

Efficient six-dimensional phase space reconstructions from experimental measurements using generative machine learning

Ryan Roussel¹, Juan Pablo Gonzalez-Aguilera², Eric Wisniewski³, Alexander Ody³,
Wanning Liu³, John Power³, Young-Kee Kim², and Auralee Edelen¹

¹SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA

²Department of Physics and Enrico Fermi Institute, University of Chicago, Chicago, Illinois 60637, USA

³Argonne National Laboratory, Lemont, Illinois 60439, USA



(Received 19 April 2024; accepted 12 August 2024; published 11 September 2024)

Next-generation accelerator concepts, which hinge on the precise shaping of beam distributions, demand equally precise diagnostic methods capable of reconstructing beam distributions within six-dimensional position-momentum spaces. However, the characterization of intricate features within six-dimensional beam distributions using current diagnostic techniques necessitates a substantial number of measurements, using many hours of valuable beam time. Novel phase space reconstruction techniques are needed to reduce the number of measurements required to reconstruct detailed, high-dimensional beam features in order to resolve complex beam phenomena and as a feedback in precision beam shaping applications. In this study, we present a novel approach to reconstructing detailed six-dimensional phase space distributions from experimental measurements using generative machine learning and differentiable beam dynamics simulations. We demonstrate that this approach can be used to resolve six-dimensional phase space distributions from scratch, using basic beam manipulations and as few as 20 two-dimensional measurements of the beam profile. We also demonstrate an application of the reconstruction method in an experimental setting at the Argonne Wakefield Accelerator, where it is able to reconstruct the beam distribution and accurately predict previously unseen measurements 75× faster than previous methods.

DOI: [10.1103/PhysRevAccelBeams.27.094601](https://doi.org/10.1103/PhysRevAccelBeams.27.094601)

I. INTRODUCTION

Current and future particle accelerators rely on the precise control of beam distributions in six-dimensional position-momentum phase space, especially as accelerators push the boundaries of intensity and brightness for collider and light source applications [1]. For example, magnetized electron beams for hadron beam cooling [2], two bunch operations in free electron lasers [3], and drive/witness bunch pairs for high efficiency wakefield acceleration [4], are enabled by the precise control of six-dimensional phase spaces beyond macroscopic (rms) beam properties. Achieving this level of control requires measurement and analysis techniques that provide detailed information about the beam distribution in all six phase space coordinates, including cross correlations between the different phase spaces. Additionally, detailed characterization of six-dimensional phase space distributions is necessary to resolve complex beam dynamics phenomena that

couple particle motion along multiple axes, such as coherent synchrotron radiation [5] or plasma wakefield accelerators [6].

A wide variety of beam manipulation and diagnostic techniques have been developed to measure and reconstruct detailed characteristics of phase space distributions. These techniques can involve rotating the beam in phase space (tomography) [7–13], using masks, meshes, or slits to isolate and observe the dynamics of portions of the transverse beam [14–16], or using specialized, nondestructive beam diagnostics, such as laser wires [17].

However, reconstructing five- or six-dimensional phase space distributions from experimental measurements has proven to be substantially more difficult than reconstructing lower dimensional spaces (four or fewer phase space coordinates). Experimental five-dimensional phase space characterization has been done in limited instances, once at the Accelerator Research Experiment at SINBAD (ARES) beamline with a polarizable X-band transverse deflecting cavity [18,19] using the simultaneous algebraic reconstruction technique (SART) algorithm [20], and once at the Spallation Neutron Source (SNS) using a set of movable masking slits [21]. Full six-dimensional reconstruction of a single beam distribution has only been performed once, by combining multiple, scanning

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

masking slits with a dipole spectrometer and a bunch shape monitor at the SNS beamline [22]. However, these measurements required a significant amount of beam time resources to carry out (960 measurements over ~ 28 h for the five-dimensional ARES case and ~ 5 million measurements over ~ 32 h for the six-dimensional SNS measurement), making these measurement procedures impractical for regular use as feedback for online accelerator tuning or understanding complex beam dynamics phenomena. There is a need for reconstruction methods that significantly reduce the number of measurements required to reconstruct the six-dimensional beam distribution in order to be used during regular accelerator operations.

Recently, machine learning methods have been used to provide up to six-dimensional phase space predictions at locations along the beamline, so-called “virtual diagnostics,” using beam dynamics simulations or experimental data to correlate downstream beam distributions with upstream measurements and/or accelerator parameters [23–31]. However, these methods require significant amounts of experimental or simulated training data that spans the full range of possible phase space distributions in a given beamline to train the machine learning model to make accurate predictions. Additionally, virtual diagnostics are not easily transferable between beamlines, in varying operating configurations, or when time dependent drifts are present (although they can be adapted with additional data, e.g., [28]). If accurate simulations of beam dynamics or virtual diagnostics are unavailable, the beam distribution must be reconstructed directly from experimental measurements.

In previous work [32], we introduced and demonstrated a novel machine learning-based method for reconstructing detailed beam distributions directly from experimental data. This method, which we refer to here as generative phase space reconstruction (GPSR), uses generative machine learning models and physics-based differentiable beam dynamics simulations to effectively reconstruct detailed, high-dimensional phase space distributions from limited datasets, without the need for prior data collection or model pretraining. This method was used to reconstruct detailed, correlated, four-dimensional phase space distributions using a single quadrupole scan (ten measurements of the beam distribution) by taking advantage of the information contained in conventional measurements of the transverse beam profile using a YAG screen [33]. More recently, this method was also used to reconstruct beams with large angular momentum components (magnetization) and large emittance ratios [34].

In this work, we extend the use of GPSR to reconstruct six-dimensional phase space distributions from experimental measurements. We demonstrate in simulation that a diagnostic beamline consisting of quadrupoles, a transverse

deflecting cavity, and a dipole spectrometer can be used to resolve detailed characteristics of six-dimensional phase space distributions using as few as 20 measurements. We show that the GPSR technique accurately reconstructs a variety of different beam distributions from simulated screen measurements, including correlated Gaussian distributions, nonlinear distributions, and beam distributions similar to those produced by the emittance exchange (EEX) beamline at the Argonne Wakefield Accelerator (AWA) [35]. We then apply the GPSR algorithm to reconstructing a beam distribution from experimental measurements at AWA and demonstrate that it makes accurate predictions of previously unseen measurements. Finally, we discuss current limitations and advantages of the reconstruction technique.

II. METHODS

The method we use to reconstruct six-dimensional phase space distributions addresses two issues encountered by conventional reconstruction techniques.

First, conventional reconstruction techniques often use one- or zero-dimensional (scalar) projections of the six-dimensional phase space to infer features of the distribution. Conventional algebraic reconstruction techniques, such as maximum entropy tomography [12] and SART [20], typically use one-dimensional projections of screen images to produce a higher dimensional reconstruction. Measuring the intensity of the beam distribution through a set of slits, as done in [22], further reduces the measurement down to a scalar quantity. As a result, these methods lose a significant amount of information about the beam distribution, requiring more individual measurements of the distribution to resolve six-dimensional features. The GPSR technique enables us to fully utilize the detailed information contained in two-dimensional images of the beam distribution without the need to project to lower dimensions.

Second, traditional methods for representing particle distributions in a beam become increasingly inefficient when extended to six-dimensional distributions. Both of the methods listed above use a high-dimensional histogramming approach to describe the beam distribution; i.e., they solve for the beam density inside a number of high-dimensional voxels along an N -dimensional mesh. While this formalism is used to effectively describe distributions in one or two dimensions, it becomes impractical to describe beam distributions in six-dimensional space this way, as the number of bins grows exponentially with the number of dimensions. For example, resolving a beam distribution with 100 bins per dimension results in 10^{12} voxels needed to describe the full six-dimensional distribution. On the other hand, alternative methods for representing beam distributions, such as distribution moments or principal component analysis techniques, may reduce reconstruction detail.

