

Article

Virtual diagnostic suite for electron beam prediction and control at FACET-II

Claudio Emma ^{1,†*}, Auralee Edelen ^{1,†}, Adi Hanuka ^{1,†}, Brendan O'Shea ^{1,†} and Alexander Scheinker^{2,†}

¹ SLAC National Accelerator Laboratory

² Los Alamos National Laboratory

* Correspondence: cemma@slac.stanford.edu

† These authors contributed equally to this work.

Abstract: We discuss the implementation of a suite of virtual diagnostics at the FACET-II facility currently under commissioning at SLAC National Accelerator Laboratory. The diagnostics will be used for prediction of the longitudinal phase space along the linac, spectral reconstruction of the bunch profile and non-destructive inference of transverse beam quality (emittance) using edge radiation at the injector dogleg and bunch compressor locations. These measurements will be folded in to adaptive feedbacks and ML-based reinforcement learning controls to improve the stability and optimize the performance of the machine for different experimental configurations. In this paper we describe each of these diagnostics with expected measurement results based on simulation data and discuss progress towards implementation in regular operations.

Keywords: machine learning; virtual diagnostics; reinforcement learning control

1. Introduction

Experiments at the forefront of e-beam accelerator R&D require increasingly finer measurement and control of the beam properties during acceleration, transport and delivery to users. For example, research planned at the Facility for Advanced Accelerator Experimental Tests II (FACET-II) [1] under commissioning at SLAC National Accelerator Laboratory aims to demonstrate ultra-high gradient plasma acceleration with preservation of beam quality as well as study the physics of extreme beams with ultra-short bunch lengths (sub μm) and ultra-high peak currents ($> 100 \text{ kA}$). These kinds of applications pose a challenge for state-of-the-art diagnostics as the high intensity and very short pulse duration limits the applicability, effectiveness and accuracy of traditional measurement techniques. This requires a re-thinking of the existing suite of diagnostics which will be used to characterize and control the e-beam properties and facilitate the success of the experimental program.

In this paper we discuss the application of Machine Learning (ML) methods for developing a suite of virtual diagnostics to be used for electron beam prediction and control at the FACET-II facility. These virtual diagnostics will provide a shot-to-shot non-destructive measurement of the electron beam longitudinal and transverse properties along the accelerator and serve as input for conventional feedbacks and optimization algorithms based on reinforcement learning that can tailor the beam properties for specific experimental applications (see Fig. 1 for schematic). These diagnostics will work in tandem with traditional measurement techniques to provide otherwise unavailable information to experimenters which will aid in both the machine setup, optimization and interpretation of experimental results during offline data analysis. We will discuss a number of examples of ML-applications of FACET-II in this work: reconstruction of e-beam Longitudinal Phase Space (LPS) along the accelerator, spectral virtual diagnostics for enhancing the accuracy and confidence of beam profile and LPS predictions, ML-based image analysis for inferring e-beam emittance using edge radiation, adaptive model-based 6D phase space predictions, and reinforcement learning controls. We describe the ML methods applied in this work in the following section and summarize key results each ML-driven diagnostic and control method in sections 3.1-3.5.

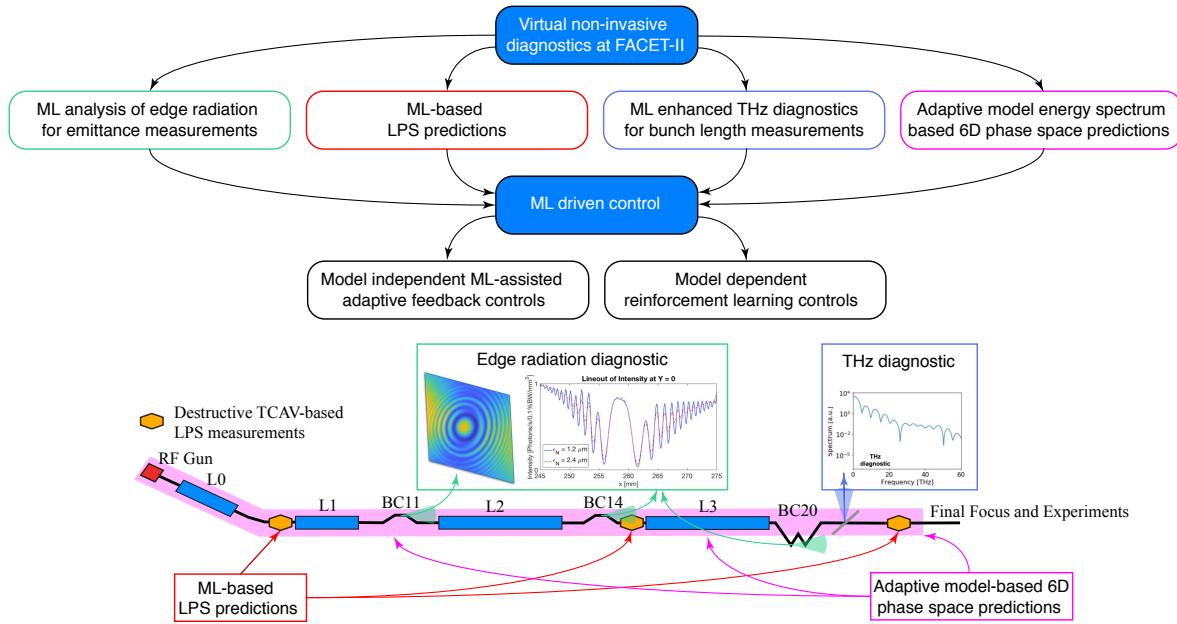


Figure 1. Schematic of the planned virtual diagnostic suite and ML driven controls at FACET-II

