Machine learning-based longitudinal phase space prediction of particle accelerators

C. Emma,[*,†] A. Edelen,[†] M. J. Hogan, B. O'Shea, G. White, and V. Yakimenko
SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA

(Received 11 September 2018; published 16 November 2018)

We report on the application of machine learning (ML) methods for predicting the longitudinal phase space (LPS) distribution of particle accelerators. Our approach consists of training a ML-based virtual diagnostic to predict the LPS using only nondestructive linac and e-beam measurements as inputs. We validate this approach with a simulation study for the FACET-II linac and with an experimental demonstration conducted at LCLS. At LCLS, the e-beam LPS images are obtained with a transverse deflecting cavity and used as training data for our ML model. In both the FACET-II and LCLS cases we find good agreement between the predicted and simulated/measured LPS profiles, an important step towards showing the feasibility of implementing such a virtual diagnostic on particle accelerators in the future.

DOI: 10.1103/PhysRevAccelBeams.21.112802

I. INTRODUCTION

Accurate nondestructive diagnostics of the electron beam longitudinal phase space (LPS) distribution can be challenging for high-intensity particle accelerators. As an example, the FACET-II accelerator is designed to deliver beams of unprecedented intensity to a suite of advanced accelerator experiments. The design parameters for the beam are 2 nC charge, 10 GeV energy, <10 μm normalized transverse emittance and up to 200 kA peak current [1]. These unique characteristics present many opportunities for scientific experiments [2], and a threefold hurdle from the diagnostic point of view. First, the high intensity of the beams limits the possibility of utilizing intercepting diagnostics due to heat-induced surface damage of the measurement devices. Second, the very short pulse duration for high-current shots ($\sigma_z \sim 1$ μm for $I > 100$ kA) is close to the resolution limit of state-of-the-art longitudinal diagnostics such as transverse deflecting cavities (TCAVs) [3]. Finally, a drawback of the aforementioned diagnostics is that they provide a destructive measurement of the electron beam properties and cannot be made in conjunction with experiments unless they are located downstream of the interaction region.

In an effort to meet these challenges, we investigate the possibility of incorporating a machine learning (ML) based virtual diagnostic to provide shot-to-shot nondestructive measurements of the LPS distribution in particle accelerators. The virtual diagnostic is a computational tool which creates a mapping between nondestructive measurements of the linac and e-beam properties and the 2D LPS distribution of the beam. The rationale behind choosing an ML-based approach for the virtual diagnostic is motivated by a number of factors. First, ML methods have made tremendous progress in the fields of image recognition and prediction in the past few years [4]. This gives us confidence that a virtual diagnostic trained on image data, e.g., from TCAV or profile monitors, can be used to reconstruct desired properties of the e-beam such as the LPS and/or the current profile. Furthermore, ML techniques are also well suited for solving problems involving large amounts of data. Large data sets can be acquired in real time on accelerators such as FACET-II or LCLS operating at nominal repetition rates of 10 and 120 Hz. ML models can also be trained off-line using simulation data from computationally expensive particle tracking codes and updated with measurements on the accelerator, as has been recently demonstrated in Ref. [5]. Finally, there is growing interest due to recent studies which highlight the versatility of ML methods used in particle accelerators as tools for prediction, control and optimization of accelerator performance [6–10]. As an example, recent work performed at Fermilab's FAST facility has been aimed at training a virtual diagnostic to predict the e-beam emittance through a combination of simulation and experimental studies [5,11].

With these potential benefits in mind, the first goal of this work is to show the feasibility of applying ML techniques to predict the LPS of the FACET-II accelerator. Our approach is to acquire training data for the ML model

[*]Corresponding author.
cemma@slac.stanford.edu
[†]C. E. and A. E. contributed equally to this work.