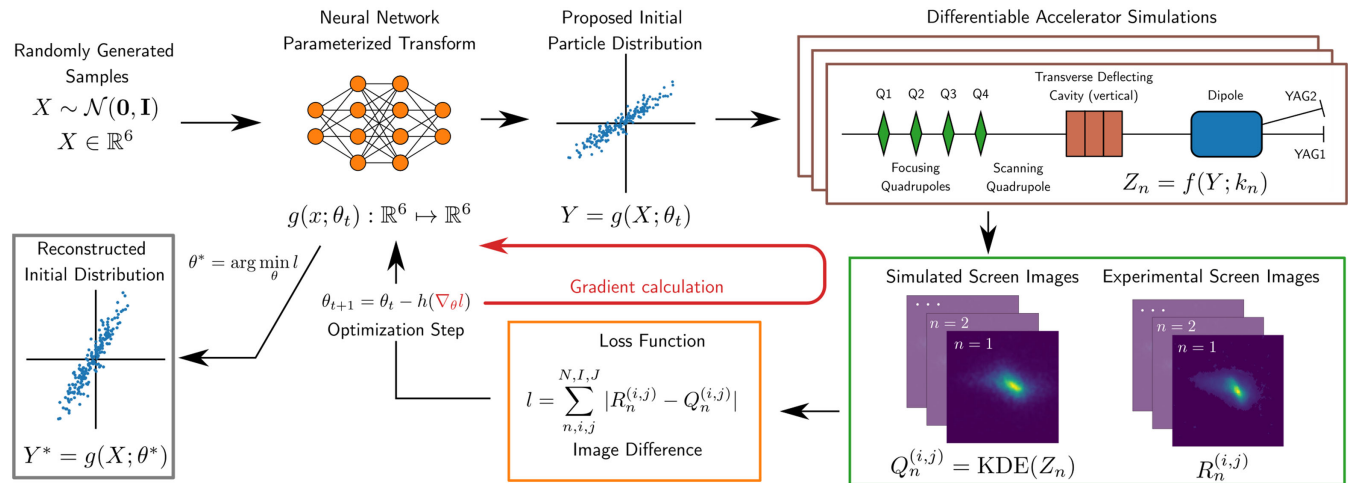


FIG. 1. Description of the GPSR approach for reconstructing phase space beam distributions from experimental data. A six-dimensional beam distribution is parameterized via a generative machine learning model, where randomly generated samples drawn from a multivariate normal distribution are transformed using a neural network into macroparticle coordinates in position-momentum space. The beam distribution is then transported through a backward differentiable simulation of the diagnostic beamline (implemented in Bmad-X [36]) to make measurement predictions at the diagnostic screens. The total per-pixel difference in intensity between simulated predictions $Q_n^{(i,j)}$ and experimental measurements $R_n^{(i,j)}$ is calculated as a loss function. The neural network parameters θ are then optimized to minimize the loss function using gradients calculated by the differentiable beam dynamics simulation. The distribution generated once the loss function has been minimized (simulation predictions match experimental measurements) is the reconstructed beam distribution.

To address these challenges, GPSR uses two concepts (shown in Fig. 1) to reconstruct detailed six-dimensional phase space distributions from experimental measurements.

First, we use a generative machine learning model to represent a distribution of macroparticles in six-dimensional phase space. Generative machine learning methods have been widely used in a variety of applications, including text [37] and image [38] generation, as well as in solving scientific problems, such as protein folding [39]. In the domain of accelerator physics, generative models have been explored for virtual diagnostics applications (see e.g., [30,31]).

In GPSR, generative modeling is used to generate samples from an arbitrary distribution of macroparticles by using a neural network to transform samples from a fixed random distribution (in this case a multivariate normal distribution) into six-dimensional position-momentum phase space coordinates. This is inspired by the development of neural radiance fields [40], which use neural networks to represent mass and color density functions in three-dimensional space. Neural networks of sufficient complexity are universal function approximators [41], enabling the generative model to produce particle distributions with nearly arbitrary structure in six-dimensional phase space. For the work here, we found that a fully connected neural network with two layers, 20 neurons each, connected by Tanh activation functions, was sufficient to represent a wide variety of beam distributions with high enough detail. With this method, the distribution of macroparticles in phase space is entirely controlled by the parameters of the neural network, namely the individual

neuron weights and biases inside the network. This results in a parameterization of phase space distributions using approximately 1000 free parameters, as opposed to the extremely large number of parameters needed for histogram-based representations of six-dimensional distributions.

The reconstruction process determines the parameters of the generative beam model by solving an optimization problem, where the goal is to minimize the total absolute error between simulated predictions and experimental measurements, in this case, the per-pixel intensity of the transverse YAG screens. Due to the number of free parameters contained inside the generative beam distribution model, solving this optimization problem in a reasonable amount of time is beyond the capabilities of black-box optimization algorithms because of the so-called “curse of dimensionality” [42]. As a result, it is necessary to use more powerful gradient-based optimization algorithms in GPSR. Unfortunately, this is also prohibitively difficult when using conventional beam dynamics simulations to predict experimental measurements, since calculating the gradients numerically in these cases requires finite-difference methods that also scale poorly with the number of free parameters used in optimization.

To address this challenge, GPSR leverages “backward-mode” [43] automatic differentiation to substantially reduce the cost of evaluating gradients of simulation outputs with respect to input parameters. Automatic differentiation is the technique of tracking derivative information alongside each computation step during physics

simulations. This, in turn, allows analytical evaluation of simulation output derivatives by using the chain rule in reverse to determine derivatives, a process commonly referred to as “backpropagation” [43] or “adjoint differentiation” [44]. Computational cost of calculating the derivatives in this manner is roughly equivalent to the cost of evaluating the simulation itself, and more importantly, it is largely independent of the number of input parameters the derivative is taken with respect to. This makes calculating derivatives substantially cheaper to compute than finite-difference methods when optimizing with respect to a large number of input parameters. It should be noted that this process is similar to, but distinct from, previous uses of automatic differentiation in accelerator physics, often referred to as “forward-mode” differentiation [44], “differential algebra” [45], or, when computing higher derivatives, “truncated power series algebra” [46]. These techniques are well suited for calculating particle transport dynamics up to arbitrary order; however, they do not scale well to calculating gradients with respect to thousands of input parameters.

Facilitating the use of backpropagation in the context of GPSR requires beam dynamics and diagnostic simulations that support the tracking of derivatives during evaluations. To this end, we have developed the simulation package $B_{\text{mad-X}}^1$ [36], which reimplements beam transport through a number of simple accelerator elements using the machine learning library *PyTorch* [48], which implements backpropagation. Additionally, we simulate the measurement of transverse beam profile intensity on a screen diagnostic by using kernel density estimation [49] as opposed to normal histogramming to preserve differentiability. By using differentiable beam dynamics simulations, we are able to cheaply compute derivatives for use in gradient descent optimization of beam distribution parameters to reconstruct the beam distribution. Differentiable representations of physics coupled to machine learning have been used for other problems in accelerator physics, such as lattice calibration, x-ray pulse reconstruction, and analysis of synchrotron-radiation-based diagnostics [47,50–52], although they were not used in this work.

III. SIX-DIMENSIONAL DIAGNOSTIC BEAMLINE DESCRIPTION

In previous work [32], we demonstrated that this reconstruction technique is able to produce accurate predictions of the four-dimensional beam distribution using images gathered from a single quadrupole scan. To extend this work toward resolving six-dimensional phase spaces, we add a transverse deflecting cavity (TDC) and a dipole spectrometer to the diagnostic

beamline, as shown in Fig. 1. In this case, quadrupoles Q1–Q3 are used to focus the beam onto YAG1 and quadrupole Q4 is scanned to measure the transverse phase space distribution. Focusing the beam on the diagnostic screen improves the measurement resolution and increases the range of Q4 strengths that can be scanned over since the transverse beam profile needs to be kept within a region of interest on both of the diagnostic screens. Pairing a TDC that kicks the beam vertically to resolve the current profile of the beam with a horizontally bending dipole magnet, which measures beam energy spread, is a common approach taken to measure the longitudinal phase space distribution of a beam [53]. The diagnostic setup used in this work is motivated by the notion that combining transverse information from the quadrupole scan with the longitudinal phase space manipulations should provide enough information to resolve the full six-dimensional phase space distribution.

We demonstrate the six-dimensional reconstruction technique using a simulation of the diagnostic beamline at AWA. Transverse diagnostic screens YAG1 and YAG2 are placed along each beam path when the dipole is off and on, respectively. These simulated screens have a region of interest that is 10×10 mm (or 200×200 pixels) in size, with a resolution of $50 \mu\text{m}/\text{px}$. The L-band transverse deflecting cavity used in simulation (1.3 GHz, $L = 0.48$ m) operated with a peak field in a range from 0 to 3 MV depending on the beam distribution, consistent with the operational range of the AWA deflecting cavity [54]. The rectangular dipole spectrometer ($L = 0.3018$ m) has a bend angle of 20 deg.




FIG. 2. Simulated screen images of the Gaussian beam distribution during the six-dimensional reconstruction scan. Brighter colors denote higher beam intensity (arbitrary scale for each image).

¹ $B_{\text{mad-X}}$ will be merged into the Cheetah [47] library for future applications of GPSR.

IV. SYNTHETIC RECONSTRUCTION EXAMPLES

In this section, we evaluate the effectiveness of the GPSR technique on several different phase space distributions. For each case described below, the following procedure was followed. The first three quadrupoles (Q1, Q2, and Q3) are tuned using Bayesian optimization [55] implemented in Xopt [56] to minimize the transverse beam size at YAG1 with the scanning quadrupole (Q4) turned-off. In practice, this optimization took less than a few minutes to conduct in simulation and experiment. This process is done in conventional longitudinal phase space measurements to reduce the betatron spot size contribution to the beam profile to increase measurement resolution [53].

We then scan the focusing strength of Q4 four times, once for each combination of TDC and dipole states. Images from the quadrupole scans are then used to reconstruct the beam distribution. In the synthetic case studies that follow, the quadrupole strength was scanned over five steps, resulting in a total dataset of 20 images. The number of quadrupole scan points was chosen empirically to provide sufficient information about the transverse phase space distribution while minimizing the number of data samples to reduce computational reconstruction costs (see Sec. VI for further discussion).

