ISBN: 978-3-95450-179-3 doi:10.18429/JACoW-FEL2017-WEP031

USING A NEURAL NETWORK CONTROL POLICY FOR RAPID SWITCHING BETWEEN BEAM PARAMETERS IN AN FEL

A. L. Edelen[†*], Colorado State University, Fort Collins, CO
S.G. Biedron[1], Element Aero, Chicago, IL
J. P. Edelen, Fermi National Accelerator Laboratory, Batavia, IL
S.V. Milton, Los Alamos National Laboratory, Los Alamos, NM
P. J. M. van der Slot, University of Twente, Enschede, The Netherlands
[1]also at University of Ljubljana, Ljubljana Slovenia and
University of New Mexico, Albuquerque, NM

Abstract

FEL user facilities often must accommodate requests for a variety of beam parameters. This usually requires skilled operators to tune the machine, reducing the amount of available time for users. In principle, a neural network control policy that is trained on a broad range of operating states could be used to quickly switch between these requests without substantial need for human intervention. We present preliminary results from an ongoing study in which a neural network control policy is investigated for rapid switching between beam parameters in a compact THz FEL.

INTRODUCTION

Free-electron laser (FEL) user facilities often must accommodate requests for a variety of beam parameters. This usually requires skilled human operators to tune the machine manually, thus reducing the amount of available up time for the users. In principle a neural network control policy that is trained on a broad range of operating states could be used to quickly switch between these requests without substantial need for human intervention [1]. Additionally, this policy could be updated concurrently to machine operation as new states are visited or as drift occurs in the system. It also provides a relatively compact way of storing the information (as compared to, e.g., a database of previous settings and measured output). We are exploring this approach using simulations of a compact THz FEL design based on the Twente/Eindhoven University FEL (TEU-FEL) [2]. This an appealing system for this study because it has a relatively small number of components, yet it exhibits non-trivial beam dynamics.

Here, we focus on an initial study: injector and beamline tuning to achieve specific electron beam parameters at the entrance of the undulator. In this case, changing the operating state consists of specifying a change in energy and then choosing the appropriate injector settings and beamline settings such that the match of the beam into the undulator is preserved and emittance is minimized.

The way the neural network controller is trained is similar to how one might try to do it for a real machine. First, we train a neural network model based on data from *a priori* simulations of the machine, with noise added. This creates a surrogate for the simulation that in principle contains the relevant physics of the machine and can execute quickly to facilitate controller training. We then train a reinforcement learning [3] controller via interaction with the learned model. Interaction directly with the simulation can then be used to fine-tune the controller.

We first give an overview of the FEL. This is followed by a discussion of our initial study. We then discuss our current efforts on a comprehensive study that includes start-to-end matching and minimizing the emittance, and we conclude with our future plans for this effort.

FEL LAYOUT AND SIMULATION

The FEL is designed to produce light with a wavelength that is tunable between 200 μm and 800 μm. It consists of a 5.5-cell, 1.3-GHz photocathode RF gun, a beam transport section, and a fixed-gap THz undulator with an undulator parameter K equal to 1. More details on this machine can be found in [4], and Fig. 1 shows the relevant components for this study.

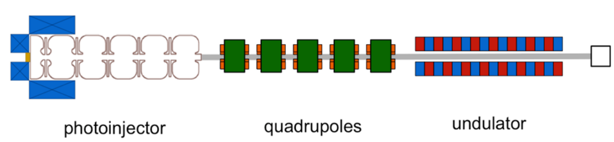

Figure 1: Layout of the accelerator showing the 5.5-cell photoinjector with its bucking coil and solenoid, the quadrupoles, the undulator, and the beam dump. There also are steering coils, but these are excluded, as the simulated electron beam is generated on-axis.

