improved by fine tuning of hyperparameters through Bayesian optimization. The BBU threshold current prediction is carried out in three cases: single HOM parameter varied, any two HOM parameters varied, and all three HOM parameters varied. It is observed that the NN's performance improves with a physics-inspired parameter rather than just the raw HOM parameters. The best prediction performance is achieved when any single HOM parameter is varied. As the number of variable parameters increases, the prediction quality decreases due to the increase in randomness. For the general case, when all the HOM parameters can vary, the

test correlation coefficient of 0.8518 is achieved with σ_{MSE} of 1.0×10^{-3} when NN is fed with $\Delta\zeta$ along with Δf_0 , ΔQ_L , $\Delta R/Q$. This study is suitable for predicting the BBU threshold current caused by the fabrication tolerances of the rf cavities. The proposed NN is also tested with extreme tolerance values. The above work is the first of its kind to integrate a physics-inspired neural network to estimate the BBU threshold current in ERLs.

ACKNOWLEDGMENTS

The research conducted within this paper is supported by the Cockcroft Core Grant No. (UKRI1887). One of the authors (S. S.) is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Contract No. DE-AC05-06OR23177. The authors would like to acknowledge high-end computing (HEC) at Lancaster University, Lancaster, United Kingdom (UK), for providing access to the HEC facility to conduct BBU analysis and neural network optimization simulations. The authors would also like to thank the PERLE collaboration, of which the authors are members of the collaboration, for allowing us to use the PERLE beam line design and cavity HOM data for the modeling in this paper.

DATA AVAILABILITY

The data that support the findings of this article are openly available [15,25,26].

