

Using convolutional neural networks to accelerate three-dimensional coherent synchrotron radiation computations

Christopher Leon^{✉,*}, Petr M. Anisimov[✉], Nikolai Yampolsky[✉], and Alexander Scheinker[✉]

Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA



(Received 5 January 2025; accepted 14 October 2025; published 4 November 2025)

Calculating the effects of coherent synchrotron radiation (CSR) is one of the most computationally expensive tasks in accelerator physics. Here, we use convolutional neural networks (CNNs), along with a latent conditional diffusion (LCD) model, trained on physics-based simulations to speed up calculations. Specifically, we produce the 3D CSR wakefields generated by electron bunches in circular orbit in the steady-state condition. Two datasets are used for training and testing the models: wakefields generated by three-dimensional Gaussian electron distributions and wakefields from a sum of up to 25 three-dimensional Gaussian distributions. The CNNs are able to accurately produce the 3D wakefields ~ 250 – 1000 times faster than the numerical calculations, while the LCD achieves a gain of a factor of ~ 34 . We also test the extrapolation and out-of-distribution generalization ability of the models. They generalize well on distributions with larger spreads than what they were trained on but struggle with smaller spreads.

DOI: [10.1103/2v1x-c97j](https://doi.org/10.1103/2v1x-c97j)

I. INTRODUCTION

Coherent synchrotron radiation (CSR) is a collective effect that can significantly distort an electron bunch undergoing circular motion. This occurs in the magnetic dipoles found in accelerator sections, such as the magnetic dipoles of storage rings and bunch compressors. Bunch lengths have continually shortened for applications like x-ray free-electron lasers (XFELs), which produce extremely bright x-ray pulses to probe matter. This has been tremendously useful for numerous fields, such as the study of protein structure, imaging of viruses, and investigations of matter at high energy density. This has led to the need to properly take into account the effects of CSR. However, calculations involving CSR can be extremely computationally expensive. Historically, 1D approximations have been widely used, but those are expected to break down eventually, specifically for large transverse-to-longitudinal ratios [1]. The need for quick, 3D calculations for CSR has become important to the accelerator community. The evolution of computational hardware along development of new techniques opens up opportunities to seek alternative routes to the traditional methods. In this work, we get the wakefields used for CSR from machine

learning (ML) surrogate models based on convolutional neural network (CNN) architectures, as well as a latent conditional diffusion model. These were all trained on the physics-based software PyCSR3D [2,3]. The goal is to create ML models that are accurate while running much faster than physics-based simulations.

A. The physics of CSR

To see the physical origins of the CSR phenomenon, first picture N_e electrons in a bunch of length σ undergoing circular motion in an external dipole magnetic field (see Fig. 1). For the radiation emitted by the electrons, consider the part of the spectrum where the wavelength, λ , is much larger than the bunch length: $\lambda \gg \sigma$. At λ resolution, the bunch looks pointlike, and thus the radiation emitted by the electrons at a given time is approximately in phase. This radiation adds up coherently and, given that in accelerators $N_e \sim 10^{10}$, the cumulative effect can be substantial. Even though the electrons are moving near the speed of light due to their small mass, they are traveling on a circular arc while the radiation travels linearly (see Fig. 1). This means it is possible for there to be significant self-interaction for the bunch, which results in serious distortions of its phase space distribution. This generally has the effect of the tail and center losing energy, the head gaining, and an unwanted increase in emittance (volume in phase space).

To calculate the effect of CSR, one can start with the electromagnetic (EM) fields generated by a single relativistic point particle undergoing arbitrary motion. These are the well-known Liénard–Wiechert (LW) fields, [4–6]:

*Contact author: cleon@lanl.gov

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.




FIG. 1. Three particles in a circular bend. Left: at time t , particle 1 absorbs radiation emitted by particles 2 and 3 at earlier points in time, t'_2 and t'_3 , respectively. Right: to get the total force a charged particle experiences at a particular point in space-time, $(t, \mathbf{r})$, one integrates the Lorentz force from the LW fields along the past light cone. This is called the wakefield, $\mathbf{W}(\mathbf{r}, t)$.

$$\mathbf{E} = -e \left(\underbrace{\frac{\mathbf{n} - \boldsymbol{\beta}}{\gamma^2(1 - \mathbf{n} \cdot \boldsymbol{\beta})^3 \rho^2}}_{\text{Coulomb term}} + \underbrace{\frac{\mathbf{n} \times (\mathbf{n} - \boldsymbol{\beta}) \times \dot{\boldsymbol{\beta}}}{c(1 - \mathbf{n} \cdot \boldsymbol{\beta})^3 \rho}}_{\text{Synchrotron radiation}} \right) \Big|_{ret.}, \quad (1)$$

$$\mathbf{B} = \mathbf{n} \times \mathbf{E}, \quad (2)$$

where Gaussian units are being used. Here, $\mathbf{E}$ is the electric field, $\mathbf{B}$ the magnetic field, $-e$ is the charge of the electron, $\mathbf{n}$ the normal vector from the source to the point of interest, $\boldsymbol{\beta}$ the velocity normalized by the speed of light, c , $\gamma = \frac{1}{\sqrt{1 - \beta^2}}$, and ρ is the distance between the source and the point of interest. As seen in Eq. (1), there are two terms: the Coulomb term, which gives the familiar Coulomb's law when the particle is at rest, and the synchrotron radiation term.

The *ret.* in Eq. (1) signifies that the expression must be evaluated at the retarded time,

$$t' = t - \rho/c, \quad (3)$$

to ensure causality. That is, the EM fields at a point in time depend on the motion of the electron at the time of emission (see Fig. 1).

From the LW fields, the wakefield, $\mathbf{W}(\mathbf{r}, t)$, the force a point particle experiences at a point in spacetime from the EM fields generated from other particles at earlier times, can be obtained. Getting the Lorentz force, $\mathbf{F}(\mathbf{r}, t) = -e[\mathbf{E}(\mathbf{r}, t) + \mathbf{v} \times \mathbf{B}(\mathbf{r}, t)]$, along with Eqs. (1) and (2), then integrate along the past light cone:

$$\mathbf{W}(\mathbf{r}, t) = \int_{\Delta} d^3\mathbf{r}' \lambda(\mathbf{r}', t') \mathbf{F}(\mathbf{r}' - \mathbf{r}, t'), \quad (4)$$

where $\lambda(\mathbf{r}', t')$ is the number density and Δ signifies the past light cone. We assume the radiation reaction force on the

emitting electron is negligible, and it will not be considered here (see Baez [7] for an overview of the thorny theoretical issues and review of the works of Abraham, Dirac, and others on the topic). Doing the calculation naively, one would go into the past light cone for each particle, sum up all LW fields produced, and use the Lorentz force to get the wakefield (see Fig. 1). This would be extremely inefficient, as it would require storing the position, velocity, and acceleration of each of the N_e particles for each time step, as well as doing computationally expensive $\mathcal{O}(N_e^2)$ calculations for each time step.

In addition, as outlined in Ref. [8], the spatial scaling of CSR goes like $\sim \frac{L}{\gamma^3}$. Thus, for the ultrarelativistic case, accurate simulations require very fine spatial and time resolution, which adds to the already significant computational burden.

