

Adaptive machine learning for time-varying systems: low dimensional latent space tuning

A. Scheinker

*Los Alamos National Laboratory,
Los Alamos, NM, U.S.A.*

E-mail: ascheink@lanl.gov

ABSTRACT: Machine learning (ML) tools such as encoder-decoder convolutional neural networks (CNN) can represent incredibly complex nonlinear functions which map between combinations of images and scalars. For example, CNNs can be used to map combinations of accelerator parameters and images which are 2D projections of the 6D phase space distributions of charged particle beams as they are transported between various particle accelerator locations. Despite their strengths, applying ML to time-varying systems, or systems with shifting distributions, is an open problem, especially for large systems for which collecting new data for re-training is impractical or interrupts operations. Particle accelerators are one example of large time-varying systems for which collecting detailed training data requires lengthy dedicated beam measurements which may no longer be available during regular operations. We present a novel method of adaptive ML for time-varying systems. Our approach is to map very high ($N \approx 100k$) dimensional inputs (a combination of scalar parameters and images) into the low dimensional ($N \approx 2$) latent space at the output of the encoder section of an encoder-decoder CNN. We then actively tune the low dimensional latent space-based representation of complex system dynamics by the addition of an adaptively tuned feedback vector directly before the decoder sections builds back up to our image-based high-dimensional phase space density representations. This method allows us to learn correlations within and to quickly tune the characteristics of incredibly large parameter space systems and to track their evolution in real time based on feedback without massive new data sets for re-training. We demonstrate that our method can accurately predict and track the phase space of charged particle beams at various locations in a particle accelerator by adaptively adjusting in real-time while the unknown input beam distribution of the accelerator is changing in shape, charge, and offset and while the RF system of the accelerator itself is also changing in an unpredictable way. For FACET-II we demonstrate that such an approach has the potential to use transverse deflecting cavity and energy spread spectrum beam measurements to accurately predict 2D projections of the 6D phase space of the electron beam at the plasma wakefield acceleration interaction point where such diagnostics are unavailable.

KEYWORDS: Accelerator Applications; Analysis and statistical methods; Beam dynamics; Data reduction methods



Contents

1	Introduction	1
1.1	Limitations of machine learning for time-varying systems	2
1.2	Adaptive feedback for time-varying systems	2
1.3	Adaptive machine learning for time-varying systems	2
1.4	Summary of main results: adaptive latent space tuning	3
2	Robust feedback control for particle accelerators	4
3	Adaptive machine learning: latent space tuning	7
4	Application for particle accelerator diagnostics	9
4.1	Latent space tuning for HiRES	10
4.2	Latent space tuning for FACET-II	11
4.3	2D latent space	12
4.4	Latent space-informed diagnostics	13
4.5	$ND, N > 2$ latent space	15
5	Conclusions	17

1 Introduction

New advanced particle accelerators are being designed and operated with the goal of finer control over the phase space of the accelerated beams than what has previously been achieved. At advanced light sources such as the Linac Coherent Light Source (LCLS) upgrade, LCLS-II two color mode operation is being developed in which the energies of two bunches separated by tens of femtoseconds are precisely controlled to generate x-rays with precisely tuned energy differences [1]. At the Swiss-FEL a two color mode is also being developed and advanced nonlinear compression techniques have enabled the production of attosecond hard x-ray FEL pulses [2]. The EuXFEL is utilizing advanced superconducting radio frequency (RF) cavities and low level RF controls to accelerate bunches at MHz repetition rates [3]. The plasma wakefield acceleration facility for advanced accelerator experimental tests (FACET) upgrade, FACET-II, is planning on providing custom tailored current profiles with fs-long bunches and hundreds to thousands of kA peak currents [4]. At CERN's Advanced Proton Driven Plasma Wake-field Acceleration Experiment (AWAKE) facility the goal is to use transversely focused high intensity (3×10^{11}), high energy (400 GeV) protons from CERN's Super Proton Synchrotron (SPS) accelerator to create a 10 meter long plasma and wake-fields into which 18.8 MeV 650 pC electron bunches will be injected and accelerated up to 2 GeV [5].

The control of intense charged particle beams is challenging due to complex collective effects such as space charge forces and coherent synchrotron radiation (CSR) which severely distort the

6D phase space (x, y, z, p_x, p_y, p_z) of intense bunches, especially during bunch compression. One of the main challenges of beam control is a lack of detailed real-time phase space diagnostics. To achieve more precise control of intense charged particle beams requires the development of new and higher resolution non-invasive beam diagnostics as well as new advanced adaptive control methods which utilize the data from these diagnostics. The vast majority of online real-time bunch-to-bunch beam measurements provide scalar data such as beam position centroids from beam position monitors or beam averaged charge from current monitors.

1.1 Limitations of machine learning for time-varying systems

Machine learning (ML) tools are being developed that can learn representations of complex accelerator dynamics directly from data. ML methods have been utilized to develop surrogate models and virtual diagnostics [6–8], neural networks are being used to represent cost functions or optimal policies in reinforcement learning [9], powerful polynomial chaos expansion-based surrogate models have been used for uncertainty quantification [10], an interesting analysis of the latent space of neural network-based surrogate models has been used for uncertainty quantification [11], convolutional neural networks have been used for time-series classification and forecasting in accelerators [12], neural network-based surrogate models have been used to speed up optimization [13, 14], Bayesian Gaussian processes utilize learned correlations in data/physics-informed kernels [15–22], and various ML methods have been used for beam dynamics studies at CERN [23–25].

One limitation of ML methods, and an active area of research in the ML community, is the problem of time-varying systems, known as distribution shift. If a system changes with time then the data that was used to train an ML-based tool will no longer provide an accurate representation of the system of interest, and the accuracy of the ML tool will degrade. Distribution drift is a challenge for all ML methods including neural networks for surrogate models, the use of neural networks to represent cost functions or optimal policies in reinforcement learning, and even for methods such as Gaussian processes which utilize learned correlations in their kernels. Incorporating methods to deal with distribution shift is a major need for the accelerator community because accelerators and their beams change unpredictably with time.

1.2 Adaptive feedback for time-varying systems

In order to handle time-varying systems one must utilize feedback which enables real-time adjustments based on measurements. Powerful model-independent feedback control techniques have been developed to automatically compensate for un-modeled disturbances and time-varying system dynamics for unknown, time-varying, high dimensional systems [26]. The limitation of most adaptive feedback methods is that they are local in nature and therefore can become stuck in local minima or may converge very slowly if their initial conditions are chosen extremely far away from the optimum in a large parameter space.

1.3 Adaptive machine learning for time-varying systems

Efforts have begun to combine the robustness of adaptive feedback with the global representations that can be learned with ML methods to develop adaptive machine learning (AML) for time-varying systems. The first such result combined neural networks and model-independent feedback to provide

fast, robust, and automatic control over the energy vs time phase space of electron bunches in the LCLS [27]. Recently, AML methods have been studied in more generality for adaptive tuning of the inputs and outputs of ML tools such as neural networks for time-varying systems [28].

1.4 Summary of main results: adaptive latent space tuning

Adaptive latent space tuning maps high dimension inputs down to extremely low dimensional latent space representations of generative encoder-decoder style convolutional neural networks (CNN) [29, 30]. Encoder-decoder CNNs are powerful ML tools that can represent complex relationships in high-dimensional systems in a low dimensional latent space [31, 32], and have been used for anomaly detection [33], for time-series data [34], for optimization of deep generative models [35]. For particle accelerators encoder-decoder CNNs have been used as surrogate models to generate longitudinal phase space (LPS) distributions [36]. Recently encoder-decoder CNNs have also been demonstrated with measured beam data at the European XFEL to provide extremely high accuracy (768×1064 pixel images) predictions of the beam’s LPS and have also demonstrated an innovative method in which once the decoder half is trained and fixed, multiple different encoders can be used for various working points without having to re-train the decoder [37].

In this paper we demonstrate an approach in which an encoder-decoder CNN’s input \mathbf{X}_{in} is a combination of a 7-dimensional vector representing particle accelerator components as well as a 128×128 pixel image which is the $\rho(x, y)$ phase space density of an electron beam entering the FACET-II accelerator. The overall input dimension of this system, $N_{in} = 16391$, is then mapped down to a latent space dimension as low as $N_L = 2$ by the encoder half of the CNN before the generative half builds back up to a 128 × 128 × 75 pixel output \mathbf{Y}_{out} , which is a representation of the beam ($N_{out} \approx 1.2e6$) which are all 15 of the unique 2D 128×128 pixel projections of the 6D phase space of the beam at 5 different locations within the FACET-II accelerator, as shown in figure 1.

By simultaneously predicting 75 phase space projections we are effectively forcing the CNN to satisfy a large number of physics constraints. This low-dimensional latent space representation can be quickly adaptively tuned based only on a comparison of a single one of the 75 output distributions, the position vs energy longitudinal phase space prediction which is available for measurement at many accelerators via transverse deflecting cavities. The network is first trained in a supervised learning approach using the input-output pairs $(\mathbf{X}_{in}, \mathbf{Y}_{out})$. After the network has been trained to learn the physics constraints within the system and the correlations between various parameters, it functions in an un-supervised feedback fashion without access to the inputs \mathbf{X}_{in} . For feedback, a subset of the output $Y_i \in \mathbf{Y}_{out}$ is compared to a measurement \hat{Y}_i which guides feedback in the latent space layer without need for measurements of an input beam distribution or knowledge of accelerator parameters which may drift. Because the network has learned to respect constraints and correlations it is able to predict the entire 6D phase space of the beam at 5 different locations. The feedback compensates for unknown time-varying input beam distributions and accelerator parameters. The low-dimensional latent space representation also informs the placement of limited numbers of diagnostics. The main results are summarized as:

1. An encoder-decoder generative convolutional neural network is trained using a supervised learning approach. The CNN takes as input the initial $\rho(x, y)$ input beam distribution and 7 accelerator parameters which are beam x and y offsets, charge, and the amplitudes and

phases of the first 2 linear acceleration sections of FACET-II. The CNN pinches down to a low dimensional latent space before the generative half of the CNN outputs all 15 unique projections of the beam's 6D phase space at 5 different accelerator locations as in figure 1.