A. Case 1: Gaussian beam reconstruction

The first case we explore is one where the beam distribution is a multivariate normal distribution that contains cross correlations between a number of the six-dimensional phase space coordinates. Simulated measurements of the beam on the two diagnostic screens for each of the quadrupole scans are shown in Fig. 2. These images were then used to reconstruct the beam distribution using the GPSR approach with 100k macroparticles. To maximize reconstruction accuracy, the reconstruction was trained for 3000 iterations of gradient descent (Adam [57], learning rate 0.01), which took roughly 10–15 min on an NVIDIA A100 GPU, although it is possible that fewer iterations could be used.

The reconstruction algorithm results in a generative model, which generates a distribution of macroparticles in six-dimensional phase space that should approximate the true beam distribution. The reconstructed beam distribution, with a comparison to the synthetic ground truth distribution, is shown in Fig. 3. We see from the 50th and 90th percentile contours (calculated by measuring the beam intensities that account for 50%/90% of the total beam distribution) that the reconstructed beam distribution closely matches the true synthetic distribution. On the other hand, the reconstruction fails to accurately predict the halo of the beam distribution (100th percentile contour). This is likely due to two factors, which include limitations on the number of macroparticles used in the reconstruction process and the finite number of iterations used to optimize the generative model, which tends to identify high-intensity features in the distribution before lower-intensity ones.

Achieving accurate reconstructions across a larger dynamic range to identify beam halo is a subject of future study, so we avoid displaying this contour in the following sections.

To provide a quantitative measure of the reconstruction accuracy, we can calculate the second order moments of the beam distribution. We observe that in this case the reconstructed beam distribution shows close agreement with the ground truth. The fractional error between the second-order components of the reconstructed beam distribution and the ground truth is frequently less than 10%. In cases where the cross covariances of the ground truth distribution are much smaller than the principal axis covariances, the error of the reconstruction is higher due to the dominance of the beam size along the principal axis over other beam features.

B. Case 2: Nonlinear beam reconstruction

Next, we examine a case where the beam has a variety of nonlinearities and correlations throughout the phase space distribution, as shown in Fig. 4 (simulated measurements of this beam distribution can be found in the Appendix). The reconstruction successfully replicates the ground truth distribution, including nonlinear features contained in cross-correlated phase spaces. The comparison between the predicted and ground truth values of the covariance matrix again shows good agreement within 10% for most elements.

C. Case 3: EEX beam reconstruction

Detailed six-dimensional phase space reconstructions are particularly important when performing complex beam manipulations that precisely shape features in six-dimensional phase space. An example of this is transverse-to-longitudinal emittance exchange (EEX) [35,58]. This process combines two doglegs and a transverse deflecting cavity to map the horizontal phase space distribution into a longitudinal distribution, allowing longitudinal profile shaping using transverse masking. For example, the EEX beamline has been used to generate ramped current profiles with a sharp drop off to improve the transformer ratio of dielectric [53] and plasma [59] wakefield acceleration. Precisely characterizing beam distributions created by the EEX beamline enables us to control accelerator and beam parameters, such as the location of leafs in a multileaf masking element [60] or focusing magnets before the EEX beamline [61], that lead to improved six-dimensional beam tailoring for accelerator applications.

As a proof-of-concept demonstration, we generated a synthetic beam distribution which mimics those created by the EEX beamline for high transformer ratio wakefield acceleration applications. In this case, EEX aims to create a drive beam with a triangular current profile and a uniform witness beam to sample the wakefield. This is achieved using a laser cut mask to shape the horizontal beam profile before EEX, which is then mapped into a current distribution during the exchange process. However, this results in a correlation between the vertical beam size and longitudinal position

within the bunch, since the vertical distribution is not exchanged inside the EEX beamline. This can be detrimental to achieving optimal matching into a wakefield structure due to time-of-flight degradation of the longitudinal profile inside strong final-focus magnets before the wakefield device [62].

Reconstruction results from simulated diagnostic measurements (seen in the Appendix) of this beam distribution are shown in Fig. 5. We see that the reconstruction algorithm can resolve important features of the beam distribution,

including the longitudinal profile and correlations between transverse and longitudinal phase spaces. However, the reconstruction algorithm has difficulty identifying large areas of uniform density within the beam, most notably in the $z - y$ phase space where the triangular head and the lower rectangular regions should both have near uniform density profiles. This leads to slightly inaccurate predictions of the longitudinal beam profile, which is critical for high transformer ratio applications. Identifying how to improve the accuracy of the reconstruction in this case, either through

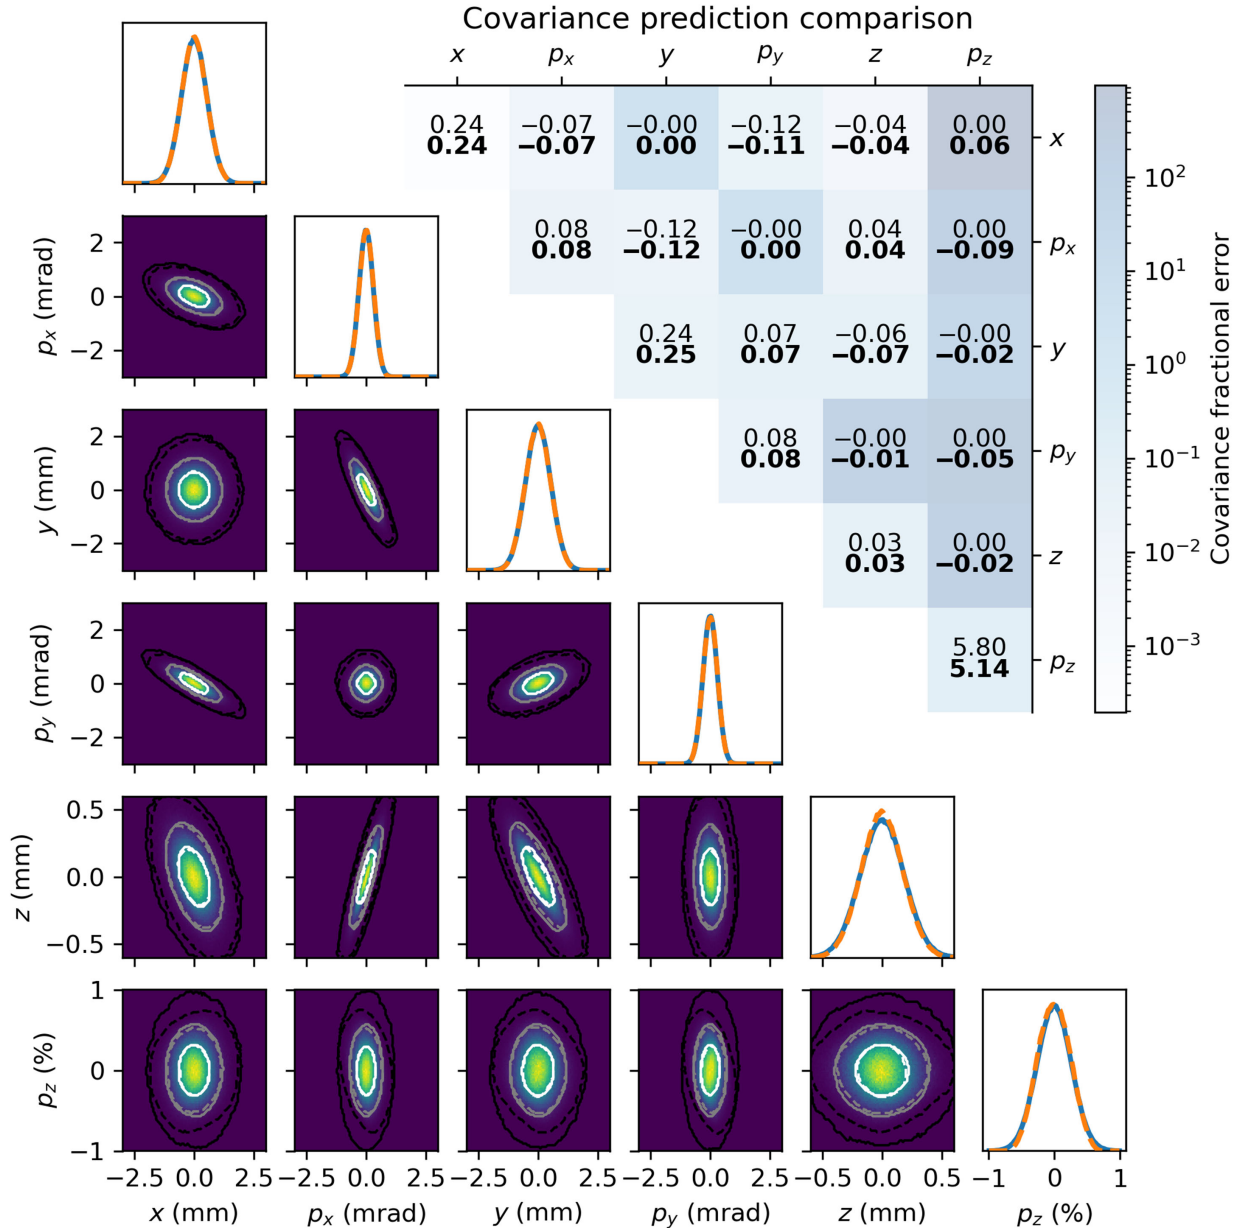


FIG. 3. Reconstruction results from Gaussian synthetic beam distribution. Lower left: comparison between two-dimensional projections of the ground truth synthetic beam distribution with the reconstructed beam distribution. Solid lines denote ground truth projections and contours, while dashed lines denote reconstruction predictions. White, gray, and black contours denote 50th, 90th, and 100th percentile intensity levels, respectively. Color map intensity denotes reconstructed prediction. Upper right: comparison between ground truth and reconstructed second order moments of 90th percentile beam particles, where regular text denotes ground truth values and bold text denotes reconstructed predictions.

algorithm modifications or changes in the diagnostic beamline, is a topic of future study.