2. Materials and Methods

One of the main running configurations for PWFA experiments at FACET-II will involve accelerating two bunches from the photocathode to the interaction point (IP) at the plasma entrance with specific longitudinal profile properties and drive-witness bunch separation. For a full description of PWFA experiments at FACET-II see Ref. [2]. The major goals for the PWFA experiments will be to demonstrate pump depletion of the 10 GeV drive beam and acceleration of the witness beam to approximately 18 GeV while preserving good beam quality. The figures of merit for the beam quality will be preservation of energy spread and emittance of the witness bunch, and these will need to be measured on a shot-to-shot basis for both the incoming distribution and the accelerated witness beam. To this end, accurate measurements of the bunch profile entering the plasma are essential for the success of the experimental campaign. Using ML based virtual diagnostics to non-destructively predict the LPS distribution at the entrance of the plasma will provide previously unavailable information which can be used to both understand experimental results from PWFA and tune the beam parameters to facilitate the PWFA interaction.

2.1. ML-enhanced diagnostics

Previous work has demonstrated the feasibility of using Machine Learning (ML) models as virtual diagnostics to non-destructively predict the LPS distribution of FACET-II single bunch operation (in simulation) and at LCLS (in experiment) [3]. These studies used neural networks to create a mapping between non-destructive diagnostic inputs (e.g. linac and e-beam diagnostics available on a single shot basis) and destructive diagnostic output which measures the beam LPS. What was previously not included the simulation study of Ref. [3] is the impact of LPS measurement resolution on the ML-based LPS reconstruction. At FACET-II the LPS distribution of the electron bunch can be measured destructively at the entrance of the plasma with an X-band TCAV operating at a peak voltage of 20 MV with a longitudinal resolution of a few μm RMS. This introduces a challenge for accurately characterizing the longitudinal bunch profile, as the accelerator is expected to produce very short bunches ($\sigma_z \leq 1 \mu\text{m}$) beyond the TCAV resolution. In section 3.1 of this work we examine the effect of the TCAV measurement on the performance of the ML-based virtual diagnostic and discuss its application in the FACET-II two-bunch operation mode. We present results from 3,125 particle tracking simulations [4] of the FACET-II linac from the exit of the injector to the end of the linac with induced

jitter of key accelerator and beam parameters described in Table 1. The simulation data is used to train a ML-based virtual diagnostic for the two-bunch LPS and results show very good agreement between the simulated LPS distribution as measured by the TCAV and the LPS distribution predicted by the ML model. Due to TCAV resolution limits there is some discrepancy when we use the projection of the measured LPS distribution to infer the current profile at the entrance of the plasma. This discrepancy affects the accuracy of the ML-based virtual diagnostic for high current shots with short bunch length. In order to flag these shots and increase the confidence and accuracy of the virtual diagnostic predictions we discuss the incorporation of spectral signals in a spectral virtual diagnostic in section 3.2.

In addition to longitudinal diagnostics, transverse diagnostics of the beam emittance are of critical importance for FACET-II's aim of achieving acceleration in PWFA while preserving beam quality. To address this need we will be implementing a series of single-shot non-destructive emittance measurement based on the interference of edge radiation at the location of bunch compressors (BC11, BC14 and BC20) as well as at the exit of the photoinjector before the first linac section (see Fig. 1). These diagnostics will provide a snapshot of the emittance at each point along the linac allowing experimenters to outline critical sources of emittance growth improving understanding of the beam dynamics in the transport, acceleration and compression from the photoinjector to the experimental area. The diagnostics will require advanced image analysis in order to obtain a real-time estimate of the beam emittance from the edge radiation interference pattern. This analysis will be conducted using Convolutional Neural Networks (CNNs) trained on image data from simulations of the edge radiation interference pattern with beams of different emittance at different points along the linac. An example of this kind of simulation is give in section 3.3. We will incorporate signals from the ML-enhanced suite of diagnostics into control methods to improve the quality and stability of the electron beams as described in the following paragraphs.

2.2. ML-enhanced control

While powerful ML methods are able to learn complex input-output relationships in large many parameter systems directly from data, their accuracy will degrade if the system for which they have been trained changes with time. One way to compensate for time variation is to repeatedly re-train several layers of an ML tool such as a CNN, but this may be problematic for particle accelerator applications such as FACET-II where, for example, acquiring new LPS training data requires re-tuning a beam line in order to transport the beam to a TCAV-based diagnostic rather than to the interaction point.

