

PAPER • OPEN ACCESS

Improving surrogate model accuracy for the LCLS-II injector frontend using convolutional neural networks and transfer learning

To cite this article: Lipi Gupta *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 045025

View the [article online](#) for updates and enhancements.

You may also like

- [Toward a protocol for quantifying the greenhouse gas balance and identifying mitigation options in smallholder farming systems](#)
T S Rosenstock, M C Rufino, K Butterbach-Bahl et al.
- [Foreword](#)
- [Advancing agricultural greenhouse gas quantification](#)
Lydia Olander, Eva Wollenberg, Francesco Tubiello et al.



PAPER

OPEN ACCESS


RECEIVED
2 April 2021REVISED
26 July 2021ACCEPTED FOR PUBLICATION
17 September 2021PUBLISHED
1 October 2021

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Improving surrogate model accuracy for the LCLS-II injector frontend using convolutional neural networks and transfer learning

Lipi Gupta^{1,*} , Auralee Edelen², Nicole Neveu², Aashwin Mishra², Christopher Mayes² and Young-Kee Kim¹¹ Physics Department, University of Chicago, Chicago, IL, United States of America² SLAC National Accelerator Laboratory, Menlo Park, United States of America

* Author to whom any correspondence should be addressed.

E-mail: lipigupta@uchicago.edu**Keywords:** convolutional neural network, transfer learning, particle accelerator physics, surrogate modeling

Abstract

Machine learning (ML) models of accelerator systems ('surrogate models') are able to provide fast, accurate predictions of accelerator physics phenomena. However, approaches to date typically do not include measured input diagnostics, such as the initial beam distributions, which are critical for accurately representing the beam evolution through the system. In addition, these inputs often vary over time, and models that can account for these changing conditions are needed. As beam time for measurements is often limited, simulations are in some cases needed to provide sufficient training data. These typically represent the designed machine before construction; however, the behavior of the installed components may be quite different due to changes over time or static differences that were not modeled. Therefore, surrogate models that can leverage both simulation and measured data successfully are needed. We introduce an approach based on convolutional neural networks that uses the drive laser distribution and scalar settings as inputs for a photoinjector system model (here, the linac coherent light source II, LCLS-II, injector frontend). The model is able to predict scalar beam parameters and the transverse beam distribution downstream, taking into account the impact of time-varying non-uniformities in the initial transverse laser distribution. We also introduce and evaluate a transfer learning procedure for adapting the surrogate model from the simulation domain to the measurement domain, to account for differences between the two. Applying this approach to our test case results in a model that can predict test sample outputs within a mean absolute percent error of 9%. This is a substantial improvement over the model trained only on simulations, which has an error of 261% when applied to measured data. While we focus on the LCLS-II Injector frontend, these approaches for improving ML-based online modeling of injector systems could be easily adapted to other accelerator facilities.

1. Introduction

Physics simulations of particle accelerators are essential tools for predicting optimal settings for different running configurations (e.g. changing the bunch charge, bunch length). Injector systems are particularly difficult to model accurately *a priori* because of nonlinear forces such as space charge at low beam energies. These simulations can also be computationally expensive, which can be prohibitive during the design stage as well as for online use in accelerators. Generating a comprehensive set of simulations of beam parameters resulting from different input settings can take several hours to complete. This is because a single simulation to compute important bulk beam parameters such as the beam size, can require minutes to hours to complete. This time-scale is too long for interactive online use in the accelerator control room. This also makes it difficult to conduct systematic comparisons with measured data and account for deviation between the idealized simulation and the as-built accelerator. In addition, obtaining machine time to characterize accelerator components can be rare, especially at large facilities with high demands on beam time.

Thus, there is a general need for fast and reliable models which can be used for online prediction, offline experiment planning, and design of new setups. Fast models would also enable more thorough investigation of differences between physics simulations and the real machine. There is also significant effort [1–8] toward using model-based control methods in real-time machine operation and tuning, with the goal to achieve faster, higher-quality tuning. Machine learning (ML) methods may help to automate tasks such as switching between standard operating schemes, or correcting small deviations that result in poor beam quality. Fast-executing, accurate machine models can aid the development and deployment of these control methods.

ML-based surrogate models are one avenue toward developing fast, reliable, and realistic models of accelerators. For injector systems, data generation and model training requires significant computational resources, but once trained, ML models offer orders of magnitude faster execution speed over classical simulation methods. Amongst the many ML algorithms available, neural network (NN) based surrogate models are being widely applied for addressing the issue of execution speed and obtaining fast, non-invasive predictions of beam parameters. Several studies have verified that ML-based models can be used to support fast optimization, particularly when trained using data that spans the operational range of the physical inputs [1, 4, 9–13].

While surrogate models trained on simulation are fast enough for use in online operation, the issue of how accurate these models are with respect to the real accelerator system also needs to be addressed. In many cases, large discrepancies are present due to simplifications assumed in the simulations or calibration differences. In addition, the physics simulation is often a static representation of the designed machine, and not an evolving representation of the physical machine. For example, in many photocathode injector systems, changes over time in the drive laser profile have a significant impact on the beam behavior. Typically, simulations are conducted with ideal initial beam distributions, or only a few example distributions from measurements. This can be a significant source of discrepancy between model predictions and measured data [21]. Overall, simulations tend to represent ideal conditions, and therefore ideal beam dynamics. ML models trained on simulation data thus reproduce these discrepancies between the ideal and as-built machine behavior.

One way to work around this issue is to train surrogate models on measured data, but in many cases there is not enough data to do so, especially as the number of input settings or output measurements being sampled increases. Typical operation of an accelerator often leaves large gaps in the parameter space, and due to limited time available for dedicated machine characterization studies, it may be challenging to collect sufficient training data to produce a surrogate model that is reliable across a broad range of inputs. Similarly, injector surrogate models to date do not include the full transverse laser distribution measurements as inputs, resulting in a loss of important time-varying information for accurately predicting the beam behavior.

The new linac coherent light source (LCLS) superconducting linac at the Stanford Linear Accelerator Center (SLAC), or LCLS-II, is one such facility that could benefit from having fast, accurate models for use in experiment planning, online prediction of beam distributions, and model-based control. Of particular importance is the injector, which sets key initial characteristics of the electron bunch, such as the overall emittance. As is the case for many injector systems, the LCLS-II injector is outfitted with a virtual cathode camera (VCC) that measures the transverse laser distribution. Changes to and non-uniformities in this laser distribution significantly impact the beam behavior.

Here, we introduce a multi-faceted, ML-based approach to address these issues by accounting for variation in the VCC image and conducting domain transfer between the simulation and measured data. First, we train a NN model based on Astra [14] simulations of our test case, the LCLS-II injector, over a wide range of the input settings. These inputs include solenoid settings and buncher phases, which are regularly scanned during tuning. Other input parameters such as the bunch charge is set for a given experiment, but can range significantly (from 1 to 100 pC). We also include changes to the initial distribution of the photocathode drive laser, as represented by measured and simulated VCC images. We demonstrate that including the VCC image as an input to the model improves the accuracy of the surrogate model with respect to the real machine, and we show that it can accurately predict beam output for unseen (out-of-distribution) VCC images. This is essential for using the model on a real accelerator, where in many cases VCC images are likely to change from day-to-day. In this case, we combine scalar setting inputs for the injector with a convolutional neural network (CNN) to do the image processing. A similar approach was taken in simulation in [9], and here we take the next step of including measured VCC images as inputs. Finally, we show we are able to compensate for the difference between the injector simulations and measurements by using transfer learning (TL) [15, 16], resulting in a surrogate model that is more representative of the real machine and can interpolate between VCC images more accurately than a model trained only on measured data.

As part of this process, we also carried out a detailed study of the sensitivity of the simulated output to changes in the initial beam distribution, as seen on VCC images, to determine whether using the images

directly would confer an advantage over scalar fits to the beam distribution. We also conducted a characterization study of the LCLS-II injector to compare measurements to simulations.

The main contribution from this work are as follows: (1) a characterization of the LCLS-II injector frontend, (2) introduction of an approach using measured VCC (laser distribution) images and a CNN to improve the accuracy of injector surrogate models, (3) demonstration of a TL method to pre-train the injector surrogate model in simulation and fine-tune on machine measurements. While the demonstration is specific to the LCLS-II injector, the approach can be used for other injector systems, and we have made our code publicly available to help facilitate this.

2. Characterization studies of the LCLS-II injector

In setting up a realistic surrogate model for the LCLS-II injector, it was important to assess how simulation results vary when using a realistic laser profile, as compared to an ideal super Gaussian (SG) or uniform profile (as is assumed in most start-to-end injector optimizations) [17], in which the intensity of the pulse is uniform within the circle of radius r , representing the transverse profile of the beam. The SG distribution, ρ , is given as function of radius r of the form:

$$\rho(r) = \frac{1}{2\pi\Gamma\left(1 + \frac{1}{p}\right)\sigma^2} \exp\left[-\left(\frac{r^2}{2\sigma^2}\right)^p\right] \quad (1)$$

where σ is the standard deviation, Γ denotes the Gamma-function, and p is the SG parameter. In the limit $p \rightarrow 1$, the SG is a standard Gaussian distribution. However, as $p \rightarrow \infty$, the SG distribution approaches a flat-top distribution. By parameterizing p by $p = 1/\alpha$, the α parameter is bounded by $[0, 1]$. This is a common choice when simulating an ideal laser distribution.

This determines whether it is necessary to include the full VCC image as an input to the model, or whether bulk metrics such as laser radius and an assumption of a Gaussian or uniform profile would be sufficient. We also characterized the LCLS-II injector with measured data scans and compared these to simulation data. This was done both to help improve the underlying physics simulation and assess both the need for and viability of using TL for this system to account for differences between the simulation and measurement domains.

2.1. The LCLS-II injector

The LCLS-II injector will produce the electron beam for the LCLS-II superconducting linac that is presently being commissioned at SLAC.

The injector, shown in figure 1, will be used for the LCLS-II project. The expected operating parameters for the injector are shown in table 1. At the time of experimentation, the injector was in early commissioning.

With such a high repetition rate, the cathode field gradient is lower than low-repetition rate photocathode injectors [18]. Thus, the kinetic energy of the electrons as they are emitted and injected into the buncher is relatively low; up to 750 keV. At this energy, the dynamics are space-charged dominated. To study the dynamics in this regime and optimize the parameters for operation, particle-in-cell simulations are necessary. As such, these calculations can take several minutes to complete.

There are several options for simulation tools, especially for accelerator injector simulations, which rely on sophisticated space charge calculations. For this study, all simulated data was generated using Astra [14], and particle generation was done using distgen [19]. The SLAC-developed Python wrapper LUME-Astra [20] was used to create, set-up, and process simulated data.

Measurements of laser input distributions and associated solenoid scans (where the electron beam size in one transverse direction was measured while the solenoid value was changed) were recorded for use in surrogate model training. These scans were taken at several different beam charges, ranging from about 1 to 25 pC. Machine values such as magnet currents and beam charge were also recorded.