Much of the previous work on CSR has focused on creating and implementing approximate models to make the challenging task of calculating its effects on electron beams computationally tractable.

B. Past work

Here, we provide a short review of past work done on CSR. For a comprehensive review of the CSR literature and the different techniques used, see Mayes [2].

Carlsten and Raubenheimer calculated the emittance growth in circular accelerators [9]. A small angle, ultra-relativistic, 1D model that treats the bunch as a line charge distribution, $\lambda_{1D}(s)$, and includes transient effects was developed by Saldin *et al.* [10]. It was later expanded in [8]. Like many examples of using only effective degrees of freedom, complications arose from infinities. The issue was dealt with by “renormalizing” the Coulomb term in Eq. (1); that is, by subtracting the field generated during uniform motion. In the 1D steady-state case, the longitudinal wakefield in Frenet-Serret coordinates (see Fig. 2) is then,




FIG. 2. Illustration of the Frenet-Serret coordinates for a position, $\mathbf{r}$. The reference path here is uniform circular motion with a circle of radius ρ in the horizontal plane.

$$W_s(s) = -\frac{e}{3^{1/3}\rho^{2/3}} \int_{-\infty}^s ds' \frac{\partial \lambda_{1D}(s')}{\partial s'} \frac{1}{(s-s')^{1/3}} \text{(1D theory)}, \quad (5)$$

with the transient case being a slight modification of Eq. (5). The commonly used software ELEGANT (ELEctron Generation AND Tracking) [11] uses this method to calculate the effects of CSR [12]. A bend is broken up into a series of segments, and the effect of Eq. (5) is calculated for each, with the approximation that the distribution used in the calculation is the one at that point.

In 2D and 3D, the infinity issue and the need for renormalization do not arise. A two-dimensional model of CSR making use of Gaussians was created, CSRTrack [13,14], which approximates a distribution as a sum of Gaussian sub-bunches. A 2D model of CSR was also created by Huang *et al.* [15]. LW3D is a result of Ryne *et al.* using the LW fields and a particle-based approach [16,17]. It has been found that the effects of the Coulomb term can be significant during transience. As the bunch enters the bend, a pancaked EM field, originating from the free relativistic electron bunch in the drift section, continues “moving” into the bend and can interact with the bunch inside [18,19].

Recently, in a series of papers [20–22], Cai created models for CSR in a bend utilizing the Hamiltonian formalism in Courant-Snyder theory [23]. This was first done for a 2D model for steady state [20], then with Ding for the 3D case in steady state [21], and finally in the 2D case with transient effects [22]. Applying the formalism to the first bunch compressor at the Linac Coherent Light Source at SLAC, Cai and Ding’s method accurately captured the change in the emittance. Furthermore, it predicted a change in the vertical emittance, whereas ELEGANT had constant vertical emittance during a bend [21].

For the problem of CSR, Frenet-Serret coordinates for a particle undergoing uniform circular motion in the horizontal plane are typically used. The longitudinal direction here is the circle arc length, s , the horizontal x corresponds to the radial distance from the reference circle, and the vertical is denoted y (see Fig. 2). In the Snyder-Courant approach [23,24], s becomes the “time” axis, and

derivatives with respect to s are denoted by a prime (e.g., x'). In this approach, the equations of motion for a single particle are [21]:

$$\delta' = \frac{r_e N_b}{\gamma} W_s, \quad (6)$$

$$x'' = \frac{r_e N_b}{\gamma} W_x, \quad (7)$$

$$y'' = \frac{r_e N_b}{\gamma} W_y, \quad (8)$$

where r_e is the classical electron radius ($r_e = \frac{e^2}{m_e c^2}$) and $\delta = (p - p_0)/p_0$, where p_0 is the reference momentum ($p_0 = \frac{e\rho B_0}{c}$). From Eqs. (7)–(9), one can see that the wakefields have units of $1/\text{m}^2$.

In Cai’s approach, shielding, where longer wavelengths are suppressed by the surrounding metallic beam pipe, was considered negligible in all cases. The paraxial approximation was also used in some instances to make the problem more tractable. To solve for the retarded time condition in Eq. (3), a fourth order approximation was used, which is computationally faster than a root finding algorithm since the fourth order solution is analytic, and it was shown to be quite accurate as well [20,21].

In practice, a combination of analytic models and numerical software can be used to minimize the effects of CSR. Xu *et al.* demonstrated the reduction of horizontal momentum shift in a bunch compressor for the proposed MaRIE FEL using a double chicane, where the momentum shifts produced by each chicane largely cancel each other out [25,26]. A coarse optimization was first done using a fast, simple analytical model, followed by refinement with a slower, more accurate simulation. The end result preserved the beam emittance well. The same approach can be used for general CSR applications, replacing the analytic model with an ML surrogate model.

C. Convolutional neural networks

CNNs, originally designed for image processing, use locality in the data to build higher-level features [27]. CNNs also possess the inductive bias of translation equivariance, which many physical systems possess; colloquially, “the laws of physics are the same everywhere.” Locality and translation equivariance make CNNs well-suited ML surrogate models for various physical systems.

While using ML surrogate models can have a significant upfront time cost for training the models on simulation data, once trained, evaluations can be performed extremely quickly. This speed arises because evaluations leverage matrix multiplication and parallelized operations optimized by modern GPUs. CNNs have been used elsewhere to replace computationally intensive physics-based simulations with ML surrogates. For example, in [28], 3D CNNs

for electrodynamics with hard physics constraints were first developed by creating neural operators that map charge and current densities directly to their associated scalar and vector potentials from which the electric and magnetic fields are generated according to Maxwell's equations. In [29], a Fourier-Helmholtz-Maxwell neural operator approach was developed, which incorporates almost all of Maxwell's equations as hard physics constraints, and can accurately simulate the electromagnetic fields generated by charged particle beams. Both of these approaches were shown to reduce evaluation time from approximately 2 days when using the General Particle Tracker simulations to just a few seconds for AI model-based inference.

Previous work has investigated the use of ML surrogate models to speed up CSR calculations. Mayes and co-workers used a dense neural network in the 1D case to speed up CSR calculations [30]. CNNs have been utilized before in the 2D transient case [31]. Gonzalez-Aguilera *et al.* used generative methods to reconstruct projections of phase space affected by CSR [32]. To our knowledge, this work is the first in-depth study employing CNN surrogate ML models to calculate the full 3D CSR wakefields.

II. TRAINING AND TEST DATA

We generated the training data for the surrogate ML model using the software package PyCSR3D [2,3], which produces the 3D wakefield of a given distribution undergoing circular motion in the steady-state case. The physics implemented in the software is based on the model by Cai and Ding [21] described in Sec. IB and includes the Coulomb term. PyCSR3D was chosen because it outperformed several other mesh-based methods in a study [2], specifically by most accurately reproducing the entire 3D wakefield, where ground truth was taken to be the results from the far more computationally expensive, particle-based LW3D [17]. Computationally, PyCSR3D uses the integrated Green function approach to get the transverse wakefields. Convolutional integrals are evaluated numerically by performing fast Fourier transforms, utilizing the convolution theorem, and then performing an inverse fast Fourier transform.

