32nd Linear Accelerator Conference (LINAC2024),Chicago, IL, USA JACoW Publishing
ISBN: 978-3-95450-219-6 ISSN: 2226-0366 doi: 10.18429/JACoW-LINAC2024-WEXA005

ADAPTIVE MACHINE LEARNING WITH HARD PHYSICS CONSTRAINTS AND GENERATIVE DIFFUSION FOR 6D PHASE SPACE DIAGNOSTICS*

A. Scheinker†, Los Alamos National Laboratory, Los Alamos, NM, USA

Abstract

Machine learning (ML) tools have been growing in popularity for accelerator applications, but still struggle with time-varying systems, for which they require lengthy brute-force re-training. LANL has developed machine learning (ML)-based tools, that utilize adaptive model independent feedback control theory together with hard physics constraints, to make the tools much more robust to distribution shift. These adaptive ML tools can extrapolate much further beyond the span of the training data and are thus much more robust for time-varying systems. This talk will give a broad overview of the challenges of various time-varying accelerator systems at various accelerator facilities (known as systems with distribution shift in the ML community) and will present adaptive ML tools for 6D phase space diagnostics of intense charged particle beams. The talk will also give a general overview of adaptive latent space tuning, which is the novel method we have developed for adaptive ML, and how we are strictly enforcing hard physics constraints in our ML tools, which traditional ML tools lack. We demonstrate our general methods for various accelerators: the 5-meter-long ultra-fast electron diffraction (UED) HiRES compact accelerator at LBNL, the kilometer long plasma wakefield accelerator FACET-II at SLAC, and the LANL ion accelerator LANSCE.

INTRODUCTION

Particle accelerators are large complex systems composed of thousands of coupled components including resonant radio frequency (RF) structures for beam acceleration and magnets for beam shaping and steering. The performance of all of these components drifts with time due to external disturbances such as environmental temperature variation, vibrations and power line fluctuations, due to ageing, and due to misalignment and shift of components during maintenance. Furthermore, the beams themselves have time-varying initial conditions due to time variation of the complex beam source properties. Changing initial conditions can strongly effect the detailed 6D phase space distributions of beams, especially for intense beams whose evolution is governed by complex collective effects such as space charge forces and coherent synchrotron radiation.

Alongside the complexity and time variation of these systems and their beams, there is a lack of detailed non-invasive beam diagnostics beyond simple 1D measurements such as beam position monitors, beam current monitors, or beam loss monitors. The result is that accelerator beams are difficult to tune up after an outage (typically requiring weeks of effort be large teams) and are difficult to re-tune between different experiments even when the machines are running. Besides the large amount of time that is consequently wasted in re-tuning of the machines, they are also typically running at sub-optimal states with operators chasing drifts to decrease beam loss by tuning tens-hundreds of parameters by hand based on experience.

Motivated by the above, there has been a strong effort to create new tools for beam diagnostics and controls. For the EuXFEL, a convolutional neural network-based approach has been developed for generating high-resolution longitudinal phase-space (LPS) (z, E) images of the electron beam [1]. Convolutional neural networks have also been combine with destructive beam measurements to develop extremely fast virtual diagnostics for 4D tomographic phase space reconstructions [2]. Neural network-based methods have also been developed for predicting the transverse emittance of space charge dominated beams [3]. Convolutional neural networks and clustering algorithms have also been developed for predicting the longitudinal phase space of FEL beams and also to cluster these images to highlight patterns within the data [4]. Recently, very interesting methods have also been studied for phase space reconstructions based on normalizing flows [5]. One benefit of all ML-based tools is that they are by default fully differentiable models, enabling easy studies of complex input-output dependencies.