2.2. Comparison between simulation and measured data

Here we show a comparison between measured data and the associated simulated values for the LCLS-II injector. In figure 2, the beam sizes measured during two solenoid scans on the injector are compared to the predictions from Astra. In order to be able to appropriately compare measurements to simulations, all beam sizes are calculated by fitting the particle distribution to a radial Gaussian distribution, and reporting the beam size at the standard deviation of the distribution. The initial particle distribution for these comparison scans were generated by sampling SG distributions in the transverse dimension, for 10 000 particles. The laser diameter was archived during the measurement and used for simulation. With first order matching of inputs (charge, radius, gradient, solenoid strength), there are clear discrepancies between simulation and

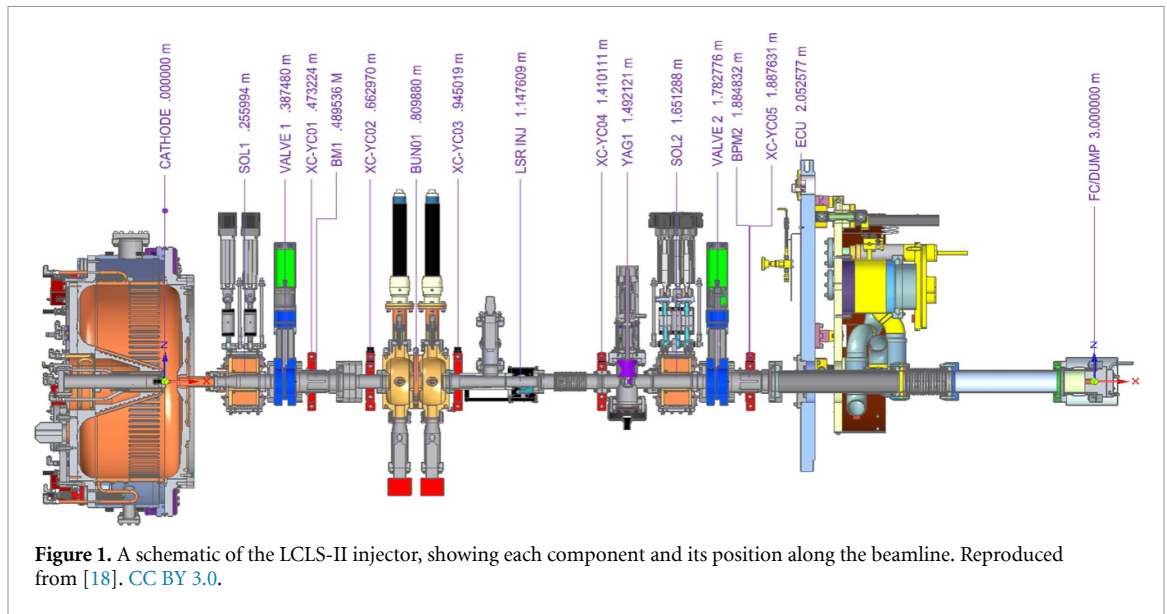


Table 1. Operating parameters for the LCLS-II injector, expected after the completion of the construction and commissioning of the machine. During this study, the injector was still in early commissioning, and had limited operational ability.

Parameter	Value	Unit
Charges	100	pC
Laser FWHM	20	ps
Laser radius	1	mm
Field on cathode	20	MV m ⁻¹
Repetition rate	1	MHz

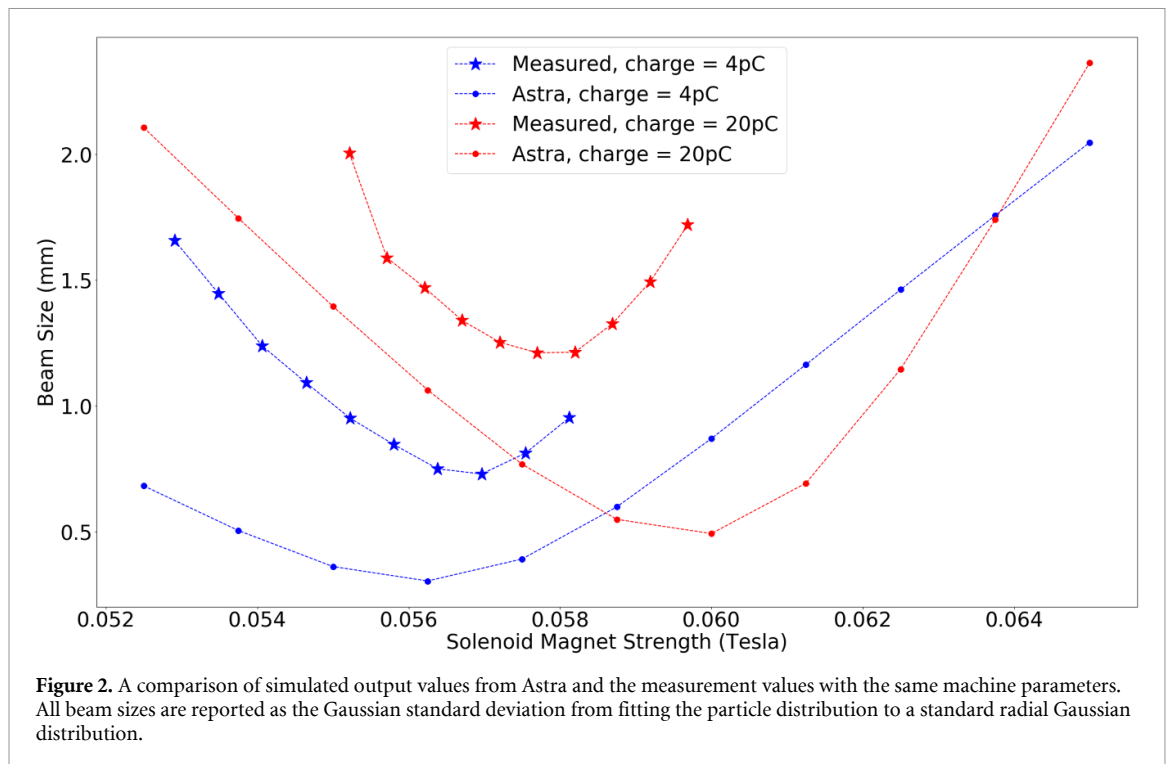
measurement, as shown in figure 2. After investigating these errors through simulation scans, it is clear that the gun gradient and laser radius can have a large effect on the focal point of the solenoid scan. While the gradient on the cathode has a significant impact on beam size and the location of the beam waist, there was not a spectrometer located in the beam line when this data was taken. Therefore, the exact beam energy at the time of measurement is not definitively known. Indirect measurements were attempted with a small corrector, but the results were inconclusive given they returned unphysical energy values (i.e. higher energies than possible given the amount of RF power supplied to the gun). This leaves the exact gun energy to be determined, and a probable cause of discrepancy in the measurement.

2.3. Simulation-based sensitivity studies

2.3.1. Generation of electron distributions from laser distributions

Using the measured laser distribution to sample particles for space charge simulation can minimize discrepancies between measurement and simulated output values [21]. Therefore, in order to attempt to create more realistic simulated data, real laser profiles were used to generate particle distributions to be tracked in LUME-Astra. These laser profiles were collected at the LCLS-II. The measurement is conducted as follows: the laser beam is passed through an optical splitter, such that approximately 5% of the beam intensity is directed toward a camera. The distance between the splitter and the camera is analogous to the distance between the splitter and the cathode, i.e. the transverse size of the laser at the camera location should be equal to the size at the cathode. The intensity at the camera is recorded, providing an image of how the laser intensity at the cathode appears.

We compared how the the beam sizes would differ when the initial particles were sampled from a realistic laser distribution and an ideal one. In order to make an idealized transverse SG distribution from measured VCC images, the following optimization procedure was completed. The measured laser profile was projected onto each transverse axis. A SG profile was then generated using an initial α value, and projected similarly into each transverse axis. An optimizer then iterated α via a Brent minimization algorithm [22], to minimize the residual between the projections. Many of the measured laser profiles are highly non-uniform or highly irregular edges. Therefore, results for which the per pixel percent error distribution standard deviation of less than 20% were selected. A nominal VCC and associated SG were chosen for the sensitivity analysis.



2.3.2. Sensitivity of LUME-Astra predictions to different laser distributions

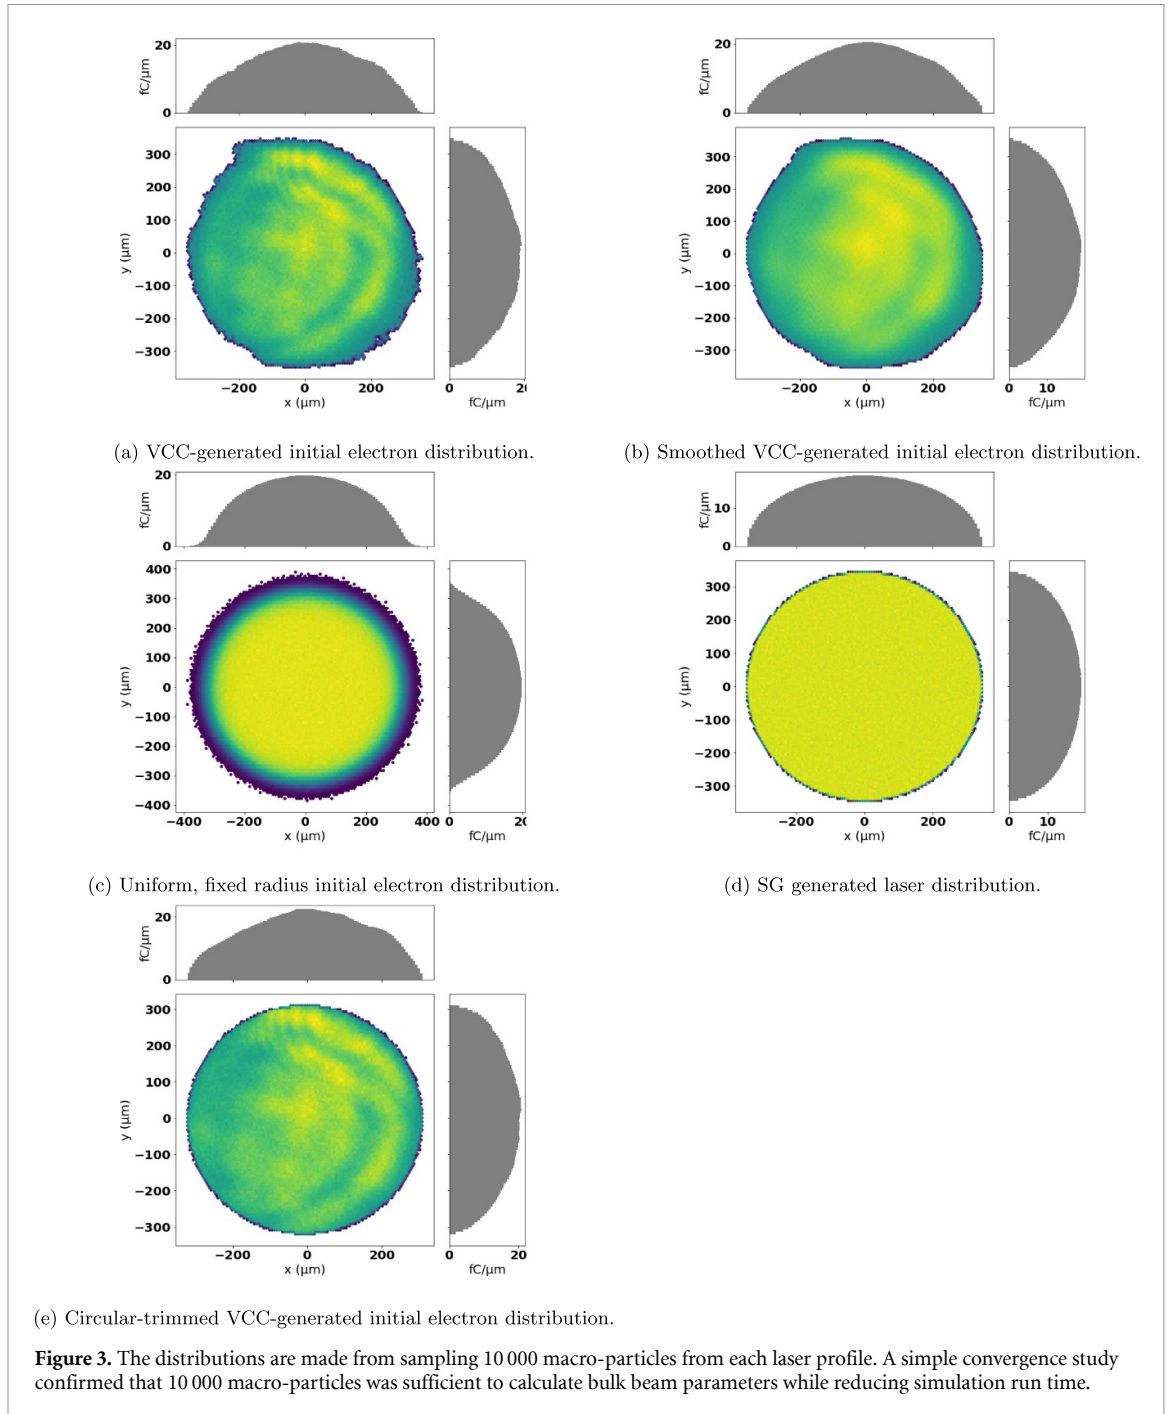
We assessed the sensitivity of the simulation results on a realistic vs. idealized laser profile in simulation. In this case, sensitivity is evaluated by whether the bulk properties of importance, normalized transverse emittance and transverse beam size, differ more than 10% from the values calculated from the uniform initial distribution. This threshold is close to the resolution of such measurements in the physical set up.