V. EXPERIMENTAL DEMONSTRATION

In addition to simulation studies, we conducted an experimental demonstration of using the six-dimensional GPSR algorithm to characterize electron beams generated at AWA. We configured the AWA drive beamline to produce 1 nC electron bunches at a 2 Hz repetition rate

with an energy of 43 MeV and transported them to the diagnostic configuration shown in Fig. 1 at the end of the AWA beamline. The beam charge was selected to provide a strong enough signal-to-noise ratio for beam imaging diagnostics while mitigating coherent synchrotron radiation effects in the dipole spectrometer.

Quadrupoles Q1–Q3 were then used to focus the beam onto YAG1 using Bayesian optimization algorithms, while the scanning quadrupole Q4 was turned-off. We then repeated four quadrupole scans (–2.9 to 2.9 T/m, nine steps) with

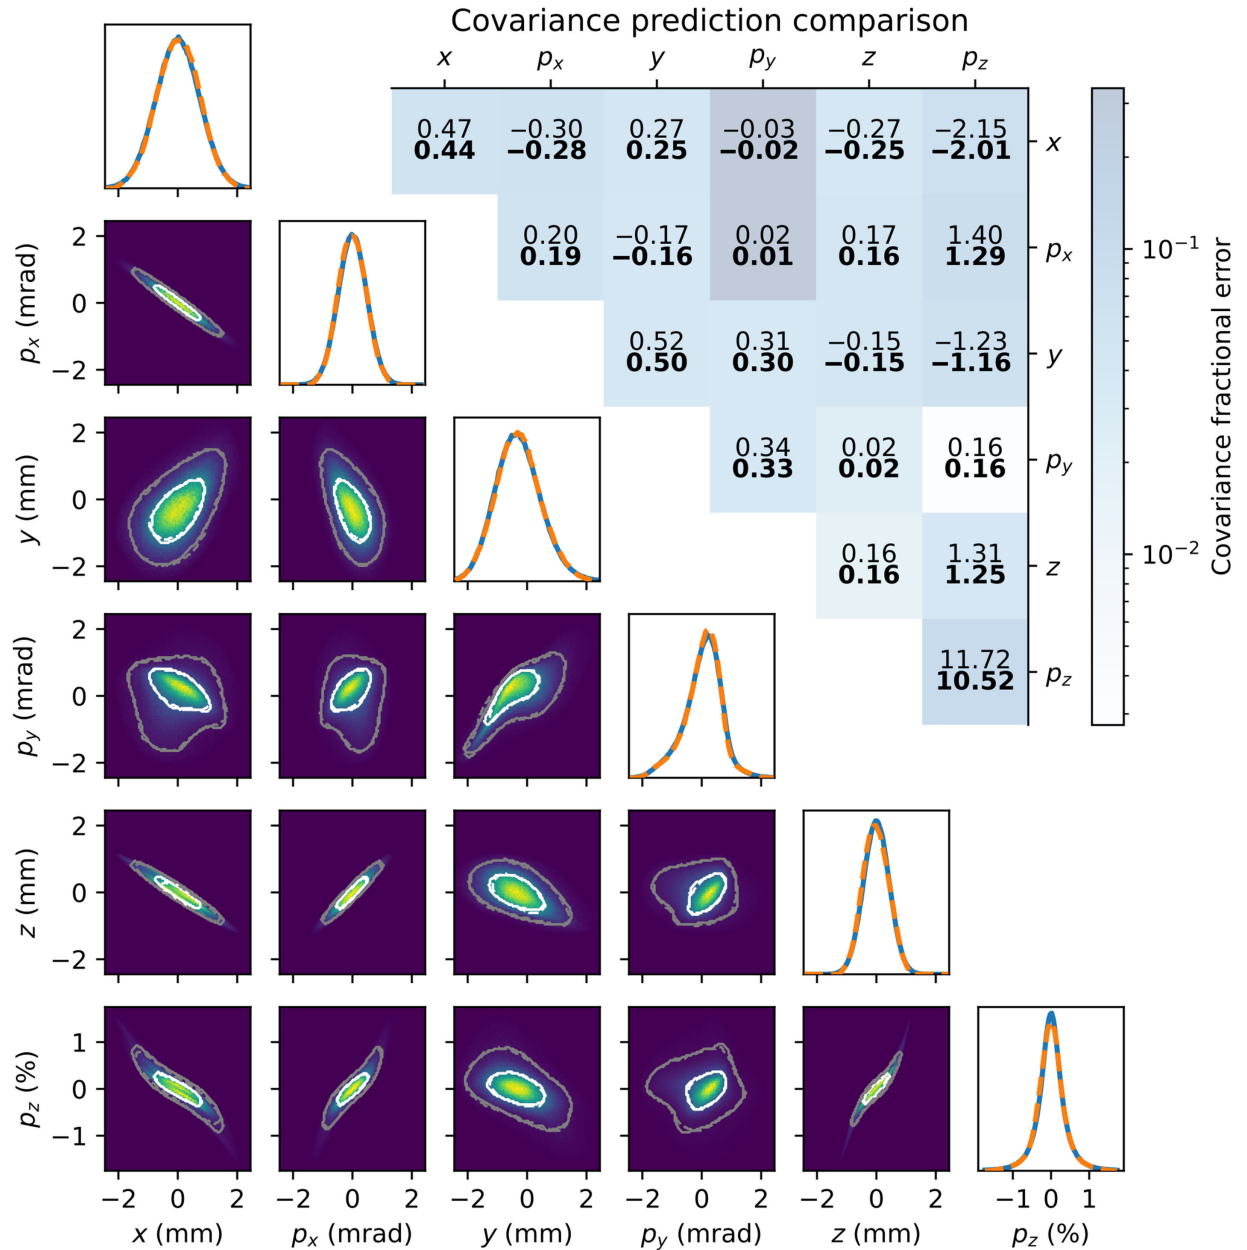


FIG. 4. Reconstruction results from nonlinear synthetic beam distribution. Lower left: comparison between two-dimensional projections of the ground truth synthetic beam distribution with the reconstructed beam distribution. Solid lines denote ground truth projections and contours, while dashed lines denote reconstruction predictions. White and gray contours denote 50th and 90th percentile intensity levels. Color map intensity denotes reconstructed prediction. Upper right: comparison between ground truth and reconstructed second order moments of 90th percentile beam particles, where regular text denotes ground truth values and bold text denotes reconstructed predictions.

different TDC and dipole settings, similar to what was done in the simulated examples. A set of five beam image shots were taken for each parameter setting with a charge window of 0.1 nC (i.e., shots where the deviation of the measured bunch charge exceeded the charge window were discarded). The four parameter scans took approximately 8 min to perform at AWA, although we estimate that given better charge stability, the scan time could be reduced by a factor of 2.

An additional complexity of using GPSR in an experimental setting was a slight difference in image resolution between the two YAG screens used to measure the beam

profile due to minor differences in camera location and focusing. These differences were incorporated into the reconstruction by defining two beamline simulations; one for the case where the dipole was off, and one case where the dipole was on, each with different definitions for the screen diagnostic. The images were cropped to a square size of 300×300 pixels, which corresponds to a side length of approximately 13 mm. Finally, the intensity of the image was clipped and set to zero at a lower bound threshold and a Gaussian smoothing filter was applied to the images to remove salt and pepper noise.