The time variation of the systems makes it clear that off-the-shelf machine learning methods cannot simply be used as they will require continuous repeated re-training based on expensive invasive beam measurements if they are applied to such systems which quickly vary with time. To overcome this difficulty, a new class of adaptive machine

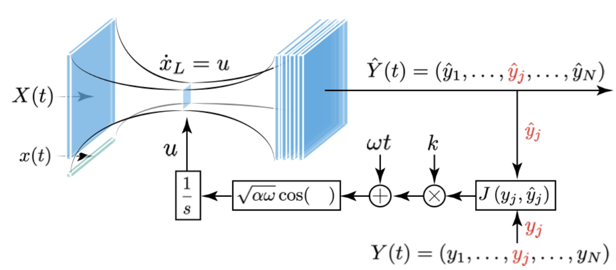

Figure 1: Adaptive latent space tuning setup in which a generative autoencoder generates all 15 projections of a beam's 6D phase space and then only those projections that can be measured non-invasively are compared to their predictions in order to track the time-varying beam with unknown and time-varying initial condtions, by adaptively tuning the low-dimensional latent embedding.

* This work was funded by the US Department of Energy (DOE), Office of Science, Office of High Energy Physics under contract number 89233218CNA000001 and the Los Alamos National Laboratory LDRD Program Directed Research (DR) project 20220074DR

† ascheink@lanl.gov

ISBN: 978-3-95450-219-6 ISSN: 2226-0366 doi: 10.18429/JACoW-LINAC2024-WEXA005

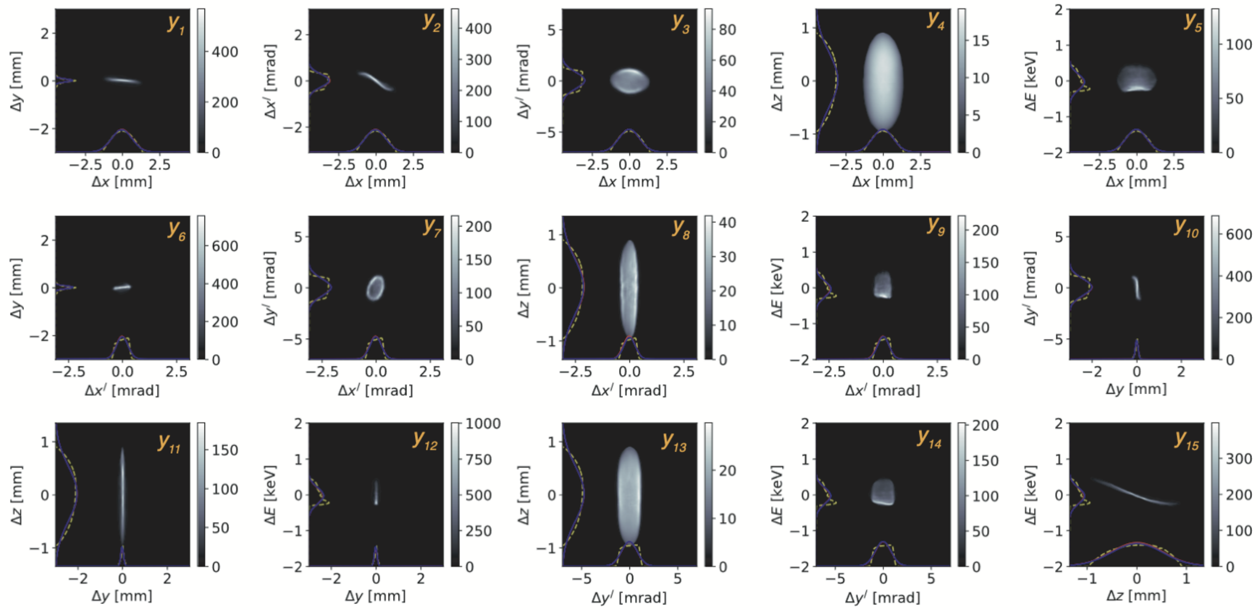

Figure 2: Generative AI creating all 15 projections of a beam's 6D phase space and then only those projections that can be measured non-invasively are compared to their predictions in order to track the time-varying beam with unknown and time-varying initial condtions, by adaptively tuning the low-dimensional latent embedding.

learning tools is being developed by the Adaptive Machine Learning Team at Los Alamos National Laboratory which incorporates adaptive feedback control and hard physics constraints directly within the architectures of generative AI in order to make these generative tools useful for time-varying systems without having to rely on continuous re-training.