Two datasets were created: the Gaussian Dataset and the Multi-Gaussian Dataset. The first dataset is used to evaluate the ML-based approach relative to how simple Gaussian distribution-based models approximate real-world conditions. The second assesses the models against more complex distributions that may also show up. Both datasets consisted of bunches in a bend with $\rho = 1$ m and $\gamma = 500$. The data are on a uniformly spaced grid in Frenet-Serret coordinates. For both the Gaussian and Multi-Gaussian Dataset, the coordinates lie between $-48 \mu\text{m} \leq x_i \leq 48 \mu\text{m}$, with grid spacings of $0.75 \mu\text{m}$ for x, y and $0.375 \mu\text{m}$ for s . Other physical characteristics can be seen in Table I.

A. Gaussian Dataset

The Gaussian Dataset consisted of 500 Gaussian distributions centered at $\mathbf{0}$, along with their standard deviations and corresponding wakefields. Field configurations are represented as an array of $(128, 128, 256, n_F)$, where $n_F = 1$ ($n_F = 3$) for a scalar (vector) field. The third index represents the s component and has twice the resolution of the transverse components, represented by the first and second indices, as the longitudinal coordinate is usually the most important.

Starting with Gaussian distributions is useful because there are a number of analytic results from approximate models to compare with, making them good sanity checks on the data and models. Real beams can also sometimes approximate Gaussians and are often modeled as such.

B. Multi-Gaussian Dataset

Each distribution in the Multi-Gaussian Dataset is a sum of Gaussians, with the exact number, N_G , being chosen randomly and uniformly between 2 and 25 for each distribution. The total number of particles is the same as in the Gaussian Dataset, and each individual Gaussian has a $1/N_G$ fraction of the particles. There are a total of 500 examples of multi-Gaussian distribution/wake pairs.

The means and standard deviations of each individual Gaussian were chosen randomly and uniformly to be inside some bounds (see Table I for exact details). Approximating

TABLE I. Physical characteristics of the two datasets. Distributions are created by being uniformly sampled in the ranges here.

	Gaussian Dataset	Multi-Gaussian Dataset
Physical range	$x_i \in (-48, 48) \mu\text{m}$ $i = x, y, s$	Same as Gaussian Dataset
Distribution attributes	$\sigma_i \in (2, 12) \mu\text{m}$ $i = x, y, s$ Centered at $\mathbf{0}$	Sum of 2 to 25 Gaussians $\mu_i \in (-12, 12) \mu\text{m}$ $\sigma_i \in (2, 12) \mu\text{m}$
Mesh details	Dimensions: $128 \times 128 \times 256$ Spacing: $0.75 \mu\text{m}(x, y), 0.375 \mu\text{m}(s)$	Same as Gaussian Dataset

a distribution as a sum of Gaussians is commonly used in kernel density estimation. More physically, something like this is seen in microbunching.

For both the Gaussian and Multi-Gaussian sets, the scale of W_s was $\sim 10^6 \frac{1}{\text{m}^3}$, the scale of W_x about an order of magnitude smaller and that of W_y about 2 orders of magnitude smaller. To avoid W_s dominating completely when training, both the Gaussian and Multi-Gaussian Datasets had each component normalized separately, by the means and standard deviations of the components of the W_i in the Gaussian Dataset.

III. MODELS

There is significant flexibility in choosing the ML architecture and the generality of the models. To explore performance across different cases, we consider three machine learning architectures. The σ -NN is the simplest of the three with only a decoder CNN architecture, while also having the most restrictive use case of just Gaussian distributions. The λ -NN uses a modified U-Net architecture, which is computationally slower than σ -NN, but allows for more general applications. Finally, the LCD uses a state-of-the-art latent diffusion and has a similar general use case as the λ -NN. The following subsections describe the models in more detail.

A. Sigma model

The σ -NN model produces the wakefield of Gaussian density distributions given the spreads in x , y , and s :

$$\sigma\text{-NN}: (\sigma_x, \sigma_y, \sigma_s) \rightarrow \mathbf{W}(\mathbf{r}). \quad (9)$$

In terms of arrays of the dataset, $(3) \rightarrow (128, 128, 256, 3)$.

The σ -NN has a decoder CNN architecture (see Fig. 3). The model has the built-in bias that the three σ_i parameters

describe the entire wakefield. This limits its applicability, but it is expected to perform better when the bunch is Gaussian. Gaussian-produced wakefields can be useful beyond Gaussian distributions. This has been demonstrated by CSRTTrack [13], which approximates a particle distribution with a superposition of multiple Gaussian sub-bunches; the exact parameters, such as charge and number of sub-bunches, are chosen by the user. The produced wakefields are then approximated as a summation of the wakefields of the Gaussians, leading to a computationally efficient means of producing general wakefields.

B. Lambda model

The λ -NN model is more general and produces the wakefield given an arbitrary 3D distribution, $\lambda(\mathbf{r})$,

$$\lambda\text{-NN}: \lambda(\mathbf{r}) \rightarrow \mathbf{W}(\mathbf{r}). \quad (10)$$

In terms of arrays, $(128, 128, 256, 1) \rightarrow (128, 128, 256, 3)$.

A U-Net architecture [33] was chosen for the λ -NN model (see Fig. 4). This builds on a CNN encoder-decoder architecture with the addition of skip connections, which concatenate layers in the encoder to same sized layers in the decoder. “Blurriness” in CNN encoder-decoder architectures is a well-known phenomenon and likely stems from the fact that much of the finer detailed information in the original input gets lost during the compression into the latent space. The skipped connections in the U-Net fix this by having the decoder also pull information from the encoder in a multiscale fashion (see Fig. 4). This gives the decoder more information while still forcing the network to identify important, salient features during the compression down to the latent space. In addition, the skipped connections also offer a benefit during the training process, as they offer a “gradient superhighway” that




FIG. 3. σ -NN architecture—A dense network is used to create a 1024 dimensional vector, which is reshaped to a $(8, 8, 16)$ array. The output is processed through several convolutional and convolutional transpose block layers. Numbers over blocks represent the number of filters.




FIG. 4. λ -NN architecture (a) A U-Net architecture [33] constructed out of residual convolution and attention blocks. A down/up sampling layer is implied for every time the block dimensions change. F is a parameter here, which was taken as $F = 4$. (b) These blocks are of a deep residual structure [34]. The input gets sent to through two paths: one part goes through several layers, while the other passes through as many $1 \times 1 \times 1$ convolutions needed to make the dimensions the same for addition. The two outputs are then summed. F here is the number of filters. (c) The Attention Blocks use the attention mechanism [35] on the feature maps. For a single input, the query, key, and value vectors, of dimension F , are obtained by multiplying the input by the matrices W^Q , W^K , and W^V , respectively. A weighted average of the value vector is taken, with the weights acquired through the use of softmax on the query and key vectors. This gets added to the group normed input. In practice, several inputs are fed simultaneously in order to train more efficiently.

allows for better optimization during gradient descent, which is particularly beneficial for parameters that would otherwise be far removed from the output.