2. Once the CNN is trained, we operate in an adaptive un-supervised manner in which we no longer need access to the 7 accelerator parameters or to the initial beam distribution. Our adaptive approach compares the CNN's prediction of the position vs energy (z, E) phase space to transverse deflecting cavity measurements and adaptively tunes the latent space in order to track unknown time-varying beams and accelerator parameters.
3. The CNN is able to predict all of the 2D phase space projections because it has learned the physics relationships and correlations within the system.

2 Robust feedback control for particle accelerators

Particle accelerators are large systems with thousands of coupled components and complex charged particle dynamics. Accelerators and beams are time varying due to external disturbances, misalignments, and time variation of both the accelerator beam sources as well as of the RF and magnet components such as phase shifts in RF cables. The time variation and complexity of accelerator systems results in time varying input-output relationships, where we refer to inputs as component settings, such as magnet power supply currents or RF accelerating voltages, and outputs as beam characteristics such as beam spill and the beam's 6D phase space distribution. Because of the complexity and variability, the performance of standard model-based controls is limited and drifts with time. Although individual components can be kept at their desired set points very precisely, the overall system dynamics drifts as it depends on the relationships between various components and the beam. Model-independent feedback methods are a major subject of research in control theory where robustness to un-modeled disturbances and changes to system dynamics is an important consideration. For example, consider the scalar linear system

$$\dot{x}(t) = ax(t) + bu(t), \quad (2.1)$$

where the values of a and b are unknown. Such a system cannot be stabilized with simple proportional integral derivative (PID)-type feedback, but if the sign of b is known, for example if $b > 0$, a stable equilibrium of (2.1) can be established at $x = 0$ by the following nonlinear controller

$$u(t) = \theta(t)x(t), \quad \dot{\theta}(t) = -kx^2(t), \quad k > 0. \quad (2.2)$$

This approach is robust as it does not depend on a detailed knowledge of system dynamics, which is important in practice because all models are only approximations. However, two major limitations of such an approach are: 1). The sign of the unknown term b must be known and cannot be time-varying. 2). The presence of an arbitrarily small un-modeled disturbance in the dynamics (2.1) can destabilize the closed loop nonlinear system. Over the last decades, many more sophisticated nonlinear and adaptive feedback control algorithms have been developed which are robust to a wide range of un-modeled disturbances and are applicable to time-varying systems [38–40].

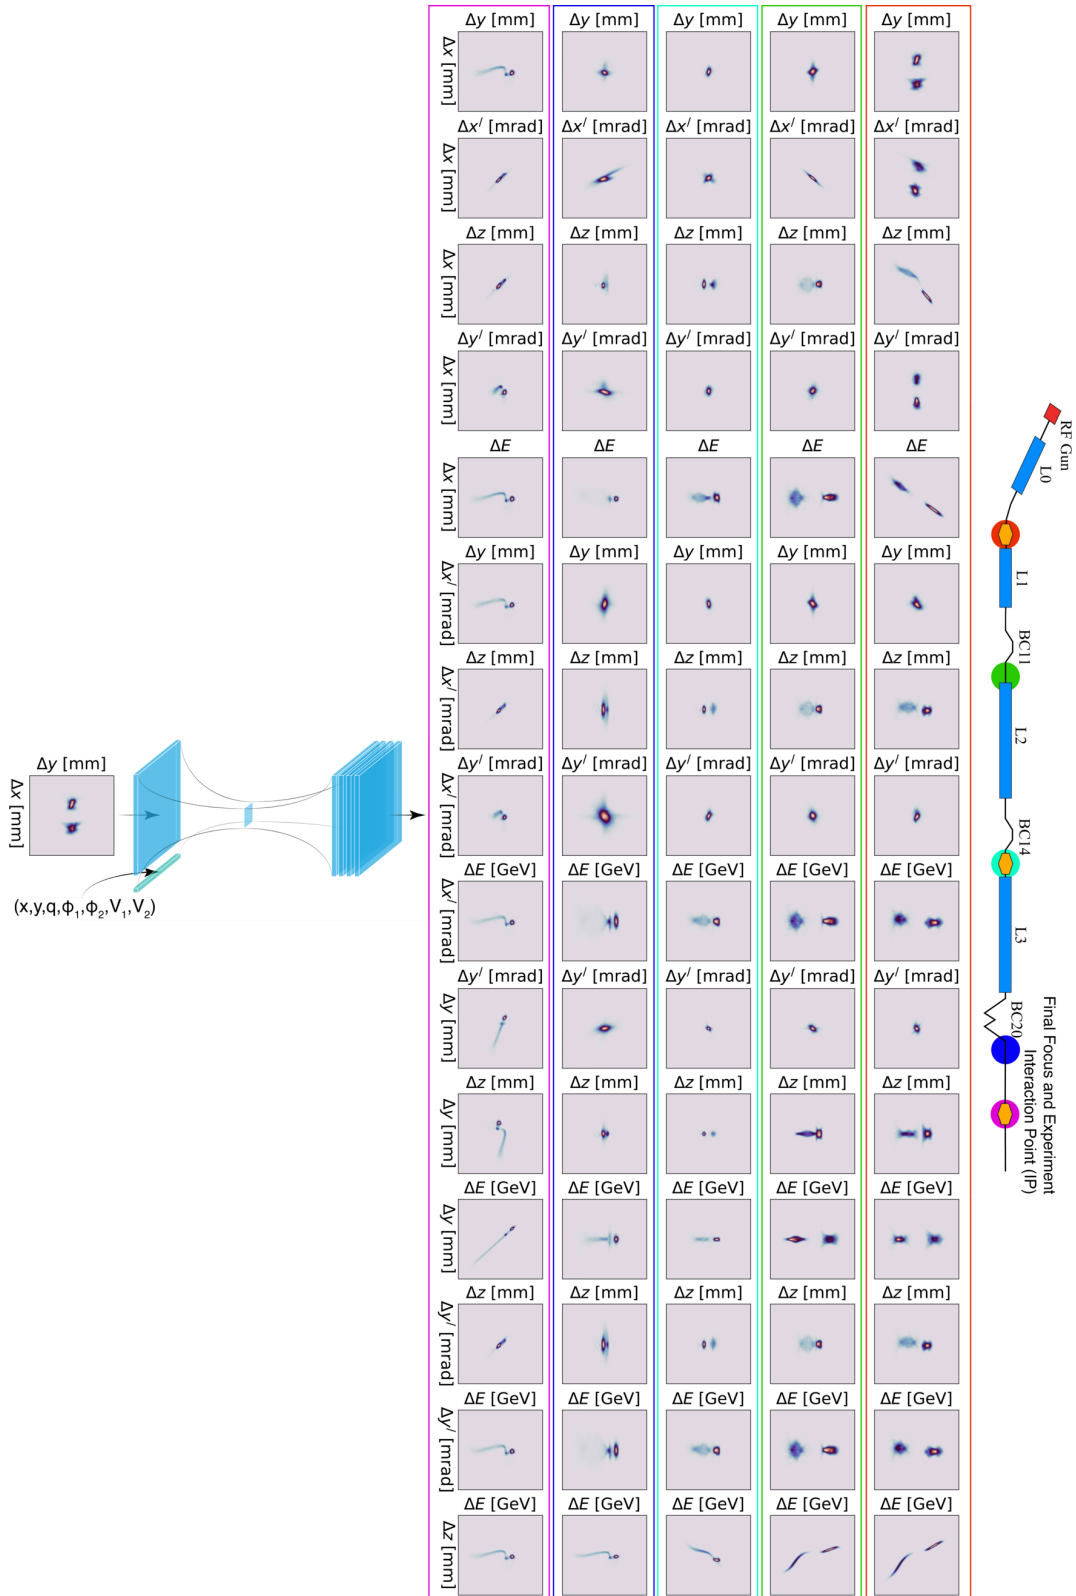


Figure 1. One example of a CNN output of 15 phase space projections at the 5 FACET-II locations grouped by colors which are highlighted by the circles on the FACET-II schematic on the right side of the image.

A limitation of nonlinear and adaptive control approaches was an inability to handle a sign-changing time varying coefficient $b(t)$ in system (2.1) which multiplies the control input $u(t)$, such as $b(t) = \cos(2\pi ft)$ which changes sign repeatedly thereby changing the effect of control input $u(t)$. For particle accelerators such variation is important. For example, consider a state $x(t)$ which describes beam loss in a particle accelerator, whose minimization is desired, which is influenced by a large collection of quadrupole magnets $\mathbf{u} = (u_1, \dots, u_m)$. The effect of a single magnet, u_m , depends on the initial beam's phase space as it enters the accelerator from the source and also on the settings of all of the other quadrupole magnets that are upstream, $u_{i < m}$ and changes with time as the upstream magnets are adjusted and as the initial beam conditions change. One day decreasing x may require decreasing the current of magnet u_m and another day it might have to be increased.