First, a candidate laser profile with features including rough edges, as well as fluctuations within the bulk of the laser spot, was chosen. The transverse profiles compared include: uniform radial distribution, SG distribution, and the candidate laser profile with a Gaussian blur applied. Shown in figure 3 are each laser profile, with nominal charge of 10 pC total with 10 000 sampled particles. The candidate VCC image is shown in figure 3(a), as well as the particles generated from a blurred version of the candidate VCC in figure 3(b). Having a similar, but slightly smoothed version of the candidate VCC will address whether the simulation is sensitive to the internal structure of the spot size, or just coarse features such as rough edges. This is further investigated by removing the edges of the candidate VCC images, and keeping internal features. This distribution and resulting particles are shown in figure 3(e).

Two highly uniform distributions (referring to the charge distribution in transverse space) were prepared for comparison as well. First, which is often used as the standard distribution for simulations, is a uniform density distribution specified, shown in figure 3(c). Next is the SG distribution, shown in figure 3(d). The time structure for all of the particle distributions generated for this sensitivity study, as well as for the surrogate model training, was a Gaussian with standard deviation of 8.5 ps. This time distribution was held constant for all simulations.

For each laser profile, a particle bunch with 10k particles was generated and tracked through the injector lattice in Astra to calculate various resulting beam outputs. It was determined that the bulk parameters in the simulation can be recovered with sufficient fidelity and speed using 10k particles. The primary quantities of interest in this study were the resulting normalized transverse emittances (95%, about two sigma, core emittance) and beam sizes. Astra simulations were completed for each laser profile at two charge settings (5 and 50 pC), with all other parameters, such as solenoid magnet gradient, held constant. The resulting transverse emittances and beam sizes at the yttrium-aluminum-garnet (YAG) screen, 1.49 m from the cathode, are shown in figure 4.

It is clear that the emittance and beam size from Astra simulations are sensitive to the realistic beam distributions, relative to a uniform beam distribution. For the SG distribution, which emulates an ideal flat-top beam distribution, the emittance in each direction is the same, however there is a difference seen in emittances calculated from a VCC generated laser profile. Clearly the asymmetry of the VCC generated laser distributions can be captured by the simulation, as shown by the difference in beam size and emittances in each transverse direction, for a given VCC laser profile. These results suggest that using realistic laser profiles



could result in simulated training data which is sufficiently different from that generated from idealized conditions.

3. Surrogate model to emulate Astra simulations

As stated, two major use cases for the LCLS-II surrogate model are: (1) to aid offline experiment planning and start-to-end optimization, and (2) to provide non-invasive predictions of the beam behavior, given measured upstream inputs. With those applications in mind, it is critical to evaluate the ability of the surrogate model to accurately emulate the behavior of the Astra simulations under optimization and in predicting the output beam distributions. For example, the model needs to be able to interpolate between different types of laser distributions (as seen on the VCC images) to be useful in online prediction under changing laser conditions. Here we evaluate the ability of a NN surrogate model to interpolate to regions of parameter space not seen during training and to reliably provide accurate predictions when used in optimization.

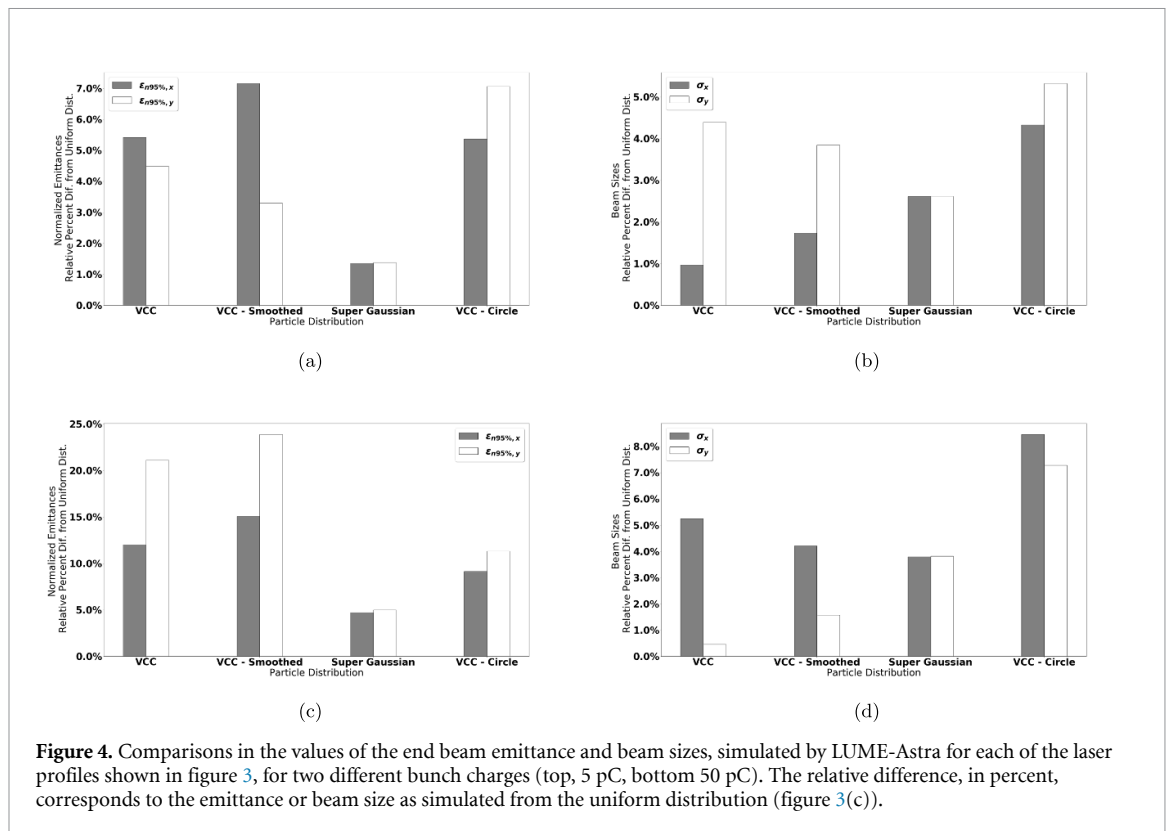


Figure 4. Comparisons in the values of the end beam emittance and beam sizes, simulated by LUME-Astra for each of the laser profiles shown in figure 3, for two different bunch charges (top, 5 pC, bottom 50 pC). The relative difference, in percent, corresponds to the emittance or beam size as simulated from the uniform distribution (figure 3(c)).

3.1. Data generation and general training procedures

Training data was generated in two ways. For scalar model development, the standard particle generator in Astra was used by supplying laser radii as inputs. A large random sample of the laser radius, cathode cavity RF phase, beam charge and solenoid strength was then simulated to create a data set which assumed standard incoming beam parameters. A sample in this data set consists of the scalar input values, and the associated bulk beam values such as emittances and beam sizes. This data was used for scalar model training.

Further training data was generated by running measured VCC laser distributions and idealized SG laser distributions through LUME-Astra while randomly sampling injector input settings. For each unique laser profile, 2000 randomly-sampled points in the input space were generated. The predicted output includes the electron beam distribution as it might be measured at a YAG screen, along with bulk statistical quantities such as normalized emittance and beam sizes. Two simulated data sets consisting of approximately 60 000 samples from SG particle distributions and 70 000 from VCC measurements were generated. Thus, a sample in this data set consisted of the scalar inputs including the dimensions of the input laser distribution and also the 50×50 binned input distribution, and the associated scalar output values and 50×50 binned output electron distribution.

All NN model development and training was done using the TensorFlow and Keras libraries [23]. Each model was trained by minimizing a mean squared error loss function, using the Adam optimization algorithm [24]. Several different NN architectures are used, as described in the following sections. During any training process, the training samples are used for fitting the model. The validation loss is calculated and monitored during training to avoid overfitting, but is not included directly in the weight updates. All testing samples are withheld from the training process entirely.

3.2. Scalar model performance in interpolation and multi-objective optimization

In this section we demonstrate a model that uses only the laser radius of a uniform laser distribution as an input rather than a laser distribution images. Typically a laser radius with a static profile is used for most multi-objective optimization studies on injectors (including the LCLS-II injector). We predict a wide variety of output scalars that are relevant for optimization studies. Figure 5 shows the basic inputs and outputs. The NN architecture itself consisted of eight layers (six hidden layers), each using a hyperbolic tangent activation function. The hidden layers each had 20 nodes; the input layer had four nodes corresponding with each input. The model output 16 scalar predictions.

The performance on the scalar predictions is similar to that of the CNN case. A selection of scalar parameters, the beam emittances $\epsilon_{x,y}$, beam sizes $\sigma_{x,y}$, and bunch charge Q are described later and shown in