We follow Ref. [36] using a U-Net with Res-Net block modules, which is based on PixelCNN++ [37] with group normalization [38]. This modifies the U-Net with the addition of five attention layers—two in the encoder, two within the decoder, and one in the latent space. These attention layers utilize the self-attention mechanism, a fundamental aspect of transformers in large language models [35]. These layers help the model learn global correlations using higher-level features.

We first trained λ -NN on the Gaussian Dataset, which we will refer to as λ -NN-Gauss. It was taken as a pretrained model and then trained on the Multi-Gaussian Dataset. This saved time compared to training from scratch and allowed us to explore how well learning transferred from going from one dataset to another. The final resulting model is labeled λ -NN-Multi.

C. Latent conditional diffusion model

Diffusion models are frequently used in modern image generation [36] and are employed in technologies such as DALL-E and Stable Diffusion [39]. They have also been

used in accelerator physics as generative virtual beam diagnostics to generate the phase space distribution of particle beams from limited data [40,41] and have been found to have connections to the stochastic quantization of quantum lattice field theories [42]. Inspired by nonequilibrium thermodynamics [43], these models draw analogies to fluid systems, which start from low entropy and evolve to pure randomness. The data are corrupted by the successive addition of noise until the image becomes effectively just pure Gaussian noise. The machine learning algorithm learns to denoise images step by step so that, starting from pure randomness, an image can be generated.

Specifically, in the forward process of diffusion, we start with an “image” of $\mathbf{Z}^0$, then define a Markov chain sequence as follows:

$$\mathbf{Z}^t = \sqrt{1 - \beta_t} \mathbf{Z}^{t-1} + \sqrt{\beta_t} \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(\mathbf{0}, I), \quad (11)$$

where $\{\beta_t\}_{t=1}^T$ is a chosen “noise schedule” and $0 < \beta_t < 1$. Noise comes to dominate in the sense that $\lim_{t \rightarrow \infty} \mathbf{Z}^t \sim \mathcal{N}(\mathbf{0}, I)$. Given the analytical properties of multidimensional normal distributions, this makes the diffusion process mathematically tractable and computationally efficient. Subsequently, the network is trained with the backward diffusion process, learning to denoise by

predicting $\mathbf{Z}^{t-1}$ given $\mathbf{Z}^t$. After training, starting from Gaussian noise, $\mathbf{Z}^0 \sim \mathcal{N}(\mathbf{0}, I)$, the model iteratively denoises to produce an image. This approach transforms the problem of sampling from a difficult, complex target distribution, such as 128×128 pixel images of cats, to a simpler one of sampling from a known prior distribution (Gaussian noise).

Latent diffusion [39] was specifically developed for cases when the data are high dimensional (e.g., high resolution images), and it becomes more computationally efficient to operate in a lower dimensional latent space representation. Conditioning is done by specifying aspects of the “image.” For example, using the text “cat” to generate an image of a cat. Here, the $\rho(\mathbf{r})$ field is compressed down to a vector $\mathbf{Z}_\rho$, using a variational autoencoder (VAE) for conditioning. Since the field configuration for a single component of $W^i(\mathbf{r})$ are for us represented by large $(128, 128, 256, 1)$ arrays, we will also work in the latent space. The structure and regularity in the data mean it can be described by far fewer than $128 \times 128 \times 256 \sim 4 \times 10^6$ parameters. This is accomplished with another convolutional autoencoder, which compresses the wakefield down to an $(8, 8, 16, 4)$ array, a compression of $\sim 0.1\%$. Although in the Gaussian case, the data can be compressed to the three σ_i parameters, we keep the latent space larger in order to better represent a more general case. We also condition based on the “time” of the diffusion, latent density distribution vector $\mathbf{Z}_\rho$, and the component we wish to focus on, i , which is represented as a one hot encoding vector. This way, all three W^i components can be calculated in parallel, but also only one or two components can be obtained if one wishes to use a 1D or 2D approximation.

In summary, the procedure of the LCD model is processed through three submodels (see Fig. 5): (i) VAE: compress $\rho(\mathbf{r})$ with the encoder of a VAE:

$$\rho(\mathbf{r}) \rightarrow \mathbf{Z}_\rho.$$

(ii) *Conditional U-Net with attention layers*: Starting with pure noise in the latent space of the wakefields, the model iteratively denoises using conditioning on time, $\mathbf{Z}_\rho$, and component i :

$$\mathbf{Z}_{W^i}^0 \sim \mathcal{N}(\mathbf{0}, I)$$

$$\text{for } t \in \{0, 1, \dots, T-1\}:$$

$$\mathbf{Z}_{W^i}^t | (t, \mathbf{Z}_\rho, i) \rightarrow \mathbf{Z}_{W^i}^{t+1}.$$

(iii) *Wake decoder*: A previously trained decoder is used to map the wakefield latent space and produce an estimate of the wakefield:

$$\hat{\mathbf{Z}}_{W^i}^T \rightarrow \hat{W}^i(\mathbf{r}).$$

For the conditional latent diffusion architecture, we also follow Ref. [36] for the denoising model, which uses a U-Net with Res-Net block modules. This model incorporates seven attention layers—three within the encoder, three within the decoder, and one in the latent space. For time embedding, we use the transformer sinusoidal position embedding [35]. The noise schedule was $\beta_1 = 10^{-4}$, $\beta_T = 0.5$, and $T = 12$. These hyperparameters were obtained after some experimentation. Smaller values of T led to a significant degradation in reconstruction, while larger T significantly increased the running time without notable improvement. In fact, at sufficiently large T , it actually produced worse results. We found these hyperparameters were a good balance of speed and accuracy.




FIG. 5. The conditional latent diffusion model.

D. Architecture hyperparameters and training

Many of the hyperparameters in the architectures and training were chosen to be powers of 2. This is recommended by Nvidia [44] for deep learning in order to most efficiently utilize its GPUs, which rely on dividing calculation operations into many parallelized processes. Filter sizes of $3 \times 3 \times 3$ are used in all the convolutional layers, as small filter sizes with more layers can provide the same coverage as larger filter sizes with fewer layers, but with fewer parameters. To smooth the predicted fields, we added a final averaging layer with a $3 \times 3 \times 3$ window and then trained for a few more epochs for σ -NN and λ -NN-Gauss. We found this significantly improved results.

For the σ and λ models, leaky ReLUs [45] with $\alpha = 0.1$ were used for the activation function, as they have the same advantages as the more commonly used ReLU but suffer less from the vanishing gradient problem. The LCD model used the Swish activation function [46].

For training of all the models, the Adam optimizer [47] was used, and the mean squared error (MSE) was chosen for the loss function. Both datasets had a 85/15% training and test split. The learning rate was decreased during training to “turn down the noise” while optimizing. That is, there is high noise in the beginning to encourage greater exploration of parameter space and avoid getting stuck in local stationary points, while the noise is decreased toward the end as the best local set of parameters is refined. There are parallels here to simulated annealing, where the temperature acts as the source of the noise and is decreased during optimization to find the global minimum [48]. TensorFlow [49] was used for the creation of the neural networks and their training. All computations were done on a workstation with an Intel Xeon Platinum 8268 processor at 2.90 GHz and two NVIDIA RTX A6000 GPUs with 48 GB of RAM.