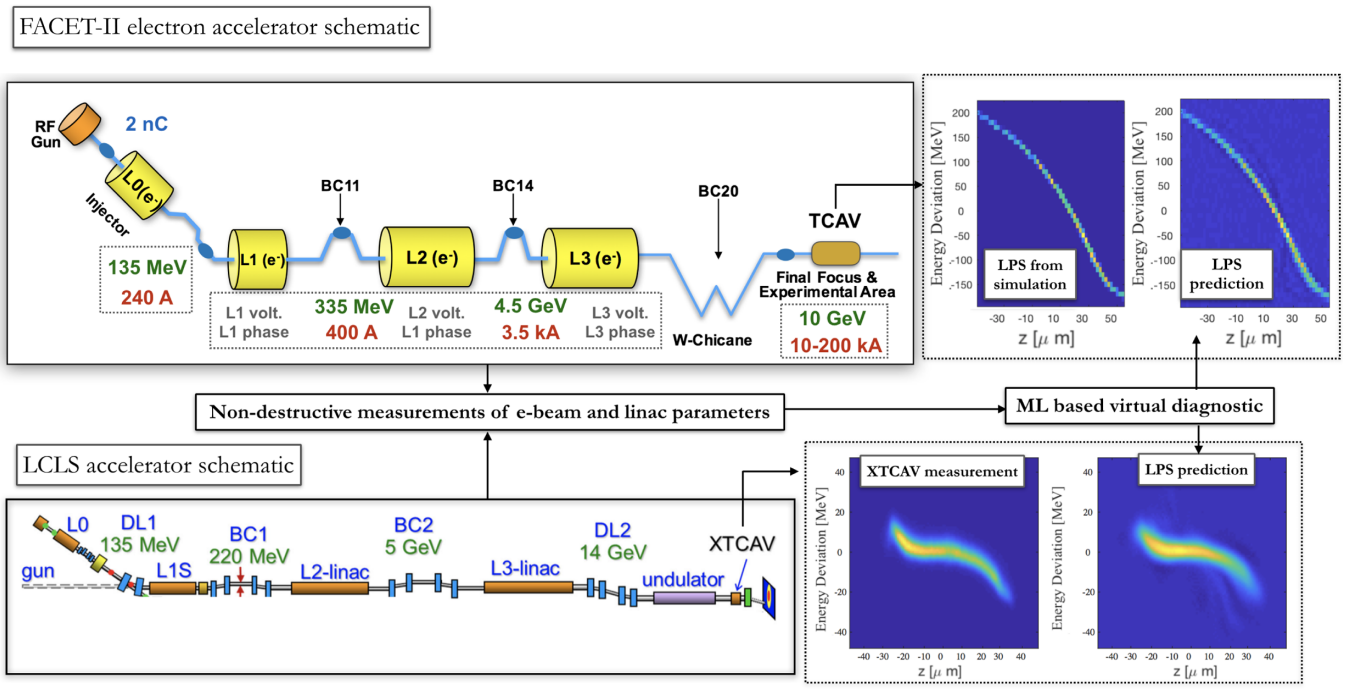

FIG. 1. Schematic of the FACET-II and LCLS electron accelerators and example LPS profiles from particle tracking simulations (FACET-II), experimental measurements (LCLS) and from the ML-based virtual diagnostic predictions. The figure highlights the similarities between the two accelerator layouts up to the BC20 chicane in FACET-II which is used to increase the current from 3–4 kA to 10–200 kA.

from a large number of simulations. These simulations represent the performance of the machine which changes as a result of several key accelerator parameters jittering around their design values. The accuracy of this ML model based on simulation data, as well as its dependence on diagnostic inputs, will inform the measurement resolution necessary for this to be successful on the actual machine. Our second goal is to test out a similar predictive ML model on the Linac Coherent Light Source (LCLS). For that part of the study we train a ML model using existing nondestructive diagnostics and images of the e-beam LPS obtained with the X-band TCAV [3]. In the following sections we present results from the simulation study of FACET-II and measurements from the LCLS, with a discussion of the steps necessary to implement this diagnostic tool on future particle accelerators.