Recently, a nonlinear extremum seeking (ES) feedback control method was developed which could stabilize and minimize the analytically unknown outputs of a wide range of dynamics systems, scalar and vector-valued which can be time-varying, nonlinear and open loop unstable with unknown control directions [26, 41]. The ES method is applicable to a wide range of systems of the form

$$\begin{aligned} \dot{x} &= a(t)x(t) + b(t)u(\hat{y}(t)), & \dot{\mathbf{x}} &= A(t)\mathbf{x}(t) + B(t)\mathbf{u}(\hat{y}(t)), \\ \dot{\hat{y}} &= \mathbf{f}(\mathbf{x}(t), \mathbf{u}(\hat{y}(t)), t), & \hat{y}(\mathbf{x}, t) &= y(\mathbf{x}, t) + n(t), \end{aligned} \quad (2.3)$$

which include scalar time-varying linear systems, vector-valued time-varying linear systems, and vector-valued nonlinear time-varying systems, where in each case the feedback control u is based only on a noise-corrupted measurement $\hat{y}(t)$ of an analytically unknown cost function $y(\mathbf{x}, t)$. For example, a measurable but analytically unknown cost function can be the sum of beam loss along a many kilometer long particle accelerator, which depends on all accelerator parameters and on the initial 6D phase space of the beam being accelerated.

Relative to accelerator applications, the ES method is used to tune groups of accelerator parameters, $\mathbf{p} = (p_1, \dots, p_m)$, which may be a subset of the state space \mathbf{x} which describes the overall accelerator dynamics. For example, tuned parameters might include RF cavity amplitude and phase set points as well as magnet power supply voltages or currents. The adaptive ES algorithm dynamically tunes parameters according to

$$\dot{p}_j = \psi_j(\omega_j t + k\hat{y}(\mathbf{x}, t)), \quad (2.4)$$

where ω_i are distinct dithering frequencies defined as $\omega_i = \omega r_i$ with $r_i \neq r_j$ for $i \neq j$, k is a feedback gain. The ψ_j may be chosen from a large class of functions which may be non-differentiable and not even continuous, such as square waves which are easily implemented in digital systems [42]. The only requirements on the ψ_j are that for a given time interval $[0, t]$ they are measurable with respect to the L^2 norm and that they are mutually orthogonal in Hilbert space in the weak sense relative to all measurable functions $f(t) \in L^2[0, t]$ in the limit as $\omega \rightarrow \infty$, which can be written as

$$\begin{aligned} \lim_{\omega \rightarrow \infty} \int_0^t \psi_i(\tau)\psi_j(\tau)d\tau &= 0, \quad \forall i \neq j, & \lim_{\omega \rightarrow \infty} \int_0^t \psi_i(\tau)f(\tau)d\tau &= 0, \quad \forall i, \quad \forall f(t) \in L^2[0, t], \\ \lim_{\omega \rightarrow \infty} \int_0^t \psi_i^2(\tau)f(\tau)d\tau &= \int_0^t c_i f(\tau)d\tau, \quad \forall i, \quad \forall f(t) \in L^2[0, t], & c_i &> 0. \end{aligned}$$

One particular implementation of the ES method is especially convenient for particle accelerator applications because the tuning functions ψ_i have analytically guaranteed bounds despite acting on

analytically unknown and noisy functions, which guaranteed known update rates and limits on all tuned parameters. This bounded form of ES was first developed for accelerator tuning [43], and then studied for dynamic systems in general [44] and takes the simple form

$$u_i = \sqrt{\alpha_i \omega_i} \cos(\omega_i t + k \hat{y}(\mathbf{x}, t)). \quad (2.5)$$

The utility of this approach is clearly demonstrated by considering a system of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(\hat{y}(t)), t), \quad \dot{p}_i = \sqrt{\alpha_i \omega_i} \cos(\omega_i t + k \hat{y}(\mathbf{x}, t)), \quad (2.6)$$

which results in average dynamics that minimize the noise-corrupted unknown function $y(\mathbf{x}, t)$:

$$\dot{p}_i = -\frac{k\alpha}{2} \frac{\partial y(\mathbf{x}, t)}{\partial p_i}. \quad (2.7)$$

3 Adaptive machine learning: latent space tuning

Convolutional neural networks are some of the most powerful machine learning tools which regularly outperform all other methods in complex high-dimensional tasks such as image classification where, for example, using 100×100 pixel photos results in a 10^4 dimensional input space [45]. The main challenge in utilizing ML approaches for particle accelerator beams is the fact that accelerator components and beam initial conditions change unpredictably with time degrading the accuracy of any data-based ML approach to virtual diagnostics.




Figure 2. An encoder-decoder AML setup is shown in (A) where the inputs are stacks of images \mathbf{X}_{nD} , vectors of parameter values \mathbf{x}_{1D} , and an adaptively tuned latent space input vector \mathbf{y}_A , while the output is another stack of images \mathbf{Y}_{mD} . A simplified diagram of the same setup is shown in (B). For adaptive tuning, a measurement of the network's output, \mathbf{Y}_m , is compared to an available measurement \mathbf{M} , based on which the adaptive update \mathbf{y}_A is created to adjust the latent space representation \mathbf{y}_L .

A new adaptive machine learning (AML) approach for time varying systems has recently been proposed in which model-independent adaptive feedback is used to tune the low dimensional latent space of encoder-decoder architecture convolutional neural networks (CNN) for systems that change with time [29, 30]. The encoder half of the CNN can use both images (initial beam distributions) and scalars (accelerator magnet and RF parameter settings) as inputs with the scalars being concatenated together with the information from the images once the images have been scaled down using 2D convolutional layers, flattened, and passed through dense fully connected layers.

The low dimensional latent space at the bottleneck of the encoder CNN is a general nonlinear transformation in which a trained CNN has found an incredibly efficient basis with which to represent

beam properties, which may outperform purely linear dimensionality reduction techniques such as principal component analysis [46]. The CNN’s decoder half can be used to generated predictions of beam phase space distributions as a function of the low-dimensional latent space whose parameters are easily adjustable due to the much lower dimension of the problem which may be orders of magnitude smaller than the many thousands of dimensions associated with large input images.

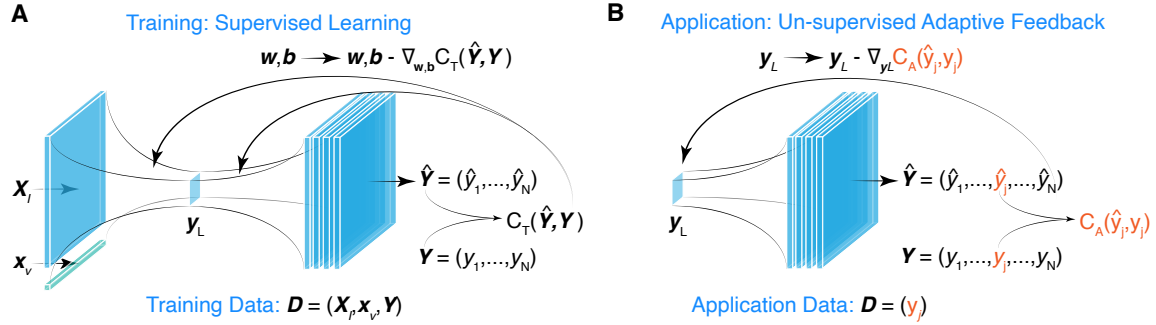


Figure 3. An overview of the adaptive machine learning latent space tuning approach.

An overview of the adaptive latent space approach is shown in figure 3. During the initial training phase a supervised learning approach is taken in which the network uses input-output pairs $(\mathbf{X}_I, \mathbf{x}_v, \mathbf{Y})$, where \mathbf{X}_I may be collections of images which represent 2D projections of a particle accelerator beam’s 6D phase space and the \mathbf{x}_v are vectors of parameters such as magnet or RF system settings. The CNN’s prediction $\hat{\mathbf{Y}}$ of phase space projections is then compared to their true values \mathbf{Y} so that a loss function $C_T(\hat{\mathbf{Y}}, \mathbf{Y})$ is minimized by adjusting the CNN weights and biases (\mathbf{w}, \mathbf{b}) in a supervised training approach as shown in figure 3A.

Once the network has been trained within its weights and activation functions are embedded the relationships and dependencies between the various components of the large output $\hat{\mathbf{Y}}$. The trained network is then utilized in an un-supervised adaptive feedback manner without access to the input data $(\mathbf{X}_I, \mathbf{x}_v)$ to make phase space predictions by looking at only those outputs which can be compared to measurements in real time, \hat{y}_j , to calculate some cost function $C_A(\hat{y}_j, y_j)$ which guides the adaptive tuning of an input to the latent space \mathbf{y}_L according to the right hand side of Equation 2.6, as shown in figure 3B.

This approach has two major strengths. The first is that we are now tuning an incredibly high dimensional system within a much smaller space which greatly speeds up the convergence of any optimization or search algorithm. The second main benefit is that once we have learned the generative half of the CNN which is capable of generating complex high dimensional outputs as a function of a low dimensional latent space vector \mathbf{y}_L we no longer need access to the original inputs that were used to train the network. This second feature is especially important for large complex systems, such as particle accelerators, where continuous operations makes it impossible to take detailed beam measurements without interrupting normal operations and therefore is only done during rare dedicated beam development times. The resulting algorithm is robust to uncertain time varying input beam distributions and accelerator parameters as they are assumed to be unknown and the application takes place as an adaptive un-supervised feedback approach which is only allowed to have access to measured quantities which may be a very small subset of all of the CNN’s

predictions. We demonstrate the approach on two very different particle accelerators: the HiRES compact ultrafast electron diffraction accelerator and the FACET-II plasma wakefield accelerator.

4 Application for particle accelerator diagnostics

One of the most advanced and detailed diagnostics is a setup using a transverse deflecting RF cavity (TCAV) together with a dipole magnet which separates a beam horizontally as a function of time while separating it as a function of energy vertically resulting in a 2D projection of the beam's (z, E) longitudinal phase space (LPS) with fs resolution [47]. Accelerators such as the LCLS, have a TCAV directly after the undulator so that once the FEL lasing process has completed the remaining electron beam can be separated from the x-rays and the beam's LPS can be measured without disrupting FEL operations. Therefore at the LCLS the TCAV is an always present real-time non-invasive LPS diagnostics. Machines, such as the EuXFEL have their TCAV near the machine's beam source, so it can be used to get an estimate of the LPS characteristics of the beam entering the accelerator. Machines such as FACET-II might have multiple TCAVs, as shown in figure 6, which can be used to periodically check the LPS properties of the beams throughout the accelerator.