IV. RESULTS

To evaluate the models, we used the relative weighted mean average error (RW-MAE) for each component $i = x, y, s$:

$$\text{RW-MAE}_i = \frac{\int d^3\mathbf{r} \lambda(\mathbf{r}) |W_{\text{pred}}^i(\mathbf{r}) - W_{\text{true}}^i(\mathbf{r})|}{\int d^3\mathbf{r} \lambda(\mathbf{r}) |W_{\text{true}}^i(\mathbf{r})|}, \quad (12)$$

where $W_{\text{true}}^i(\mathbf{r})$ and $W_{\text{pred}}^i(\mathbf{r})$ are the true and predicted wakefields, respectively. The idea behind RW-MAE is that we care most about being correct where the particle density is highest, since $\mathbf{W}(\mathbf{r})$ is used to update the phase space distribution. Moreover, it differs significantly enough from the MSE loss function used in training to serve as a complementary metric.

A. Gaussian Dataset

We compare the results of the ML models to the approximate 1D theory [10], which can be obtained by using a Gaussian distribution in Eq. (5):

$$W_s(q) = \frac{2^{11/6}}{3^{7/6} \rho^{2/3} \sigma_s^{4/3}} \left[\frac{\sqrt{\pi} q {}_1F_1\left(\frac{7}{6}, \frac{3}{2}, -\frac{q^2}{2}\right)}{\sqrt{3} \Gamma\left(\frac{5}{3}\right)} - \frac{2^{5/6} \Gamma\left(\frac{2}{3}\right) {}_1F_1\left(\frac{2}{3}, \frac{1}{2}, -\frac{q^2}{2}\right)}{\Gamma\left(\frac{7}{3}\right)} \right], \quad (13)$$

where $q = z/\sigma_z$ and ${}_1F_1(a, b, z)$ is a confluence hypergeometric function of the first kind. While the networks do not produce a perfect match, both were generally closer to the W_s produced by PyCSR3D than the 1D theory frequently used in practice, as can be seen in Fig. 10.

For the horizontal wake in the 1D approximation, various authors have expressed $W_x(z) = \frac{\Lambda}{\rho} \lambda(z)$, with different sources arriving at different values for Λ , depending on the assumptions used. For a Gaussian distribution:

$$W_x(z) = \frac{\Lambda}{\rho} \frac{1}{\sqrt{2\pi\sigma_s^2}} e^{-z^2/2\sigma_s^2}, \quad (14)$$

While some authors assume constant values, such as $\Lambda = 2$ and $\Lambda = 4$ [21,50], we found that

$$\Lambda = \log \left[\left(\frac{\rho \sigma_s^2}{\sigma_\perp^3} \right)^{2/3} \left(1 + \frac{\sigma_\perp}{\sigma_s} \right) \right], \quad (15)$$

from Ref. [51] with $\sigma_\perp = \sqrt{\sigma_x^2 + \sigma_y^2}$ worked the best at approximating the PyCSR3D wakefields. Again, the predictions of σ -NN and λ -NN were both generally closer to PyCSR3D than the 1D theory, an example of which can be seen in Fig. 6. It was observed that the tails generally drop much slower than Eq. (15) suggests.

As shown in Table II, the σ -NN model and λ -NN-Gauss performed comparably across the RW-MAE_{*i*} metric. Examples of the models performance can be seen in Figs. 7 and 8. Thus, despite σ -NN having an easier task of just predicting the wakefields from the given σ_i 's, its accuracy and that of λ -NN were similar. The addition of attention layers and residual connections was invaluable, as testing showed that a simpler CNN architecture without these features increased the error of λ -NN by a factor of ~ 3 . However, the performance of λ -NN-Gauss comes at the cost of taking ~ 4 times longer to evaluate than σ -NN (see Table III), likely due to the greater model complexity.

The LCD performed worse than the other two ML models, both in terms of accuracy and evaluation time. Examples of how it works can be seen in Fig. 9. The accuracy result is likely due to having more sources of potential error: encoding the distribution, running the backward diffusion and then decoding to get $W^i(\mathbf{r})$. Similarly, the use of several networks and having to perform backward diffusion make the process slower than the other models.

Nonetheless, all the machine learning models achieved their predictions much faster than PyCSR3D. For the time to




FIG. 6. Gaussian Dataset: The performance of σ -NN and λ -NN-Gauss on a distribution in the test set. This is along the z axis ($x = y = 0 \mu\text{m}$). The electron distribution is shown in the shaded blue area. The 1D theory for W_s is given by Eq. (5), while for W_x it uses Eqs. (14) and (15).

predict $\mathbf{W}$ from a single distribution, σ -NN achieved a speedup of ~ 1000 , λ -NN-Gauss ~ 250 , and LCD ~ 34 (see Table III). The most computational time-consuming part for PyCSR3D was calculating grid values for the Green's functions.

B. Multi-Gaussian Dataset

The λ -NN-Gauss model was able to qualitatively predict the wakes in the Multi-Gaussian Dataset, as can be seen in Fig. 10, despite only being trained on single Gaussian




FIG. 7. Gaussian Dataset: σ -NN performance on a distribution in the test set in the $y = 0 \text{ } \mu\text{m}$ plane. Wake units: $10^6/\text{m}^2$.

distributions. Given this result, researchers using CNN models for CSR calculations can benefit greatly from pretraining their models on easily producible, large datasets using analytic results from approximate models,

before fine-tuning them on more realistic and general data from numerically intensive simulations. Since much learning transferred, training for λ -NN-Multi took half the time of training λ -NN-Gauss from scratch. Nonetheless,




FIG. 8. Gaussian Dataset: λ -NN-Gauss performance on a distribution in the test set in the $y = 0 \text{ } \mu\text{m}$ plane. Wake units: $10^6/\text{m}^2$.

λ -NN-Multi benefited greatly from training on Multi-Gaussian Dataset, as can be seen in Fig. 11.

Evaluation times for the λ -NN-Multi and PyCSR3D were similar to what was found with the Gaussian Dataset in Table III.

The performance of σ -NN and λ -NN-Multi can be seen in Table IV. The σ -NN model was fed the σ_i 's of the various Gaussians and then its wake predictions were centered on the appropriate μ_i 's.




FIG. 9. Gaussian Dataset: Latent conditional diffusion. Top: a cross section of latent space during the denoising process at certain time steps. Bottom: a cross section of the corresponding wakefield during denoising. Wake units: $10^6/\text{m}^2$.

C. Extrapolation and out-of-distribution generalization

To see how well λ -NN generalizes to distributions very different from those in the training and test data, we performed three different tests.

1. Gaussian extrapolation

First, we evaluated how well λ -NN-Gauss and λ -NN-Multi extrapolate by using Gaussian distributions with widths smaller or larger than those the models had previously encountered. Since our lattice spacing for

TABLE II. Gaussian Dataset: The RW-MAE_i over the training and test set.