Another approach to dealing with time-varying systems is the use of model-independent adaptive feedback. Adaptive feedback is by design applicable to unknown and changing systems. Recently an adaptive tuning method known as extremum seeking (ES) has been developed for the ptimization and stablization of unknown, time-varying, nonlinear dynamic systems which is able to tune many parameters simultaneously based only on noisy measurement data [5]. This ES method has been implemented for various particle accelerator applications including virtual diagnostics and beam optimization. In [6] an adaptive ES feedback-based virtual LPS diagnostic was developed which was able to non-invasively predict TCAV LPS measurements and track changing beam parameters over a wide range of bunch lengths and bunch-to-bunch separation at FACET. ES was also recently demonstrated for online multi-objective optimization for simultaneous trajectory control and transverse emittance growth minimization at the electron beam line of the AWAKE plasma wakefield accelerator at CERN by adaptively simultaneously tuning 15 parameters: 2 solenoids, 3 quadrupole magnets, and 10 steering magnets [7]. One possible limitation of adaptive feedback is that it is usually based on local iterative methods whose convergence speed may be lengthy or which may become stuck at local extrema in very large parameter spaces. Recently, a first of its kind adaptive ML approach was demonstrated for the automatic control of the LPS of the electron beam in the LCLS FEL [8]. In [8] a neural network was trained to map TCAV-based LPS measurements directly to the accelerator

parameters required for those beam properties, the NN's predictions were able to find the correct neighborhood of parameter space, after which ES was able to adaptively tune all of the parameters to zoom in on and track their optimal settings despite noise and time-variation of both the beam and accelerator parameters.

In contrast to doing adaptive feedback and system modeling separately, a Reinforcement Learning (RL) approach can be used which combines both model learning and control. In RL an "agent" (i.e. the controller) learns how to interact with an environment over time in order to achieve the highest long-term reward. In the context of accelerator tuning, the "environment" is the accelerator and the reward could be, for example, specific beam shapes one wants to achieve. Critically, RL takes the present system state (e.g. system control settings and observable outputs) into account when choosing the next action to take. Over the course of many interactions with the environment the RL algorithm learns to improve its overall control strategy while retaining information about previously-visited environmental states. Some RL algorithms directly learn a map from system states to actions (a learned policy). Others use a learned model that estimates the likely future reward that will be obtained when specific actions are taken in various observed system states. These predictions can then be combined with simple policies to determine actions to take. RL has been applied to the problem FEL tuning at LCLS [9] and FERMI@Elettra [10], and it is at present being developed for a variety of other online optimization and control tasks in accelerators (for example, round-to-flat beam transforms at UCLA [11], beam size control at AWAKE).

Deep RL leveraging neural networks is appealing for the task of LPS tuning in part because it can directly learn policies from images. Deep RL has been used for end-to-end visuomotor control tasks in robotics cite and game-playing tasks where the state of the system is given as an image. In the case of FACET-II, both LPS image and upstream diagnostics, such as the virtual cathode camera (VCC) could be used to inform the present system state. By directly using images, features that otherwise would not be captured by bulk scalar metrics derived from the images may be learned from and exploited in tuning, potentially leading to finer control over the LPS. In section 3.4 we discuss the use of combined adaptive feedback and ML for virtual 6D diagnostics of the FACET-II beam's phase space and for active feedback control of the beam properties.

3. Results

3.1. Longitudinal Phase space reconstruction

We present three examples of the simulated LPS profiles at the FACET-II experimental area as measured by the TCAV in Fig. 2 with corresponding current profiles and prediction from the ML-based virtual diagnostic. This ML tool is a NN that takes scalar inputs from accelerator and electron beam diagnostics and outputs a prediction of the 2D LPS image. The diagnostic inputs can be measured

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]	± 0.25
L1 & L2 voltage [%]	± 0.05
I_{pk} at BC (11,14,20) [kA]	± (0.25,1,5)
Beam centroid BC (11,14) [m]	N/A

Table 1. 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 TDR [12]. The diagnostics fed to the ML model include random errors introduced artificially to approximate the measurement accuracy present in the accelerator.

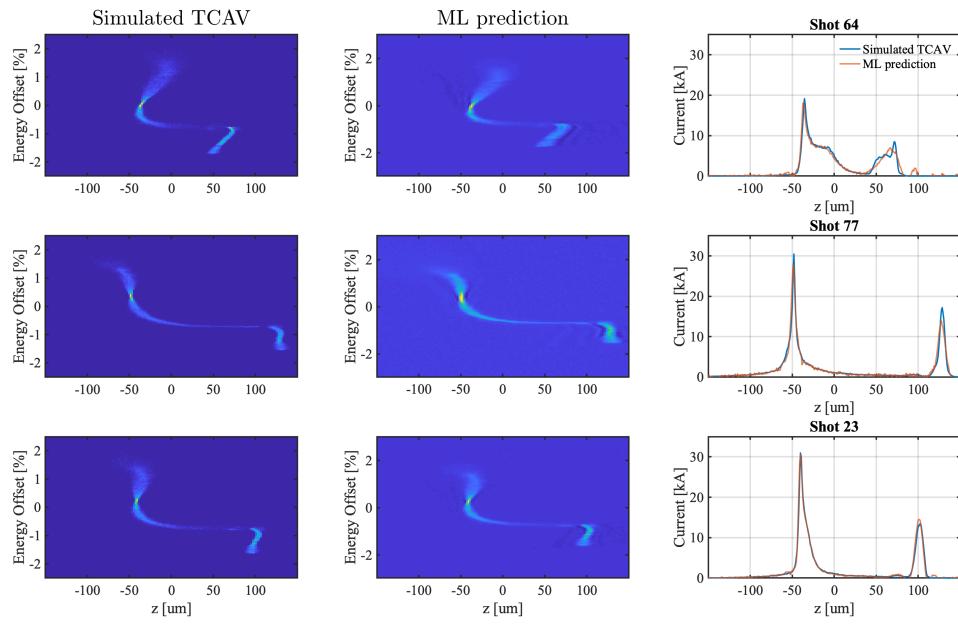


Figure 2. Example shots from numerical simulation of the FACET-II two-bunch operation mode with nominal jitter values given in Table 1. The ML model accurately predicts the LPS distribution including chirp, time separation and bunch charge ratio. The current profile matches well with what is measured on the TCAV. As shown in Fig. 3 this may deviate from the true current profile at the IP due to resolution limits of the TCAV for some high current shots.