Beam Dynamics Considerations

While this machine is comparatively simple in terms of the number of components and interacting systems, the beam dynamics exhibit some subtleties that if not addressed properly will result in decreased performance of the FEL. Because the bunch charge can be as large as 5 nC and the beam energy is low (3–6 MeV), space-charge effects will be significant both in the photoinjector and in the beamline. A solenoid is typically used to compensate for space-charge effects by matching the beam to the invariant envelope [5]. This can easily be done in long (more than ~10 cells) accelerating structures or with in-

*A.L.E. is partially supported by a scholarship from the Directed Energy Professional Society
† auralee.l.morin@gmail.com

ISBN: 978-3-95450-179-3 doi:10.18429/JACoW-FEL2017-WEP031

jectors that have a short (e.g. 1.5-cell) gun followed by a drift and a secondary accelerating stage. The TEU-FEL photoinjector, with 5.5 cells, is neither a long gun nor a short gun with a secondary accelerating structure. This decreases the ability to perform adequate compensation of space-charge-induced emittance growth. Additionally, due to physical space constraints there is not enough room to allow the emittance to fully damp before the matching section into the undulator, and because the beam is low-energy, space-charge forces will continue to impact the emittance along the beamline. This complicates the task of focusing the beam even after initial space-charge compensation is performed by the solenoid, which in turn complicates the task of matching the beam into the FEL. Thus, a full study where the gun parameters and the beamline parameters are optimized simultaneously is needed to achieve optimal FEL performance.

PARMELA Simulation

A simulation model of the TEU-FEL injector and beamline was constructed using SUPERFISH [6] for the RF fields and PARMELA [7] for the beam dynamics. Because the coupling between the cells of the cavity is not axially symmetric, a full simulation of the gun geometry could not be conducted using SUPERFISH. Therefore, we approximated the geometry by simulating each of the individual cell types and splicing them together in PARMELA to create the proper field geometry. When compared with measurements of the field geometry this approach performs reasonably well [8]. The solenoid and bucking coil assembly were modelled using PANDRIA [6] using the nominal current settings in the solenoid. The bucking coil was then scaled to cancel the magnetic field on the cathode. The combined field map is then scaled in PARMELA in order to tune the space-charge compensation. The PARMELA simulations were performed using 5000 macro-particles with a 0.1º phase integration step. This was determined to be well within the stable region for reasonable estimation of bulk parameters [8]. In order to achieve proper matching into the undulator, the beam is focused to a waist at the entrance, which means the alpha parameter $\alpha_{x,y} = 0$. The good magnetic field region for equal focusing in both planes requires the beam size to be less than 4 mm [9]. We chose to optimize the beamline, which contains 5 quadrupoles (Q1–Q5), with the same beta function ($\beta_{x,y}$) value at the entrance of the undulator for each electron beam energy: $\beta_{x,y} = 0.106$ [m/rad]. This corresponds to a beam size that is within the optimal field region of the undulator for all energies.

INITIAL BEAMLINE STUDY

For the initial study the aim was to set the quadrupole settings such that specific $\alpha_{x,y}$ and $\beta_{x,y}$ are achieved at the entrance of the undulator for a given electron beam energy. Limited adjustment of the RF power, phase, and solenoid strength is also included.

Neural Network Model

The model was trained to predict the Twiss parameters at the entrance of the undulator, given the RF power, the RF phase, the solenoid strength, and the quadrupole settings. The general setup is shown in Figure 2, and in this case the extraneous outputs are not used during training.

The training set consisted of the output from each iteration of simplex optimization of the quadrupole settings for 12 different beam energies. This looks similar to what one might see in the data archive of an operational accelerator: a lot of tuning around roughly optimal settings. In order to reduce the overall optimization time and simplify the initial problem for the neural network, the gun and the beamline were optimized separately, with the knowledge that the match may not be fully optimized for FEL performance. This training data includes the following variable ranges: -2.14 to 2.11 rad for α_x, -5.76 to 1.45 for α_y, 0.058 to 1.86 m/rad for β_x, 0.074 to 3.65 m/rad for β_y, -0.98 to 0.83 T/m for Q1, 0.65 to 1.98 T/m for Q2, -2.24 to -1.07 T/m for Q3, 0.89 to 2.26 T/m for Q4, -1.90 to -0.23 T/m for Q5, 0.67 to 1.05 for solenoid strength (normalized units), and 10.3º to 21.4 º for RF phase.

For the validation set, the optimization data for the 5.7-MeV electron beam was used. The performance of the model in terms of mean absolute error (MAE) and standard deviation (STD) is shown in Table 1. A representative plot from the validation set is shown in Figure 3.

The neural network architecture consists of four hidden layers containing 50, 50, 30, and 30 nodes, respectively. Each node in the hidden layers uses a hyperbolic tangent activation function and a dropout [10] probability of 10%.