The LPS of beams has a strong influence over their characteristics in terms of FEL lasing, bunch compression, and wakefield acceleration and therefore many recent efforts have been made for machine learning (ML)-based virtual LPS predictions to complement TCAV-based measurements by providing LPS estimates where TCAVs are not available or by providing LPS estimates without the use of destructive TCAVs. Tools for virtual phase space estimates have been developed which use model-independent adaptive feedback to adaptively tune physics models based on non-invasive beam measurements such as energy spread spectra [48]. Once an accurate estimate of a beam's LPS is obtained, whether by using a TCAV-based measurement or a virtual diagnostic, powerful adaptive feedback methods can be used to automatically control the LPS distribution [49], including combinations of ML and adaptive feedback for LPS control which is both global and robust to time-variation of accelerator and beam parameters [27].

A latent space tuning AML approach is developed for predicting all of the 2D projections of the beams entire 6D phase space given by the 15 distributions:

$$\begin{aligned}
 &\rho_{12}(x, x'), \quad \rho_{13}(x, y), \quad \rho_{14}(x, y'), \quad \rho_{15}(x, z), \quad \rho_{16}(x, E), \\
 &\rho_{23}(x', y), \quad \rho_{24}(x', y'), \quad \rho_{25}(x', z), \quad \rho_{26}(x', E), \\
 &\rho_{34}(y, y'), \quad \rho_{35}(y, z), \quad \rho_{36}(y, E), \\
 &\rho_{45}(y', z), \quad \rho_{46}(y', E), \\
 &\rho_{56}(z, E),
 \end{aligned} \tag{4.1}$$

where we are now using the coordinates

$$\mathbf{p} = (p_1, \dots, p_6) = (x, x', y, y', z, E), \quad x' = \frac{p_x}{p_z}, \quad y' = \frac{p_y}{p_z}, \tag{4.2}$$

and the distributions are defined as

$$\rho_{ij}(p_i, p_j) = \iiint \rho(\mathbf{p}) dp_{i_1} dp_{i_2} dp_{i_3} dp_{i_4}, \quad i_1, i_2, i_3, i_4 \neq i, j. \tag{4.3}$$

It is usually difficult or impossible to measure the transverse phase space distributions $\rho_{12}(x, x')$ and $\rho_{34}(y, y')$ of the beam quickly and without interrupting operations. Our latent space approach provides an estimate of all 15 2D phase space projections in (4.1) by comparing the CNN’s predictions of the LPS $\rho_{56}(z, E)$ to available TCAV measurements and when possible also projecting $\rho_{16}(x, E)$ distributions down to 1D energy spread spectra when such measurements are available, as was done in [48] to adaptive tune an online model to track the $\rho_{56}(z, E)$ phase space non-invasively. This approach relies on the fact that the encoder-decoder CNN has learned correlations between the various phase space projections for a given accelerator lattice, as demonstrated below.

By training the CNN over a wide range of beam and parameter characteristics, and by forcing the CNN to satisfy the many (15) physics constraints of the 2D phase space projections, the encoder-decoder did learn enough of the underlying beam physics so that it could be used in situations where the input beam distribution is no longer available, but only output beam TCAV-based measurements and possibly energy spread spectra measurements are available. Adaptively tuning the latent space allows us to use the encoder-decoder CNN without needing access to the original input parameter and distribution measurements with which the CNN was trained. This is a major advance for ML-based virtual diagnostics because outside of dedicated beam development times it is often very difficult to gather detailed measurements on input beam distributions in a non-invasive manner that does not interrupt accelerator operations.

Figure 4 shows the results of training the encoder-decoder CNN with a 2D latent space layer. The figure shows the 2D latent space positions (y_1, y_2) of 500 test data sets which the CNN was not shown during training which are colored by the difference between the $\rho_{56}(z, E)$ phase space projection at each point and that of the point labeled as “goal” in the top left corner, based on the beam at the first TCAV1 as shown in the schematic of FACET-II and defined by the equation:

$$\text{cost} = \int_z \int_E |\hat{\rho}_{56}(z, E) - \rho_{56}(z, E)| dE dz. \quad (4.4)$$

Two points at opposite positions in the latent space are shown in red and black. As we move from the start to the goal point we see the evolution of the $\rho_{13}(x, y)$ and $\rho_{14}(x, y')$ projections at the location of TCAV1 and the $\rho_{56}(z, E)$ projection at the interaction point (IP) location. We see that the encoder section of the CNN has sorted the beam in the latent space so that as we move along the green dashed line the transverse beam size is increasing, the beam divergence is flattening, and the bunch-to-bunch separation is squeezed down to a single bunch.

4.1 Latent space tuning for HiRES

As first demonstrated in [30] an adaptive latent space tuning approach was applied at the HiRES ultra fast electron diffraction compact accelerator whose schematic is shown in figure 5. At HiRES the input was the (x, y) phase space projection at the electron gun and the accelerator parameters included magnet strengths and beam energy. The latent space tuning approach enabled measurements of the (z, E) LPS to be mapped to the transverse phase space (x, x') , (y, y') of the beam as shown in figure 5 by minimizing the difference between a target and guess (z, E) longitudinal phase spaces by applying adaptive latent space feedback tuning as defined in (2.5).

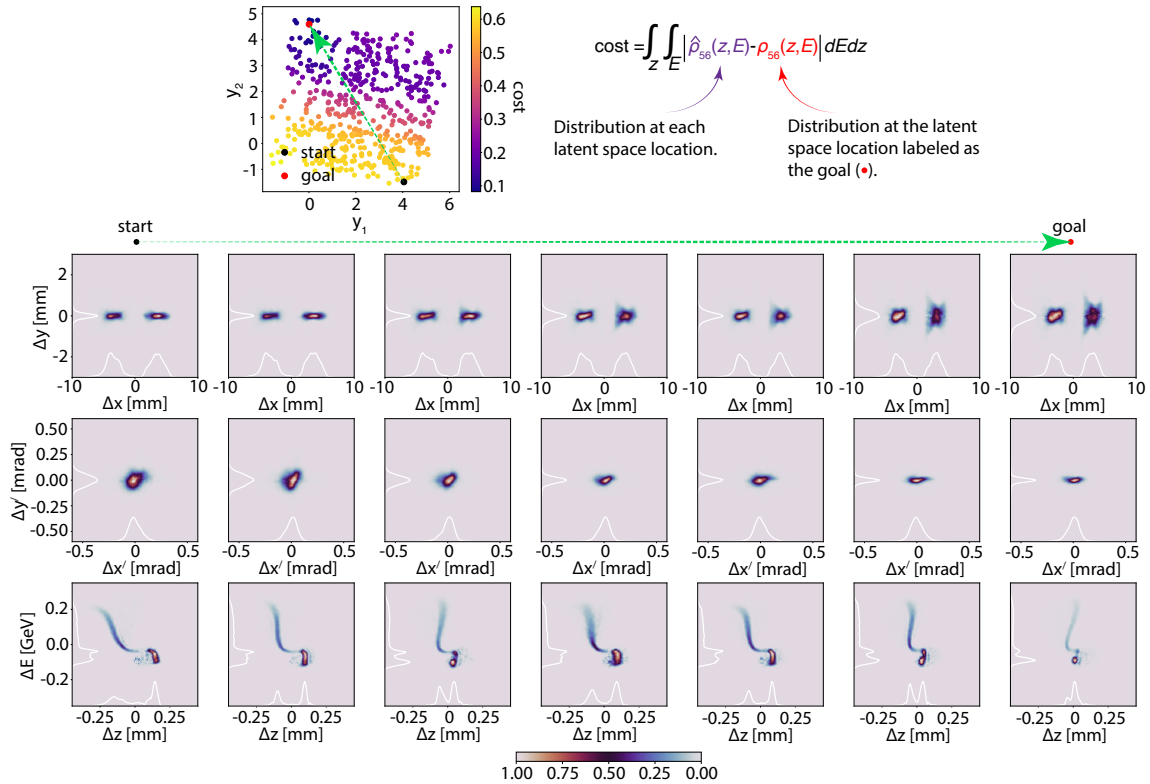


Figure 4. Beam representation as organized by the encoder CNN is shown when moving in the 2D latent space from the start (bottom right corner of the latent space) to the goal (top left corner of the latent space) point. Seven equidistant $\rho_{13}(x, y)$ and $\rho_{14}(x, y')$ 2D phase space projections are shown at the location of TCAV1 and $\rho_{56}(z, E)$ is shown at the interaction point (IP) location. The encoder-decoder CNN has naturally organized the various distributions in the 2D latent space. The cost function value, as defined in Equation (4.4), is used to color all latent space points relative to the goal point in the top left corner.

4.2 Latent space tuning for FACET-II

For application to FACET-II, the \mathbf{X}_I input in figure 3 is a 128×128 pixel 2D $\rho_{13}(x, y)$ distribution of the (x, y) phase space projection of the beam entering the particle accelerator and \mathbf{x}_v is a vector of 7 accelerator and beam parameters which are the RF phase and amplitude settings, transverse beam offsets, and beam charge variation. The 16391 dimensional inputs ($\mathbf{X}_I, \mathbf{x}_v$) are then combined into a single low-dimensional vector \mathbf{y}_L of dimension between 2 and 100 before the generative half of the encoder-decoder CNN builds back up to a high dimensional output $\hat{\mathbf{Y}}$ which is a $128 \times 128 \times 75$ output with the 75 channels representing a set of 15 2D projections of the 6D phase space: (x, y) , (x, z) , (x, x') , (x, y') , (x, E) , (x', y) , (x', z) , (x', y') , (x', E) , (y, z) , (y, y') , (y, E) , (y', z) , (y', E) , (z, E) at 5 different accelerator locations in FACET-II as shown in figure 1. Our simulation study of FACET-II uses the particle tracking code LUCRETIA for a ~ 2 nC beam (this was adjusted to make training and test data sets) using $2e5$ macroparticles and tracking from the exit of the FACET-II injector to the end of the linac [50]. The simulations include the effects of longitudinal space charge, wakefields, and coherent and incoherent synchrotron radiation on the electron beam distribution and has been reported on in detail elsewhere [49].