non-destructively on a single shot basis and include random offsets to the readings to simulate expected measurement accuracy (see Table 1). The three distributions shown represent an under-compressed, over-compressed and nearly fully-compressed (nominal) beam respectively. Note that the head of the bunch is on the left of the images. The ML model we used was a three-layer fully-connected neural network with (500,200,100) neurons in each successive hidden layer and a rectified linear unit activation function for each neuron. The network was trained using the open source ML library Tensorflow, and two separate models with the same architecture were trained for the 2D LPS prediction and the 1d current profile prediction. As evidenced Fig. 1, we see very good agreement between the LPS profiles measured by the TCAV and those predicted by the ML model. There is also good agreement between the ML-predicted current profiles and those extracted from the TCAV image. The variety of LPS images input to the ML model for training result from the expected shot-to-shot-jitter of linac and e-beam parameters outlined in the FACET-II Technical Design Report (TDR, see Table 1) [12]. The nominal settings produce a $\sim 150 \mu\text{m}$ bunch spacing, a 2:1 ratio between the peak currents and a 3:1 ratio in the bunch charge between the drive and witness beam. The variation in bunch profile from shot-to-shot jitter results in a 9 % RMS variation in the drive-witness charge ratio, a 30 μm RMS variation in the bunch separation and a 36 % RMS variation in the ratio of the peak current from the nominal settings. These parameter variations are well predicted by the ML model.

As discussed above, the FACET-II two-bunch configuration operates at near full compression and will generate very short bunches with RMS sizes of a few μm putting them at the limit of the TCAV resolution. This means the values measured for the peak current on the TCAV sometimes differ from the values at the IP and therefore so will the prediction from the ML model, which is trained using TCAV measurements as inputs. We examine this discrepancy in detail in Fig. 2 where we show the same current profiles from the three example shots in Fig. 1 as measured on the TCAV and compare this with the current profile we calculate from the distribution at the IP binned at $0.25 \mu\text{m}$ per pixel. There are a few observations we can make by looking at Fig. 2 (a)-(c). The first is that the peak current

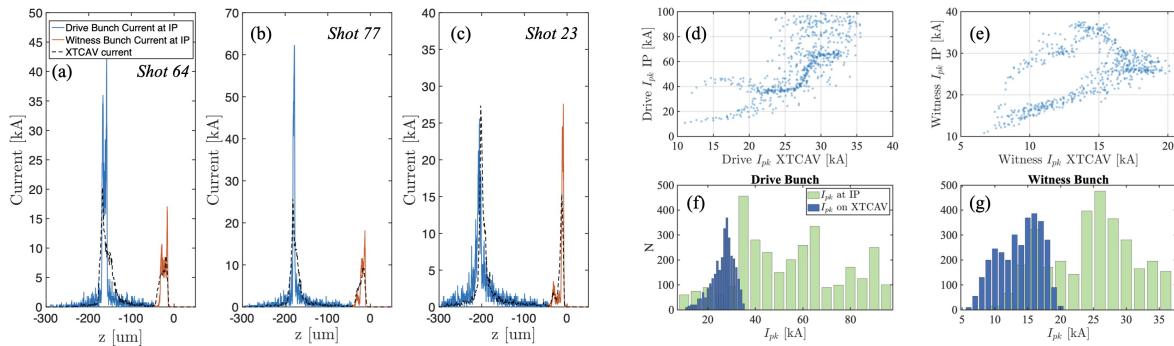


Figure 3. (a)-(e) Single shot comparisons of the current profiles measured by the TCAV and calculated from the macroparticle distribution in simulation dumped at the IP. The shots match those displayed in Fig. 1. (d)-(f) Comparison of the peak current measured by the TCAV vs. at the IP for the drive and witness beam. The plots show that the TCAV measurement underestimates the peak current value and smears out some details of the current distribution for the very short bunches which will be produced at FACET-II.

values measured by the TCAV under-estimate the true value for shots with peak current greater than ~ 35 kA. We note that these high peak currents are greater than those which we plan to deliver for the two-bunch pump depletion experiments outlined in Ref. [2]. Nonetheless, close to the nominal settings (as shown in Fig. 2c) the correct value of the ratio of the peak currents may be under-estimated if the witness bunch current profile is poorly resolved by the TCAV measurement. To quantitatively understand the limits imposed by the TCAV measurement we can estimate the longitudinal resolution as follows:

$$\sigma_z = \frac{E_e}{eV_{rf}k_{rf}|\sin \Delta\psi|} \frac{\sqrt{\sigma_S^2 + \beta_S \epsilon}}{\sqrt{\beta_T \beta_S}} \quad (1)$$