II. FACET-II SIMULATION STUDY

A key performance feature for the success of advanced acceleration experiments is knowledge and control of the e-beam LPS and current profile [12]. We therefore train two separate ML models to predict the current profile and LPS of the bunch using some key nondestructive diagnostics as input to the models (see Table I). We consider the nominal operation of the FACET-II accelerator in single-bunch mode, with the machine set up to deliver a beam of 10 GeV energy, 25 kA peak current, and <10 μm transverse emittance (see Fig. 1 for a schematic). In order to capture the performance of the machine we perform 5^5 LUCRETIA [13] simulations starting from the exit of the injector, with key linac parameters and the bunch charge jittering around their nominal values. The simulations include longitudinal space charge and incoherent and coherent synchrotron radiation. The mean value and the range for each simulation parameter scanned was set using the values from the FACET-II technical design report (TDR) [1]. The output of these simulations is a $6 \times 6 \times N_p$ distribution of $N_p = 2 \times 10^5$ macroparticles

TABLE I. Linac and e-beam parameters scanned in the 5^5 simulations of the FACET-II accelerator. The ranges are chosen closely based on the jitter parameters from the FACET-II technical design report (TDR) [1]. The diagnostics fed to the ML model include random errors introduced artificially to approximate the measurement accuracy present in the accelerator.

Simulation parameter scanned	Range
L1 & L2 phase [deg]	± 0.25
L1 & L2 voltage [%]	± 0.1
Bunch charge [%]	± 1
Input to ML model	Accuracy
L1 & L2 phase [deg]	± 0.1
L1 & L2 voltage [%]	± 0.05
I_{pk} at BC (11,14,20) [kA]	$\pm(0.25, 1, 5)$
ϵ_n at BC (11,14) [μm]	± 1
Beam centroid BC (11,14) [m]	

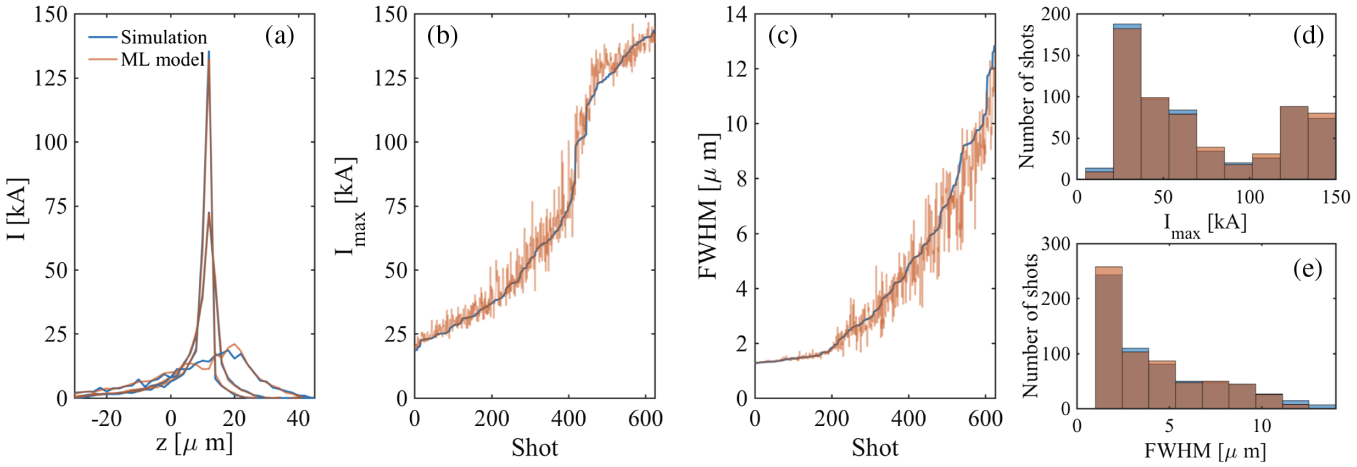

FIG. 2. (a) Examples of reconstructed current profiles from a ML model compared to the true data from simulation at the exit of the FACET-II accelerator. The profiles are taken from a set of 625 test cases which the NN is not trained on. (b)–(e) Sorted plots of the peak current and full width half maximum and histograms of the distributions for the ML prediction and the simulation, showing good agreement between the ML model and simulation data.

