NEURAL NETWORK-BASED AMPLITUDE FEEDFORWARD CONTROL ALGORITHM FOR LLRF SYSTEMS*

Yuchen Wang, Xiaofang Hu†, Shenghua Yang, Jian Pang,
Fangfang Wu, Kai Zhang, Shancai Zhang
National Synchrotron Radiation Laboratory, University of Science and Technology of China,
Hefei, Anhui, 230029, China

Abstract

In free-electron laser facilities, the amplitude-phase stability of the microwave pulses driving the electron beam is a key factor determining beam energy spread. Aiming at the long bunch train operation mode, this paper proposes a neural network-based amplitude feedforward control for low-level radio frequency (LLRF) systems to suppress intra-pulse amplitude fluctuations. The algorithm has been validated at the output of a solid-state amplifier (SSA): under four randomly selected vector modulator (VM) output configurations, the average intra-pulse amplitude flatness (RMS) was reduced from 1.208 % to 0.398 %, and the average peak-to-peak variation was reduced from 4.683 % to 1.353 %, demonstrating a significant compensation effect.

INTRODUCTION

For linear accelerators with relatively high average beam currents, the beam loading effect within the accelerating tubes is significant, leading to inconsistent energy gains for particles at different times within the pulse, thereby substantially degrading the uniformity of particle acceleration [1, 2]. Furthermore, as critical components of the accelerator's microwave high-frequency system, the nonlinear characteristics of solid-state amplifiers (SSA) and klystrons introduce significant waveform distortion. This results in issues such as non-flat tops and rising edge distortions in the ideally flat-top pulses generated by the low-level radio frequency (LLRF) vector modulator (VM) after transmission and amplification, upon reaching the entrance of the accelerating tube, also impacting beam quality. To ensure uniformity in the accelerating fields experienced by each bunch, the LLRF control system and the microwave high-frequency system must provide RF power with stable intra-pulse and pulse-to-pulse amplitude and phase characteristics to the accelerating structure. In engineering practice, feedback control is often employed to enhance pulse-to-pulse amplitude and phase stability [3–5]. However, due to inherent delays in the feedback loop, achieving rapid, point-by-point intra-pulse amplitude and phase flattening control is challenging. Consequently, precise and fast waveform control is typically addressed using feedforward compensation. For compensating this nonlinear distortion, traditional fixed-model-based feedforward methods often prove complex to implement under conditions where amplitude settings vary over a wide range and pulse widths and delays are adjusted [6, 7]. U-Net is a deep learning architecture that is data-efficient and excels at capturing local details within complex data and performing nonlinear fitting, making it widely applied in the field of image processing [8, 9]. This paper proposes a data-driven deep learning method by introducing U-Net into the LLRF controller domain. By constructing a one-dimensional U-Net neural network, the inverse response characteristic from the distorted amplitude output of the microwave high-frequency system to the ideal amplitude input is directly learned, serving as a pre-distortion model to implement feedforward control.

* Supported by National Natural Science Foundation of China (12505182)
Supported by National Key Research and Development Program of China (2024YFA1612200)
† Corresponding author, email: huxiaofang@ustc.edu.cn

MODEL AND DATASET DESIGN

The core objective of this study is to construct a deep neural network model capable of accurately learning the nonlinear inverse response of an RF power system, thereby achieving waveform pre-distortion compensation based on feedforward control. The overall application framework comprises two phases: offline training and online application. In the offline phase, we first need to acquire a large-scale dataset consisting of measured VM output waveforms and the corresponding waveforms at the controlled location of the RF power system. Subsequently, using this dataset, we employ the waveforms from the controlled location as the network input and the measured VM output waveforms as the target output. This training process yields a deep neural network that learns the mapping from the distorted output of the RF power system to the ideal low-level input (i.e., the inverse model of the power chain). During the online application phase, this deep neural network module, serving as a pre-distortion module, is deployed on a computer equipped with a GPU. The desired waveform for the controlled location is fed into the neural network, which generates the corresponding pre-distorted low-level excitation waveform data. This data is then written directly into the memory of the LLRF system's DAC via Direct Memory Access (DMA) and output through the VM, thereby enabling real-time compensation at any controlled location within the RF power system.

Network Architecture: 1D U-Net Convolutional Neural Network

To effectively process long-sequence time-domain waveforms and capture their local features along with global context, we adopted a one-dimensional U-Net architecture. This network adopts an encoder-decoder structure as its

main body, specifically optimized for the characteristics of RF pulse waveforms.

In the encoder-decoder backbone, the network employs symmetric contracting and expanding paths. The encoder progressively extracts multi-level features through a series of down-sampling blocks while compressing the sequence length. Conversely, the decoder gradually restores the spatial resolution and fuses high-resolution features from the encoder through up-sampling and skip connections, ultimately reconstructing the complete 2048-point pre-distortion waveform. The encoder part of this one-dimensional U-Net architecture is shown in Fig. 1, and the decoder part is shown in Fig. 2.