where E_e is the electron beam energy V_{rf} , k_{rf} is the TCAV voltage and wavenumber, $\Delta\psi$ is the phase advance between the TCAV and the measurement screen, σ_S is the resolution of the screen (we assume $4\mu\text{m}$ for a transition radiation target), β_S is the beta function at the screen ϵ is the beam emittance and β_T is the beta function at the TCAV. The ~ 35 kA max resolvable peak current come from the constrained optimization of the beta function at the screen and at the TCAV while meeting the beam stay-clear constraints in the experimental area and mitigating loss of resolution from chromatic errors and emittance growth in the transport. For a $10\mu\text{m}$ normalized emittance at 10 GeV with a phase advance of $3\pi/2$ between TCAV and screen, the optimized values of β_T and β_S are 107 and 6.5 m give a resolution of $\sigma_{z,min} = 4.58\mu\text{m}$. Given a Gaussian drive bunch at 1.5 nC charge this corresponds to $I_{max} = 39.2$ kA which is in reasonable agreement with the trend shown in the scatter plot in Fig. 2 (d).

For shots that are not beyond the TCAV resolution we can see from Fig. 2 (d)-(e) that we can correlate the TCAV measured peak current with the peak current at the IP. These shots are mostly in the region defined by $I_{pk,drive} < 30$ kA and $I_{pk,wit} < 16$ kA as measured on the TCAV. Some shots in this region still show large discrepancy between the TCAV current profile and that measured at the IP and these represent the spiky ‘double-horn’ type distributions in the drive and witness beam exemplified in Fig. 2 (a). One of the challenges this particular virtual diagnostic faces is to flag whether or not a single shot falls within the ‘high-current’ region beyond the TCAV resolution. Accurately determining this on a shot-to-shot basis will provide added assurance that the current profiles predicted by the ML model map to the electron beam current profile at the IP and we discuss efforts to address this issue in the following section.

3.2. Spectral virtual diagnostics

One potential method to address the current profile resolution limits of virtual diagnostics trained on TCAV data would be to use a secondary non-destructive diagnostic in tandem with the ML prediction that is sensitive to changes in the peak current beyond the TCAV resolution [13]. This would help identify the region in which a given shot falls. The secondary diagnostic may be a mid-IR and/or Thz spectrometer similar to those described in Ref. [14,15] and could use diffraction or bend radiation as a non-destructive radiation source. It may also be possible to implement a simple upgrade (adding an appropriate set of spectral filters) to the existing radiation-based bunch length monitor at the exit of the final bunch compressor (see Ref. [16]) to mimic a more complicated spectroscopic measurement. This would allow us to measure the integrated radiation signal over a given frequency band proportional to the bunch length for the high peak current shots. Figure 4(a) shows example of high peak current shot (blue) that would be smeared out on the TCAV, thus appearing similar to a lower current shot (red) on the TCAV. However, the corresponding shots' spectrum is clearly different - see 4(b). We use this spectrum in two ways: to train a spectral virtual diagnostic and to flag high current (bad) shots. First, we train a VD using the spectrum as an input to a NN rather than scalars as described in subsection 3.1. As shown in Ref. [13], the spectral VD predicts the current profile more accurately than the scalar VD. We can quantify the discrepancy between the predictions of the two VDs on a single shot basis and if this is greater than some pre-determined threshold, we may supplement the prediction of that shot with a low-confidence label.

Second, we use an integrated spectrum signal in some frequency band (shown by grey interval in 4(b) to veto suspect shots beyond the TCAV peak current resolution. We optimize the boundaries of this band pass filter by maximizing the difference between low and high peak current shots. Higher peak current shots would have more spectral content at higher frequencies. Figure 4(c) shows the fraction of shots that could be trusted with high reliability to be within the TCAV resolution. The measured TCAV current should be correlated with the current at the IP for shots that are within the TCAV resolution - as shown in Figure 4(d) for the optimized frequency band. Shots with spectral intensity smaller than the a pre-defined threshold (shown in black line) will be flagged. Determining on a shot-to-shot basis if the predicted TCAV current profile is valid will be complementary to the spectral VD, increasing the confidence in that prediction.

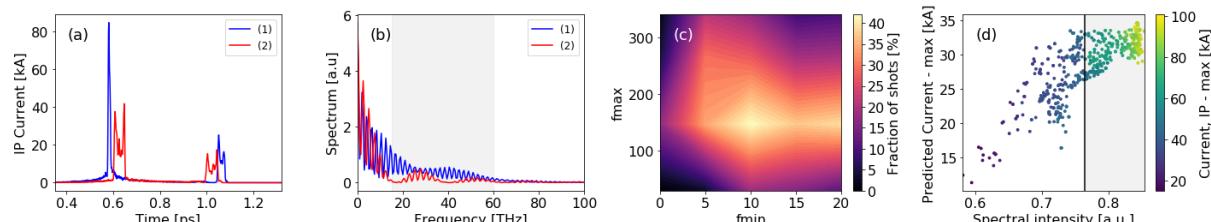


Figure 4. (a) Single shot current profile at the interaction point (IP) of high and low peak current profile that would be appear similar on the TCAV and (b) their matching spectrum. (c) Fraction of shots that are within the TCAV resolution for various frequency bands. (d) Maximum predicted TCAV current vs integrated spectral intensity that was optimized to maximize the fraction of shots in (c). Corresponding maximum IP current is shown in the colorbar.

Using the spectrum, we obtain increased confidence in the overall VD prediction, especially in the cases of high current shots. In addition, spectral VD is able to resolve shot-to-shot features of the electron beam (such as microbunching in the LPS) in cases wherein scalar VD isn't applicable at all since integrated scalar beam diagnostics cannot capture them.