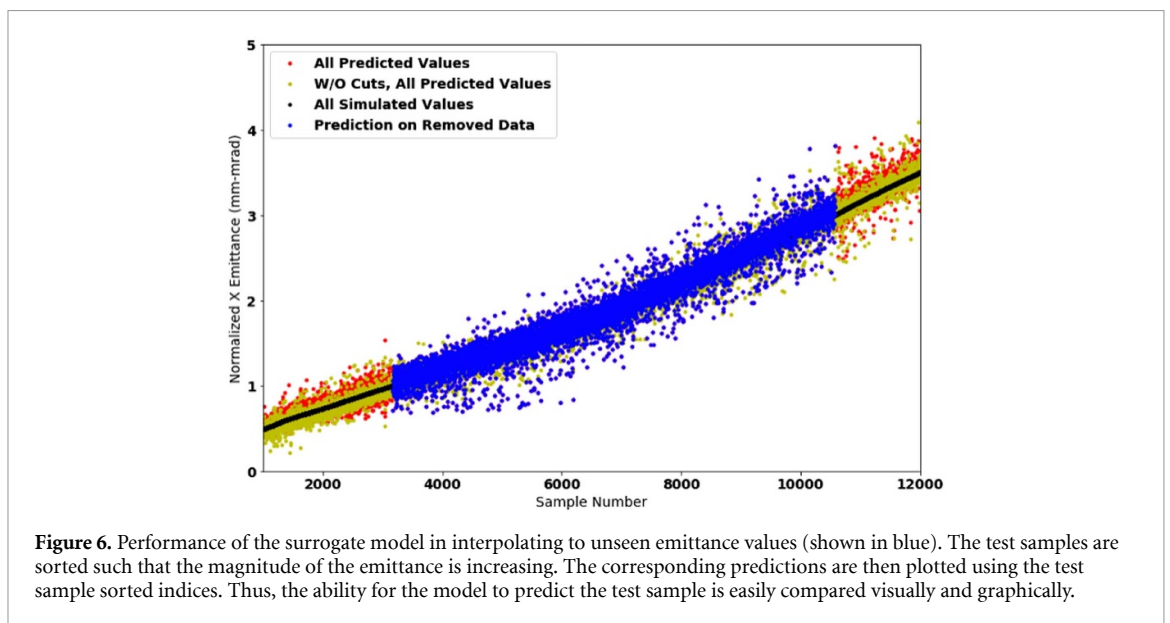
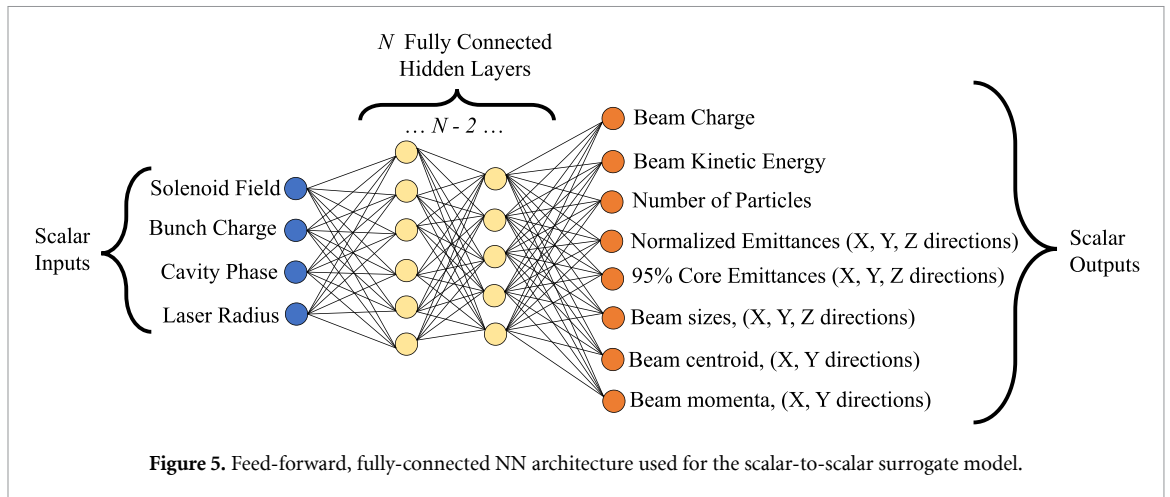
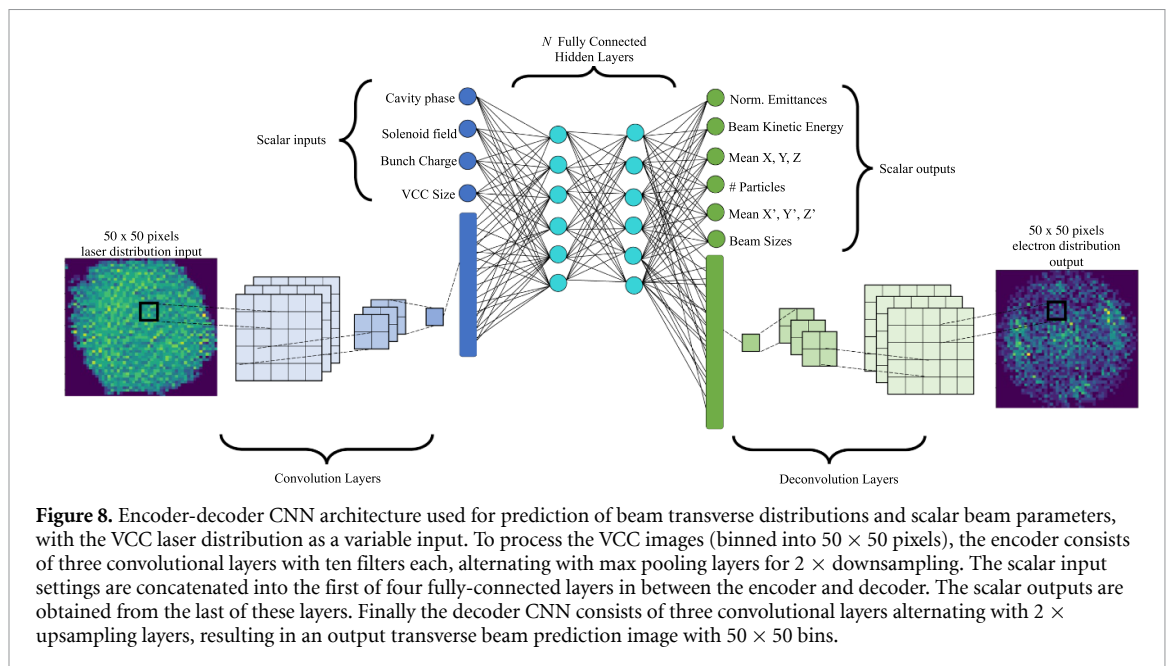
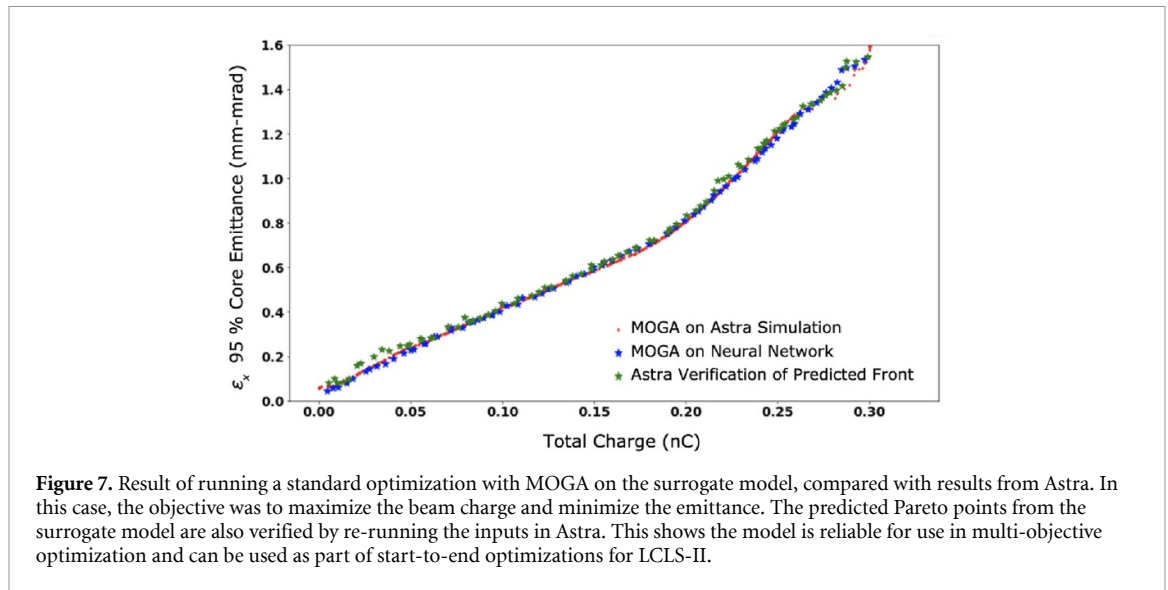


figure 11. To ensure the model can be used for optimization studies, we first left out sections of the parameter space from both the training and the validation set to verify the model can interpolate accurately (see figure 6). Next, we also verified that optimization with a standard multi-objective genetic algorithm (MOGA) on the model produces an accurate Pareto front (see figure 7). We use a similar MOGA setup to that described in [4]. For the verification, we run the input settings from the predicted front in Astra. Based on these results, this scalar version of the injector surrogate model can already be used as a component in start-to-end optimization of new setups for LCLS-II (e.g. by replacing the simulation of the gun).

3.3. Performance for beam distribution predictions and interpolation to new (out-of-distribution) VCC images

For prediction of the transverse beam distributions and scalar outputs and taking into account the VCC image, we introduced an encoder-decoder style CNN architecture (shown in figure 8); this approach has not been taken before for injector surrogate modeling. Each output transverse distribution is binned into a 50×50 image. Sixteen scalar beam parameter outputs and the scalar extents of the beam distributions are predicted as well. For each of the 62 unique VCC images, a random sample of the input settings was conducted in Astra. Thus, for each VCC image, several thousand Astra simulations with unique scalar inputs were completed. The model was then trained on 60 320 samples, with 7540 samples held out for training and testing.

To assess the ability of the model to interpolate between different laser distributions (so that it can provide accurate predictions on new VCC images as the laser distribution shifts over time), we selected a set of VCC images that had patches of intensity within the bulk of the laser spot missing, and held this data out of the training and validation set. We find good agreement between simulation results and the surrogate



model predictions, even for cases with irregular beam distributions (see figures 9 and 10). This indicates that the model can be used online with the running accelerator to provide non-invasive estimates of the transverse beam profile (i.e. as both an online model of the injector and a virtual diagnostic), similar to how an online physics simulator could, but with much faster-to-execute predictions. The performance of the model on bulk scalar predictions is shown in figure 11.

4. Transfer learning

The previous sections demonstrated the ability of the surrogate model to reliably emulate predictions from Astra simulations. However, the issue of how these predictions compare to measured beam parameters remain. Because we have very little measured data, we generated an initial model trained on simulation data and then modify it to be consistent with measured data afterward. Here, we develop and demonstrate a TL procedure to accomplish this.

TL encompasses a broad class of ML approaches wherein the performance of a model at a particular task or domain may be improved by transferring information from another related but different task or domain [25]. In traditional approaches to ML, the distribution over feature space and the distribution in target space must be identical during training and deployment. If any such differences, termed distribution shifts, exist, the performance of the trained model is severely degraded [26]. TL is thus one approach to handle