which approximates the six-dimensional phase space of the e-beam. We artificially simulate measurement error by inserting random deviations in the diagnostic readings to approximate the measurement accuracy on the real accelerator. The value for each shot was set by adding Gaussian random errors to the mean with the $\pm\sigma$ ranges shown in Table I. These values are determined from operational experience on the FACET linac and estimated performance of diagnostics for FACET-II [1]. Note that we assume the nondestructive emittance measurements will be made using a coherent edge radiation diagnostic currently under design. The prediction accuracy of the ML model is not critically sensitive to these inputs. We also assume the peak current after BC-20 can be measured nondestructively either using coherent edge radiation or coherent undulator radiation [14,15]. For all the examples presented we use the open source ML library scikit-learn, specifically the multilayer perceptron (MLP) regressor from the library's neural network (NN) module. Note that in the plots of the current profile or 2D LPS we use the convention that the head of the beam is on the left.

The results for the current profile prediction are shown in Fig. 2, where the shots displayed are not used in the training of the NN and are 625 randomly selected cases (20% of the total data set). The specific NN architecture for the current profile prediction is a three hidden layer (200,100,50)-neuron fully connected feed-forward NN with a relu as the activation function for each neuron in the hidden layers. There is very good agreement between the NN prediction and the current profile from simulation as shown by the example profiles in Fig. 2(a). A comparison of the peak current and FWHM of the actual distribution vs the prediction also shows good overlap between the two [Figs. 2(b) and 2(c)]. The difference in charge between the predicted and actual profiles, integrated from the current profile, is below 3% in all cases. As shown in Figs. 2(d) and 2(e), the distributions of predicted and simulated peak current and FWHM values are also very well matched.

We use the same MLP regressor and the same diagnostic inputs as for the current profile to predict the 2D LPS distribution. The ability to combine the prediction from both models—one for the current profile and one for the 2D LPS—will provide valuable information for commissioning the accelerator as well as tailoring specific beam properties for different experiments. It is important to note that the LPS reconstruction accuracy depends critically on defining a suitable region of interest for each image which has to be done in a preprocessing step. For the cases shown we crop each 2D LPS picture to a 52×42 pixel image with a 2 μm and 10 MeV/pixel resolution in (z,E) respectively. The resolution values are obtained from estimates of the FACET-II TCAV performance [1]. A quantitative measure of the accuracy of the prediction for each shot is given by the score, defined as

$$\text{score} \equiv R^2 = 1 - \frac{\Sigma_{i,j}(x_{ij}^{\text{true}} - x_{ij}^{\text{predicted}})^2}{\Sigma_{i,j}(x_{ij}^{\text{true}} - \bar{x}^{\text{true}})^2}, \qquad (1)$$

where x_{ij} are the pixel values of the ith row and jth column in the image and $\bar{x}$ denotes the mean over those values for each image. In most cases, as shown in Fig. 3, the 2D LPS reconstruction is quite accurate and provides a good indication of the LPS shape and chirp. In the worst-performing examples, the reconstruction is subject to some numerical artifacts and blurring, which smear out the phase space density. The mean ($\pm$rms) score for the whole test set is 0.68 ± 0.16. The majority of the shots (85%) have a good reconstruction with a score above 0.5. It is interesting to note that shots with worse reconstruction [see Fig. 3(a)] are

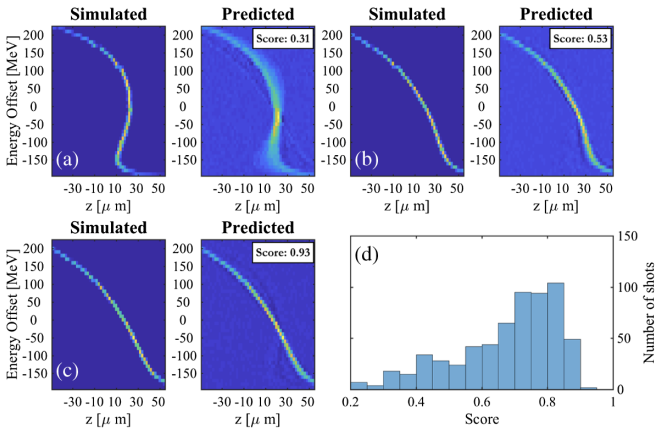

FIG. 3. (a)–(c) Simulated and predicted 2D LPS distributions for FACET-II. The predictions are outputs of a NN with 11 scalar valued nondestructive diagnostic inputs (see Table I). Case (a) represents an imperfect prediction with noticeable blurring and some visual artifacts. Cases (b) and (c) represent average to good predictions and show good agreement with the LPS distribution. (d) Histogram of the prediction accuracy in the test set quantified by the score (R^2 coefficient of the LPS prediction). The mean score is 0.68, and 85% of the shots are above a score of 0.5.