ADAPTIVE LATENT SPACE TUNING FOR TIME-VARYING SYSTEMS

The idea of adaptively tuning the low-dimensional latent embedding of a generative model to quickly track the time-varying 6D phase space of a charged particle beam was first introduced in [6]. It was shown that this adaptive generative AI method enables generative models to track time-varying accelerator beams and makes them robust much further beyond the span of the training data than standard ML approaches which fail to extrapolate [7]. The setup of the adaptive approach is shown in Fig. 1 alongside one example of the 15 generated projections of a beam's 6D phase space (Fig. 2). Such virtual phase space diagnostics are very useful as having an ability to non-invasively measure details of a beam's phase space allows us to control the phase space via adaptive ML-based feedback methods. This was demonstrated in the first combination of an adaptive feedback algorithm [8] together with deep neural networks for automatic shaping of the longitudinal (z, E) phase space of intense short electron beams in the LCLS FEL [9]. Adaptive convolutional neural networks have also been designed for inverse problems that map downstream beam measurements back to the initial beam distribution [10].

GENERATIVE DEEP LEARNING WITH HARD PHYSICS CONSTRAINTS FOR ELECTRODYNAMICS

The predictive abilities of most generative models are limited by a lack of physics constraints. The models are biased towards generating physical data because they are trained on data coming from physics simulations or physical experiments, however the models are prone to hallucination especially when they extrapolate beyond the training data. The adaptive generative approach in [7] mitigates ML-driven hallucination by closely coupling the generative model with an online physics model, but this slows down the process because physics based simulations of intense beams are computationally expensive. For example, considering a 2 nC bunch of $N \approx 1.25 \times 10^{10}$ electrons, calculating individual particle to particle SC or CSR interactions is a computationally expensive $\mathcal{O}(N^2)$ process orders of magnitude slower than the inference speed (10-100 microseconds) of trained generative neural networks running on modern GPUs. Although the adaptive feedback in the generative approaches for 6D phase space diagnostics developed in [6, 11] has been demonstrated to make the models more robust and more physically consistent than generative models without adaptive feedback, but they still lack hard physics constraints.

Physics informed neural networks (PINNs) are the standard approach to trying to get neural networks to respect physics. PINNs include an extra term in a neural network's training cost penalizing the violation of a PDE which describes the desired physics [12]. PINNs are incredibly flexible and easy to use because any PDE can easily be added to a cost function by utilizing the automatic differentiation

ISBN: 978-3-95450-219-6 ISSN: 2226-0366 doi: 10.18429/JACoW-LINAC2024-WEXA005

built in to all ML frameworks without having to modify the neural network itself. PINNs also provide a powerful way to include boundary conditions and sparse observations to solve inverse problems [13]. However, PINNs also face several major challenges. The first problem is that PINNs do not guarantee any hard physics constraints, but rather gently nudge a neural network toward respecting a PDE. The second problem is that training PINNs can be incredibly difficult as the PDE constraint and reconstruction accuracy of the PINN are usually at odds because simply setting everything to 0 usually perfectly satisfies any PDE solution. How to actually set up and train PINNs successfully is its own research area. Intuitively this problem can be easily understood by the following accelerator physics-relevant example.