3.3. Emittance reconstruction using edge radiation

Non-destructive single-shot monitoring of beam emittance with ML-based image analysis will be carried out at FACET-II by training CNNs to predict the beam emittance given a 2D interference

pattern of edge radiation emitted by the electron beam. This technique was previously studied in applications for the Siberia-1 electron storage ring [17] and the FERMI free electron laser [18]. We have carried out simulations of the edge radiation interference process at FACET-II using the code Synchrotron Radiation Workshop (SRW) [19]. Simulation results for two different emittance beams are shown in Fig. 5. The blurring of the interference fringes at the larger emittance value of $1.2 \mu\text{m}$ is visible to the naked eye and is also evidenced in the lineouts displayed in Fig. 5 c. We will plan on training CNNs in a supervised learning paradigm using multiple simulations of the edge radiation for different electron beams and train the ML-based image analysis software with simulation and experimental data. This will be accomplished starting from the photoinjector and progressively moving down the accelerator where the alignment tolerances for the radiation pattern generated at the magnet edges becomes tighter as the electron beam energy increases.

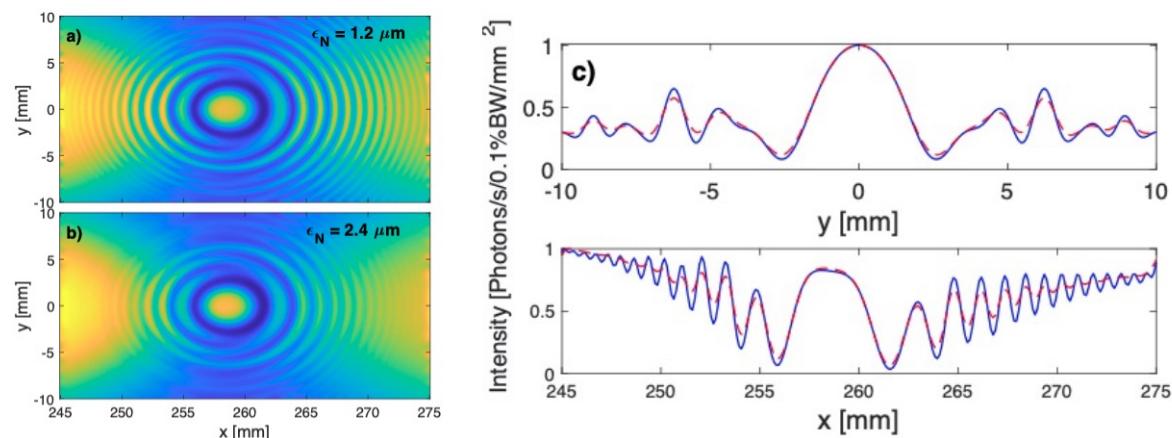


Figure 5. Examples of edge radiation interference patterns for beams of different emittance at FACET-II with corresponding lineouts. The interference pattern will be used as input to a CNN based image analysis which will determine the beam emittance from the 2D images in real-time.

3.4. Adaptive feedback with ML for virtual 6D diagnostics and control

A proof of principle adaptive model-based virtual diagnostic was demonstrated at FACET [6] and shown to track the TCAV measurement-based LPS non-invasively based only on accelerator readouts and an energy spread spectrum of the beam. In [6] the focus was on tracking time-varying accelerator parameters and predicting 1D current profiles. Recently, this approach was studied in simulation for FACET-II in which the 2D LPS was predicted and tracked based on energy spread spectrum measurements alone [20]. Figure 6 shows one simulation-based example of adaptive tuning that adjusts parameters to match two initially different energy spread spectra, such as those that will be measured non-invasively at FACET-II, and the result is a match of the LPS and its projections. The method works by repeatedly comparing the measured and simulated energy spread spectra and adjusting multiple components of the model in order to get a close match. Once the spectra are in agreement the physics model's constraints result in a unique reconstruction of the electron bunch's LPS.

Such an adaptive model tuning-base approach in which an online model is adjusted in real time based on beam and accelerator component measurements in order to give more accurate predictions of the beam's 6D phase space and to track both beam and accelerator parameters as they drift with time.

Once such an adaptive model-based diagnostic is running, it not only provides a prediction of the 2D LPS, but of the entire 6D phase space of the beam since it is adjusting the 6D LUCRETIA tracking code [4]. Figure 7 shows one example of a simulation of one section of bunch compression in FACET-II for a 2 nC beam using 2e5 macroparticles. If this model had been tuned online to match all possible real time diagnostics from FACET-II, then based on previous results [6,20], it is expected to provide a virtual diagnostic of the actual beam's 6D phase space. Furthermore, such an adaptively tuned