mostly fully compressed or overcompressed (low energy electrons arrive at the same time as high energy electrons on the head of the beam). This is due to the fact that there are fewer shots with these characteristics in the data set since this configuration is the furthest from the nominal output LPS distribution [shown in Figs. 3(b) and 3(c)]. These shots occur due to larger excursions in L1 and L2 phase ($|\Delta\Phi_{L_{1,2}}| = 0.2$–$0.25$ deg).

We have considered a limited set of diagnostic inputs in our simulation study of FACET-II as a conservative approach for determining the feasibility of applying the ML-based virtual diagnostic on the real machine whilst retaining satisfactory prediction accuracy. For the diagnostics considered, the ML model is least sensitive to the emittance and centroid measurements in BC 11–14, achieving a mean score of 0.53 for the LPS reconstruction with those diagnostics removed. A more critical input is the nondestructive peak current measurement after BC20, without which the model's mean prediction score drops to 0.33. We note that a reduction in the measurement error associated with the BC20 current reading (from ± 5 to ± 1 kA) increases the mean score to 0.73. Additional more advanced diagnostics such as current profile monitors after each bunch compressor may also increase the prediction accuracy of the model and are improvements which will be considered in future optimization studies.

III. LCLS EXPERIMENTAL DEMONSTRATION

In order to validate our simulation study of FACET-II, we apply the same NN approach to predict the LPS at the exit of the LCLS linac. The linac was set at a nominal operating energy of 13.4 GeV and 180 pC charge. To collect a data set with a large variety of LPS profiles we scan the values of the L1S phase between -27.8 and -21 degrees and the BC2 peak current between 1–7 kA generating LPS profiles with multiple different features (see Figs. 1 and 4). The diagnostics we use as inputs to the ML model are amplitude and phase readings from L1s and amplitude readings from the L1x accelerator sections as well as nondestructive current measurements (coherent radiation monitors [14]) after BC1 and BC2. The XTCAV was used to measure the LPS at the exit of the accelerator with a resolution of ~ 1.2 μm and 0.92 MeV/pixel [3]. As for the FACET-II case, the prediction accuracy is critically sensitive to preprocessing the LPS image, specifically normalization, centering and cropping of the distribution, with the cases presented cropped to 100×100 pixels. This kind of preprocessing has to be done on-the-fly if such a virtual diagnostic is to be applied to an accelerator during run-times and should be adapted depending on the expected output LPS distribution.

As shown in Figs. 4(a)–4(c), the reconstruction has high fidelity with respect to the current profile and LPS shape. These examples are not used in the training of the ML model and are taken from a test set of 808 shots (20% of the size of total data set). As in the case for the FACET-II simulations, the LPS reconstruction suffers from some numerical artifacts which for some shots smears out the phase space [see Fig. 4(c)]. Nonetheless, for a limited set of input diagnostics (five scalar inputs) and a data set with large variations in the LPS and current profile, the NN is fairly successful in predicting the bunch profile. Note that the current profiles shown are normalized using an independent measurement of the beam charge. As shown in the distribution of peak current vs FWHM [Fig. 4(d)] the measured values and those predicted by the ML model mostly overlap. The model fails to predict some current profiles, mostly with high peak current ($I > 4$ kA), which are furthest away from the nominal LCLS settings at this energy ($I \sim 1$–4 kA). The low prediction accuracy for these shots [labeled "bad shots" in Fig. 4(e)] is correlated to the large discrepancy between the BC2 peak current reading and the peak current measured on the XTCAV. The result is the ML model predicts a profile with lower peak current in-line with the BC2 measurement. This is an important point as it underscores the fact that a diagnostic input error may result in a prediction error for a ML-based virtual diagnostic. One potential way to address this issue would be to have built-in redundancy in the diagnostic inputs fed to the ML model. This would facilitate flagging suspect shots for which there is significant discrepancy between two separate measurements of the same quantity.