Consider the goal of generating a magnetic field $\hat{\mathbf{B}}$ associated with some kind of electrodynamic problem, maybe this is the self-field of a relativistic electron bunch. The PINN approach is to train a neural network with a cost function defined as

$$\begin{aligned} C &= w_{\mathbf{B}} \iiint |\mathbf{B} - \hat{\mathbf{B}}|^2 dV + w_{\Delta} \iiint |\nabla \cdot \hat{\mathbf{B}}|^2 dV \\ &= w_{\mathbf{B}} \|\mathbf{B} - \hat{\mathbf{B}}\|_2 + w_{\Delta} \|\nabla \cdot \hat{\mathbf{B}}\|_2, \end{aligned} \tag{1}$$

where the first term is related to magnetic field prediction accuracy and boundary conditions and the second term penalizes violation of the physics constraint $\nabla \cdot \hat{\mathbf{B}} = 0$. With soft PINN-type constraints there is a trade-off between the minimization of the two terms in Eq. 1 based on the choice of weights $w_{\mathbf{B}}$ and w_{Δ}. This trade-off can be understood by the fact that the most simple way for a neural network to satisfy $\nabla \cdot \hat{\mathbf{B}} = 0$ is to generate $\hat{\mathbf{B}} \equiv C$ for any constant C, which has absolutely nothing to do with the correct physical field $\mathbf{B}$.

Recently a new method was developed for enforcing hard physics constraints in generative ML, respecting physics and improving generalization via physics constrained neural networks (PCNNs) [14]. In [14], 3D convolutional neural operators were used to generate the vector and scalar potential fields $\mathbf{A}(\mathbf{r}, t)$ and $\varphi(\mathbf{r}, t)$, respectively, from which the electromagnetic fields were generated as $\mathbf{B}(\mathbf{r}, t) = \nabla \times \mathbf{A}(\mathbf{r}, t)$, $\mathbf{E}(\mathbf{r}, t) = -\nabla\varphi(\mathbf{r}, t) - \partial\mathbf{A}(\mathbf{r}, t)/\partial t$, as shown in Fig. 3.

The PCNN approach generating fields based on potentials as in [14] builds hard physics constraints (Maxwell's equations), such as $\nabla \cdot \mathbf{B}(\mathbf{r}, t) = 0$ and $\nabla \times \mathbf{E} + \partial\mathbf{B}/\partial t = 0$, directly into the architecture of generative neural networks without the trade-offs or training difficulties of PINNs.

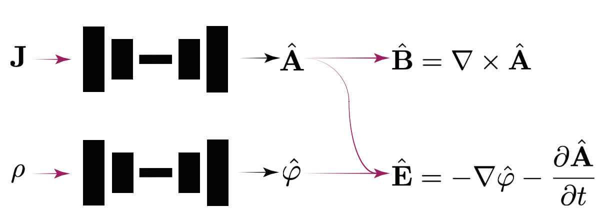

Figure 3: A method has been developed to incorporate Maxwell's Equations as hard physics within the structure of generative models by generating potentials rather than electromagnetic fields directly [14].

DIFFUSION-BASED MODELS

In the field of generative machine learning, diffusion-based models are the current state-of-the-art for generating high resolution, highly complex representations of widely varying objects. Generative models based on diffusion utilize a gradual denoising approach inspired by statistical thermodynamics for modeling complex distributions [15]. This approach was then further developed for the generation of high resolution images [16–19]. The generative ability of diffusion-based models has made them powerful tools for a wide range of scientific applications [20], such as conditional generation of hypothetical new familes of superconductotrs [21], for brain imaging [22], for various bioengineering applications [23], for protein structure generation [24].

The first application of diffusion to particle accelerators was recently developed for providing a virtual XTCAV to give virtual non-invasive views of the 2D longitudinal phase space (LPS) (z, E) of the European XFEL electron beam [25]. In this approach a conditional vector, $\mathbf{c}$, is created out of a combination of accelerator settings (such as RF cavity phase and amplitude set points) together with non-invasive beam measurements (such as current monitors and beam position monitors).