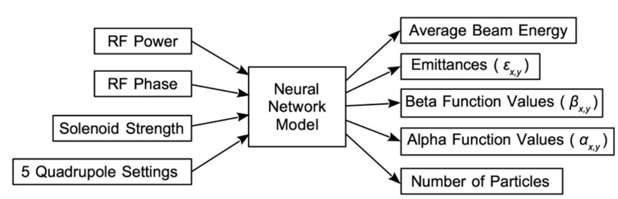

Figure 2: General setup for the neural network model. Output parameters are at the entrance of the undulator.

ISBN: 978-3-95450-179-3 doi:10.18429/JACoW-FEL2017-WEP031

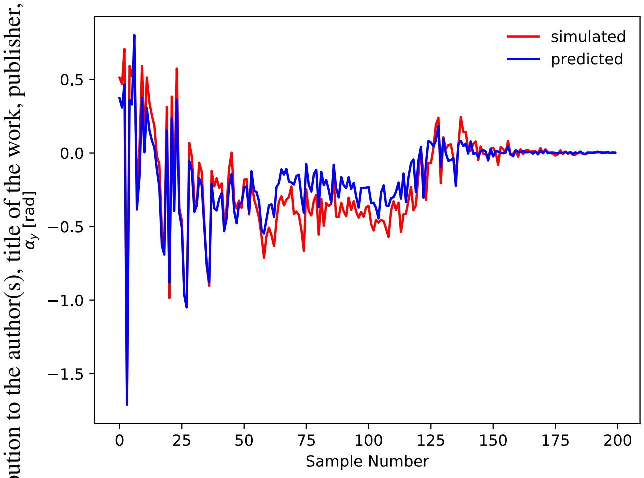

Figure 3: Neural network model predictions and simulated values for α_x on the validation data set for the initial study, with a beam energy of 5.7 MeV.

Table 1: Model Performance

Variable	Train MAE	Val. MAE	Train STD	Val. STD
α_x [rad]	0.018	0.067	0.042	0.091
α_y [rad]	0.022	0.070	0.037	0.079
β_x [m/rad]	0.004	0.008	0.009	0.012
β_y [m/rad]	0.005	0.012	0.011	0.017

Neural Network Controller

The controller was trained via reinforcement learning by allowing it to interface with the model. For the controller, random desired beam energy values between 3.1 and 6.2 MeV were specified, along with a target set of $\alpha_{x,y}$ and $\beta_{x,y}$ values. Larger quadrupole settings were penalized proportionally. This time, energy values between 4.8 and 5.2 MeV were excluded from training and used exclusively for the validation set. The controller network architecture consists of three hidden layers, with 30, 30, 20, and 20 nodes, respectively. As before, a hyperbolic tangent activation function and a dropout probability of 10% for each node was used.

Given random requested energy values within 3–6 MeV, Table 2 shows the performance in reaching the desired Twiss parameters ($\alpha_{x,y} = 0$ rad, $\beta_{x,y} = 0.106$ m/rad) in one iteration. This shows that for a given energy, the controller will immediately reach the desired beam size to within about 10% and the beam will be close to a waist, requiring minimal further tuning to reach the target values. The maximum absolute errors for the energy range seen in the training data set were 0.063 rad, 0.023 m/rad, 0.067 rad, and 0.041 m/rad for α_x, β_x, α_y, and β_y respectively. The maximum absolute errors for the validation set were 0.141 rad, 0.140 m/rad, 0.008 rad, and 0.038 m/rad for α_x, β_x, α_y, and β_y respectively.

Table 2: Ability to achieve $\alpha_{x,y} = 0$ rad and $\beta_{x,y} = 0.106$ m/rad for 3–6 MeV electron beams in one iteration.