To ensure the reliability of the ML prediction a long-term study of the prediction accuracy under the influence of linac drifts and long-term parameter variations is under

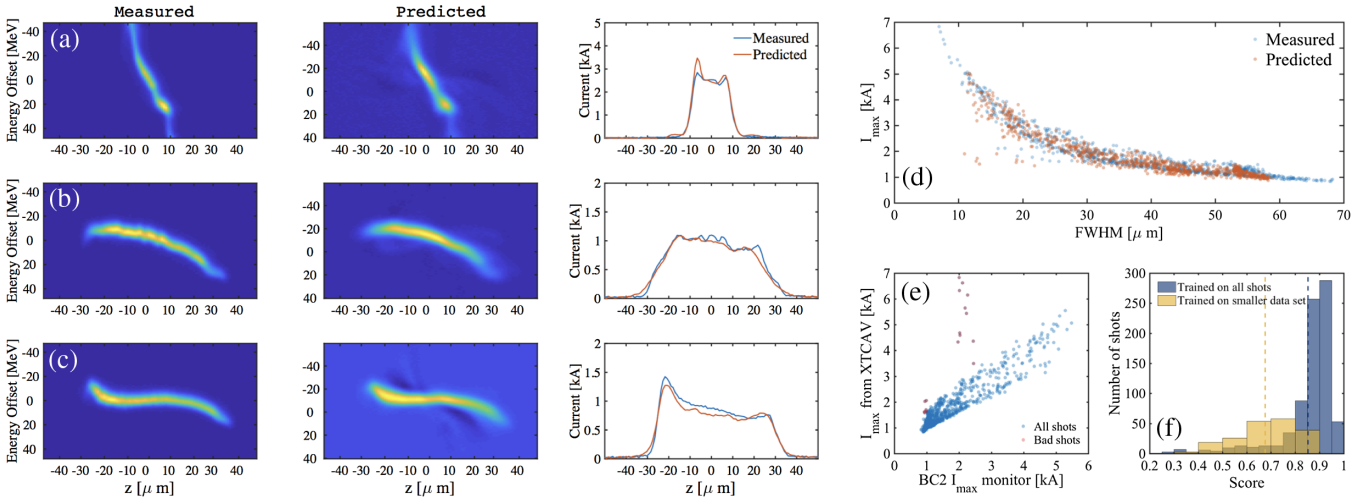

FIG. 4. (a)–(c) Examples of reconstructed LPS and current profiles from the LCLS accelerator. The measured data is collected using the XTCAV and the prediction is made using two separate NNs for the LPS and the current profile. The plots show good agreement in predicting both the LPS and the current profile. For some shots the LPS reconstruction suffers from numerical artifacts [see (c)] which lead to an imperfect reconstruction. (d) Measured and predicted values for the peak and FWHM of the current profile. (e) Correlation between peak current from the XTCAV and the BC2 current monitor highlighting a number of bad shots (2% of the total) where the difference between the two values is large and the prediction accuracy is low. (f) Score for the 2D LPS prediction model trained on 3236 shots and tested on 808 shots randomly selected from the entire data set (grey). Score for the 2D LPS prediction with model trained on 410 shots from the start of the data set and tested on 200 shots recorded at the end of the data set two hours later (yellow).

consideration. Preliminary considerations can be made by examining changes in the prediction accuracy by training the ML models using data from the start of our data set and making predictions on shots taken at the end of the data set (two hours later). For this case, there is no reduction in performance for the prediction of the current profile. The 2D LPS reconstruction however does suffer from a small but noticeable decrease in prediction accuracy as shown in Fig. 4(f). The mean (±rms) score for the 2D LPS reconstruction drops from 0.85 ± 0.14 to 0.68 ± 0.14 which may in part be due to the smaller size of the training set (410 shots compared to 3236). A detailed study of the robustness and reliability of the ML model for longer drift times (one day, one week etc.) and with larger data sets will be conducted and the results reported in future work. Following such long-term prediction accuracy studies, the ML model could be deployed as a virtual diagnostic for predicting the LPS at LCLS when the XTCAV is off.