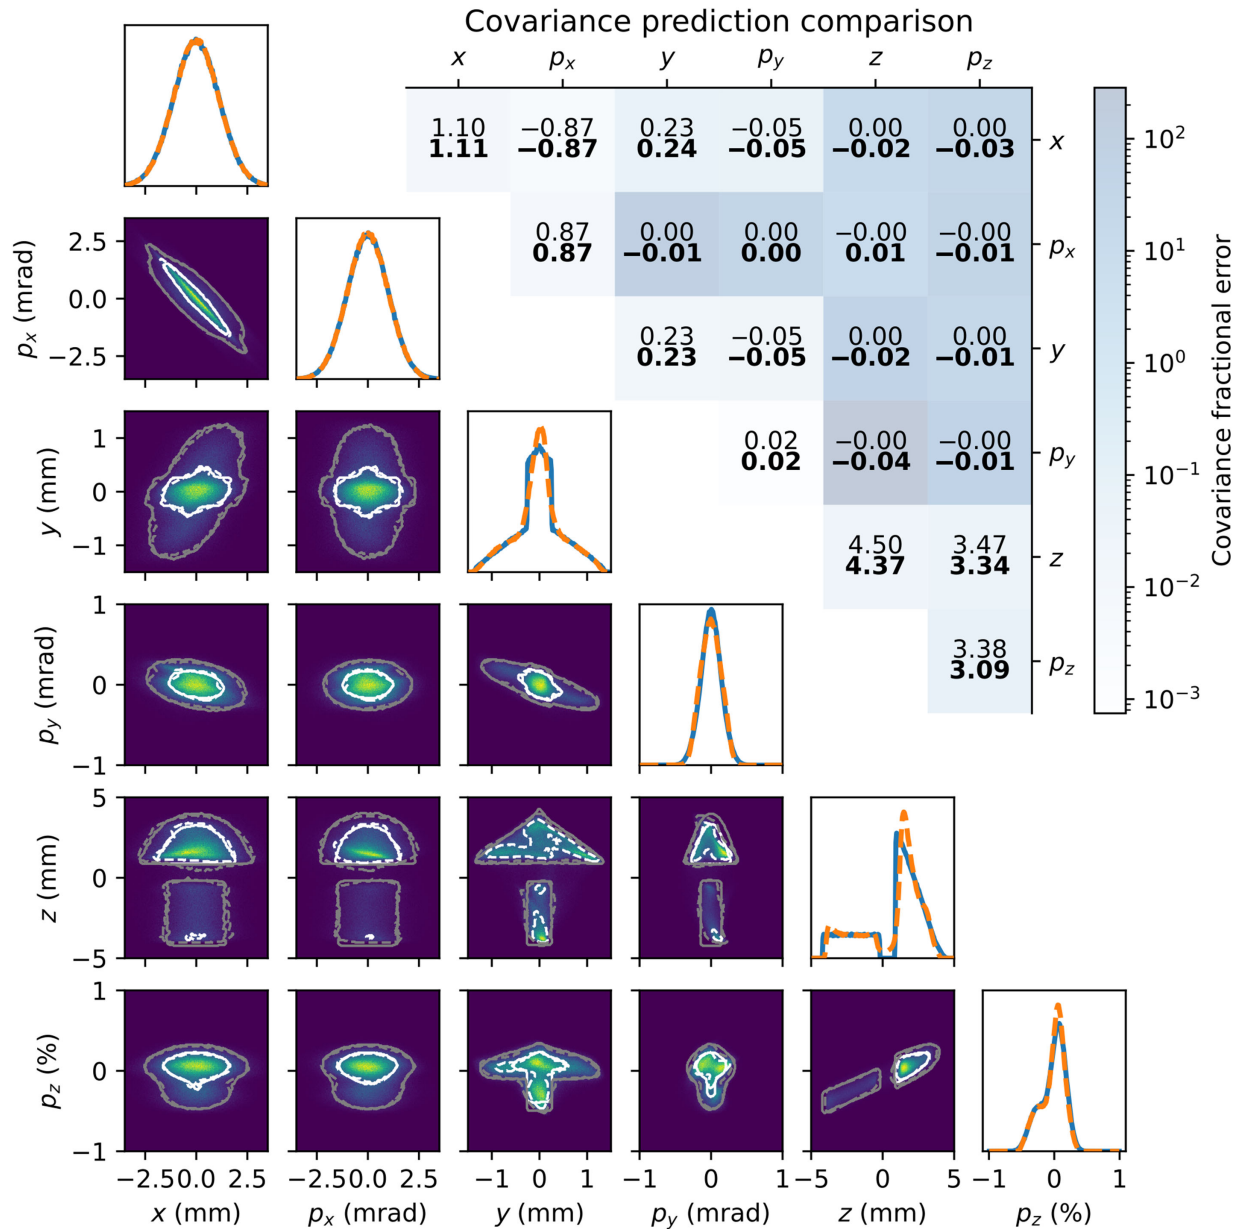


FIG. 5. Reconstruction results from the EEX synthetic beam distribution. Lower left: comparison between two-dimensional projections of the ground truth synthetic beam distribution with the reconstructed beam distribution. Solid lines denote ground truth projections and contours, while dashed lines denote reconstruction predictions. White contours denote 90th percentile intensity levels. Color map intensity denotes reconstructed prediction. Upper right: comparison between ground truth and reconstructed second order moments of 90th percentile beam particles, where regular text denotes ground truth values and bold text denotes reconstructed predictions.

While comparing individual pixel intensities on the diagnostic screen to simulation predictions provides more information about the beam distribution, it does make phase space reconstruction sensitive to shot-to-shot jitter in the beam centroid. Detailed features of the beam distribution can be obscured when averaging over images that contain shot-to-shot variations in the beam centroid. To mitigate this effect, we shifted the distribution of each image such that the beam centroid was at the center of the region of interest and averaged the intensity profile over the five shots. As a result, our reconstruction predicts the structure of the distribution, but not its offset with respect to magnetic element centers, mean energy, or timing relative to the zero crossing of the TDC cavity. A more robust treatment of the shot-to-shot jitter would allow us to

identify these aspects of the beam distribution, as was done in [32], and is a topic of future study.

In order to validate the accuracy of the reconstruction in the absence of a ground truth beam distribution, we compared model predictions to a subset of data that was not included in determining the phase space distribution (referred to here as the test dataset). We can be confident that the reconstructed distribution is accurate if the generative model can accurately predict measurements inside the test set. We selected every other quadrupole strength in the four quadrupole scans to be test data, resulting in a training dataset consisting of 20 images (5 quadrupole settings \times 4 LPS diagnostic settings) and a test dataset consisting of 16 images (4 quadrupole settings \times 4 LPS diagnostic settings). The entire dataset can be viewed in the Appendix. With this




FIG. 6. Comparison between averaged, experimentally measured test images with predictions from propagating the reconstructed beam distribution to the screen diagnostics. White and gray contours denote 50th and 90th percentile intensity levels, with solid lines representing measurements and dashed lines denoting predictions. Color map intensity denotes reconstructed prediction on an arbitrary scale. Blue lines denote measured one-dimensional projections, while orange dashed lines denote projection predictions.

training dataset, the reconstruction took approximately 17 min on an A100 NVIDIA GPU due to the increased image size compared to simulated examples.

A comparison between predictions from the reconstruction and the experimentally measured test data is shown in Fig. 6. We see that the reconstruction accurately reproduces the one- and two-dimensional beam structure (with minor discrepancies in some cases) seen in the test images, including nonlinear and beamlet features. The reconstructed six-dimensional phase space from experimental measurements is shown in Fig. 7. The reconstruction predicts normalized transverse emittances of $\varepsilon_{x,n} = 19$ mm mrad, $\varepsilon_{y,n} = 8.5$ mm mrad, which are consistent with nonoptimized AWA beamline parameters, the bunch charge of

1 nC, and the observed asymmetric beam size growth as a function of quadrupole strength when the dipole and deflecting cavity are off. One prominent feature of the distribution is the beamlet structure observed in the longitudinal current profile. We believe that this structure results from a set of 5 alpha-BBO crystals used at AWA to produce longer, flat-top laser pulse profiles by stacking 32 laser subpulses together [63]; however, it is unclear how this leads to variations in beam energy along the bunch.

VI. DISCUSSION

While the examples shown here demonstrate the effectiveness of GPSR in these cases, the reliability and