- [1] L. Merminga, Energy Recovery Linacs, in *Synchrotron Light Sources and Free-Electron Lasers*, edited by E. Jaeschke, S. Khan, J. Schneider, and J. Hastings (Springer, Cham, 2016), .
- [2] A. Hutton, Energy-recovery linacs for energy-efficient particle acceleration, *Nat. Rev. Phys.* **5**, 708 (2023).
- [3] S. A. Bogacz *et al.* (PERLE Collaboration), Beam dynamics driven design of powerful energy recovery linac for experiments, *Phys. Rev. Accel. Beams* **27**, 031603 (2024).
- [4] D. Angal-Kalinin *et al.*, PERLE. Powerful energy recovery linac for experiments. Conceptual design report, *J. Phys. G* **45**, 065003 (2018).
- [5] P. Agostini *et al.* (LHeC and FCC-he Study Group), The Large Hadron–Electron Collider at the HL-LHC, *J. Phys. G* **48**, 110501 (2021).
- [6] O. Brüning, A. Seryi, and S. Verdú-Andrés, Electron-Hadron Colliders: EIC, LHeC and FCC-eh, *Front. Phys.* **10**, 886473 (2022).
- [7] European Strategy Group, Update of the European strategy for particle physics, CERN, Tech. Rep. CERN-ESU-013, 2020.
- [8] E. Pozdeyev, Regenerative multipass beam breakup in two dimensions, *Phys. Rev. ST Accel. Beams* **8**, 054401 (2005).
- [9] C. D. Tennant, K. B. Beard, D. R. Douglas, K. C. Jordan, L. Merminga, E. G. Pozdeyev, and T. I. Smith, First observations and suppression of multipass, multibunch beam breakup in the Jefferson Laboratory free electron laser upgrade, *Phys. Rev. ST Accel. Beams* **8**, 074403 (2005).
- [10] E. Pozdeyev, C. Tennant, J. J. Bisognano, M. Sawamura, R. Hajima, and T. I. Smith, Multipass beam breakup in energy recovery linacs, *Nucl. Instrum. Methods Phys. Res., Sect. A* **557**, 176 (2006).
- [11] B. C. Yunn, Expressions for the threshold current of multipass beam breakup in recirculating linacs from single cavity models, *Phys. Rev. ST Accel. Beams* **8**, 104401 (2005).
- [12] S. Setiniyaz, R. Apsimon, and P. H. Williams, Implications of beam filling patterns on the design of recirculating energy recovery linacs, *Phys. Rev. Accel. Beams* **23**, 072002 (2020).
- [13] S. Setiniyaz, R. Apsimon, and P. H. Williams, Filling pattern dependence of regenerative beam breakup instability in energy recovery linacs, *Phys. Rev. Accel. Beams* **24**, 061003 (2021).
- [14] S. Setiniyaz, R. Apsimon, M. Southerby, and P. H. Williams, in *Proceedings of the 31st International Linear Accelerator Conference (LINAC2022)* (JACoW Publishing, Geneva, Switzerland, 2022), p. TUPORI20.
- [15] S. Setiniyaz, R. Apsimon, S. A. Bogacz, R. M. Bodenstein, K. Deitrick, P. H. Williams, I. Bailey *et al.*, Beam breakup instability studies of powerful energy recovery linac for experiments, *Phys. Rev. Accel. Beams* **28**, 011003 (2025).
- [16] A. Edelen and X. Huang, Machine learning for design and control of particle accelerators: A look backward and forward, *Annu. Rev. Nucl. Part. Sci.* **74**, 557 (2024).
- [17] A. Scheinker, F. Cropp, S. Paigua, and D. Filippetto, An adaptive approach to machine learning for compact particle accelerators, *Sci. Rep.* **11**, 19187 (2021).
- [18] A. L. Edelen, S. G. Biedron, B. E. Chase, D. Edstrom, S. V. Milton, and P. Stabile, Neural networks for modeling and control of particle accelerators, *IEEE Trans. Nucl. Sci.* **63**, 878 (2016).
- [19] A. Hanuka, X. Huang, J. Shtalenkova, D. Kennedy, A. Edelen, Z. Zhang, V. R. Lalchand, D. Ratner, and J. Duris, Physics model-informed Gaussian process for online optimization of particle accelerators, *Phys. Rev. Accel. Beams* **24**, 072802 (2021).
- [20] O. Convery, L. Smith, Y. Gal, and A. Hanuka, Uncertainty quantification for virtual diagnostic of particle accelerators, *Phys. Rev. Accel. Beams* **24**, 074602 (2021).
- [21] K. Fujita, Physics-informed neural network method for space charge effect in particle accelerators, *IEEE Access* **9**, 164017 (2021).
- [22] A. Ivanov and I. Agapov, Physics-based deep neural networks for beam dynamics in charged particle accelerators, *Phys. Rev. Accel. Beams* **23**, 074601 (2020).
- [23] C. D. Tennant, Studies of energy recovery linacs at Jefferson Laboratory: 1 GeV demonstration of energy recovery at CEBAF and studies of the multibunch, multipass beam breakup instability in the 10 kW FEL upgrade driver, Ph. D. thesis, College of William and Mary, Williamsburg, VA, USA, 2006.
- [24] N. Valles, Pushing the frontiers of superconducting radio frequency science: From the temperature dependence of the superheating field of niobium to higher-order mode

- damping in very high quality factor accelerating structures, Ph. D. thesis, Cornell University, Ithaca, NY, USA, 2014.
- [25] C. Barbagallo, P. Duchesne, W. Kaabi, G. Olry, F. Zomer, R. A. Rimmer, H. Wang, R. Apsimon, and S. Setiniyaz, HOM-damping studies in a multi-cell elliptical superconducting RF cavity for the multi-turn energy recovery linac PERLE, [arXiv:2409.13798](https://arxiv.org/abs/2409.13798).
- [26] C. Barbagallo, Design and optimization of higher order mode couplers for the superconducting cavities of the PERLE energy recovery linac, Ph.d. thesis, University Paris-Saclay, Geneva, Switzerland, 2024.
- [27] The MathWorks, Inc., MATLAB and Deep Learning Toolbox, version: R2025b (The MathWorks, Inc., Natick, Massachusetts, United States, 2025), <https://www.mathworks.com>.