IV. CONCLUSION

Accurate measurement and control of the LPS distribution is often critically important for applications of high brightness electron beams, ranging from free electron lasers to beam-driven plasma wakefield accelerators. We have explored the feasibility of training a ML-based virtual diagnostic for predicting the LPS distribution of particle accelerators. The study was divided into two parts: a first section using particle tracking simulations of the FACET-II linac as training data for the ML model, and a second using experimental data from the LCLS accelerator. The simulation study explored the single bunch operation mode of FACET-II for which we trained two separate neural networks to predict the current profile and the 2D LPS image based on the input from a number of non-intercepting diagnostics (e.g., beam position monitors, bunch length monitors, emittance measurements). The experimental study performed on the LCLS linac used five measurements from nondestructive diagnostics as well as the XTCAV to measure the electron beam LPS and train the ML models. The results showed close agreement between the predicted current and 2D LPS profiles and those obtained from both simulation and experiment.

It is important to note that the accuracy of a predictive virtual diagnostic based on this kind of supervised learning, in which the neural network generates a mapping between input-output pairs of data, depends critically on the accuracy and resolution of diagnostic inputs. In the experimental study for LCLS, the temporal resolution of the TCAV was ~ 1.2 μm, much smaller than the typical bunch length which ranged from 6–60 μm. In our FACET-II simulation example, the training data fed to the ML model assumes a 2 μm resolution for the LPS images which may present a challenge for the current FACET-II TCAV design. The ability to resolve fine features in the LPS will be challenging, especially in the longitudinal direction due to the very short bunches ($\sigma_z \sim 1$ μm) which are at or beyond the resolution limit of the existing TCAV diagnostic. While the temporal reconstruction may be subject to experimental challenges, the simulation study gives us confidence in the ability of the virtual diagnostic to accurately resolve and predict the energy distribution with ~ 10 MeV/pix resolution.

We expect that we will be able to obtain this resolution or better experimentally using the energy spectrometer downstream of the FACET-II experimental area [1]. We also note that these methods can be used to predict transverse phase space properties of e-beams, such as the emittance, using single shot emittance reconstruction techniques [16,17].

As a next step we plan to include the realistic effect of TCAV measurements in the LUCRETIA tracking code for both single and two-bunch operation, and use the simulated LPS profile on the TCAV rather than the actual LPS distribution to train the ML model. This, together with a more accurate simulation of the diagnostic inputs, will more closely approximate the actual implementation of the virtual diagnostic in the real accelerator. Sensitivity studies related to the accuracy of the neural network prediction based on varying degrees of error for the nondestructive inputs and LPS outputs are currently under consideration. Finally, we plan to use this virtual diagnostic in tandem with optimization methods such at extremum seeking (ES) [18,19], to not only predict the phase space distribution, but to tailor it specifically for different experimental setups. Recent results from LCLS [9] have shown that applying a neural network inverse model to predict machine settings based on LPS images as inputs can improve the convergence and accuracy of an ES-based feedback for customization of the 2D LPS distribution.

ACKNOWLEDGMENTS

We would like to thank L. Alsberg, D. Bohler, A. Lutman and A. Scheinker for useful discussions and assistance in data collection. This work was supported by the U.S. Department of Energy under Contract No. DE-AC02-76SF00515.

[1] Technical design report for the FACET-II project at SLAC National Accelerator Laboratory, http://www.slac.stanford.edu/pubs/slacreports/reports19/slac-r-1072.pdf.

[2] V. Yakimenko, Y. Cai, C. Clarke, S. Green, C. Hast, M. Hogan, N. Lipkowitz, N. Phinney, G. White, and G. Yocky, FACET-II accelerator research with beams of extreme intensities, in *Proceedings of the International Particle Accelerator Conference, Busan, Korea* (JACoW, Busan, 2016).