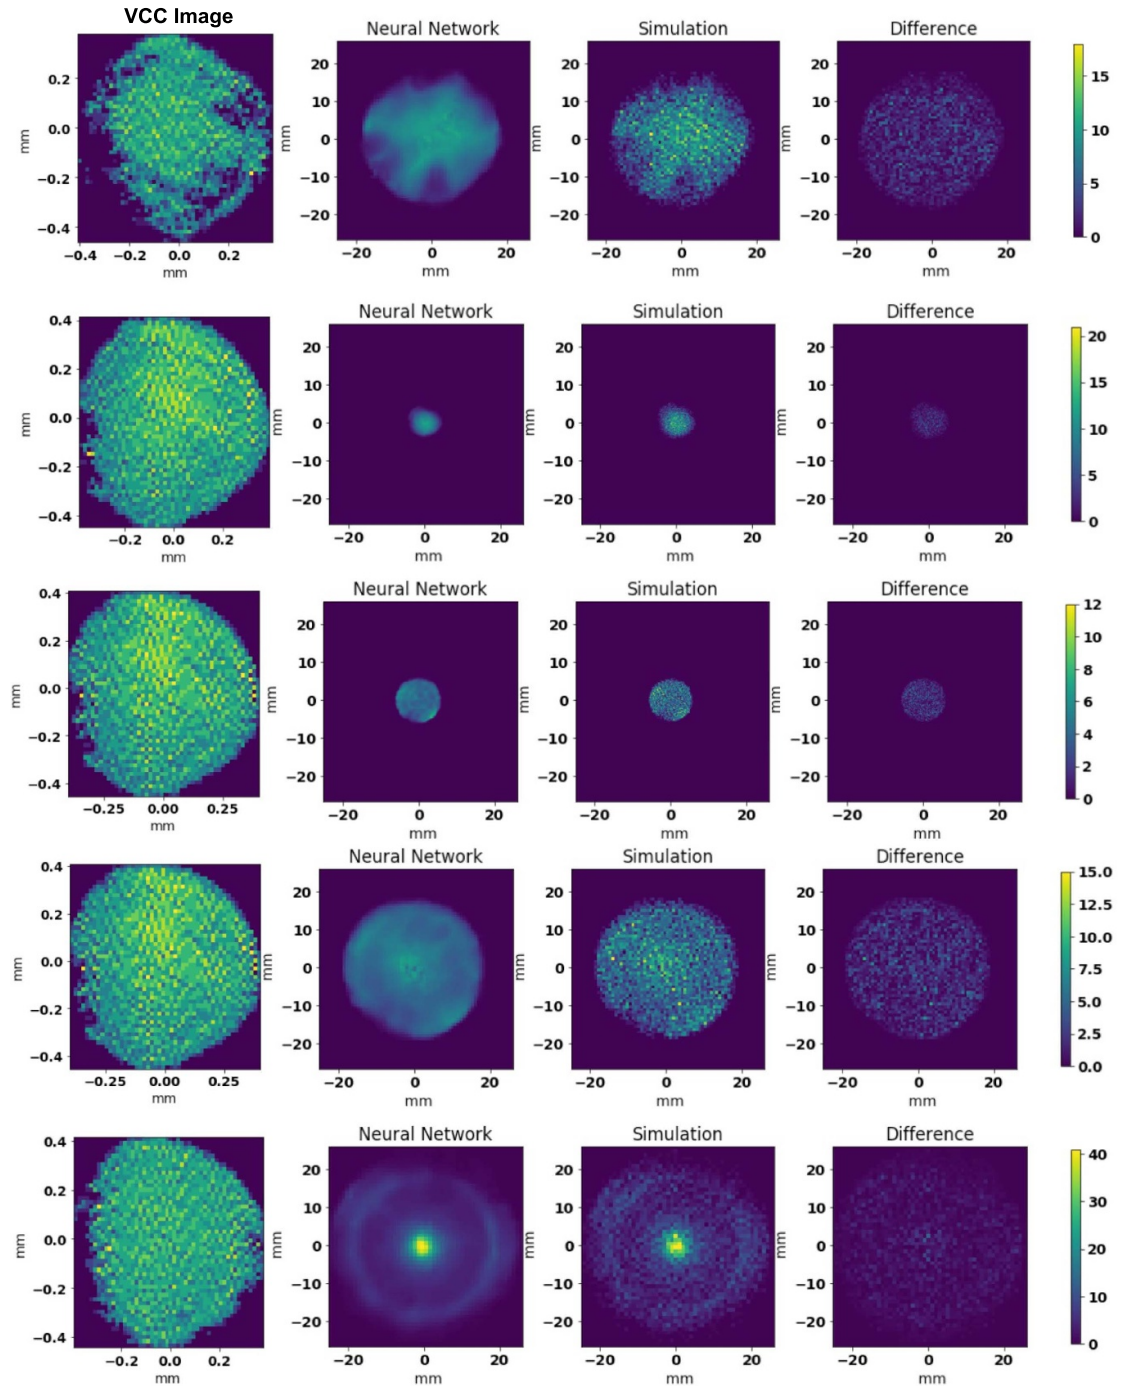
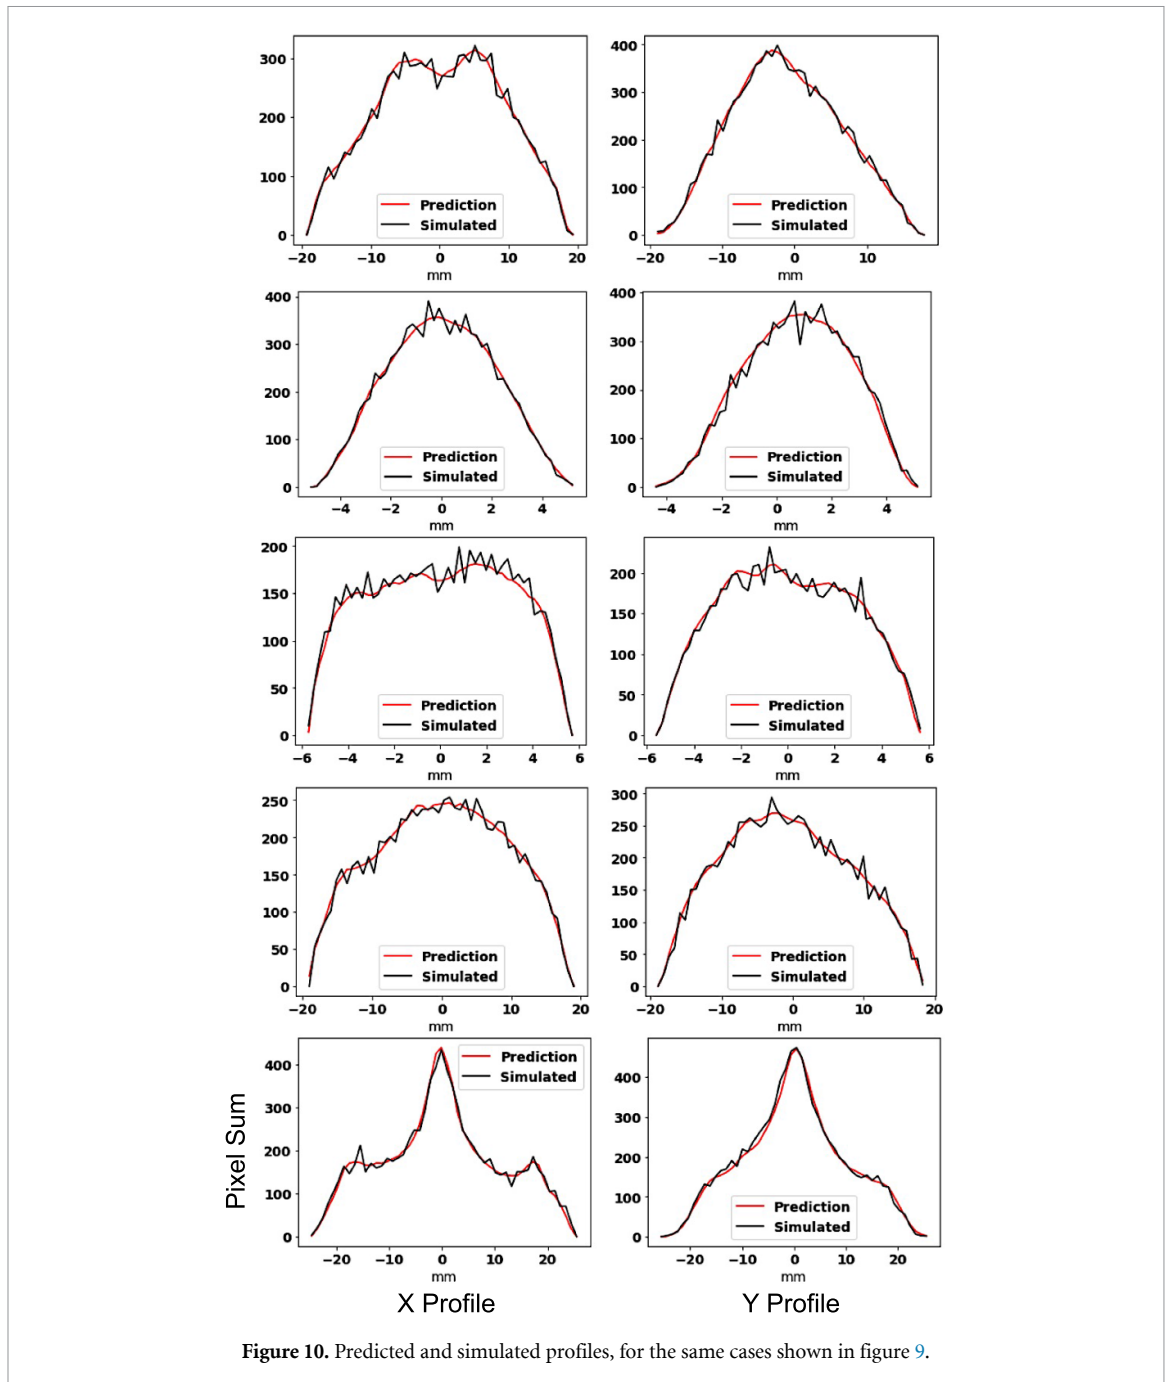


Figure 9. Examples of the NN predictions and Astra simulation results for the transverse beam distributions. The corresponding VCC inputs used in each case are shown at left. The agreement is good, even for cases with irregular beam distributions. This demonstrates that the model can interpolate between measured input laser distributions (as seen on the VCC) and provide realistic predictions of the expected transverse beam distribution from simulation. This is important for using this model online in the accelerator, as the initial beam distribution will vary with time.

distribution shifts between target domains (e.g. simulation to measurements, idealized laser beam shapes to non-idealized ones), and it has been successfully applied to diverse applications including image classification [27], anomaly detection [28], text sentiment analysis [29], etc.

To find a suitable TL approach for this class of accelerator surrogate model, we first prototyped the approach on simulation data. We started with a primary model trained on SG beam distributions and then expand the model training to include the simulation results obtained from measured VCC images. We then applied the same procedure to the measured data.

A base model was trained on simulation data set until the mean-squared error (MSE) loss did not decrease over several training epochs. Then all of the model weights except the weights and bias values for the