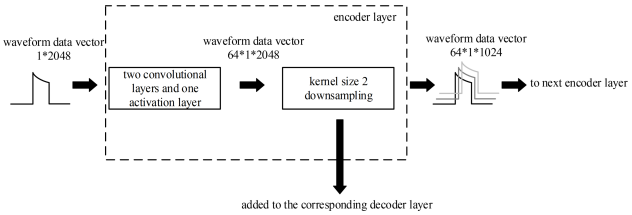

Figure 1: The encoder part of one-dimensional U-Net architecture.

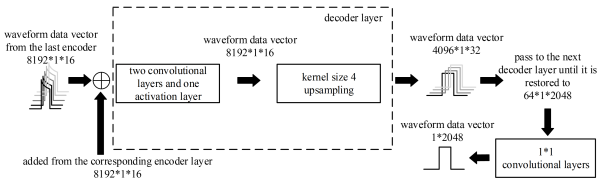

Figure 2: The decoder part of one-dimensional U-Net architecture.

Dataset Processing and Construction

The dataset utilized in this study originates from a test platform comprising a physically implemented LLRF control system and an SSA. By systematically acquiring the system's responses under various excitation parameters, a large-scale collection of waveform pairs was constructed for training and validating the deep inverse model, with the aim of comprehensively covering its amplitude nonlinear characteristics.

The hardware for data acquisition consists of an LLRF controller, an SSA, an attenuator, a trigger generator, an RF signal generator, a frequency synthesizer, and a trigger level converter. During system operation, the LLRF controller generates and outputs the initial RF pulse, which is then power-amplified by the SSA. The signal is subsequently attenuated back to an appropriate level via the attenuator and routed back to the acquisition channel of the LLRF controller. Finally, the Process Variables (PVs) of the corresponding channels are recorded by the host computer's Input-Output Controller (IOC) based on the EPICS framework. The setup of the data acquisition system is illustrated in Fig. 3.

The waveforms actually collected after distortion by the SSA are used as the input to the neural network. The corresponding ideal waveforms originally sent by the LLRF

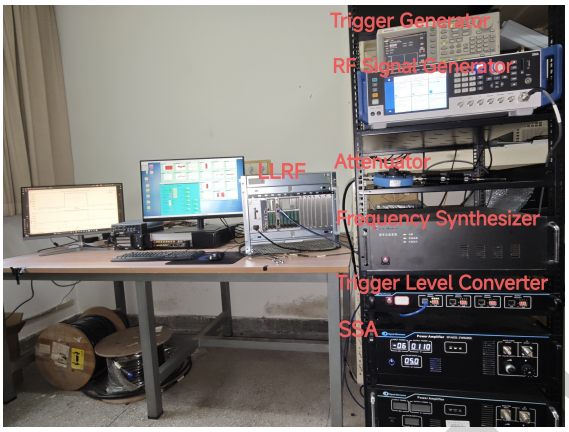

Figure 3: The setup of the data acquisition system.

system serve as the target output for the network. The model training adopts a supervised learning paradigm, with the objective of enabling the pre-distortion signal predicted by the network to produce the desired ideal waveform after passing through the actual physical system.

The model is trained using the AdamW optimizer, combined with a learning rate scheduler that dynamically adjusts the learning rate upon plateaus. To prevent overfitting, in addition to monitoring with an independent validation set, we employ an early stopping strategy, terminating training when the validation loss ceases to decrease for a consecutive number of epochs. The network hyperparameters are determined via grid search in conjunction with the loss on the test set.

EXPERIMENTAL SETUP AND RESULTS

To evaluate the feedforward compensation performance of the aforementioned model, we conducted a bench-top test. The hardware devices and their interconnections in this test remained identical to those used during the data acquisition process. The model was implemented on an additional computer equipped with a GPU and trained offline using the PyTorch framework. The performance of the model training was characterized by the value of the same loss function employed during training, calculated based on the difference between the model's predictions for the waveform data in the test set and the corresponding ground truth labels of the test set. Upon completion of the model training, the LLRF controller was used once again to acquire waveform data from both the VM readback channel and the SSA output channel. The waveform from the re-acquired VM readback channel was then fed as input into the model. Subsequently, the pre-distortion waveform generated by the model was written into the DAC, thereby enabling the output of the pre-distorted excitation waveform.

With the selected set of optimal training parameters, validation experiments were conducted to evaluate the feedforward compensation performance. Four excitation signals with different VM output amplitudes, pulse delays, and pulse widths were randomly selected for comparison before and after feedforward compensation. The distorted waveforms

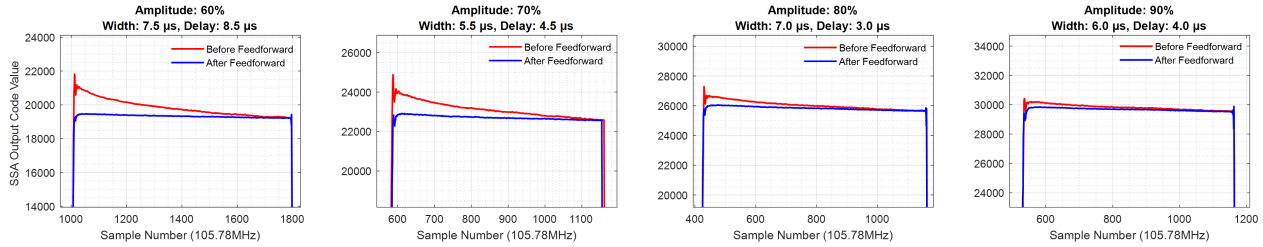

Figure 4: U-Net feedforward performance comparison under different VM pulse settings.