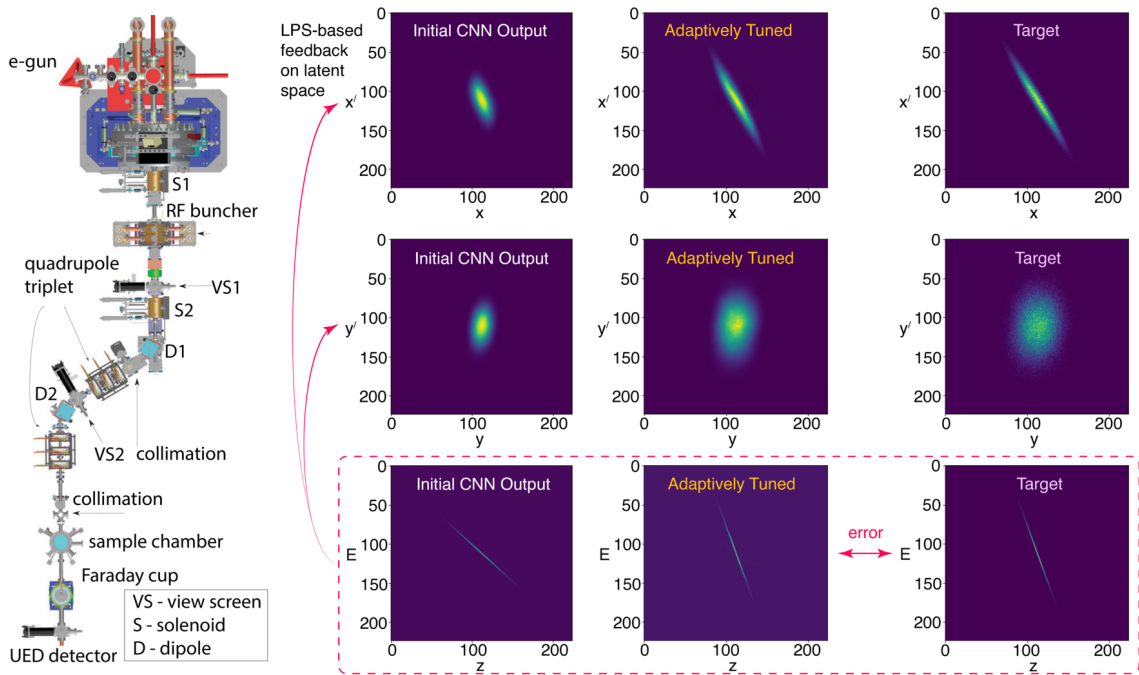


Figure 5. Predicting the transverse phase space (x, x') , (y, y') based on LPS measurements (z, E) at HiRES.




Figure 6. A simplified schematic of the FACET-II beamline is shown where $L0-L3$ are linac sections and $BC11-BC20$ are bunch compressors. Multiple TCAV locations are shown in yellow, as well as an energy spread spectrum measurement location within bunch compressor $BC20$. The AML approach discussed in this paper can be trained by and provide non-invasive estimates of the beam’s phase space at the TCAV locations as implied by the black arrows which represent data that may be used as inputs to the CNN and red arrows that represent possible CNN-based virtual diagnostics once the network has been trained.

4.3 2D latent space

For FACET-II the CNN was trained with twenty thousand randomly generated input-output pairs, over a wide range which included ± 1 degree variation of the L1 and L2 linac phase set points and $\pm 1\%$ variation over their accelerating voltages, $\pm 1\%$ variation in the input beam’s 2.5 nC charge, $\pm 20 \mu\text{m}$ offsets of the input beam, and a $\pm 10\%$ change of the input beam’s transverse (x, y) size. By forcing the CNN to simultaneously predict all 15 2D distribution projections which must satisfy all of the physics constraints of both general relativistic electrodynamics as well as the specific constraints of the FACET-II setup, the correlations between the various phase spaces at FACET-II were learned. Various latent space dimensions were tested (2,3,5,7,50,100), and various statistics of

test-set errors are shown in figure 7. Clearly as the number of latent space dimensions is increased the accuracy of the CNN grows because the network size is bigger and the nonlinear representation is more general. However it is clear that after a dimension of 7 is reached the increase in accuracy is only incremental, which is understood by the fact that only 7 accelerator parameters were changed to create the training and test data for this demonstration. The CNN trained with Adam optimizer and mean square error, used 2D convolutional layers with selu activation functions to squeeze images down, reshape, and flatten before passing through dense layers. Vectors inputs passed through dense layers before concatenation with the dense layers of the image resulting in a dense layer of length N_L which is the latent space dimension. Dense layers are then re-shaped into an image, which is up-sampled repeatedly using transposed 2D convolutional layers with selu activation functions.

The adaptive approach at FACET-II was to use the mean of the training set as a guess for the input beam distribution and then iteratively adjust the low dimensional latent space to minimize a cost function which measured the difference between the predicted LPS $\hat{\rho}_{56}(z, E)$ and the correct LPS $\rho_{56}(z, E)$ according to several candidate cost functions defined as weighted combinations of:

$$C_{T,j} = \underbrace{\int_z \int_E |\hat{\rho}_{56}(z, E) - \rho_{56}(z, E)| dE dz}_{\text{at TCAVN}}, \quad N \in \{1, 2, 3\},$$

$$C_{B,j} = \underbrace{\int_x |\hat{\rho}_x(E) - \rho_x(E)| dE}_{\text{at BCN}}, \quad N \in \{11, 14, 20\}, \quad \rho_x(E) = \int_E \rho_{16}(x, E) dE,$$

where $C_{T,j}$ are costs based on TCAV measurements at the three TCAV locations considered for FACET-II and the $C_{B,j}$ are costs based on non-invasive energy spread spectra measurements which could possibly be collected at each of the three bunch compressors.

4.4 Latent space-informed diagnostics

For the 2D latent space CNN a pair of adaptive tuning results are shown in figure 8. In figure 8A-D, the cost being adaptively minimized is the sum $\sum_{j=1}^3 w_j C_{T,j}$, and we see that the resulting cost function has a local minimum near the starting point in the 2D latent space, as shown in more detail in 8D. The result is that the adaptive scheme becomes stuck in this local minimum and is unable to reach the target. This result can be understood by looking at figure 4 where the IP $\hat{\rho}_{56}(z, E)$ phase space is shown, which is almost identical to the TCAV3 LPS. We can see that the way the CNN has sorted the 2D latent space, moving along the green dashed line from the starting to the goal position the bunch length is initially compressed and then spread out again before finally being compressed to a single bunch. This oscillating variation caused by including the TCAV3 measurement $\hat{\rho}_{56}(z, E)$ creates the local minimum in the cost function and a similar result happened with TCAV2. If we instead define the cost function to simply include only the TCAV1 LPS measurement, then we get the result shown in figure 8E-H, the cost function is convex and the adaptive optimizer makes its way across the entire range of the latent space in order to adaptively match the unknown beam.

This result is somewhat counter intuitive: using less information gives better convergence. One explanation is that FACET-II is such a large and complex machine that a given final beam phase space distribution can be created from many different initial distributions in a non-unique