Variable	Train MAE	Val. MAE	Train STD	Val. STD
α_x [rad]	0.012	0.075	0.011	0.046
α_y [rad]	0.013	0.079	0.012	0.045
β_x [m/rad]	0.008	0.004	0.006	0.002
β_y [m/rad]	0.014	0.011	0.011	0.011

FULL PHOTOINJECTOR AND BEAMLINE STUDY

Based on the encouraging results from the initial study, we have been conducting a more comprehensive version of it that also includes the task of minimizing the emittance at the entrance of the undulator. As this requires finer adjustment of the RF phase, RF power, and solenoid strength in conjunction with the quadrupole settings, additional data sets consisting of the output from start-to-end optimizations and parameter scans were collected that include fine adjustment of these parameters. In this case, the neural network model inputs and outputs are as shown in Figure 2. This time the controller takes in desired beam energy, Twiss parameters, emittances, and transmission at the entrance of the undulator and sets the RF power, RF phase, solenoid strength, and quadrupole settings needed to achieve them. Figure 4 shows an example of the model's performance on the validation set, which consisted of start-to-end optimization data for 3.5 MeV

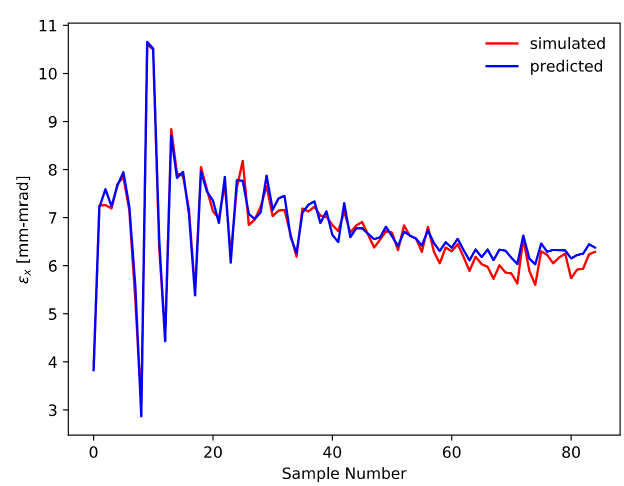

Figure 4: Neural network model predictions and simulated values for normalized ε_x on the new validation data set at 3.5 MeV for full the photoinjector and beamline study.

ISBN: 978-3-95450-179-3 doi:10.18429/JACoW-FEL2017-WEP031

CONCLUSIONS AND FUTURE WORK

We have shown encouraging results from an initial beamline tuning study, indicating that in one iteration the controller can set up the machine to achieve close to the correct Twiss parameters for arbitrary beam energies between 3–6 MeV. This is a successful first step toward the development of a neural network reinforcement learning controller that will facilitate fast switching between operational parameters along with fine-tuning. The TEU-FEL presents a good platform to explore this technique because of its relative simplicity in terms of the number of control parameters and its nonlinear beam dynamics. Our next steps are to (1) finish the full photoinjector and beamline study and (2) incorporate FEL simulations and train the controller based on FEL output. While the true merit of the approach won't be clear until it is tested experimentally, we can optimistically say that with some further R&D a neural network control policy may well be an expedient way of switching between operating states in an FEL.

REFERENCES

[1] A. L. Edelen *et al.*, "Neural networks for modeling and control of particle accelerators," *IEEE Transactions on Nuclear Science*, vol. 63, no. 2, pp. 878-897, Apr. 2016.

[2] G. J. Ernst *et al.*, "Status of the TEU-FEL project," *Nuclear Instruments and Methods in Physics Research A*, vol. 318, no. 1-3, 1992.

[3] R. S. Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1988.

[4] S. Milton *et al.*, "Colorado State University Accelerator and FEL," *Proc. IPAC 2014*, Dresden, Germany, THPRI074.

[5] L. Serafini and J. B. Rosenzweig, "Envelope analysis of intense relativistic quasilaminar beams in RF photoinjectors: a theory of emittance compensation," *Phys. Rev. E*, vol. 55, pp. 7565, 1997.

[6] J. Billen, "Poisson Superfish Codes," Los Alamos National Laboratory, Copyright 1985-2005.

[7] L. Young, "PARMELA Codes," Los Alamos National Laboratory, Copyright 1985-2005.

[8] J. P. Edelen, "Design and simulation of the Colorado State University linear accelerator system," M.S. Thesis, Colorado State University, Fort Collins, CO, 2014.

[9] J. Verschuur, G. J. Ernst, and W. J. Witteman, "The TEU-FEL undulator," *Nuclear Instruments and Methods in Physics Research A*, vol. 318, pp. 847-852, 1992.

[10] N. Srivastav *et al.*, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, pp. 1929-1958, 2014.