[3] C. Behrens, F.-J. Decker, Y. Ding, V. A. Dolgashev, J. Frisch, Z. Huang, P. Krejcik, H. Loos, A. Lutman, T. J. Maxwell *et al.*, Few-femtosecond time-resolved measurements of X-ray free-electron lasers, Nat. Commun. **5**, 3762 (2014).

[4] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, Nature (London) **521**, 436 (2015).

[5] A. Edelen, D. R. Erdstrom, P. R. Piot, A. Halavanau, J. Edelen, and S. Biedron, Neural network virtual diagnostic and tuning for the FAST low energy beamline, in *Proceedings of the International Particle Accelerator Conference, Vancouver, Canada* (JACoW, Vancouver, 2018).

[6] A. Sanchez-Gonzalez, P. Micaelli, C. Olivier, T. R. Barillot, M. Ilchen, A. A. Lutman, A. Marinelli, T. Maxwell, A. Achner, M. Agker *et al.*, Accurate prediction of X-ray pulse properties from a free-electron laser using machine learning, Nat. Commun. **8**, 15461 (2017).

[7] A. Edelen *et al.*, First steps toward incorporating image based diagnostics into particle accelerator control systems using convolutional neural networks, in *Proceedings of the North American Particle Accelerator Conference, Chicago, USA* (JACoW, Chicago, 2016).

[8] 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).

[9] A. Scheinker, A. Edelen, D. Bohler, C. Emma, and A. Lutman, Demonstration of Model-Independent Control of the Longitudinal Phase Space of Electron Beams in the Linac-Coherent Light Source with Femtosecond Resolution, Phys. Rev. Lett. **121**, 044801 (2018).

[10] A. Scheinker, D. Rees, B. Garnett, S. Milton, A. L. Edelen, and D. Bohler, Applying artificial intelligence to accelerators, in *Proceedings of the International Particle Accelerator Conference, Vancouver, Canada* (JACoW, Vancouver, 2018).

[11] A. Edelen, S. Biedron, D. L. Bowring, B. E. Chase, D. R. Erdstron, J. Steimel, J. P. Edelen, and P. J. M. van der Slot, Recent applications of neural network-based approaches to modeling & control of accelerators, in *Proceedings of the International Particle Accelerator Conference, Vancouver, Canada* (JACoW, Vancouver, 2018).

[12] C. Joshi, E. Adli, W. An, C. E. Clayton, S. Corde, S. Gessner, M. J. Hogan, M. Litos, W. Lu, K. A. Marsh *et al.*, Plasma wakefield acceleration experiments at FACET II, Plasma Phys. Controlled Fusion **60**, 034001 (2018).

[13] P. Tennenbaum, in *Proceedings of the 21st Particle Accelerator Conference, Knoxville, TN, 2005* (IEEE, Piscataway, NJ, 2005), pp. 4197–4199.

[14] H. Loos, T. Borden, P. Emma, J. Frisch, and J. Wu, Report No. SLAC-PUB 12619, 2007.

[15] C. P. Neumann, W. S. Graves, and P. G. O'Shea, Coherent off-axis undulator radiation from short electron bunches, Phys. Rev. ST Accel. Beams **3**, 030701 (2000).

[16] R. Weingartner, S. Raith, A. Popp, S. Chou, J. Wenz, K. Khrennikov, M. Heigoldt, A. R. Maier, N. Kajumba, M. Fuchs *et al.*, Ultralow emittance electron beams from a laser-wakefield accelerator, Phys. Rev. ST Accel. Beams **15**, 111302 (2012).

[17] S. K. Barber, J. van Tilborg, C. B. Schroeder, R. Lehe, H.-E. Tsai, K. K. Swanson, S. Steinke, K. Nakamura, C. G. R. Geddes, C. Benedetti *et al.*, Measured Emittance Dependence on the Injection Method in Laser Plasma Accelerators, Phys. Rev. Lett. **119**, 104801 (2017).

[18] A. Scheinker and M. Krstic, *Model-Free Stabilization by Extremum Seeking* (Springer, New York, 2017).

[19] A. Scheinker and S. Gessner, Adaptive method for electron bunch profile prediction, Phys. Rev. ST Accel. Beams **18**, 102801 (2015).