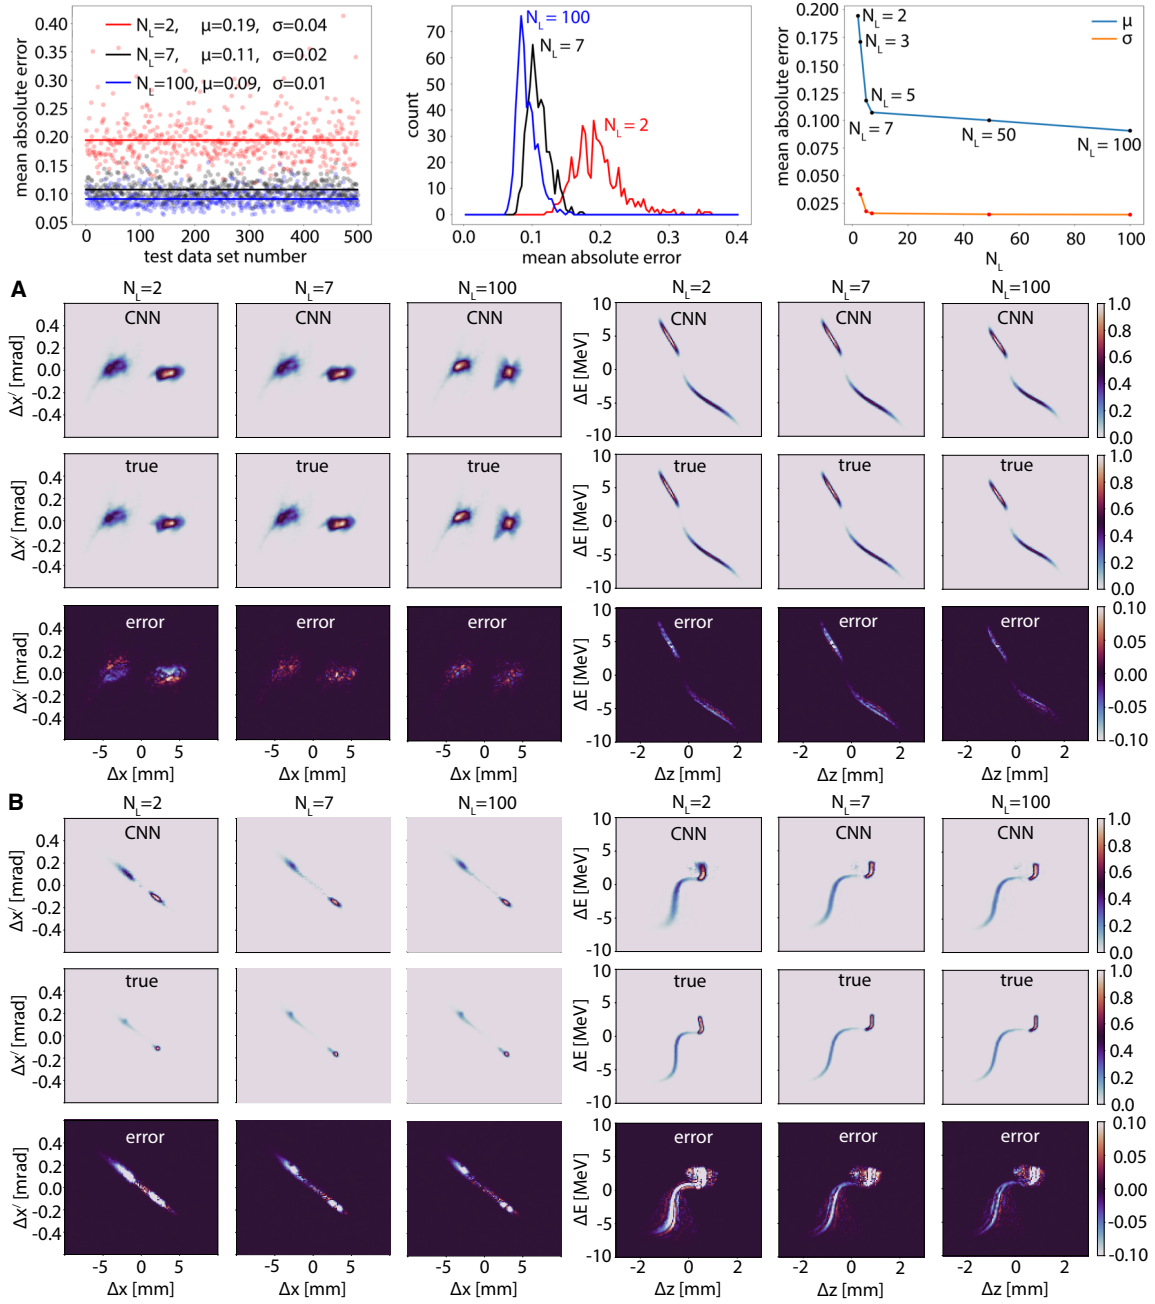


Figure 7. Error statistics for 500 test data sets when using networks with latent space dimensions $N_L \in \{2, 3, 5, 7, 50, \text{and } 100\}$ (Top). Comparison of prediction error for 2 projections of phase space at the first TCAV (A) and at the interaction point (B). To give a representative understanding of each latent space size CNN, for the three cases shown, of $N_L = 2, 7, 100$ for each network the test example chosen was one whose mean absolute error was closest to the average for that network’s performance over the entire test set.

way based on all of the accelerator settings between the start and end of the machine. With this in mind it is understandable that the earlier diagnostics give a better chance of uniquely distinguishing between phase space distributions. This analysis shows that if a limited number of diagnostics must

be chosen, such as placing a single TCAV at some location of an accelerator, it may be beneficial to place it earlier, as the measurements can be used to uniquely identify different beam characteristics. There is an obvious downside to placing a diagnostic such as a TCAV too early in an accelerator configuration as one of the most useful applications of a TCAV is to measure the influence of collective effects such as space charge forces and CSR which grow in intensity as the beam reaches higher power and is compressed. Although figure 7 shows that using a lower dimensional latent space gives less accurate results, figure 8 shows the benefit of a low-dimensional latent space in extremely fast convergence averaging 40 steps for 300 random initial conditions.

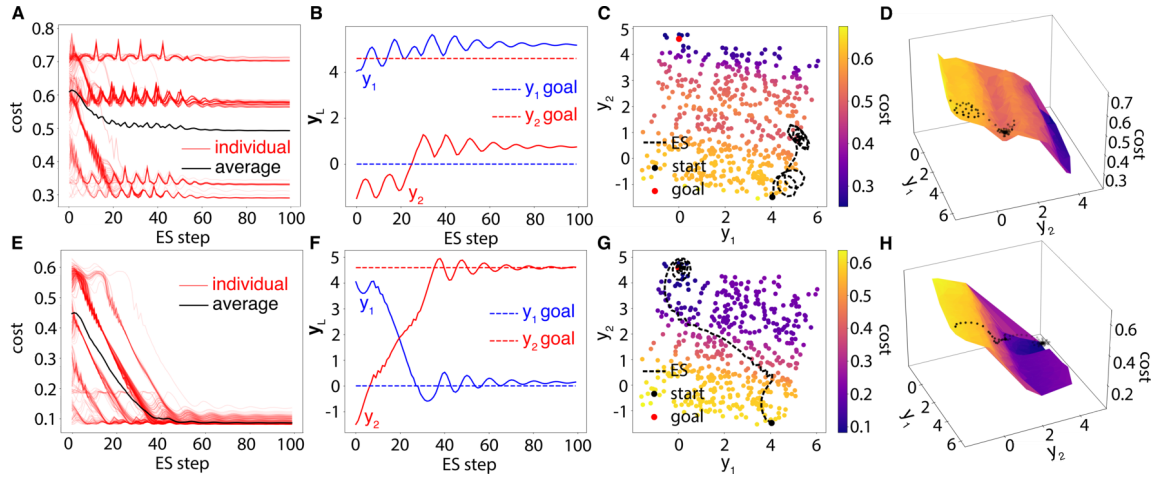


Figure 8. The cost being adaptively minimized for 300 random starting points is the sum $\sum_{j=1}^3 w_j C_{T,j}$ (A), the resulting cost function has a local minimum valley near the bottom left corner in the 2D latent space as seen with one example trajectory (B-D). The cost only includes the TCAV1 LPS measurement and is convex (E). One example of trajectories is shown resulting in convergence to the global minimum (F-H).

4.5 ND, $N > 2$ latent space

Higher dimensional latent spaces result in more accurate CNN predictions, they also make it less likely to get stuck in a local minimum as it is easier to avoid local minima in a higher dimensional space. Figure 9 shows the results of tuning in a 3D latent space when using a cost which includes all of the TCAV and BC terms. The global minimum is reached by taking a wider path that avoids local minima. A downside of increasing the latent space dimension is clearly seen in the growing number of steps required for convergence to take place. In practice there is a tradeoff between simplicity and accuracy/generalizability that must be made when choosing the latent space dimension for a given problem. Figure 10 shows 7D latent space convergence in 600 steps. To calculate prediction accuracy the (z, E) , (x_p, y_p) , and (x, y) distributions are projected onto the axes and fit with Gaussians of the form $f(\xi) = A \exp\left[-(\xi - \xi_c)^2 / 2\sigma_\xi^2\right]$, $\xi \in \{x, y, x_p, y_p, z, E\}$. The differences between target and predicted $\Delta\sigma_\xi = \sigma_{\xi,\text{tar}} - \sigma_{\xi,\text{pred}}$ fits are then calculated and the mean and standard deviation for 450 cases is shown in figure 11 at the interaction point which the AML method does not have access to for feedback. Figure 12 demonstrates that despite using only the projections measured for the cost functions $C_{T,j}$ and $C_{B,j}$ as defined above, the other 2D

projections are also accurately predicted. The CNN is predicting all 75 projections which are 5 sets of 15 projections at 5 locations in FACET-II, but to save space we just show a few of the projections.

Furthermore, we emphasize that both the input beam distributions and the accelerator parameters in the test set were unseen during the CNN's training process and the CNN is now predicting the correct 2D phase space distributions without needing access to accurate input beam distributions or to accelerator parameters. This is the major strength of this approach which makes it robust to time-varying systems for which invasive input beam distributions are no longer available after an initial dedicated training period which might take place during a dedicated beam development time.

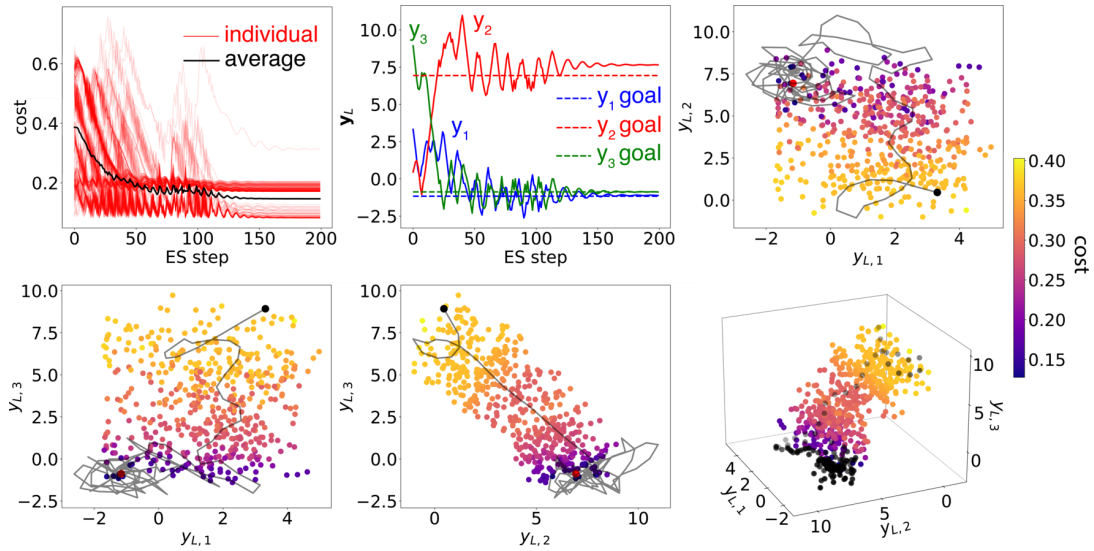


Figure 9. Convergence in the 3D latent space is shown, which takes approximately 150 steps. We show all 3 2D projections as well as the evolution of the trajectory (black) in the full 3D latent space with the underlying test set-based latent space colored by the value of the cost function.