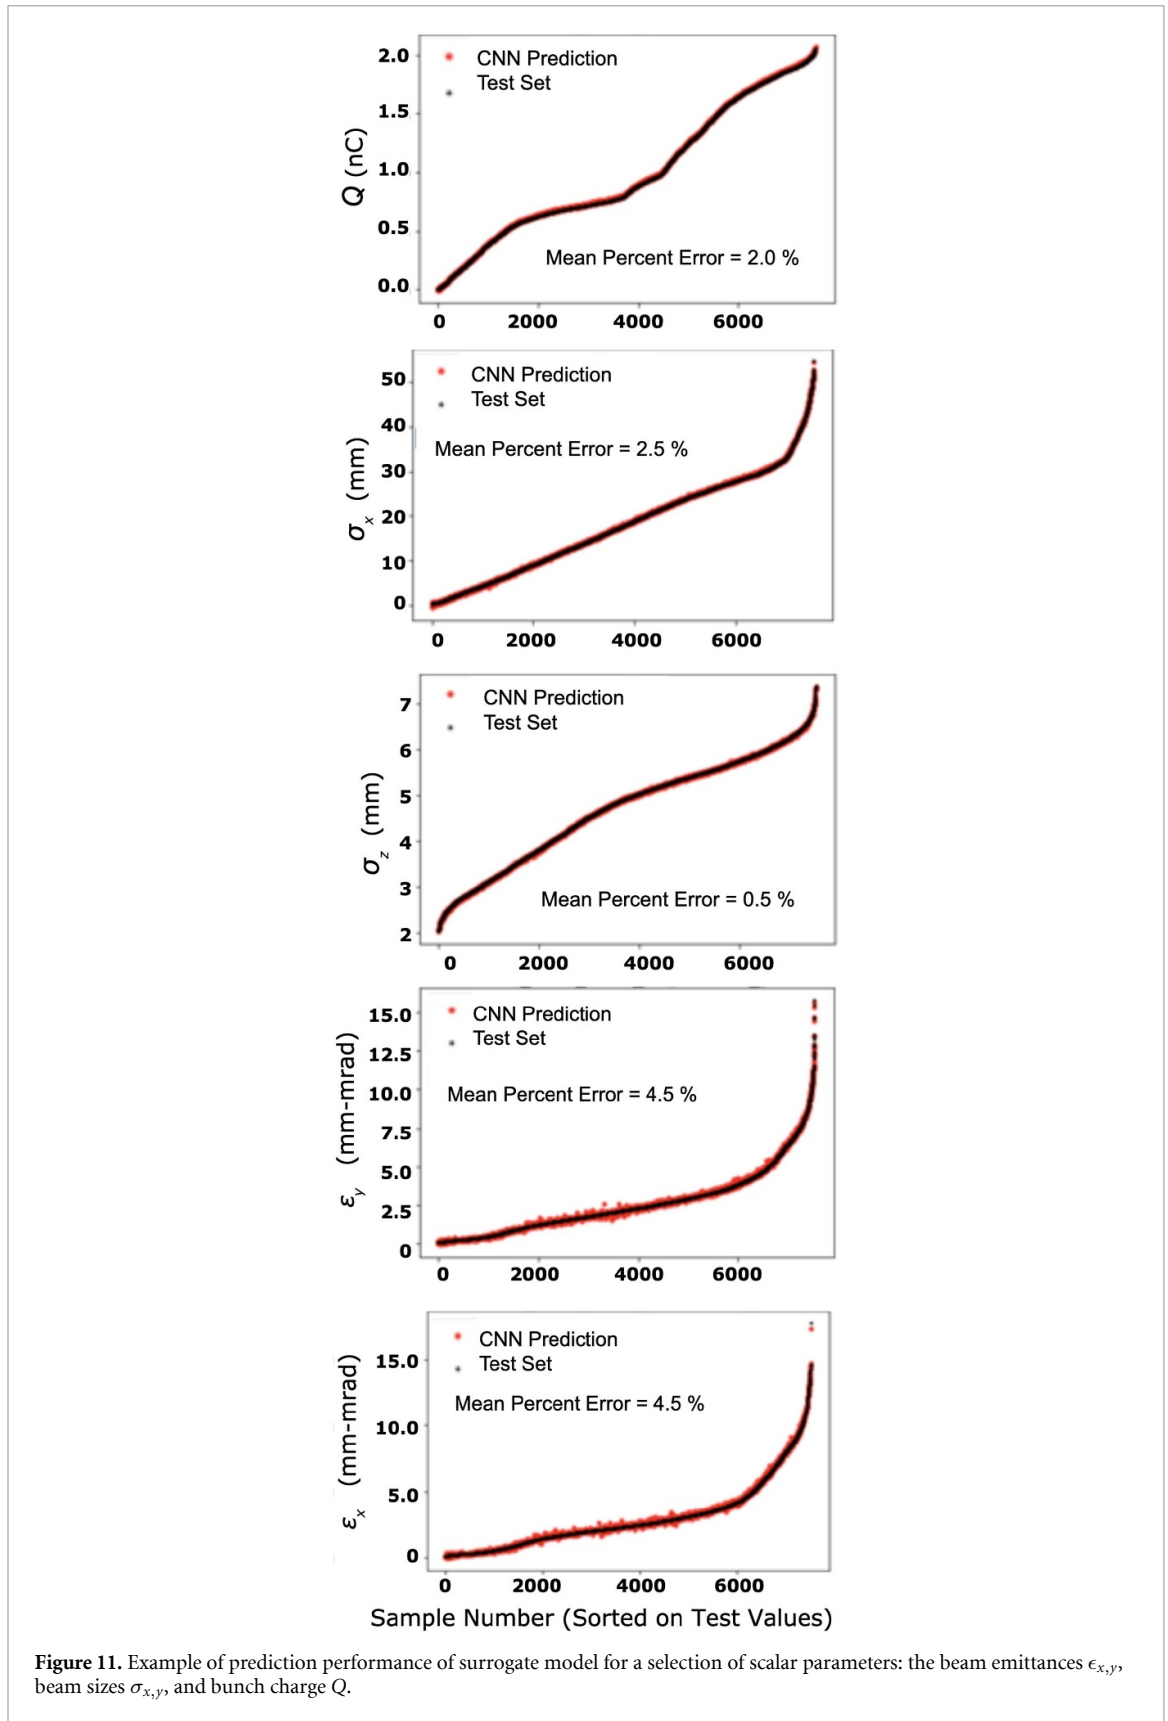


last two layers were frozen. A dropout layer [30] was also included between the frozen and unfrozen layers to reduce over-fitting. The initial learning rate was then set at 5×10^{-4} , and decreased on an exponential schedule. The training was terminated using an early stopping method. After this, the learning rate was reduced by two orders of magnitude, and all of the weights and bias values in the model were trained, referred to as fine-tuning. Models after the TL procedure are referred to as TL models.

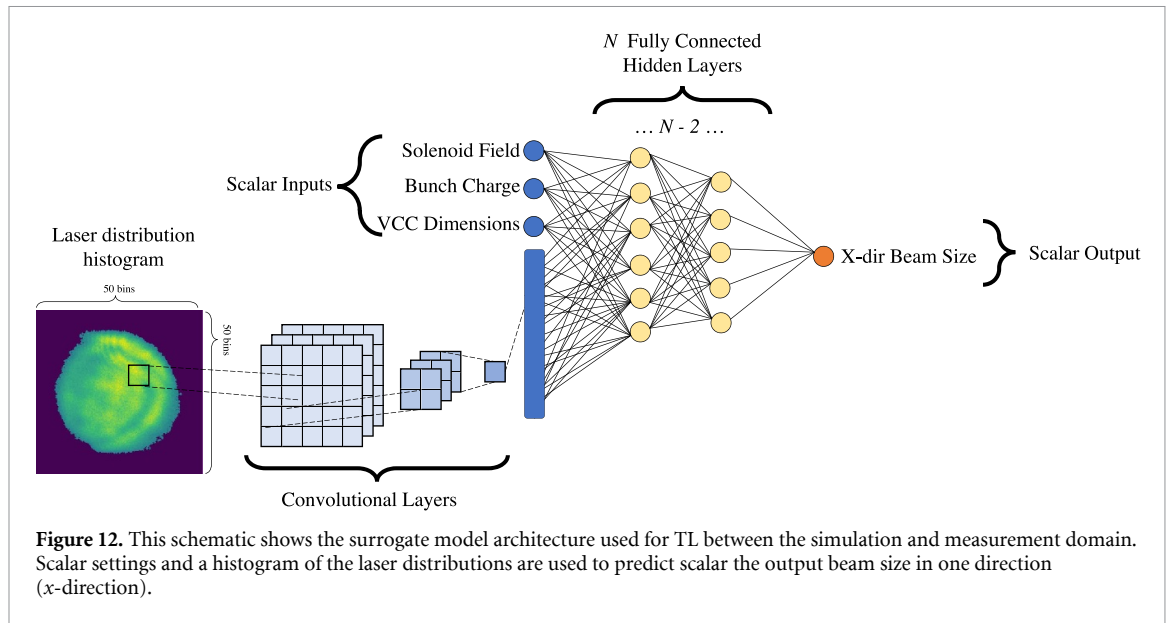
4.1. Data sampling and NN architecture

As mentioned, many simulated data samples were generated, but a sparse sub-sample was used for surrogate model training. This can ensure the model is able to interpolate well, and minimize over-fitting. Sub-sampling was also used to emulate small amounts of data for re-training (as one would typically have for measured data sets on an accelerator).

The model architecture described, and shown in figure 12, depicts the general architecture of the NN surrogate models. All models took scalar settings for the solenoid value and charge as inputs, along with the two-dimensional histogram representation of the laser intensity on the cathode. The laser distributions were 50×50 bins. The size of the laser distribution were given as the horizontal and vertical extents of the



histogram, relative to the center of the histogram. These six scalar values and the 50×50 bin laser distribution are considered inputs to the models. The binned images were input into convolutional layers. Three convolutional layers with ten 4×4 filters each are applied to the image inputs. The resulting nodes are then fed to densely-connected layers. The densely-connected part of the network consists of six hidden layers with 1024, 512, 256, 64, 32, 16, 6 neurons respectively. The output layer consists of one node predicting the transverse beam size in x.



As before, all NN model development and training was done using the TensorFlow and Keras libraries [23]. Each model was trained by minimizing a MSE loss function, using the Adam optimization algorithm [24]. To evaluate the accuracy of all models, the mean absolute percent error (MAPE) was calculated as shown:

$$\text{MAPE} = \text{mean} \left(\frac{|y_{\text{true}} - y_{\text{pred}}|}{y_{\text{true}}} \right). \quad (2)$$

4.2. TL in simulation: idealized laser distributions to measured laser distributions

We prototyped the TL approach by training a model on SG laser distributions, then retraining the NN model to predict VCC-based simulation results. Since this is a major potential source of disagreement between idealized simulations and the as-built injector, it enabled us to refine the approach prior to applying it to the measured data.

The SG generated data was used as the primary training data set, shown in figure 13. The data set was down-sampled to 700 training, 150 validation, and 151 test samples, which provided sufficiently sparse coverage to ensure we were not oversampling the parameter space. The resulting predictions are shown in figure 14.

The final MAPE on the test values was 5%. Then, the VCC-based data set was down-sampled randomly to represent the small amount of measured data available. The down-sampled data is shown in figure 15.

In this case, after training the base model, we combine the training data sets such that the NN is trained on both simulation data sets simultaneously (but, as described earlier, with limited model adaption enabled). Because there is five times more training data for the base model, the smaller data set was repeated five times in order to create proportionally equal representation in the data set (a standard practice when dealing with imbalanced data sets). The performance on the combined test set (the test samples from the SG data set, and the VCC-generated data) is shown in figure 16. In this case, we see that the model trained only on VCC data cannot predict the combined distribution as well as the model which underwent the TL procedure. The model which underwent TL is able to predict test samples from the combined data set with a MAPE of 12%, compared to a model trained on a single data set which had a MAPE of 113%.

4.3. TL to measured data

In order to assess model performance when interpolating to new combinations of input settings we evaluated TL from simulation in a case where the full operational range is present, to measurements where the model would need to interpolate to new setting ranges. This could be an effective method for producing a surrogate model, particularly when only limited measured data is available.

This scenario (missing ranges of parameter space) was emulated by withholding measurement values with charge between 20 and 22 pC, with the data distribution shown in figure 17. The previously prototyped procedure was attempted, but we found we needed to adjust the TL procedure to accommodate the large

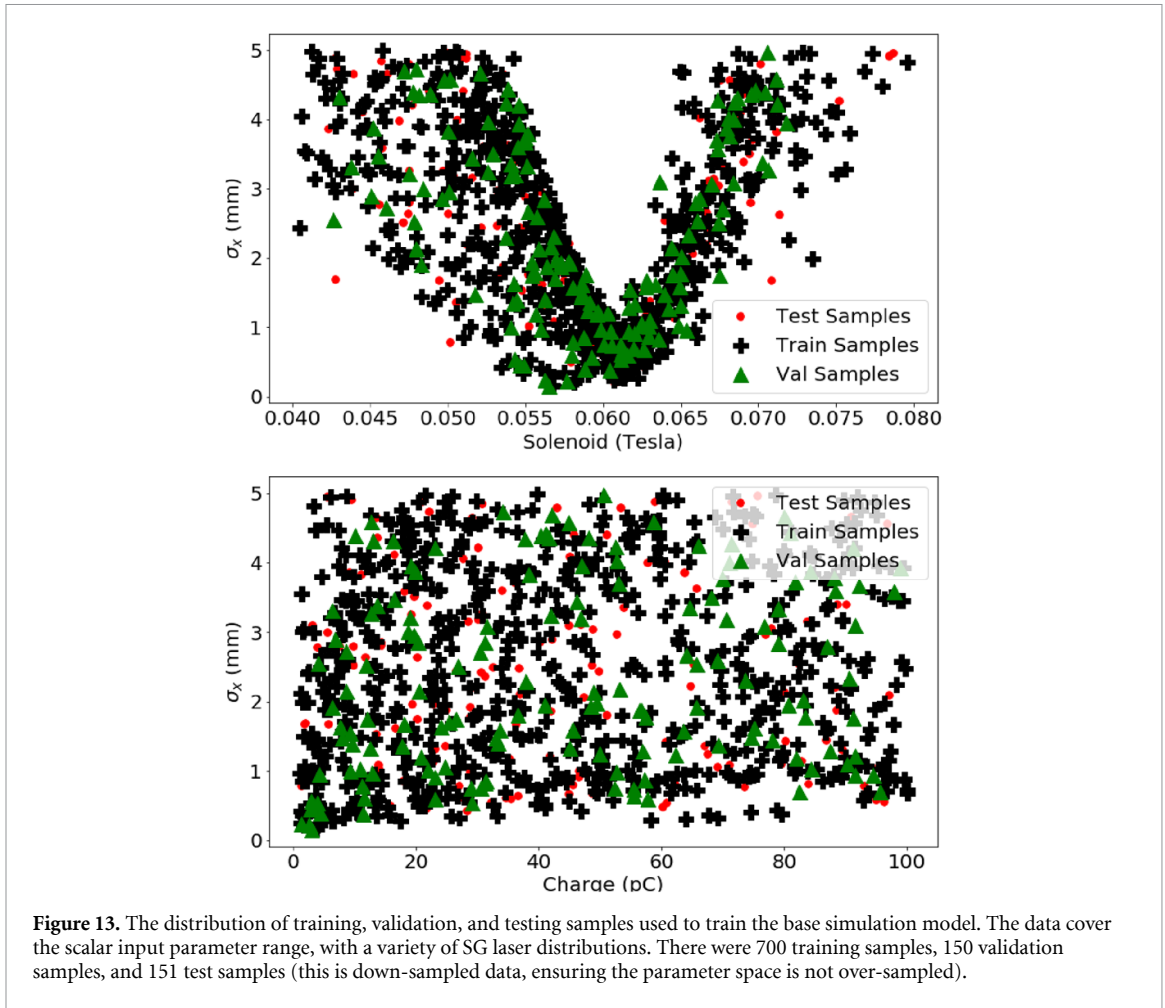


Figure 13. The distribution of training, validation, and testing samples used to train the base simulation model. The data cover the scalar input parameter range, with a variety of SG laser distributions. There were 700 training samples, 150 validation samples, and 151 test samples (this is down-sampled data, ensuring the parameter space is not over-sampled).