The vector is then used to condition the generative diffusion process $D(\mathbf{c}, t, \theta)$ to accurately predict the LPS of electron beams at the EuXFEL injector, where t is the time-step of the diffusion process and θ represents the weights of the diffusion U-Net. This approach provides a non-invasive virtual view of the initial conditions of the beam at 150 MeV before it is accelerated through the remaining 5 km of the EuXFEL up to 17 GeV. By using time-varying measurements, such as BPM readings and loss/current monitors, the method was shown to be robust for predicting unseen beams over an extended period of time (36 hours). For predictions over longer periods of time, over which much larger distribution shift is expected, adaptive feedback on the conditional vector would be required to track the time-varying beam. The setup is shown in Fig. 4.

A second application of diffusion for particle accelerators was recently published in which a variational autoencoder (VAE) is combined with a diffusion model in an adaptive and multi-modal generative diffusion process. This method is used to generate all 15 $((x, y), (x, z), \ldots, (z, E))$ unique 2D projections of a charged particle beam's 6D phase space distribution at the HiRES compact accelerator UED facility at LBNL [26]. In this setup a variational autoencoder (VAE) embeds input beam images into a latent vector $\mathbf{z}$ which is combined with a vector of accelerator and beam measurements $\mathbf{p}$ as well as a physics-informed encoding of the projection choice $\mathbf{d}$ to define the overall conditional vector $\mathbf{c} = [\mathbf{z}, \mathbf{p}, \mathbf{d}]$.

The process is applied in an unsupervised adaptive approach by first fixing a best guess value of $\mathbf{c}$ because the VAE's time-varying input beam image cannot be measured without interrupting operations and the accelerator and beam parameters $\mathbf{p}(t)$ are time-varying and their measurements are

ISBN: 978-3-95450-219-6 ISSN: 2226-0366 doi: 10.18429/JACoW-LINAC2024-WEXA005

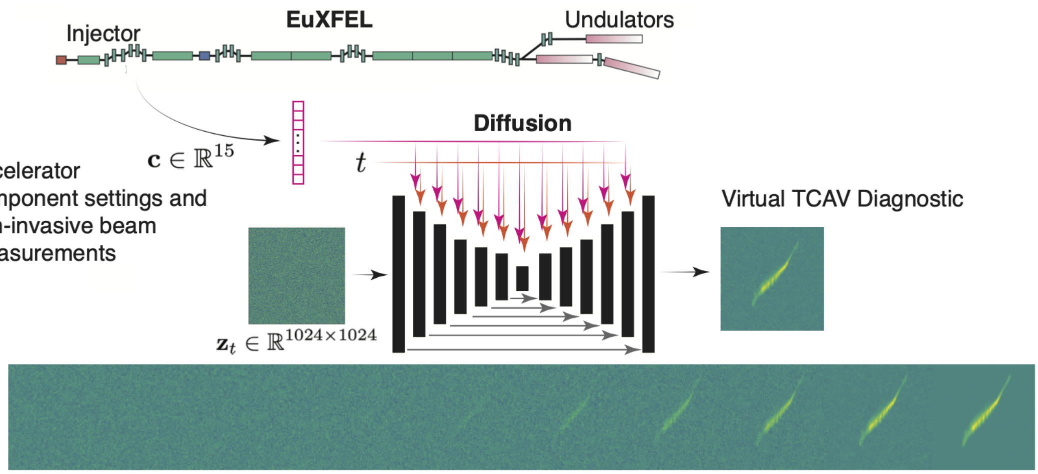

Figure 4: Set up at the EuXFEL for use of conditional diffusion as a virtual TCAV/longitudinal phase space diagnostic based only on accelerator parameter set-points and non-invasive beam measurements.