Table 1: Comparison of Intra-Pulse Amplitude Flatness (RMS) and Peak-to-Peak Values Before and After Feedforward Compensation Under Different VM Excitation Signals

Amplitude (%)	RMS (before/after)	peak-to-peak (before/after)
60	1.72 %/0.47 %	6.39 %/1.18 %
70	1.54 %/0.40 %	6.25 %/1.52 %
80	1.04 %/0.45 %	3.97 %/1.62 %
90	0.53 %/0.28 %	2.12 %/1.08 %

at the SSA output corresponding to these VM outputs were recorded, and the VM output waveforms were used as inputs to the surrogate model to generate the pre-distorted outputs of the LLRF system's VM, thereby demonstrating the effectiveness of the proposed method in flattening the waveform amplitude. The waveforms before and after feedforward compensation are shown in Fig. 4, and the corresponding intra-pulse amplitude flatness (RMS) and peak-to-peak values are presented in Table 1. Since the data for these VM excitation signals are not included in the training set or validation set, it can be demonstrated that the feedforward compensation effect is consistently significant over a wide range of amplitude settings, as well as under arbitrary adjustments of pulse widths and delays. However, as shown in Table 1, there remains room for improvement in the feedforward compensation effect corresponding to specific VM output amplitudes. Therefore, future work will focus on further optimizing the model and algorithm to address this issue.

CONCLUSION

This study presents a novel neural network-based technical pathway for intra-pulse amplitude nonlinear feedforward compensation in accelerator LLRF systems. The proposed algorithm has been validated at the output of an SSA: under four randomly selected VM output configurations, the average intra-pulse amplitude flatness (RMS) was reduced from 1.208 % to 0.398 %, and the average peak-to-peak variation was reduced from 4.683 % to 1.353 %, demonstrating a significant compensation effect. Future work will focus on advancing this technology toward practical application, extending the ultimate evaluation metric to the amplitude flatness of the electric field at the accelerating tube entrance, and validating its improvement on beam quality through beam diagnostics.

REFERENCES

[1] Z. G. He, Q. K. Jia, L. Wang, W. Xu, and S. C. Zhang, "Linac Design of the IR-FEL Project in CHINA", in *Proc. FEL'15*, Daejeon, Korea, Aug. 2015, pp. 46–48. doi:10.18429/JACoW-FEL2015-MOP010

[2] N. Mesbah, S. Doebert, M. Dayyani Kelisani, A. Latina, Y. Zhao, and J. Olivares Herrador, "Optimisation of the CLIC positron capture LINAC taking into account Beam Loading effects", *EPJ Web Conf.*, vol. 315, p. 02001, 2024. doi:10.1051/epjconf/202431502001

[3] M. Kurata *et al.*, "High current, long-pulse operation at the KEK superconducting rf test facility", *Phys. Rev. Accel. Beams*, vol. 28, no. 2, Feb. 2025. doi:10.1103/physrevaccelbeams.28.021001

[4] Z. T. Ang *et al.*, "TRIUMF ISAC LINAC Developments and Upgrades", in *Proc. LINAC'18*, Beijing, China, Sep. 2018, pp. 355–357. doi:10.18429/JACoW-LINAC2018-TUPO014

[5] V. Ayvazyan *et al.*, "LLRF Control System Upgrade at FLASH", in *Proc. PCaPAC'10*, Saskatoon, Canada, Oct. 2010, paper THPL012, pp. 150–152.

[6] Z. Gao *et al.*, "Study on transient beam loading compensation for China ADS proton linac injector II", *Chin. Phys. C*, vol. 40, no. 5, p. 057005, May 2016. doi:10.1088/1674-1137/40/5/057005

[7] X. C. Xu and J. Y. Ma, "Simulation Analysis of LLRF Feedforward Compensation to Beam Loading for CiADS LINAC", in *Proc. IPAC'19*, Melbourne, Australia, May 2019, pp. 2027–2029. doi:10.18429/JACoW-IPAC2019-TUPTS045

[8] C.-H. Chuang *et al.*, "IC-U-Net: A U-Net-based Denoising Autoencoder Using Mixtures of Independent Components for Automatic EEG Artifact Removal", *NeuroImage*, vol. 263, p. 119586, Aug. 2022. doi:10.1016/j.neuroimage.2022.119586

[9] L. Qiu *et al.*, "A lightweight U-net for ECG denoising using knowledge distillation", *Physiol. Meas.*, vol. 43, p. 115004, Nov. 2022. doi:10.1088/1361-6579/ac96cd