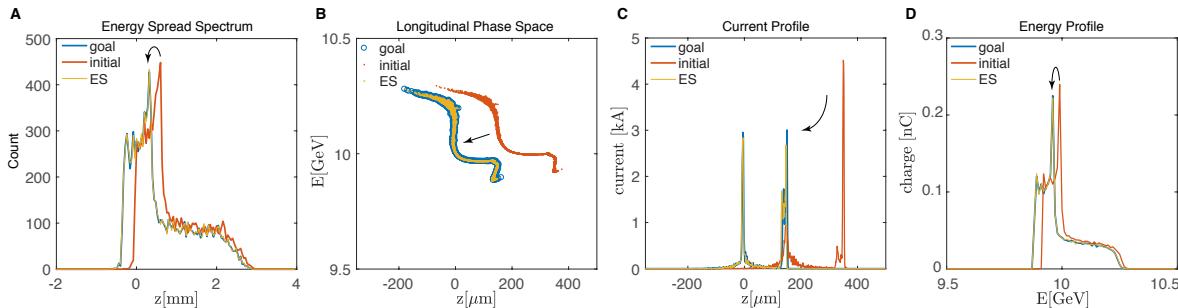


Figure 6. Simulation-based demonstration of adaptive tuning-based phase space diagnostic. Adaptively tuning the model to match a measured energy spread spectrum (A). Once the spectrum matches the LPS is uniquely reconstructed (B). Projection onto the z-axis gives an accurate diagnostic of the current profile (C). Projection onto the E-axis gives an accurate diagnostic of the energy profile (D).

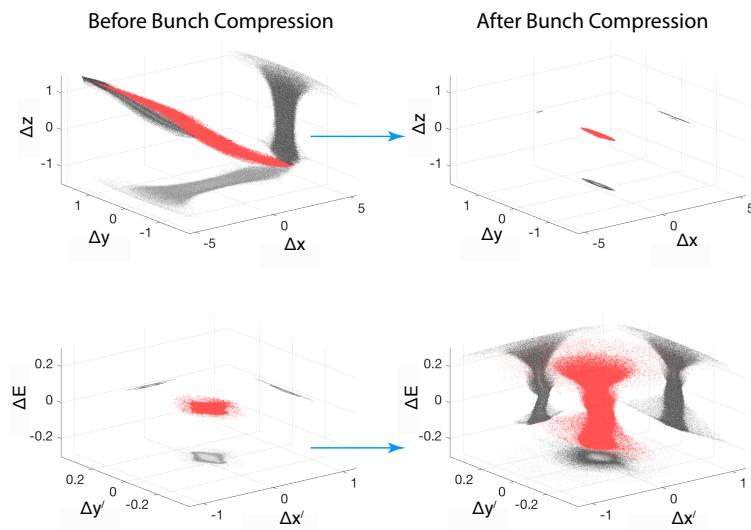


Figure 7. Simulation results showing the 6D phase space evolution of electron bunch compression in FACET-II. Axes are shown with arbitrary units with the same ranges before and after compression to show the difference.

model-based diagnostic approach will utilize all of the non-invasive diagnostics described in this paper as inputs to help more accurately match model predictions with the actual beam. Finally, the real time beam data provided by an adaptive diagnostic can be used to perform real time feedback control such as maintaining a desired phase space or to tune the beam to achieve desired custom current profiles and phase space distributions, as demonstrated in simulation in [20].

3.5. Reinforcement learning controls

For FACET-II, we plan to implement and test deep RL for the injector and linac. The RL algorithm takes present images of the LPS, external state information like the present settings and virtual cathode camera images, and a target LPS image to decide what the next setting changes should be. This is then repeated until the target phase space is achieved. Initial studies of this RL approach, and comparisons with approaches outlined above, will be conducted during the commissioning stage of the FACET-II accelerator which is currently underway and results will be reported on in future publications.

4. Discussion

In this paper we have described the approach for implementing a suite of ML-based virtual diagnostics and ML based controls to be used in regular operations at the FACET-II accelerator

facility. These ML based tools will be used to aid machine setup, optimize beam delivery for different experiments, on-the-fly data analysis to rapidly extract beam parameters, and offline data analysis/interpretation of experimental results. Moving from proof-of-concept demonstrations of to regular deployment of these diagnostics in accelerator operations will require addressing the major challenge of how to obtain a reliable and accurate measures of the uncertainty associated with ML model predictions. To this end we are planning on employing redundancy in our ML-based predictions of beam properties e.g. using spectral data and scalar linac parameters to independently predict the beam current profile. Part of the challenge of transitioning these methods to operation will also require understanding how to effectively re-train ML models in the presence of machine drifts. Progress in this area will begin by quantifying the ML-model accuracy over time and retraining the model as the observed prediction error increases beyond some acceptable threshold.

Author Contributions: LPS reconstruction, C.E.; Spectral Virtual Diagnostics, A.H.; edge radiation, B.O.; reinforcement learning controls, A.E.; adaptive feedback virtual diagnostics, A.S.

Funding: This research was funded in part by the U.S. Department of Energy under contract number DE-AC02-76SF00515.

Acknowledgments: The authors would like to acknowledge G. White for his contributions in setting up the start-to-end simulations of the FACET-II linac.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

LPS	Longitudinal Phase Space
SVD	Spectral Virtual Diagnostic
RL	Reinforcement learning
CER	Coherent Edge Radiation