Model	Set	s mean (STD), %	x mean (STD), %	y mean (STD), %
σ -NN	Training	1.24 (0.54)	0.59 (0.10)	1.09 (0.18)
	Test	1.26 (0.65)	0.60 (0.13)	1.12 (0.17)
λ -NN-Gauss	Training	1.39 (1.18)	0.48 (0.33)	0.97 (0.62)
	Test	1.43 (1.33)	0.49 (0.39)	0.97 (0.61)
LCD	Training	6.59 (4.32)	2.78 (1.82)	5.35 (3.27)
	Test	6.68 (4.22)	2.93 (1.92)	5.43 (2.45)

TABLE III. Gaussian Dataset: Time to predict $\mathbf{W}(\mathbf{r})$ for a single distribution, from highest to lowest. The mean and standard deviation are derived from seven trials.

Model	Evaluation time mean (STD), seconds
PyCSR3D	99.2 (2.2)
LCD	2.94 (0.09)
λ -NN-Gauss	0.410 (0.051)
σ -NN	0.106 (0.002)




FIG. 10. Multi-Gaussian dataset: The performance of λ -NN-Gauss and λ -NN-Multi on a distribution in the test set. This is along the z axis ($x = y = 0 \mu\text{m}$), with the electron distribution along the axis shown in the shaded blue area.




FIG. 11. Multi-Gaussian Dataset: λ -NN-Multi performance on a distribution in the testset on the $y = 0 \mu\text{m}$ plane. Wake units: $10^6/\text{m}^2$.

TABLE IV. Multi-Gaussian Dataset: The average of RW-MAE_i over the training and test set.

Model	Set	s mean (STD), %	x mean (STD), %	y mean (STD), %
σ -NN	Training + test	5.13 (1.91)	3.71 (0.81)	6.16 (1.68)
	Training	3.60 (1.51)	0.96 (0.41)	2.56 (0.76)
	Test	3.60 (1.51)	0.99 (0.56)	2.74 (0.92)

the transverse (longitudinal) component is $0.75 \mu\text{m}$ ($0.375 \mu\text{m}$) and the lowest width in the training data is $\sigma_i = 2 \mu\text{m}$, it would be difficult to tell if failure to extrapolate to smaller σ_i stems from the model being unable to extrapolate or from the small number of samples due to the lattice spacing. Looking over the wakefields, the results were not especially noisy, and we ruled out this coming from the ultrarelativistic limit discussed earlier. We therefore modified the training and test data to the range $6 \leq \sigma_i \leq 12 \mu\text{m}$ ($i = x, y, s$) for both the Gaussian and Multi-Gaussian datasets, keeping the number of samples in each set the same. New λ -NNs were then trained from scratch on the modified datasets, using the same procedure as before.

Next, a new dataset was created, which consisted of Gaussians with spreads that went outside the training and test set. Specifically, the longitudinal spread went from $\sigma_s = 1$ to $18 \mu\text{m}$, while the transverse spreads were kept fixed at $\sigma_x = \sigma_y = 9 \mu\text{m}$.

As can be seen in Fig. 12, the λ -NN-Multi model performed well at extrapolating to larger σ_s with the RW-MAE error for each component increasing only slowly outside training and test data. In contrast, the predictions for smaller σ_s did much worse and rose rapidly as σ_s was decreased. This occurred well before σ_s approached the z lattice length of $0.375 \mu\text{m}$, suggesting that the decline was due to the model's failure to extrapolate to lower σ_s .




FIG. 12. Extrapolation to Gaussians with different σ_s from the Gaussian training/test data.

2. Generalization: Smoothed uniform cubical distribution

A 1D uniform distribution between two points, a and b , can be considered the difference of two step functions, each centered at a and b . This can be smoothed by replacing the step function with sigmoids: $f(x) = \frac{1}{b-a} \left(\frac{1}{1+e^{x-a}} - \frac{1}{1+e^{x-b}} \right)$. We call a smoothed uniform cubical distribution one that uses $f(x)$ for all three components with the same widths (note: it is not exactly a cube in physical space due to curvilinear coordinates being used, but due to its small size relative to ρ , it is only approximately cubical). Smoothing improves the numerics in the simulation, which depend on derivatives in the longitudinal component. It also provides a more realistic smooth drop-off in density for particle beams compared to a discontinuous one. While the sampling in the Multi-Gaussian Dataset was performed from a cubical distribution, the relatively low number of samples, 25, the smoothed edges, and the varying cube length make the smoothed uniform cubical distributions quite different from those the model had seen.

The λ -NN-Multi model from the previous subsection was used. Going from side lengths of $l = 4 \mu\text{m}$ to $l = 50 \mu\text{m}$, the RW-MAE can be seen in Fig. 13. The results are similar to the Gaussian case, where the models fail at small widths, but become almost level at larger ones. Based on this and the previous results, it seems the λ -NN-Multi should be applied only to distributions with widths larger than $\sim 10 \mu\text{m}$.




FIG. 13. Smooth uniform cube—RW-MAE error for each component as a function of cubic side length, l .




FIG. 14. 1D profile of double horn distribution.

3. Generalization: Double horn distribution

Nonlinearities on the longitudinal phase space can lead to the undesirable formation of a “double horn” on distributions after bunch compressors. This is especially important to XFELs [52,53].

We tested λ -NN-Multi on a double horn distribution. Its 1D longitudinal profile can be seen in Fig. 14, with two peaks at the front and back of the distribution. It was formed from two Gaussians and the smoothed rectangular function, $f(x)$, used for the




FIG. 15. Double horn distribution: Performance of λ -NN-Multi on a double horn distribution. Wake units: $10^6/\text{m}^2$.

smoothed uniform cubical distribution. For the transverse part of the distribution, it was made from a 2D Gaussian with $\sigma_x = \sigma_y = 9 \mu\text{m}$ with means $\mu_x = \mu_y = 0 \mu\text{m}$.

The results can be seen in Fig. 15. Qualitatively, it matches fairly well and even predicts the large wakes produced from the “horns” at around $z = \pm 18 \mu\text{m}$ for W_s , which could be predicted by using Eq. (5) and approximating the horns with a Dirac delta. Its performance comes close to how it does on the Multi-Gaussian Dataset, as seen in Figs. 16 and 17. One can compare this to the smoothed uniform cubical distribution, which lacks the horns, in Figs. 18 and 19 in the Appendix. It did, however, noticeably worse on W_s compared to the other two components. This matches with the relative weighted MAE error metric, which were 16.8%, 3.1%, and 5.6% for W_s , W_x , and W_y , respectively. Since W_s depends on the z derivative of the distribution, the sudden drop that was quite different from what the model was trained on had the most effect on this component. Finding other realistic, non-Gaussian distributions encountered in accelerators to train on can further improve the model.

V. CONCLUSION

We have created surrogate ML models to find the 3D CSR wakes. For the λ -NN and σ -NN models, at the cost of a percent level error, we achieve a increase in speed ranging from ~ 250 and ~ 1000 , respectively.

Now that it has been shown these ML models have the ability to produce the wakefields accurately and quickly, it would be interesting to apply to real accelerators. One

interesting case would be emittance exchangers, where CSR plays a large role.

A long-term goal would be to make a computationally efficient, self-consistent general CSR code for the full 6D phase space for arbitrary motion.