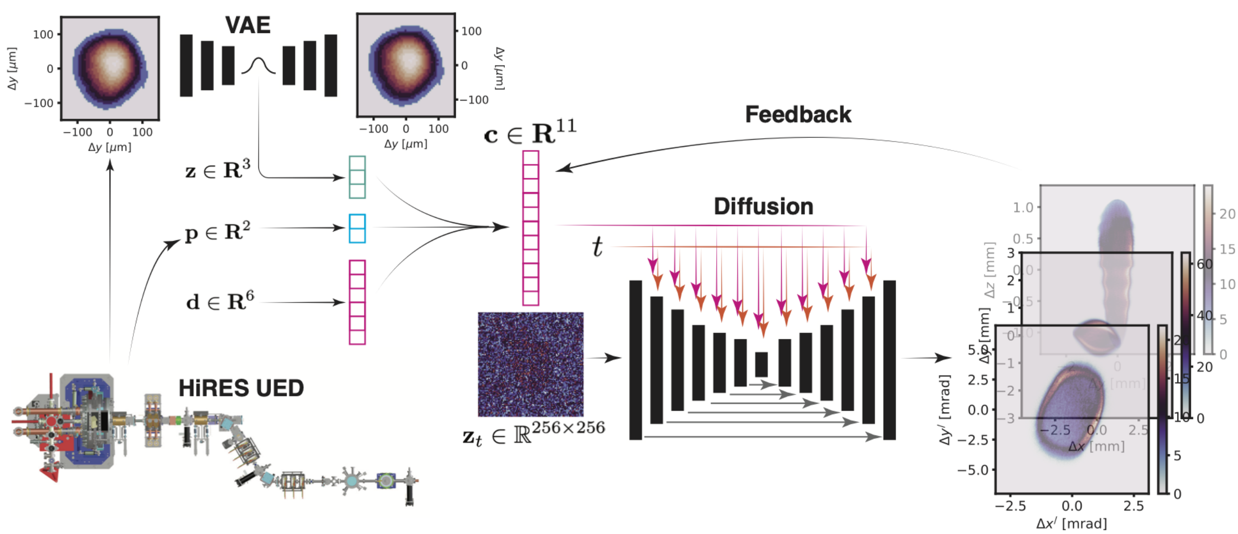

Figure 5: Set up at HiRES for use of adaptive conditional diffusion as a virtual diagnostic of all 15 unique projections of the beam's 6D phase space.

noisy and have arbitrary offsets. The next step is to generate only a single projection from the 15, one which is available for measurement. For example, at the LCLS free electron laser there is a (z, E) LPS projection always available at the end of the accelerator which is non-invasive in the sense that the TCAV and dipole magnet are after the undulator and so all of the X-rays for an experiment have already been created by the electron bunch and they go on to the experimental halls unaffected by the TCAV or dipole that are then used to image the electron beam. In various accelerators (such as FACET-II and LANSCE) there are also non-invasive (x, y) beam measurements available in which a relatively high energy beam is passed through a scintillating screen and measured with the beam continuing down the accelerator with minimal impact on the beam's characteristics.

The model's prediction $\hat{\mathbf{Y}}$ of some beam projection (p_1, p_2) is then compared to the non-invasive measurement $\mathbf{Y}$ of that same projection in order to calculate a metric of the form

$$C = \iint \left(\mathbf{Y} - \hat{\mathbf{Y}}\right)^2 dp_1 dp_2, \tag{2}$$

which is then iteratively minimized by adaptive tuning of the conditional embedding $\mathbf{c}(t)$ which represents the unknown accelerator settings and beam initial conditions. The result is that the model can track all 15 projections of the 6D phsae space of the time-varying beam for unknown and time-varying accelarator and beam parameters and beam initial conditions. The overall setup is shown in Fig. 5.