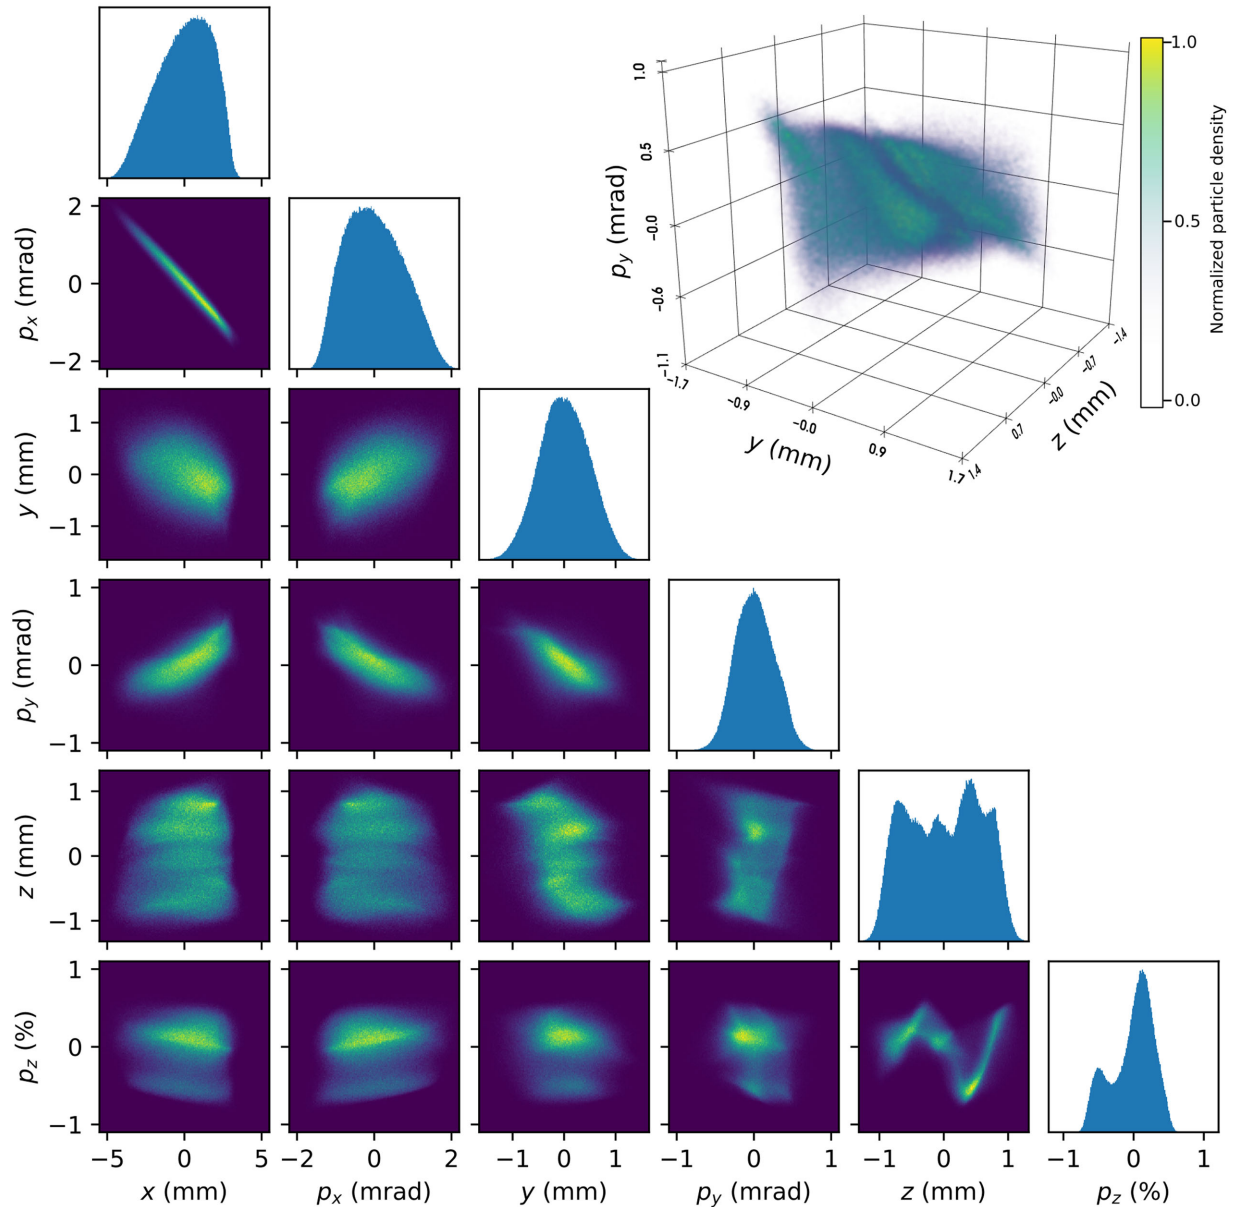


FIG. 7. Projections of the reconstructed six-dimensional beam distributions where the color map intensity denotes beam density (independent scaling for each projection). (Top right) Three-dimensional density map of beam particles in $z - y - p_y$ space.

robustness of the algorithm can be improved for general applications.

The most critical next step for this work is to enable the algorithm to quantify uncertainties in the reconstructed phase space distribution from limited experimental measurements. Conventional wisdom and empirical evidence [64] suggest that evenly spaced angular rotations of the distribution over 360 deg is necessary to accurately reconstruct two-dimensional phase space distributions from one-dimensional projections. However, it is not clear how this rule-of-thumb translates to reconstructing beam distributions in six-dimensional phase space from two-dimensional images or how this is dependent on variations in the beam distribution. Developing a calibrated estimation of the reconstruction uncertainty would enable a greater understanding of the reconstruction accuracy by tracking the reduction in reconstruction uncertainty as measurements are added. This in-turn would reduce the number of redundant measurements that do not contribute significant information about the distribution, reducing the experimental and computational effort needed to perform the reconstruction. Additionally, uncertainty quantification methods can provide information about the minimum resolution of the reconstruction, while also providing potential pathways to improve diagnostic resolution.

The flexibility of the GPSR algorithm can be used to develop novel diagnostic techniques that were previously impractical to analyze using conventional methods. Unlike other analysis methods, the GPSR algorithm does not require simplifications or approximations of beam dynamics or measurement signals needed to enable analytical tractability. By utilizing high-performance numerical optimization techniques, GPSR can solve extremely complex optimization problems that can incorporate exact measurement information into the reconstruction process, regardless of analytical tractability. This allows for the use of nonlinear phase space manipulations, which may provide more information about the beam distribution than linear manipulations. Additionally, GPSR can also easily combine data from multiple, potentially heterogeneous sources of information (as done here, where different diagnostics have different resolutions) about the beam distribution by adding additional terms into the training loss function shown in Fig. 1. For example, screen measurements of the two-dimensional beam profile can be easily combined with upstream measurements of the beam profile, nondestructive measurements of the beam distribution, such as edge radiation in bends [65], coherent optical transition radiation [66], or machine learning-based predictions of the longitudinal phase space [23] into a single, self-consistent description of the beam distribution.

Another advantage of using GPSR is that the generative model creates these macroparticles by transforming samples from a simple random distribution. As a result, predicted particle distributions can contain any number of macroparticles regardless of the number of particles used during training the model. For example, even though the

generative model in the cases demonstrated here was trained with 100k particles, the figures shown here use predictive distributions containing one million macroparticles. This method for representing beam distributions is advantageous in a number of ways for use in simulations. Saving and transporting high-fidelity beam distributions containing a large number of six-dimensional macroparticle coordinates is a memory intensive process. Generative models, on the other hand, use substantially fewer scalar quantities (the weights and offsets of neural network parameters) and a description of the model structure to represent the beam distribution at any level of fidelity, significantly reducing the memory needed to share the beam distribution between simulations.

Finally, it is important to note that the GPSR algorithm can utilize arbitrary generative model architectures to generate particle distributions for reconstruction purposes. As stated previously, the work done here uses a basic, densely connected neural network to transform samples from a random distribution into a distribution of macroparticles. However, alternative methods for generating particle distributions can be used in this workflow, provided they are backward differentiable. For example, normalizing flows [67] have also shown promise in representing complex beam distributions [68]. Normalizing flows have the additional benefit of providing the full probability distribution of the beam in six-dimensional space (not just samples from that distribution), which can then be used with maximum entropy-based loss functions to solve for the most likely beam distribution given a small set of measurements. Additionally, lower fidelity reconstructions can be performed using reduced models that improve reconstruction speed at the cost of accuracy for online control applications. Comparing the performance of different generative models for use in phase space reconstruction is the subject of future work.

VII. CONCLUSION

In this work, we have demonstrated in several simulated and experimental case studies that detailed characterizations of six-dimensional beam distributions can be achieved rapidly using the GPSR algorithm. This analysis method leverages the detail contained in two-dimensional screen images of the beam distribution and knowledge of beam dynamics in the accelerator to significantly reduce (by up to a factor of 75 \times) the time needed to produce detailed predictions of the beam six-dimensional phase space distribution. Furthermore, the generative machine learning representation of the beam distribution is trained from scratch on experimental data, requiring no previous data collection or pretraining needed by other applications of machine learning in accelerator physics. As a result, the GPSR algorithm can be used to provide six-dimensional phase space information during accelerator operations and in a wide variety of contexts. This technique has major implications for allowing six-dimensional information to be used to inform accelerator control and understand complex physical phenomena.

Software and processed data are available at [69]. Raw data can be made available upon request.

ACKNOWLEDGMENTS

This work is supported by the U.S. Department of Energy, Office of Science under Contract No. DE-AC02-76SF00515 and the Center for Bright Beams, NSF Award No. PHY-1549132. The work at AWA was funded through the U.S. Department of Energy, Office of Science under Contract No. DE-AC02-06CH11357. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC Award No. ERCAP0020725. U.S. Department of Energy, Office of Science under Contract No. DE-AC02-06CH11357.

Conceptualization, R. R., J. P. G. A., and A. L. E.; data curation, R. R. and J. P. G. A.; formal analysis, R. R., and J. P. G. A.; funding acquisition, A. L. E., J. P., and Y. K. K.; investigation, R. R. and J. P. G. A.; methodology, R. R., J. P. G. A., and A. L. E.; software, R. R. and J. P. G. A.; experiment, R. R., J. P. G. A., W. L., E. W., and A. O.; supervision, A. L. E., R. R., Y. K. K., and J. P.; validation, R. R., and J. P. G. A.; visualization, R. R., and J. P. G. A.; writing—original draft, R. R., and J. P. G. A.; and writing—review and editing, R. R., J. P. G. A., and A. L. E. All authors have read and agreed to the published version of the manuscript.