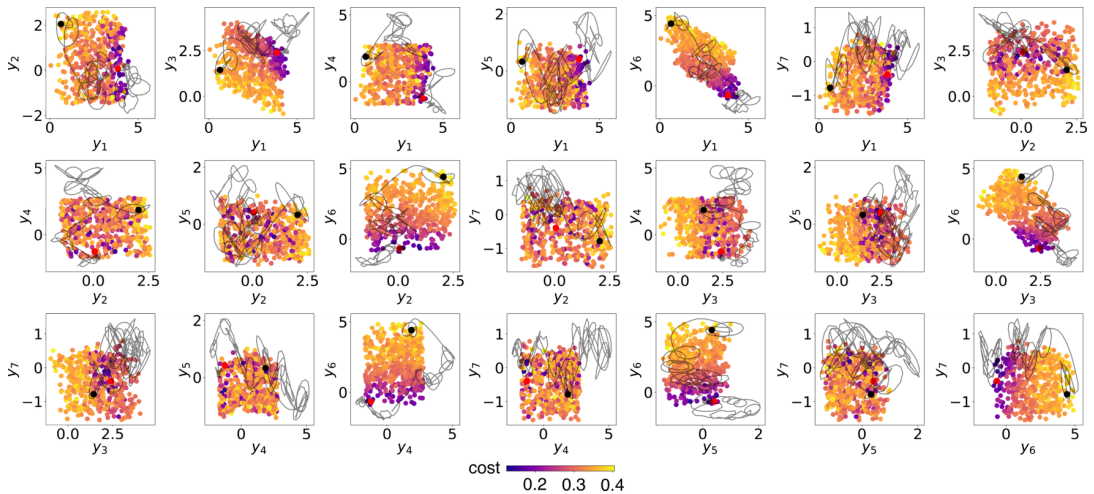


Figure 10. The adaptive feedback trajectory in 7D latent space is shown (black) on top of all 21 projections of 2D views of the 7D latent space, which is colored by the value of the cost function being minimized.

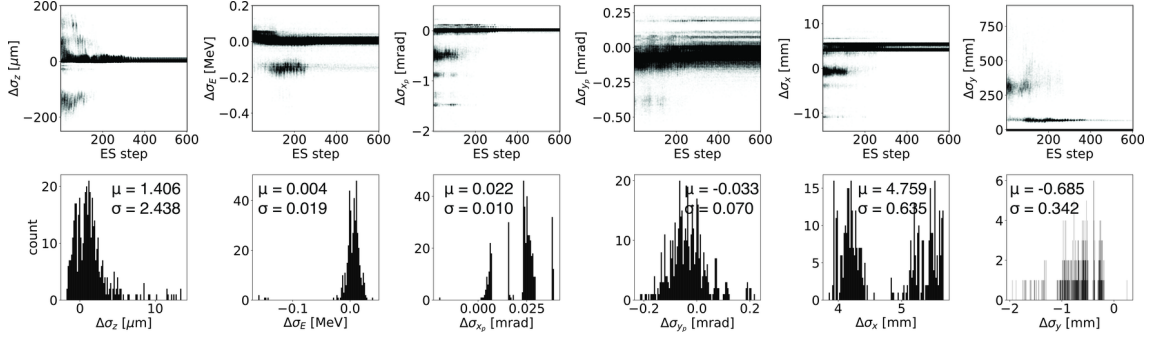


Figure 11. The top row shows the convergence for each case and the bottom row shows histograms and statistics of the final mean and standard deviations of all of the errors found by fitting gaussians to the target and final distributions . For example, for the (z, E) phase space the standard deviation and mean of the error of all of the final fits are $(\sigma_z, \sigma_E) = (2.438 \mu\text{m}, 0.019 \text{ MeV})$ and $(\mu_z, \mu_E) = (1.406 \mu\text{m}, 0.004 \text{ MeV})$.

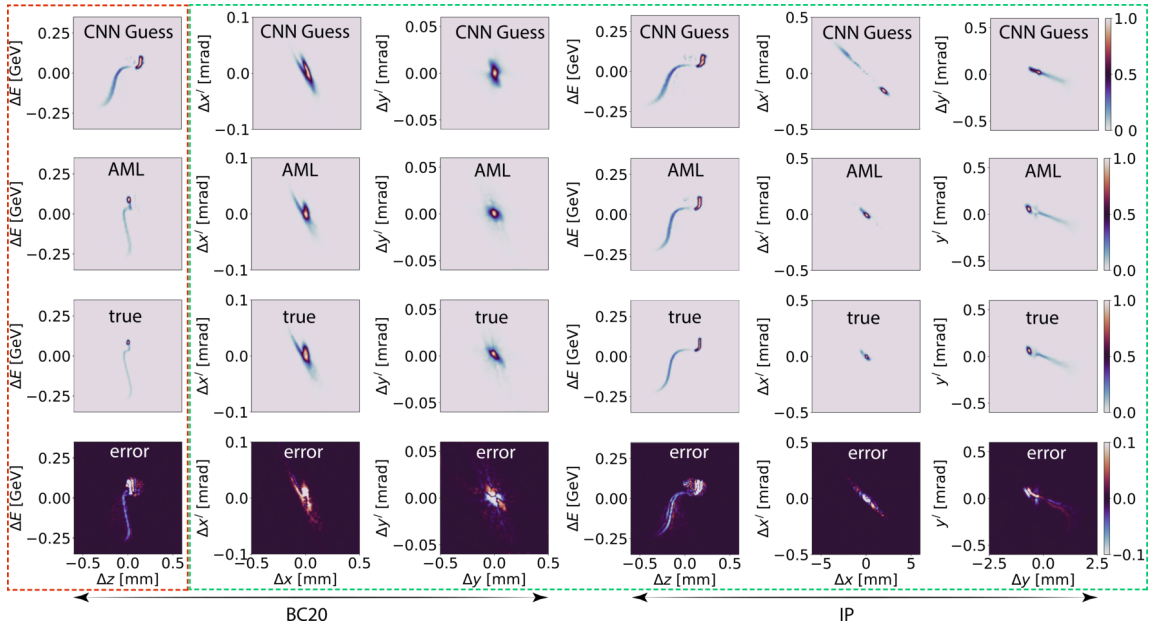


Figure 12. 7D latent space model predictions at BC20 and at the IP. Top row is average-based guess, second row is AML prediction. Red box shows a LPS diagnostic used for the cost function, others predicted.

5 Conclusions

An adaptive latent space tuning method has been developed and demonstrated on FACET-II and HiRES based on encoder-decoder style CNNs whose generative halves predict various 2D phase space projections of charged particle beams based only on a single 2D phase space slice measurement such as the LPS provided by TCAVs. Unlike traditional ML-based methods this is robust to unknown time-varying input beam distributions which are represented by the low dimensional latent space of the CNN which is quickly and automatically adjusted based on model-independent feedback.

Acknowledgments

Work was funded by Los Alamos National Laboratory LDRD Program project 20210160DR and Department of Energy Office of Science Basic Energy Sciences High Energy Physics Program on Data, Artificial Intelligence, and Machine Learning at DOE Scientific User Facilities, LAB 20-2261, project “Leveraging ML/AI techniques to enable a breakthrough in ultra-short-bunch paradigm.”

References

- [1] A.A. Lutman, T.J. Maxwell, J.P. MacArthur, M.W. Guetg, N. Berrah, R.N. Coffee et al., *Fresh-slice multicolour x-ray free-electron lasers*, *Nat. Photon.* **10** (2016) 745.
- [2] A. Malyzhenkov, Y.P. Arbelo, P. Craievich, P. Dijkstal, E. Ferrari, S. Reiche et al., *Single- and two-color attosecond hard x-ray free-electron laser pulses with nonlinear compression*, *Phys. Rev. Res.* **2** (2020) .
- [3] W. Decking, S. Abeghyan, P. Abramian, A. Abramsky, A. Aguirre, C. Albrecht et al., *A MHz-repetition-rate hard x-ray free-electron laser driven by a superconducting linear accelerator*, *Nat. Photon.* **14** (2020) 391.
- [4] V. Yakimenko et al., *FACET-II facility for advanced accelerator experimental tests*, *Phys. Rev. Accel. Beams* **22** (2019) 101301.
- [5] AWAKE collaboration, *Acceleration of electrons in the plasma wakefield of a proton bunch*, *Nature* **561** (2018) 363 [[arXiv:1808.09759](https://arxiv.org/abs/1808.09759)].
- [6] C. Emma, A. Edelen, M.J. Hogan, B. O’Shea, G. White and V. Yakimenko, *Machine learning-based longitudinal phase space prediction of particle accelerators*, *Phys. Rev. Accel. Beams* **21** (2018) 112802.
- [7] A. Hanuka, C. Emma, T. Maxwell, A. Fisher, B. Jacobson, M.J. Hogan et al., *Accurate and confident prediction of electron beam longitudinal properties using spectral virtual diagnostics*, *Sci. Rep.* **11** (2021) 2945 [[arXiv:2009.12835](https://arxiv.org/abs/2009.12835)].
- [8] L. Gupta, A. Edelen, N. Neveu, A. Mishra, C. Mayes and Y.-K. Kim, *Improving Surrogate Model Accuracy for the LCLS-II Injector Frontend Using Convolutional Neural Networks and Transfer Learning*, [arXiv:2103.07540](https://arxiv.org/abs/2103.07540).
- [9] S. Hirlaender and N. Bruchon, *Model-free and Bayesian Ensembling Model-based Deep Reinforcement Learning for Particle Accelerator Control Demonstrated on the FERMI FEL*, [arXiv:2012.09737](https://arxiv.org/abs/2012.09737).
- [10] A. Adelmann, *On nonintrusive uncertainty quantification and surrogate model construction in particle accelerator modeling*, *SIAM/ASA J. Uncert. Quant.* **7** (2019) 383.
- [11] O. Convery, L. Smith, Y. Gal and A. Hanuka, *Uncertainty quantification for virtual diagnostic of particle accelerators*, *Phys. Rev. Accel. Beams* **24** (2021) 074602 [[arXiv:2105.04654](https://arxiv.org/abs/2105.04654)].
- [12] S. Li, M. Zacharias, J. Snuverink, J.C. de Portugal, F. Perez-Cruz, D. Reggiani et al., *A novel approach for classification and forecasting of time series in particle accelerators*, *Information* **12** (2021) 121.
- [13] A. Edelen, N. Neveu, Y. Huber, M. Frey, C. Mayes and A. Adelmann, *Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems*, *Phys. Rev. Accel. Beams* **23** (2020) 044601 [[arXiv:1903.07759](https://arxiv.org/abs/1903.07759)].