A number of modifications can be made to make the models to further improve performance and applicability of the models. One issue is that by using a CNN architecture, we have introduced a mesh dependence. Finding ways to embed the data, such as is done in Fourier neural operators [54], can get rid of the mesh dependence. We have neglected shielding effects from the surrounding metallic beam pipe. In some cases, these may be important and future work can look into including these. Transience has also been neglected here. The 2D theory already exists [22] and ML has been applied to that case [31], so a fully 3D transient case would be the next goal.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX: PERFORMANCE OF λ -NN-MULTI ON DIFFERENT DISTRIBUTIONS

To further illustrate the performance of λ -NN-Multi, this section shows how its predictions compared to ground truth on the more exotic distributions.




FIG. 16. Multi-Gaussian Dataset: λ -NN-Multi performance on a distribution in the test set on the $y = 0 \mu\text{m}$ plane. Contour lines in red on the left column plots show the distribution at different level sets. Wake units: $10^6/\text{m}^2$.




FIG. 17. Multi-Gaussian Dataset: λ -NN-Multi performance on a distribution in the test set on the $y = 0 \mu\text{m}$ plane. Contour lines in red on the left column plot show the distribution at different level sets. Wake units: $10^6/\text{m}^2$.




FIG. 18. Smoothed uniform cube: λ -NN-Multi performance on a smooth uniform cube distribution in the $y = 12.47 \mu\text{m}$ plane. Wake units: $10^6/\text{m}^2$.




FIG. 19. Smoothed uniform cube: λ -NN-Multi performance on a smooth uniform cube distribution, larger than in Fig. 18, in the $y = 24.57 \mu\text{m}$ plane. Wake units: $10^6/\text{m}^2$.

- [1] Y. S. Derbenev, J. Rossbach, E. L. Saldin, and V. D. Shiltsev, Microbunch radiative tail—head interaction, DESY, Report No. PRINT-98-023.
- [2] C. E. Mayes, Computational approaches to coherent synchrotron radiation in two and three dimensions, *J. Instrum.* **16**, P10010 (2021).
- [3] Christopher Mayes, ChristopherMayes/PyCSR3D: PyCSR3D version 0.2.0, Zenodo, 2021, [10.5281/zenodo.5496096](https://doi.org/10.5281/zenodo.5496096).
- [4] Alfred Liénard, *Champ électrique et Magnétique Produit Par Une Charge électrique Concentrée en Un Point et Animée d'un Mouvement Quelconque* (G. Carré et C. Naud, Paris, France, 1898).
- [5] Emil Wiechert, Elektrodynamische elementargesetze, *Ann. Phys. (Berlin)*, **309**, 667 (1901).
- [6] John David Jackson, *Classical Electrodynamics* (American Association of Physics Teachers, New York, NY, 1999).
- [7] J. C. Baez, Struggles with the continuum, [arXiv:1609.01421](https://arxiv.org/abs/1609.01421).
- [8] G. Stupakov and P. Emma, CSR wake for a short magnet in ultrarelativistic limit, in *Proceedings of the 8th European Particle Accelerator Conference, Paris, France* (2002), Report No. SLAC-PUB-9242.
- [9] B. E. Carlsten and T. O. Raubenheimer, Emittance growth of bunched beams in bends, *Phys. Rev. E* **51**, 1453 (1995).
- [10] E. L. Saldin, E. A. Schneidmiller, and M. V. Yurkov, On the coherent radiation of an electron bunch moving in an arc of a circle, *Nucl. Instrum. Methods Phys. Res., Sect. A* **398**, 373 (1997).
- [11] Michael Borland, ELEGANT: A flexible SDDS-compliant code for accelerator simulation, Technical Report, Argonne National Laboratory, Argonne, IL, 2000, [10.2172/761286](https://doi.org/10.2172/761286).
- [12] M. Borland, Simple method for particle tracking with coherent synchrotron radiation, *Phys. Rev. ST Accel. Beams* **4**, 070701 (2001).
- [13] M. Dohlus and T. Limberg, CSRTrack: Faster calculation of 3-D CSR effects, in *Proceedings of the 2004 FEL Conference* (2004), <https://proceedings.jacow.org/f04/papers/MOCOS05/MOCOS05.PDF>.
- [14] M. Dohlus and T. Limberg, *CSRTrack Version 1.2 User's Manual* (DESY, Hamburg, Germany, 2007).
- [15] C. Huang, T. J. T. Kwan, and B. E. Carlsten, Two dimensional model for coherent synchrotron radiation, *Phys. Rev. ST Accel. Beams* **16**, 010701 (2013).
- [16] C. E. Mitchell, J. Qiang, R. D. Ryne, B. E. Carlsten, and N. A. Yampolsky, Large-scale simulation of synchrotron radiation using a Lienard-Wiechert approach, *Conf. Proc. C* **1205201**, 1689 (2012), <https://proceedings.jacow.org/IPAC2012/papers/tuppp036.pdf?n=IPAC2012/papers/TUPPP036.pdf>.
- [17] R. Ryne, B. Carlsten, C. Mitchell, and J. Qiang, Self-consistent modeling using a Lienard-Wiechert particle-mesh method, in *Proceedings of the 9th International Particle Accelerator Conference (IPAC2018), Vancouver, BC, Canada* (JACoW, Geneva, Switzerland, 2018), [10.18429/JACoW-IPAC2018-THPAK044](https://doi.org/10.18429/JACoW-IPAC2018-THPAK044).
- [18] Alexander Novokhatski, *Coherent Synchrotron Radiation: Theory and Simulations* (Citeseer, Menlo Park, CA, 2012).
- [19] A. D. Brynes, P. Smorenburg, I. Akkermans, E. Allaria, L. Badano, S. Brussaard, M. Danailov, A. Demidovich, G. De Ninno, D. Gauthier *et al.*, Beyond the limits of 1D coherent synchrotron radiation, *New J. Phys.* **20**, 073035 (2018).
- [20] Y. Cai, Coherent synchrotron radiation by electrons moving on circular orbits, *Phys. Rev. Accel. Beams* **20**, 064402 (2017).
- [21] Y. Cai and Y. Ding, Three-dimensional effects of coherent synchrotron radiation by electrons in a bunch compressor, *Phys. Rev. Accel. Beams* **23**, 014402 (2020).
- [22] Y. Cai, Two-dimensional theory of coherent synchrotron radiation with transient effects, *Phys. Rev. Accel. Beams* **24**, 064402 (2021).
- [23] E. D. Courant and H. S. Snyder, Theory of the alternating gradient synchrotron, *Ann. Phys. (N.Y.)* **3**, 1 (1958).
- [24] G. Stupakov and G. Penn, *Classical Mechanics and Electromagnetism in Accelerator Physics* (Springer, New York, 2018), ISBN 978-3-319-90187-9, [10.1007/978-3-319-90188-6](https://doi.org/10.1007/978-3-319-90188-6).
- [25] H. Xu, P. M. Anisimov, L. D. Duffy, and Q. R. Marksteiner, Comparison study on first bunch compressor schemes by conventional and double C-chicane for MaRIE XFEL, in *Proceedings of the 5th International Particle Accelerator Conference (NAPAC2022), Albuquerque, NM, USA* (JACoW, Geneva, Switzerland, 2022), TUPA72, [10.18429/JACoW-NAPAC2022-TUPA72](https://doi.org/10.18429/JACoW-NAPAC2022-TUPA72).
- [26] Haoran Xu, Petr M. Anisimov, Bruce E. Carlsten, Leanne D. Duffy, Quinn R. Marksteiner, and River R. Robles, X-ray free electron laser accelerator lattice design using laser-assisted bunch compression, *Appl. Sci.* **13**, 2285 (2023).
- [27] Yann LeCun, Bernhard Boser, John S. Denker, Donnie Henderson, Richard E. Howard, Wayne Hubbard, and Lawrence D. Jackel, Backpropagation applied to handwritten zip code recognition, in *Neural Computation* (MIT Press, Cambridge, MA, 1989), Vol. 1, pp. 541–551.
- [28] A. Scheinker and R. Pokharel, Physics-constrained 3D convolutional neural networks for electrodynamics, *APL Mach. Learn.* **1**, 026109 (2023).
- [29] C. Leon and A. Scheinker, Physics-constrained machine learning for electrodynamics without gauge ambiguity based on Fourier transformed Maxwell's equations, *Sci. Rep.* **14**, 14809 (2024).
- [30] A. Edelen, C. Emma, C. Mayes, and R. Roussel, Neural network solver for coherent synchrotron radiation wake-field calculations in accelerator-based charged particle beams, in *Proceedings of the 13th International Particle Accelerator Conference (IPAC2022), Bangkok, Thailand* (JACoW, Geneva, Switzerland, 2022), WEPOMS013, [10.18429/JACoW-IPAC2022-WEPOMS013](https://doi.org/10.18429/JACoW-IPAC2022-WEPOMS013).
- [31] R. Robles, A. Edelen, W. Lou, C. Mayes, and R. Roussel, Efficient computation of two-dimensional coherent synchrotron radiation with neural networks, in *Proceedings of the 14th International Particle Accelerator Conference (IPAC2023), Venice, Italy* (JACoW, Geneva, Switzerland, 2023), WEP105, [10.18429/JACoW-IPAC2023-WEP105](https://doi.org/10.18429/JACoW-IPAC2023-WEP105).
- [32] J. P. Gonzalez-Aguilera, R. Roussel, A. Edelen, and Y. K. Kim, Detailed characterization of coherent synchrotron radiation effects using generative phase space reconstruction, in *Proceedings of the 15th International Particle Accelerator Conference (IPAC2024), Nashville,*