1. Yakimenko, V.; Alsberg, L.; Bong, E.; Bouchard, G.; Clarke, C.; Emma, C.; Green, S.; Hast, C.; Hogan, M.J.; Seabury, J.; Lipkowitz, N.; O'Shea, B.; Storey, D.; White, G.; Yocky, G. FACET-II facility for advanced accelerator experimental tests. *Phys. Rev. Accel. Beams* **2019**, *22*, 101301. doi:10.1103/PhysRevAccelBeams.22.101301.
2. Joshi, C.; Adli, E.; An, W.; Clayton, C.E.; Corde, S.; Gessner, S.; Hogan, M.J.; Litos, M.; Lu, W.; Marsh, K.A.; Mori, W.B.; Vafaei-Najafabadi, N.; O'shea, B.; Xu, X.; White, G.; Yakimenko, V. Plasma wakefield acceleration experiments at FACET II. *Plasma Physics and Controlled Fusion* **2018**, *60*, 034001. doi:10.1088/1361-6587/aaa2e3.
3. Emma, C.; Edelen, A.; Hogan, M.J.; O'Shea, B.; White, G.; Yakimenko, V. Machine learning-based longitudinal phase space prediction of particle accelerators. *Phys. Rev. Accel. Beams* **2018**, *21*, 112802. doi:10.1103/PhysRevAccelBeams.21.112802.
4. Tenenbaum, P. Lucretia: A Matlab-based toolbox for the modeling and simulation of single pass electron beam transport systems. Proceedings of the Particle Accelerator Conference, Knoxville, USA, 2005, pp. 4197–4199.
5. Scheinker, A. Simultaneous stabilization and optimization of unknown, time-varying systems. 2013 American Control Conference. IEEE, 2013, pp. 2637–2642. doi:10.1109/ACC.2013.6580232.
6. Scheinker, A.; Gessner, S. Adaptive method for electron bunch profile prediction. *Physical Review Special Topics-Accelerators and Beams* **2015**, *18*, 102801.
7. Scheinker, A.; Hirlaender, S.; Velotti, F.M.; Gessner, S.; Della Porta, G.Z.; Kain, V.; Goddard, B.; Ramjiawan, R. Online multi-objective particle accelerator optimization of the AWAKE electron beam line for simultaneous emittance and orbit control. *AIP Advances* **2020**, *10*, 055320.

8. Scheinker, A.; Edelen, A.; Bohler, D.; Emma, C.; Lutman, A. Demonstration of model-independent control of the longitudinal phase space of electron beams in the Linac-coherent light source with Femtosecond resolution. *Physical review letters* **2018**, *121*, 044801.
9. Wu, J.; Huang, X.; Raubenheimer, T.; Scheinker, A.; others. Recent On-Line Taper Optimization on LCLS. 38th Int. Free Electron Laser Conf.(FEL'17), Santa Fe, NM, USA, August 20-25, 2017. JACOW, Geneva, Switzerland, 2018, pp. 229–234.
10. Bruchon, N.; Fenu, G.; Gaio, G.; Lonza, M.; O'Shea, F.H.; Pellegrino, F.A.; Salvato, E. Basic Reinforcement Learning Techniques to Control the Intensity of a Seeded Free-Electron Laser. *Electronics* **2020**, *9*, 781. doi:10.3390/electronics9050781.
11. Edelen, A.; E., C.; Emma, C.; A., H.; others. Machine Learning-Based Tuning of the Round-to-Flat Beam Transform at the UCLA Pegasus Photoinjector. in preparation.
12. Technical design report for the FACET-II Project at SLAC National Accelerator Laboratory.
13. Hanuka, A.; Emma, C.; Maxwell, T.; Fisher, A.; Jacobson, B.; Hogan, M.J.; Huang, Z. Accurate and confident prediction of electron beam longitudinal properties using spectral virtual diagnostics, 2020, [arXiv:physics.acc-ph/2009.12835].
14. Wesch, S.; Schmidt, B.; Behrens, C.; Delsim-Hashemi, H.; Schmüser, P. A multi-channel THz and infrared spectrometer for femtosecond electron bunch diagnostics by single-shot spectroscopy of coherent radiation. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **2011**, *665*, 40–47. doi:10.1016/j.nima.2011.11.037.
15. Maxwell, T.J.; Behrens, C.; Ding, Y.; Fisher, A.S.; Frisch, J.; Huang, Z.; Loos, H. Coherent-Radiation Spectroscopy of Few-Femtosecond Electron Bunches Using a Middle-Infrared Prism Spectrometer. *Phys. Rev. Lett.* **2013**, *111*, 184801. doi:10.1103/PhysRevLett.111.184801.
16. Loos, H.; Borden, T.; Emma, P.; Frisch, J.; Wu, J. Relative bunch length monitor for the Linac Coherent Light Source (LCLS) Using Coherent Edge Radiation. *SLAC-PUB 12619* **2007**.
17. Chubar, O. Determining electron beam parameters from edge radiation measurement results on Siberia-1 storage ring. Proceedings Particle Accelerator Conference. IEEE, 1995. doi:10.1109/pac.1995.505564.
18. Fiorito, R.; Shkvarunets, A.; Castronovo, D.; Cornacchia, M.; Di Mitri, S.; Kishek, R.; Tschalaer, C.; Veronese, M. Noninvasive emittance and energy spread monitor using optical synchrotron radiation. *Phys. Rev. ST Accel. Beams* **2014**, *17*, 122803. doi:10.1103/PhysRevSTAB.17.122803.
19. Chubar, O.; Elleaume, P. Accurate and efficient computation of synchrotron radiation in the near field region. Proc. of EPAC, 1998, p. 1177.
20. Scheinker, A.; Gessner, S.; Emma, C.; Edelen, A.L. Adaptive model tuning studies for non-invasive diagnostics and feedback control of plasma wakefield acceleration at FACET-II. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **2020**, p. 163902.