APPENDIX: SYNTHETIC AND EXPERIMENTAL DATASETS

Figures 8–10 contain training datasets used to reconstruct the beam distributions described in Secs. IV and V.




FIG. 8. Simulated screen images of the nonlinear beam distribution during the six-dimensional reconstruction scan. Brighter colors denote higher beam intensity (arbitrary scale for each image).




FIG. 9. Simulated screen images of the emittance exchange generated beam distribution during the six-dimensional reconstruction scan. Brighter colors denote higher beam intensity (arbitrary scale for each image).




FIG. 10. Comparison between all averaged, experimentally measured images with predictions from propagating the reconstructed beam distribution to the screen diagnostics. White and gray contours denote 50th and 90th percentile intensity levels, with solid lines representing measurements and dashed lines denoting predictions. Color map intensity denotes reconstructed prediction on an arbitrary scale. Blue lines denote measured one-dimensional projections, while orange dashed lines denote projection predictions. Images with red borders denote test images not included in determining the phase space distribution.

-
- [1] J. Blazey, D.L. Bruhwiler, J. Byrd, S. Cousineau, Z. Huang, S. Nagaitsev, M. Palmer, P. Piot, J. Power, J. Rosenzweig, V. Shiltsev, L. Spentzouris, and J.-L. Vay, in *Proceedings of the Accelerator and Beam Physics Roadmap Workshop* (2023), https://science.osti.gov/hep-/media/hep/pdf/2022/ABP_Roadmap_2023_final.pdf.
- [2] Y. S. Derbenev and A. Skrinskii, Magnetization effects in electron cooling, AN SSSR, Technical Report No. BNL-TR-663, 1977.
- [3] A. Marinelli, D. Ratner, A. A. Lutman, J. Turner, J. Welch, F.J. Decker, H. Loos, C. Behrens, S. Gilevich, A. A. Miahnahri, S. Vetter, T.J. Maxwell, Y. Ding, R. Coffee, S. Wakatsuki, and Z. Huang, High-intensity double-pulse X-ray free-electron laser, *Nat. Commun.* **6**, 6369 (2015).
- [4] C. Lindström, J. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D’Arcy, P. Gonzalez, A. Knetsch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, and J. Osterhoff, Energy-spread preservation and high efficiency in a plasma-wakefield accelerator, *Phys. Rev. Lett.* **126**, 014801 (2021).
- [5] H. H. Braun, R. Corsini, L. Groening, F. Zhou, A. Kabel, T. O. Raubenheimer, R. Li, and T. Limberg, Emittance growth and energy loss due to coherent synchrotron radiation in a bunch compressor, *Phys. Rev. ST Accel. Beams* **3**, 124402 (2000).
- [6] C. A. Lindström and M. Thévenet, Emittance preservation in advanced accelerators, *J. Instrum.* **17**, P05016 (2022).
- [7] V. Yakimenko, M. Babzien, I. Ben-Zvi, R. Malone, and X.-J. Wang, Electron beam phase-space measurement using a high-precision tomography technique, *Phys. Rev. ST Accel. Beams* **6**, 122801 (2003).
- [8] O.R. Sander, D.D. Chamberlin, R. A. Jameson, and G. N. Minerbo, Beam tomography in two and four dimensions, in *Proc. LINAC’79, Montauk, NY, USA* (1979), paper S5-10, pp. 314–318.
- [9] B. Hermann, V. A. Guzenko, O. R. Hürzeler, A. Kirchner, G. L. Orlandi, E. Prat, and R. Ischebeck, Electron beam transverse phase space tomography using nanofabricated wire scanners with submicrometer resolution, *Phys. Rev. Accel. Beams* **24**, 022802 (2021).
- [10] C. B. McKee, P. G. O’Shea, and J. M. J. Madey, Phase space tomography of relativistic electron beams, *Nucl. Instrum. Methods Phys. Res., Sect. A* **358**, 264 (1995).
- [11] J. J. Scheins, Tomographic reconstruction of transverse and longitudinal phase space distribution using the maximum entropy algorithm, Deutsches Elektronen-Synchrotron DESY, TESLA Report No. 2004-08, 2004.
- [12] K. M. Hock and M. G. Ibson, A study of the maximum entropy technique for phase space tomography, *J. Instrum.* **8**, P02003 (2013).
- [13] A. Hoover, Towards the production of a self-consistent phase space distribution, Ph.D. diss., University of Tennessee, 2022.
- [14] J. G. Wang, D. X. Wang, and M. Reiser, Beam emittance measurement by the pepper-pot method, *Nucl. Instrum. Methods Phys. Res., Sect. A* **307**, 190 (1991).
- [15] D. Marx, J. Giner Navarro, D. Cesar, J. Maxson, B. Marchetti, R. Assmann, and P. Musumeci, Single-shot reconstruction of core 4D phase space of high-brightness

- electron beams using metal grids, *Phys. Rev. Accel. Beams* **21**, 102802 (2018).
- [16] M. Zhang, Emittance formula for slits and pepper-pot measurement, Fermi National Accelerator Laboratory, Batavia, IL, Technical Report No. FNAL-TM-1988, 1996.
- [17] J. C. Wong, A. Shishlo, A. Aleksandrov, Y. Liu, and C. Long, 4D transverse phase space tomography of an operational hydrogen ion beam via noninvasive 2D measurements using laser wires, *Phys. Rev. Accel. Beams* **25**, 042801 (2022).
- [18] B. Marchetti *et al.*, Experimental demonstration of novel beam characterization using a polarizable X-band transverse deflection structure, *Sci. Rep.* **11**, 3560 (2021).
- [19] S. Jaster-Merz, R. Assmann, R. Brinkmann, F. Burkart, W. Hillert, M. Stanitzki, and T. Vinatier, 5D tomographic phase-space reconstruction of particle bunches, *Phys. Rev. Accel. Beams* **27**, 072801 (2024).
- [20] A. H. Andersen and A. C. Kak, Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm, *Ultrasonic Imaging* **6**, 81 (1984).
- [21] A. Hoover, K. Ruisard, A. Aleksandrov, A. Zhukov, and S. Cousineau, Analysis of a hadron beam in five-dimensional phase space, *Phys. Rev. Accel. Beams* **26**, 064202 (2023).
- [22] B. Cathey, S. Cousineau, A. Aleksandrov, and A. Zhukov, First six dimensional phase space measurement of an accelerator beam, *Phys. Rev. Lett.* **121**, 064804 (2018).
- [23] C. Emma, A. Edelen, M. Hogan, B. O'Shea, G. White, and V. Yakimenko, Machine learning-based longitudinal phase space prediction of particle accelerators, *Phys. Rev. Accel. Beams* **21**, 112802 (2018).
- [24] C. Emma, A. Edelen, A. Hanuka, B. O'shea, and A. Scheinker, Virtual diagnostic suite for electron beam prediction and control at FACET-II, *Information (Switzerland)* **12**, 61 (2021).
- [25] A. Hanuka, C. Emma, T. Maxwell, A. S. Fisher, B. Jacobson, M. J. Hogan, and Z. Huang, Accurate and confident prediction of electron beam longitudinal properties using spectral virtual diagnostics, *Sci. Rep.* **11**, 2945 (2021).
- [26] A. Wolski, M. A. Johnson, M. King, B. L. Militsyn, and P. H. Williams, Transverse phase space tomography in an accelerator test facility using image compression and machine learning, *Phys. Rev. Accel. Beams* **25**, 122803 (2022).
- [27] O. Convery, L. Smith, Y. Gal, and A. Hanuka, Uncertainty quantification for virtual diagnostic of particle accelerators, *Phys. Rev. Accel. Beams* **24**, 074602 (2021).
- [28] A. Scheinker, F. Cropp, S. Paiguá, and D. Filippetto, An adaptive approach to machine learning for compact particle accelerators, *Sci. Rep.* **11**, 19187 (2021).
- [29] A. Scheinker, F. Cropp, and D. Filippetto, Adaptive autoencoder latent space tuning for more robust machine learning beyond the training set for six-dimensional phase space diagnostics of a time-varying ultrafast electron-diffraction compact accelerator, *Phys. Rev. E* **107**, 045302 (2023).
- [30] A. Scheinker, cDVAE: Multimodal generative conditional diffusion guided by variational autoencoder latent embedding for virtual 6D phase space diagnostics, [arXiv:2407.20218](https://arxiv.org/abs/2407.20218).
- [31] M. Rautela, A. Williams, and A. Scheinker, Towards latent space evolution of spatiotemporal dynamics of six-dimensional phase space of charged particle beams, [arXiv:2406.01535](https://arxiv.org/abs/2406.01535).
- [32] R. Roussel, A. Edelen, C. Mayes, D. Ratner, J. P. Gonzalez-Aguilera, S. Kim, E. Wisniewski, and J. Power, Phase space reconstruction from accelerator beam measurements using neural networks and differentiable simulations, *Phys. Rev. Lett.* **130**, 145001 (2023).
- [33] Noting that a quadrupole scan does not provide information about the longitudinal phase space distribution.
- [34] S. Kim, J. P. Gonzalez-Aguilera, P. Piot, G. Chen, S. Doran, Y.-K. Kim, W. Liu, C. Whiteford, E. Wisniewski, A. Edelen, R. Roussel, and J. Power, Four-dimensional phase-space reconstruction of flat and magnetized beams using neural networks and differentiable simulations, *Phys. Rev. Accel. Beams* **27**, 074601 (2024).
- [35] G. Ha, M. Cho, W. Namkung, J. Power, D. Doran, E. Wisniewski, M. Conde, W. Gai, W. Liu, C. Whiteford, Q. Gao, K.-J. Kim, A. Zholents, Y.-E. Sun, C. Jing, and P. Piot, Precision control of the electron longitudinal bunch shape using an emittance-exchange beam line, *Phys. Rev. Lett.* **118**, 104801 (2017).
- [36] J.-P. Gonzalez-Aguilera, Y.-K. Kim, R. Roussel, A. Edelen, and C. Mayes, Towards fully differentiable accelerator modeling, in *Proceedings of the 14th International Particle Accelerator Conference, IPAC-2023, Venice, Italy (JACoW, Geneva, Switzerland, 2023)*, pp. 2797–2800.
- [37] T. Iqbal and S. Qureshi, The survey: Text generation models in deep learning, *J. King Saud Univ. Comput. Inform. Sci.* **34**, 2515 (2022).
- [38] C. G. Turhan and H. S. Bilge, Recent trends in deep generative models: A review, in *Proceedings of the 2018 3rd International Conference on Computer Science and Engineering (UBMK)* (IEEE, New York, 2018), pp. 574–579.
- [39] J. Ingraham, V. Garg, R. Barzilay, and T. Jaakkola, Generative models for graph-based protein design, in *Proceedings of the Advances in Neural Information Processing Systems 32* (2019), https://papers.nips.cc/paper_files/paper/2019/hash/f3a4ff4839c56a5f460c88cce3666a2b-Abstract.html.
- [40] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, NeRF: Representing scenes as neural radiance fields for view synthesis, [arXiv:2003.08934](https://arxiv.org/abs/2003.08934).
- [41] K. Hornik, M. Stinchcombe, and H. White, Multilayer feedforward networks are universal approximators, *Neural Netw.* **2**, 359 (1989).
- [42] R. Bellman, Dynamic programming, *Science* **153**, 34 (1966).
- [43] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller, Efficient BackProp, in *Neural Networks: Tricks of the Trade*, 2nd ed., edited by G. Montavon, G. B. Orr, and K.-R. Müller, Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2012), pp. 9–48.
- [44] T. Dorigo *et al.*, Toward the end-to-end optimization of particle physics instruments with differentiable programming, *Rev. Phys.* **10**, 100085 (2023).