- TV (JACoW, Geneva, Switzerland, 2024), WEPG94, [10.18429/JACoW-IPAC2024-WEPG94](https://doi.org/10.18429/JACoW-IPAC2024-WEPG94).
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, U-Net: Convolutional networks for biomedical image segmentation, in *Proceedings of the Part III Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany* (Springer, New York, 2015), pp. 234–241.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA* (2016), [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90).
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, Attention is all you need, [arXiv:1706.03762](https://arxiv.org/abs/1706.03762).
- [36] Jonathan Ho, Ajay Jain, and Pieter Abbeel, Denoising diffusion probabilistic models, *Adv. Neural Inf. Process. Syst.* **33**, 6840 (2020), <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>.
- [37] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma, PixelCNN++: Improving the pixelCNN with discretized logistic mixture likelihood and other modifications, [arXiv:1701.05517](https://arxiv.org/abs/1701.05517).
- [38] Yuxin Wu and Kaiming He, Group normalization, in *Proceedings of the European Conference on Computer Vision (ECCV)* (2018), pp. 3–19, https://openaccess.thecvf.com/content_ECCV_2018/html/Yuxin_Wu_Group_Normalization_ECCV_2018_paper.html.
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, High-resolution image synthesis with latent diffusion models, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 10684–10695, https://openaccess.thecvf.com/content/CVPR2022/html/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper.html.
- [40] A. Scheinker, cDVAE: VAE-guided diffusion for particle accelerator beam 6D phase space projection diagnostics, *Sci. Rep.* **14**, 29303 (2024).
- [41] A. Scheinker, Conditional guided generative diffusion for particle accelerator beam diagnostics, *Sci. Rep.* **14**, 19210 (2024).
- [42] L. Wang, G. Aarts, and K. Zhou, Diffusion models as stochastic quantization in lattice field theory, *J. High Energy Phys.* **05** (2024) 060.
- [43] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, [arXiv:1503.03585](https://arxiv.org/abs/1503.03585).
- [44] Nvidia, Get started with deep learning performance, <https://docs.nvidia.com/deeplearning/performance/dl-performance-getting-started/index.html> [Accessed: July 19, 2024].
- [45] A. L. Maas, A. Y. Hannun, A. Y. Ng *et al.*, Rectifier nonlinearities improve neural network acoustic models, in *Proceedings of the ICML, Atlanta, GA* (2013), Vol. 30, p. 3, https://robotics.stanford.edu/~amaas/papers/relu_hybrid_icml2013_final.pdf.
- [46] P. Ramachandran, B. Zoph, and Q. V. Le, Searching for activation functions, [arXiv:1710.05941](https://arxiv.org/abs/1710.05941).
- [47] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [48] S. L. Smith, P. J. Kindermans, C. Ying, and Q. V. Le, Don't decay the learning rate, increase the batch size, [arXiv:1711.00489](https://arxiv.org/abs/1711.00489).
- [49] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, <https://tensorflow.org> (2015).
- [50] Y. S. Derbenev, J. Rossbach, E. L. Saldin, and V. D. Shiltsev, Microbunch radiative tail—head interaction, DESY, Report No. TESLA-FEL 1995-05, 1995, [10.3204/PUBDB-2018-04128](https://pubdb-2018-04128.slac.stanford.edu/).
- [51] Y. S. Derbenev and V. D. Shiltsev, Transverse effects of microbunch radiative interaction, SLAC National Accelerator Laboratory (SLAC), Reports No. SLAC-PUB-7181, No. FERMILAB-TM-1974, [10.2172/251656](https://pubdb-2018-04128.slac.stanford.edu/).
- [52] J. Wu, P. Emma, and Z. Huang, Coherent synchrotron radiation as a diagnostic tool for the LCLS longitudinal feedback system, *Conf. Proc. C* **0505161**, 428 (2005), <https://proceedings.jacow.org/p05/PAPERS/TOAC005.PDF>.
- [53] J. Seok, M. Chung, M. Conde, G. Ha, and J. Power, Double-horn suppression in EEX based bunch compression, in *Proceedings of the 4th North American Particle Accelerator Conference, Lansing, MI, USA* (JACoW, Geneva, Switzerland, 2019), [10.18429/JACoW-NAPAC2019-TUPLM16](https://doi.org/10.18429/JACoW-NAPAC2019-TUPLM16).
- [54] Z. Li, N. B. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, Fourier neural operator for parametric partial differential equations, in *Proceedings of the International Conference on Learning Representations* (2021), <https://iclr.cc/virtual/2021/poster/3281>.