- [14] M. Kranjcevic, B. Riemann, A. Adelmann and A. Streun, *Multiobjective optimization of the dynamic aperture using surrogate models based on artificial neural networks*, *Phys. Rev. Accel. Beams* **24** (2021) 014601 [[arXiv:2008.04151](#)].
- [15] M. McIntire et al., *Bayesian optimization of FEL performance at LCLS*, in *Proceedings of the 7th International Particle Accelerator Conference*, Busan, Korea, 8–13 May 2016, pp. 2972–2975.
- [16] Y. Li, R. Rainer and W. Cheng, *Bayesian approach for linear optics correction*, *Phys. Rev. Accel. Beams* **22** (2019) 012804 [[arXiv:1904.08214](#)].
- [17] Y. Hao, Y. Li, M. Balcewicz, L. Neufcourt and W. Cheng, *Reconstruction of Storage Ring’s Linear Optics with Bayesian Inference*, [arXiv:1902.11157](#).
- [18] J. Duris et al., *Bayesian optimization of a free-electron laser*, *Phys. Rev. Lett.* **124** (2020) 124801 [[arXiv:1909.05963](#)].
- [19] A. Hanuka, J. Duris, J. Shtalenkova, D. Kennedy, A. Edelen, D. Ratner et al., *Online tuning and light source control using a physics-informed Gaussian process*, [arXiv:1911.01538](#).
- [20] Y. Li, R.S. Rainer, W. Cheng and Y. Hao, *Analysis of Beam Position Monitor Requirements with Bayesian Gaussian Regression*, in *3rd North American Particle Accelerator Conference (NAPAC2019)*, 4, 2019, DOI [[arXiv:1904.05683](#)].
- [21] R.J. Shalloo et al., *Automation and control of laser wakefield accelerators using Bayesian optimization*, *Nature Commun.* **11** (2020) 6355 [[arXiv:2007.14340](#)].
- [22] R. Roussel, A. Hanuka and A. Edelen, *Multi-Objective Bayesian Optimization for Accelerator Tuning*, [arXiv:2010.09824](#).
- [23] E. Fol, J.C. de Portugal, G. Franchetti and R. Tomas, *Optics corrections using machine learning in the LHC*, in *Proceedings of the 10th International Particle Accelerator Conference*, Melbourne, Australia, 19–24 May 2019, pp. 3990–3993.
- [24] E. Fol et al., *Unsupervised Machine Learning for Detection of Faulty Beam Position Monitors*, in *Proceedings of the 10th International Particle Accelerator Conference*, Melbourne, Australia, 19–24 May 2019, pp. 2668–2671.
- [25] P. Arpaia et al., *Machine learning for beam dynamics studies at the CERN Large Hadron Collider*, *Nucl. Instrum. Meth. A* **985** (2021) 164652 [[arXiv:2009.08109](#)].
- [26] A. Scheinker and M. Krstic, *Minimum-seeking for CLFs: Universal semiglobally stabilizing feedback under unknown control directions*, *IEEE Trans. Automat. Contr.* **58** (2013) 1107.
- [27] A. Scheinker, A. Edelen, D. Bohler, C. Emma and A. Lutman, *Demonstration of Model-Independent Control of the Longitudinal Phase Space of Electron Beams in the Linac-Coherent Light Source with Femtosecond Resolution*, *Phys. Rev. Lett.* **121** (2018) 044801.
- [28] A. Scheinker, *Adaptive machine learning for robust diagnostics and control of time-varying particle accelerator components and beams*, *Information* **12** (2021) 161.
- [29] A. Scheinker, F. Cropp, S. Paiagua and D. Filippetto, *Adaptive deep learning for time-varying systems with hidden parameters: Predicting changing input beam distributions of compact particle accelerators*, [arXiv:2102.10510](#).
- [30] A. Scheinker, F. Cropp, S. Paiagua and D. Filippetto, *Adaptive Latent Space Tuning for Non-Stationary Distributions*, [arXiv:2105.03584](#).
- [31] A. L. Caterini, A. Doucet and D. Sejdinovic, *Hamiltonian variational auto-encoder*, [arXiv:1805.11328](#) (2018).

- [32] Y. Pu, Z. Gan, R. Henao, X. Yuan, C. Li, C.A. Stevens and L. Carin, *Variational autoencoder for deep learning of images, labels and captions*, in *Advances in neural information processing systems*, vol. 29, 2016, pp. 2360–2368.
- [33] J. An and S. Cho, *Variational autoencoder based anomaly detection using reconstruction probability*, *Special lecture on IE*, 2015, vol. 2.1, pp. 1–18.
- [34] R. Maulik, A. Mohan, B. Lusch, S. Madireddy, P. Balaprakash and D. Livescu, *Time-series learning of latent-space dynamics for reduced-order model closure*, *Physica D* **405** (2020) 132368.
- [35] A. Tripp, E. Daxberger and J.M. Hernández-Lobato, *Sample-efficient optimization in the latent space of deep generative models via weighted retraining*, in *Advances in Neural Information Processing Systems*, vol. 33, 2020 [[arXiv:2006.0919](https://arxiv.org/abs/2006.0919)].
- [36] A. Edelen, N. Neveu, C. Mayes, C.Emma and D. Ratner, *Machine learning models for optimization and control of x-ray free electron lasers*, talk given at *NeurIPS Machine Learning for the Physical Sciences Workshop*, Vancouver, BC, Canada, 8–14 December 2019.
- [37] J. Zhu, Y. Chen, F. Brinker, W. Decking, S. Tomin and H. Schlarb, *High-fidelity Prediction of Megapixel Longitudinal Phase-space Images of Electron Beams using Encoder-Decoder Neural Networks*, *Phys. Rev. Applied* **16** (2021) 024005 [[arXiv:2101.10437](https://arxiv.org/abs/2101.10437)].
- [38] H.K. Khalil, *Nonlinear Systems*, Prentice Hall (2002).
- [39] P. Ioannou and B. Fidan, *Adaptive control tutorial*, Society for Industrial and Applied Mathematics (2006).
- [40] P.A. Ioannou and J. Sun, *Robust adaptive control*, Courier Corporation (2012).
- [41] A. Scheinker, *Simultaneous stabilization and optimization of unknown, time-varying systems*, in *American Control Conference*, Washington, DC, U.S.A., 17–19 June 2013, pp. 2637–2642.
- [42] A. Scheinker and D. Scheinker, *Bounded extremum seeking with discontinuous dithers*, *Automatica* **69** (2016) 250.
- [43] A. Scheinker, *Model independent beam tuning*, in *Proceedings of the 2013 International Particle Accelerator Conference*, Shanghai, China, 12–17 May 2013, pp. 1862–1864.
- [44] A. Scheinker and M. Krstić, *Extremum seeking with bounded update rates*, *Syst. Control Lett.* **63** (2014) 25.
- [45] A. Krizhevsky, I. Sutskever and G.E. Hinton, *Imagenet classification with deep convolutional neural networks*, in *Advances in neural information processing systems*, vol. 25, 2012, pp. 1106–1114.
- [46] S. Wold, K. Esbensen and P. Geladi, *Principal component analysis*, *Chemom. Intell. Lab. Systems* **2** (1987) 37.
- [47] C. Behrens, F.-J. Decker, Y. Ding, V.A. Dolgashev, J. Frisch, Z. Huang et al., *Few-femtosecond time-resolved measurements of x-ray free-electron lasers*, *Nature Commun.* **5** (2014) 3762.
- [48] A. Scheinker and S. Gessner, *Adaptive method for electron bunch profile prediction*, *Phys. Rev. ST Accel. Beams* **18** (2015) 102801.
- [49] A. Scheinker, *Adaptive Machine Learning for Time-Varying Systems: Low Dimensional Latent Space Tuning*, *Nucl. Instrum. Meth. A* **967** (2020) 163902 [[arXiv:2107.06207](https://arxiv.org/abs/2107.06207)].
- [50] P. Tenenbaum, *Lucretia: A matlab-based toolbox for the modelling and simulation of single-pass electron beam transport systems*, in *Proceedings of the 2005 Particle Accelerator Conference*, Knoxville, Tennessee, U.S.A., 16–20 May 2005, pp. 4197–4199.