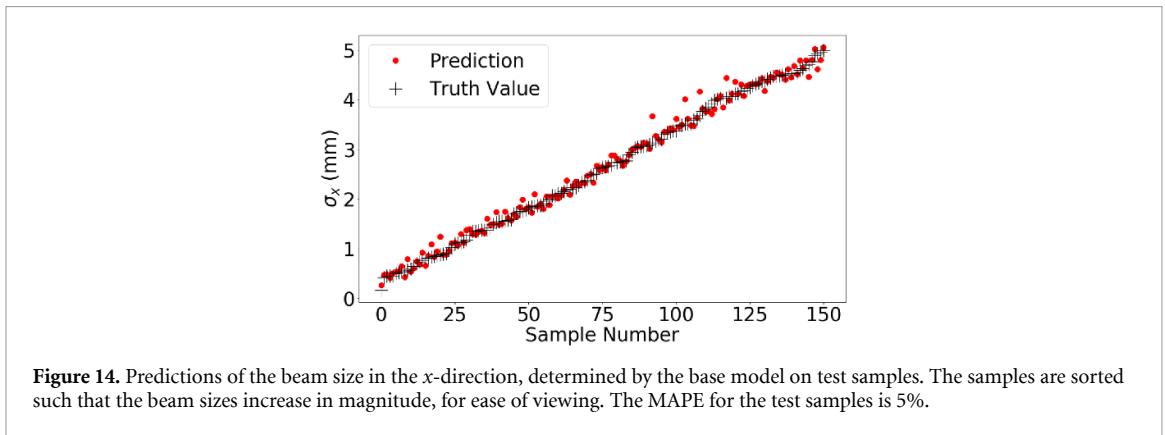


Figure 14. Predictions of the beam size in the x -direction, determined by the base model on test samples. The samples are sorted such that the beam sizes increase in magnitude, for ease of viewing. The MAPE for the test samples is 5%.

systematic differences between the simulated and measure data. For the case with TL to measured data, the main difference is the large systematic differences in the scalar output parameters, rather than in the types of input laser distributions seen during training. Thus, allowing more of the fully-connected layers to adapt to the new data was warranted.

The procedure was modified in the following ways. The base model (trained only on simulation data) for this procedure is the same as that produced previously during the simulation prototype. In this new case, the fully-connected portion (i.e. excluding the CNN layers of the NN) were allowed to train with a reduced learning rate starting at 5×10^{-5} and decreasing every ten epochs for 2000 epochs. The final learning rate is then used while fine-tuning the model (with all layers trained) for another 2000 epochs. The results are shown in figure 18.

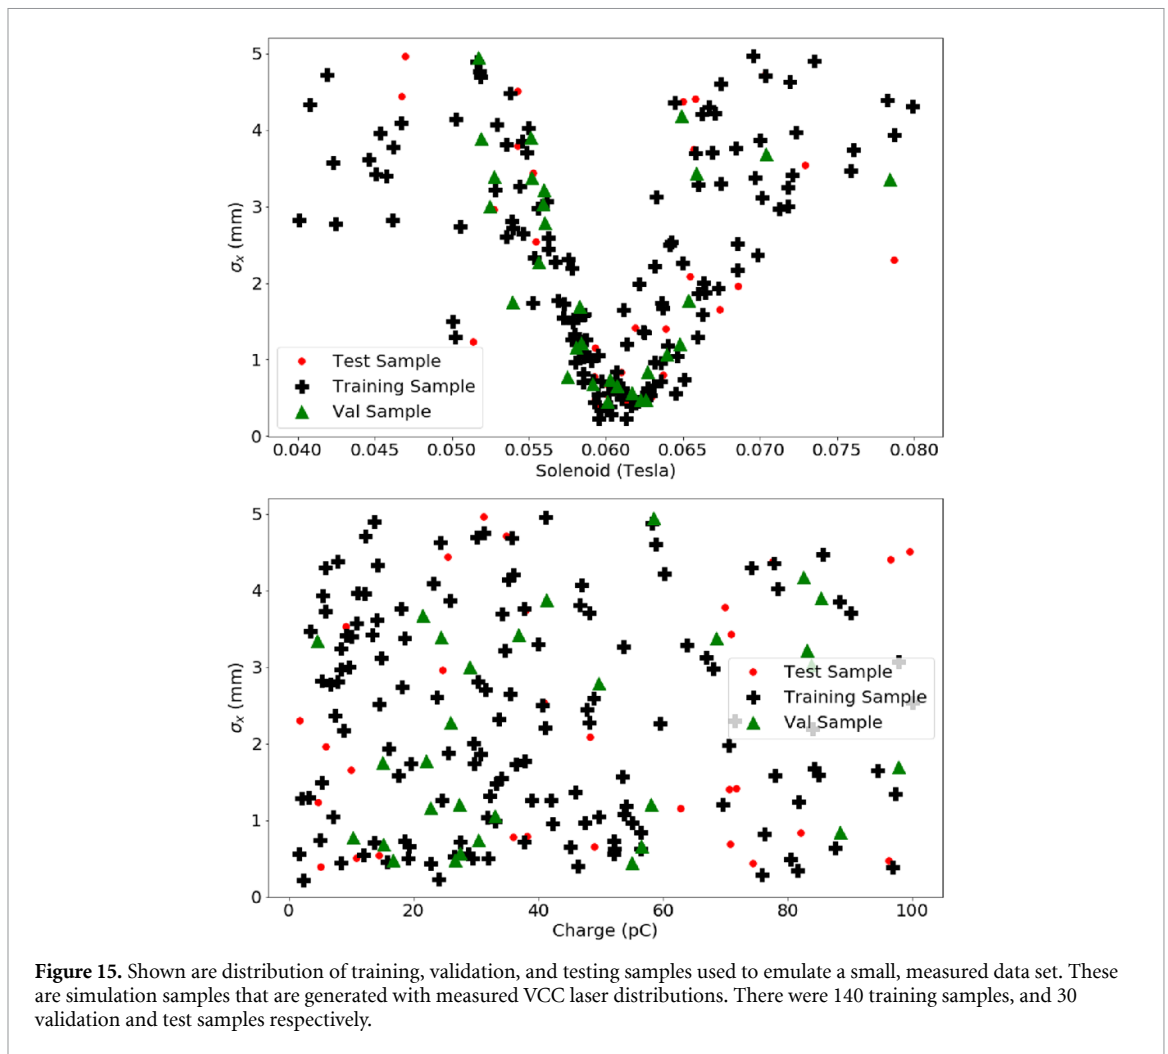


Figure 15. Shown are distribution of training, validation, and testing samples used to emulate a small, measured data set. These are simulation samples that are generated with measured VCC laser distributions. There were 140 training samples, and 30 validation and test samples respectively.

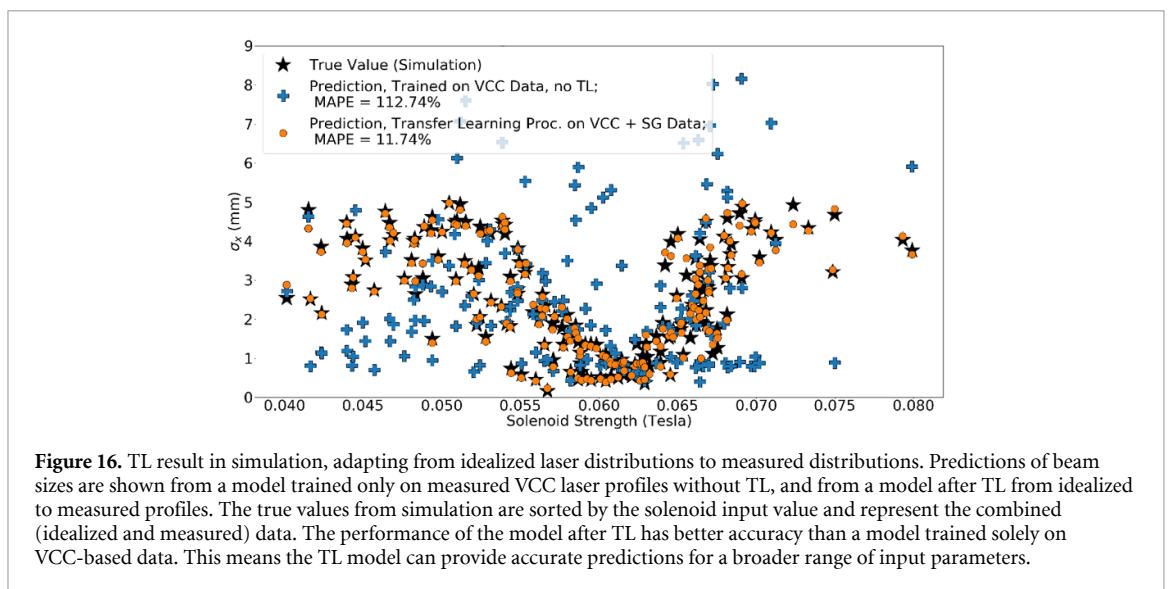
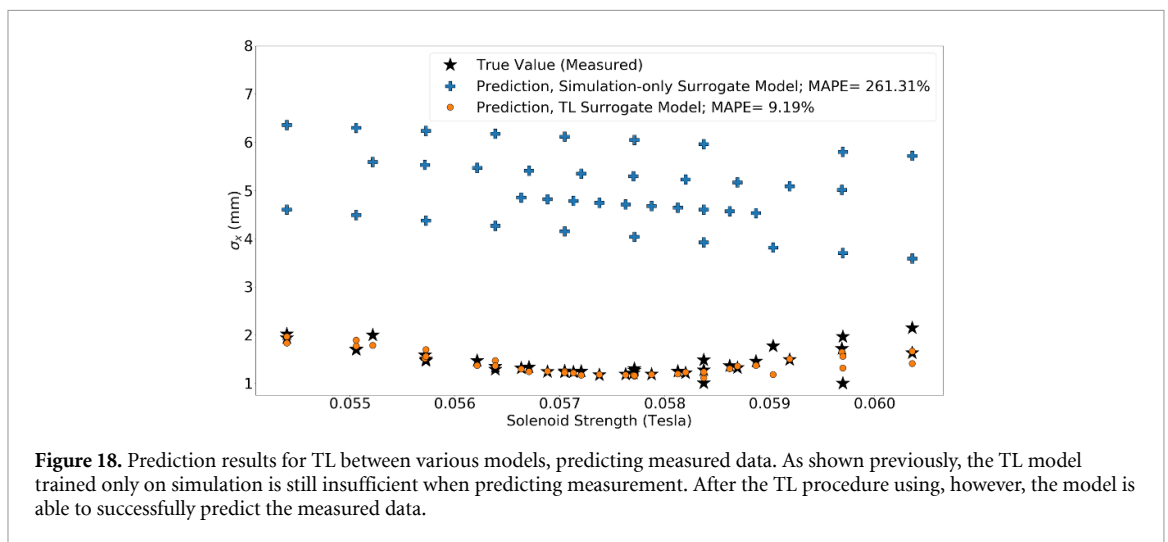
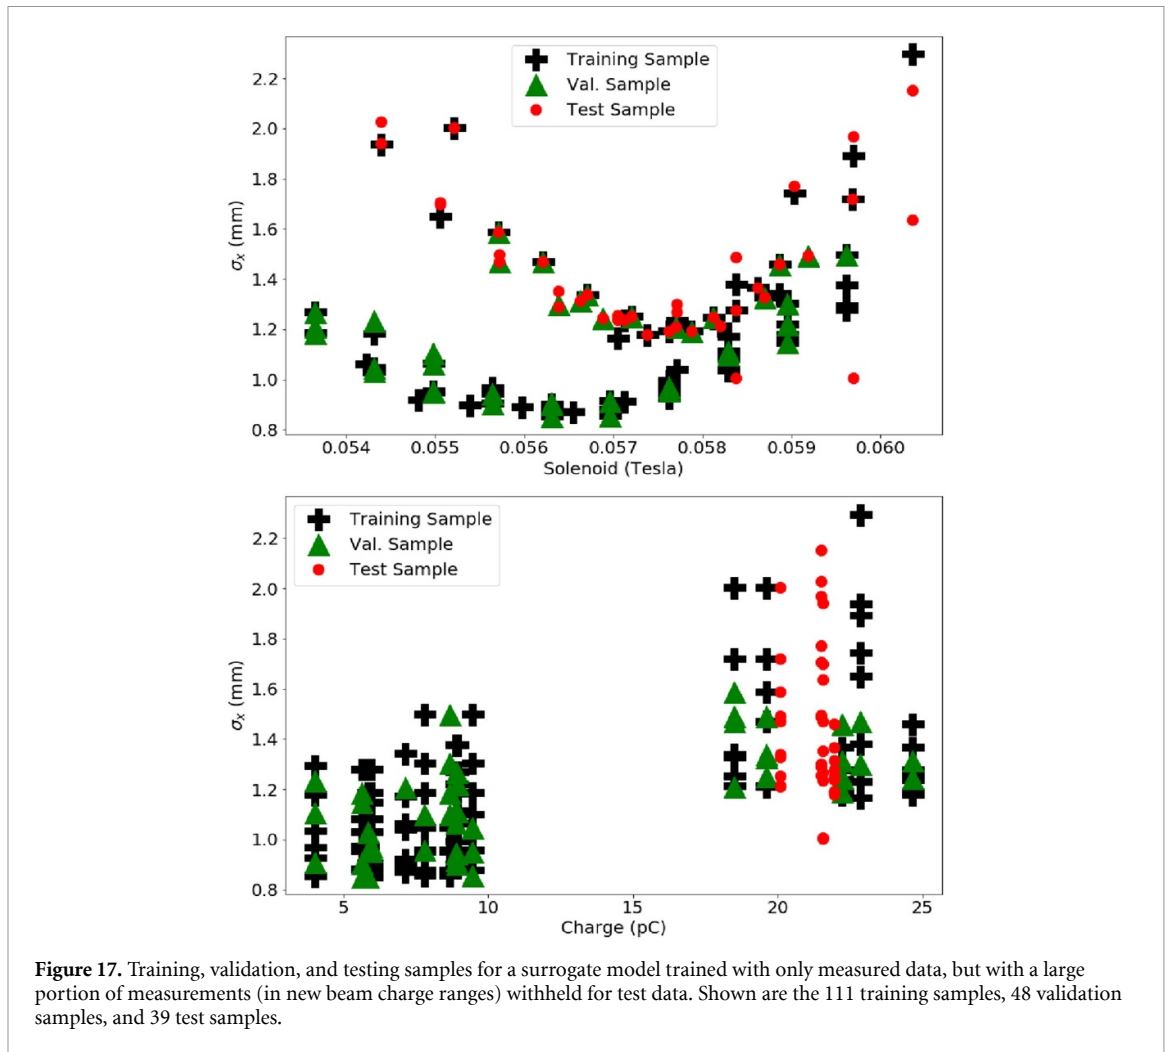