- [45] B. Martin, Differential algebraic techniques, *Advances in Imaging and Electron Physics* (Elsevier, Amsterdam, 1999), Vol. 108, pp. 81–117.
- [46] L. Deniau and C. I. Tomoiagă, Generalised truncated power series algebra for fast particle accelerator transport maps, in *Proceedings of the 6th International Particle Accelerator Conference, IPAC2015, Richmond, VA (JACoW, Geneva, Switzerland, 2015)*, pp. 374–377.
- [47] J. Kaiser, C. Xu, A. Eichler, and A. S. Garcia, Bridging the gap between machine learning and particle accelerator physics with high-speed, differentiable simulations, *Phys. Rev. Accel. Beams* **27**, 054601 (2024).
- [48] A. Paszke *et al.*, PyTorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d. Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., Red Hook, NY, 2019), pp. 8024–8035.
- [49] S. Weglarczyk, Kernel density estimation and its application, in *Proceedings of the ITM Web of Conferences (EDP Sciences, Les Ulis Cedex France, 2018)*, Vol. 23, p. 00037.
- [50] A. Ivanov and I. Agapov, Physics-based deep neural networks for beam dynamics in charged particle accelerators, *Phys. Rev. Accel. Beams* **23**, 074601 (2020).
- [51] D. Ratner, F. Christie, J. Cryan, A. Edelen, A. Lutman, and X. Zhang, Recovering the phase and amplitude of X-ray FEL pulses using neural networks and differentiable models, *Opt. Express* **29**, 20336 (2021).
- [52] R. Watt and B. O’Shea, A differentiable simulation package for performing inference of synchrotron-radiation-based diagnostics, *J. Synchrotron Radiat.* **31**, 409 (2024).
- [53] Q. Gao, J. Shi, H. Chen, G. Ha, J. G. Power, M. Conde, and W. Gai, Single-shot wakefield measurement system, *Phys. Rev. Accel. Beams* **21**, 062801 (2018).
- [54] P. Piot *et al.*, Commissioning of a 1.3-GHz deflecting cavity for phase space exchange at the Argonne Wakefield Accelerator, in *Proc. IPAC’12, New Orleans, LA, USA (2012)*, paper THPPC031, pp. 3350–3352.
- [55] R. Roussel *et al.*, Bayesian optimization algorithms for accelerator physics, *Phys. Rev. Accel. Beams* **27**, 084801 (2024).
- [56] R. Roussel, C. Mayes, A. Edelen, and A. Bartnik, Xopt: A simplified framework for optimization of accelerator problems using advanced algorithms, in *Proceedings of the 14th International Particle Accelerator Conference, IPAC-2023, Venice, Italy (JACoW, Geneva, Switzerland, 2023)*, pp. 4796–4799.
- [57] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [58] M. Cornacchia and P. Emma, Transverse to longitudinal emittance exchange, *Phys. Rev. ST Accel. Beams* **5**, 084001 (2002).
- [59] R. Roussel, G. Andonian, W. Lynn, K. Sanwalka, R. Robles, C. Hansel, A. Deng, G. Lawler, J. Rosenzweig, G. Ha, J. Seok, J. Power, M. Conde, E. Wisniewski, D. Doran, and C. Whiteford, Single shot characterization of high transformer ratio wakefields in nonlinear plasma acceleration, *Phys. Rev. Lett.* **124**, 044802 (2020).
- [60] N. Majernik, G. Andonian, W. Lynn, S. Kim, C. Lorch, R. Roussel, S. Doran, E. Wisniewski, C. Whiteford, P. Piot, J. Power, and J. Rosenzweig, Beam shaping using an ultrahigh vacuum multileaf collimator and emittance exchange beamline, *Phys. Rev. Accel. Beams* **26**, 022801 (2023).
- [61] G. Ha, M. Cho, W. Gai, K.-J. Kim, W. Namkung, and J. Power, Perturbation-minimized triangular bunch for high-transformer ratio using a double dogleg emittance exchange beam line, *Phys. Rev. Accel. Beams* **19**, 121301 (2016).
- [62] R. J. Roussel *et al.*, Transformer ratio measurements from ramped beams in the plasma blowout regime using emittance exchange, in *Proc. IPAC’19, Melbourne, Australia, 2019 (2019)*, pp. 3778–3781, [10.18429/JACoW-IPAC2019-THPGW088](https://doi.org/10.18429/JACoW-IPAC2019-THPGW088).
- [63] J. G. Power and C. Jing, Temporal laser pulse shaping for RF photocathode guns: The cheap and easy way using UV birefringent crystals, *AIP Conf. Proc.* **1086**, 689 (2009).
- [64] K. Hock, M. Ibson, D. Holder, A. Wolski, and B. Muratori, Beam tomography in transverse normalised phase space, *Nucl. Instrum. Methods Phys. Res., Sect. A* **642**, 36 (2011).
- [65] R. Fiorito, A. Shkvarunets, D. Castronovo, M. Cornacchia, S. Di Mitri, R. Kishek, C. Tschalaer, and M. Veronese, Noninvasive emittance and energy spread monitor using optical synchrotron radiation, *Phys. Rev. ST Accel. Beams* **17**, 122803 (2014).
- [66] M. LaBerge, B. Bowers, Y.-Y. Chang, J. Cabadağ, A. Debus, A. Hannasch, R. Pausch, S. Schöbel, J. Tiebel, P. Ufer, A. Willmann, O. Zarini, R. Zgadzaj, A. Lumpkin, U. Schramm, A. Irman, and M. Downer, Revealing the three-dimensional structure of microbunched plasma-wakefield-accelerated electron beams, *Nat. Photonics*, **18**, 952 (2024).
- [67] I. Kobayev, S. J. D. Prince, and M. A. Brubaker, Normalizing flows: An introduction and review of current methods, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964 (2021).
- [68] A. Hoover and J. C. Wong, High-dimensional maximum-entropy phase space tomography using normalizing flows, *Phys. Rev. Res.* **6**, 033163 (2024).
- [69] https://github.com/roussel-ryan/gpsr_6d_paper