Figure 16. TL result in simulation, adapting from idealized laser distributions to measured distributions. Predictions of beam sizes are shown from a model trained only on measured VCC laser profiles without TL, and from a model after TL from idealized to measured profiles. The true values from simulation are sorted by the solenoid input value and represent the combined (idealized and measured) data. The performance of the model after TL has better accuracy than a model trained solely on VCC-based data. This means the TL model can provide accurate predictions for a broader range of input parameters.

Here, the TL procedure resulted in a model that can predict the measured data well. The MAPE of the best simulation-only model, with the TL, performed very poorly on the measured data. Once the TL with measured data is applied, the results improve drastically; the MAPE is now 9%. All of the MAPEs as they have progressed from simulation-only to simulation and measured via TL, are shown in table 2. However,



there was some degradation in ability to predict the simulated data. There is a clear trade off between the accuracy reached on the target data set (measured data) versus the base data set (simulation data). Iteration on this procedure may further improve the agreement on both data sets as needed for experimental use. Further, the TL model is able to predict on a broader range of laser input distributions, is expected to generalize better to new beam distributions.

Table 2. Comparisons of the MAPE between models before and after TL has been applied. The percent error is gauging how well the model predicts a test samples within the target data set.

Model	Target data	MAPE (%)
SG-Only Training	SG test samples	5
SG-Only Training	SG + VCC test samples	130
VCC-Only Training	SG + VCC test samples	113
SG + VCC (TL Proc.)	SG + VCC test samples	12
SG + VCC (TL Proc.)	Measured test samples	261
SG + VCC + Meas. (TL Proc.)	Measured test samples	9

5. Conclusions and future study

Surrogate models are a viable solution for many challenges faced while designing and operating particle accelerators. They can be used for real-time feedback in the form of virtual diagnostics, for offline experiment planning, and many other applications. In our study, we demonstrated novel methods for designing and training more comprehensive injector surrogate models.

First, a scalar surrogate model based on a wide range of simulated data was demonstrated, and we verified that it can be used for offline multi-objective optimization. Next, we showed that incorporating measured fluctuations in the initial laser distribution can improve the surrogate model. Specifically, by including measured laser inputs during the training process, the model can more accurately predict beam outputs for out-of-distribution laser inputs. Previous injector surrogate models have not leveraged measured laser input fluctuations. Therefore, we showcase how important this inclusion is toward improving long term surrogate model viability during operation.

Then, to train a simulation-based surrogate model trained on idealized laser distributions to be more representative of the real machine, a simple data augmentation technique and a TL procedure was able to successfully learn both output distributions (ideal and VCC-generated). As the LCLS-II injector is operated, additional measured VCC images could be incorporated into the model using this approach. Other methods for data augmentation for improving disparity in sampling such as the synthetic minority oversampling technique [31] could also be tried and compared.

Finally, we developed and applied a TL procedure for transferring from simulation to measured data, which successfully reduced the model prediction error on a held out range of beam charges from 112.7% to 7.6%. Further iteration of the TL process will likely improve the surrogate model training on both simulated and measured training data.

Further, the simulated data can be expanded to include more operational ranges such as gun gradient values, which may help resolve the shift in beam waist seen in the measured data. The ability of the surrogate model to successfully interpolate predictions within the range of possible input parameters (demonstrated in this case for previously unseen charges) can be very helpful for quickly estimating output parameters without needing experimental data. Our study shows that this is possible with a comprehensive machine-learning based surrogate model for the LCLS-II injector frontend. While we demonstrate this only for the LCLS-II injector frontend, these approaches for improving online modeling of injector systems could be easily adapted to other accelerator facilities.

Data availability statement

The data that support the findings of this study will be openly available following an embargo at the following URL/DOI: https://dataverse.harvard.edu/dataverse/SLAC_CNN_TransferLearning. Data will be available from 26 July 2021.

Acknowledgments

The authors would like to acknowledge Jane Shtalenkova and Alex Saad for their invaluable help and expertise, and the METSD department at SLAC for CAD drawings. Funding for this work was provided in part by the U.S. Department of Energy, Office of Science, under Contracts No. DE-AC02-76SF00515 and No. DE-AC02-06CH11357, FWP 100494 through BES on B&R code KC0406020. Funding for L Gupta provided by the U.S. Department of Energy, Office of Science Graduate Student Research (SCGSR) Program, the U.S. National Science Foundation under Award No. PHY-1549132, the Center for Bright Beams, and the University of Chicago Department of Physics.

ORCID iD

Lipi Gupta  <https://orcid.org/0000-0002-5179-7692>

References

- [1] Edelen A, Edelen J, Biedron S, Milton S and van der Slot P 2018 Recent applications of neural networks to modeling and control of particle accelerators *Proc. IPAC* p THYGBE2
- [2] Edelen A L, Biedron S G, Chase B E, Edstrom D, Milton S V and Stabile P 2016 Neural networks for modeling and control of particle accelerators *IEEE Trans. Nucl. Sci.* **63** 878–97
- [3] Edelen A et al 2018 Opportunities in machine learning for particle accelerators *White Paper From the 1st ICFA Machine Learning Workshop for Particle Accelerators*
- [4] Edelen A, Neveu N, Frey M, Huber Y, Mayes C and Adelmann A 2020 Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems *Phys. Rev. Accel. Beams* **23** 044601
- [5] Scheinker A, Emma C, Edelen A L and Gessner S 2020 Advanced control methods for particle accelerators (ACM4PA) 2019 *Workshop Report*
- [6] Scheinker A, Edelen A, Bohler D, Emma C and Lutman A 2018 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
- [7] Duris J et al 2020 Bayesian optimization of a free-electron laser *Phys. Rev. Lett.* **124** 124801
- [8] Hanuka A, Duris J, Shtalenkova J, Kennedy D, Edelen A, Ratner D and Huang X 2019 Online tuning and light source control using a physics-informed Gaussian process *Machine Learning for the Physical Sciences Workshop, NeurIPS (Vancouver, BC, Canada)*
- [9] Edelen A, Edelen J, Biedron S and Milton S 2016 First steps toward incorporating image based diagnostics into particle accelerator control systems using convolutional neural networks *Proc. NAPAC* p TUOA51
- [10] Emma C, Edelen A, Hogan M, O'Shea B, White G and Yakimenko V 2018 Machine learning-based longitudinal phase space prediction of particle accelerators *Phys. Rev. Accel. Beams* **21** 112802
- [11] Edelen A, Neveu N, Emma C, Ratner D and Mayes C 2019 Machine learning models for optimization and control of x-ray free electron lasers *NeurIPS Machine Learning for the Physical Sciences Workshop*
- [12] Komkov H, Dovlatyan L, Perevalov A and Lathrop D 2019 Reservoir computing for prediction of beamevolution in particle accelerators *NeurIPS Machine Learning for the Physical Sciences Workshop*
- [13] Emma C, Alverson M, Edelen A, Hanuka A, Hogan M, O'Shea B, Storey D, White G and Yakimenko V 2019 Machine learning based longitudinal phase space prediction of two-bunch operation at FACET-II *Proc. IBIC*
- [14] Floettmann K 1997–2021 *ASTRA: a space charge tracking algorithm*
- [15] Pratt L Y 1993 Discriminability-based transfer between neural networks *Advances in Neural Information Processing Systems* p 204
- [16] Thrun S and Pratt L 2012 *Learning to Learn* (New York: Springer Science & Business Media) (<https://doi.org/10.1007/978-1-4615-5529-2>)
- [17] Zhou F, Brachmann A, Emma P, Gilevich S and Huang Z 2012 Impact of the spatial laser distribution on photocathode gun operation *Phys. Rev. ST Accel. Beams* **15** 090701
- [18] Sannibale F et al 2012 Advanced photoinjector experiment photogun commissioning results *Phys. Rev. ST Accel. Beams* **15** 103501
- [19] Gulliford C 2019 DistGen: particle distribution generator (available at: <https://github.com/ColwynGulliford/distgen>)
- [20] Mayes C LUME-astra (available at: <https://github.com/ChristopherMayes/lume-astra>)
- [21] Bartnik A, Gulliford C, Bazarov I, Cultera L and Dunham B 2015 Operational experience with nanocoulomb bunch charges in the cornell photoinjector *Phys. Rev. ST Accel. Beams* **18** 083401
- [22] Brent R P 1973 *Algorithms for Minimization Without Derivatives* (Englewood Cliffs, NJ: Prentice-Hall)
- [23] Chollet F et al 2015 Keras (available at: <https://keras.io>)
- [24] Kingma D P and Ba J 2015 Adam: a method for stochastic optimization *ICLR (Poster)*
- [25] Pan S J and Yang Q 2009 A survey on transfer learning *IEEE Trans. Knowl. Data Eng.* **22** 1345–59
- [26] Shimodaira H 2000 Improving predictive inference under covariate shift by weighting the log-likelihood function *J. Stat. Plan. Inference* **90** 227–44
- [27] Kulis B, Saenko K and Darrell T 2011 What you saw is not what you get: domain adaptation using asymmetric kernel transforms *CVPR 2011 (IEEE)* pp 1785–92
- [28] Nam J, Fu W, Kim S, Menzies T and Tan L 2017 Heterogeneous defect prediction *IEEE Trans. Softw. Eng.* **44** 874–96
- [29] Wang C and Mahadevan S 2011 Heterogeneous domain adaptation using manifold alignment *IJCAI Proc.-Int. Conf. on Artificial Intelligence* vol 22 p 1541
- [30] Srivastava N, Hinton G, Krizhevsky A, Sutskever I and Salakhutdinov R 2014 Dropout: a simple way to prevent neural networks from overfitting *J. Mach. Learn. Res.* **15** 1929–58
- [31] Chawla N V, Bowyer K W, Hall L O and Kegelmeyer W P 2002 Smote: synthetic minority over-sampling technique *J. Artif. Intell. Res.* **16** 321–57