CONCLUSIONS

The development of generative high-resolution non-invasive virtual phase space diagnostics is important for enabling real-time tuning and optimization and precise control of charged particle beam phase space distributions. We

ISBN: 978-3-95450-219-6 ISSN: 2226-0366 doi: 10.18429/JACoW-LINAC2024-WEXA005

have demonstrated adaptive generative AI methods which can track beams with unknown and time-varying initial conditions, we have developed adaptive AI-based inverse models for estimating time-varying initial beam conditions, we have developed methods for incorporating hard electrodynamics physics constraints within generative models, and we have developed the first conditional diffusion-based beam diagnostics with unprecedented resolution for widely varying beam distributions. Furthermore, all of these AI tools are by default fully differentiable models, enabling easy studies of complex input-output dependencies.

REFERENCES

[1] 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", *Physical Review Applied*, vol. 16, no. 2, p. 024 005, 2021. doi:10.1103/PhysRevApplied.16.024005

[2] A. Wolski, M. A. Johnson, M. King, B. L. Militsyn, and P. H. Williams, "Transverse phase space tomography in an accelerator test facility using image compression and machine learning", *Physical Review Accelerators and Beams*, vol. 25, no. 12, p. 122 803, 2022. doi:10.1103/PhysRevAccelBeams.25.122803

[3] F. Mayet *et al.*, "Predicting the transverse emittance of space charge dominated beams using the phase advance scan technique and a fully connected neural network", *Physical Review Accelerators and Beams*, vol. 25, no. 9, p. 094 601, 2022. doi:10.1103/PhysRevAccelBeams.25.094601

[4] M. Maheshwari, D. Dunning, J. Jones, M. King, H. Kockelbergh, A. Pollard, *et al.*, "Prediction and clustering of longitudinal phase space images and machine parameters using neural networks and k-means algorithm", *Proc. IPAC 2021*, pp. 3417–3420, 2021. doi:doi:10.18429/JACoW-IPAC2021-WEPAB318

[5] A. Hoover and J. C. Wong, "High-dimensional maximum-entropy phase space tomography using normalizing flows", *Physical Review Research*, vol. 6, no. 3, p. 033 163, 2024. doi:10.1103/PhysRevResearch.6.033163

[6] A. Scheinker, "Adaptive machine learning for time-varying systems: Low dimensional latent space tuning", *Journal of Instrumentation*, vol. 16, no. 10, p. P10008, 2021. doi:10.1088/1748-0221/16/10/P10008

[7] A. Scheinker, F. Cropp, S. Paiagua, and D. Filippetto, "An adaptive approach to machine learning for compact particle accelerators", *Scientific Reports*, vol. 11, no. 1, p. 19 187, 2021. doi:10.1038/s41598-021-98785-0

[8] A. Scheinker and M. Krstić, "Minimum-seeking for CLFs: Universal semiglobally stabilizing feedback under unknown control directions", *IEEE Transactions on Automatic Control*, vol. 58, no. 5, pp. 1107–1122, May 2013. doi:10.1109/tac.2012.2225514

[9] 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", *Physical Review Letters*, vol. 121, no. 4, p. 044 801, Jul. 2018. doi:10.1103/physrevlett.121.044801

[10] A. Scheinker, F. Cropp, S. Paiagua, and D. Filippetto, "An adaptive approach to machine learning for compact particle accelerators", *Scientific Reports*, vol. 11, no. 1, p. 19 187, 2021. doi:10.1038/s41598-021-98785-0

[11] A. Scheinker, F. Cropp, and D. Filippetto, "Adaptive autoencoder latent space tuning for more robust machine learning beyond the training set for six-dimensional phase space diagnostics of a time-varying ultrafast electron-diffraction compact accelerator", *Physical Review E*, vol. 107, no. 4, p. 045 302, 2023. doi:10.1103/PhysRevE.107.045302

[12] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning", *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021. doi:10.1038/s42254-021-00314-5

[13] A. D. Jagtap, Z. Mao, N. Adams, and G. E. Karniadakis, "Physics-informed neural networks for inverse problems in supersonic flows", *Journal of Computational Physics*, vol. 466, p. 111 402, 2022. doi:10.1016/j.jcp.2022.111402

[14] A. Scheinker and R. Pokharel, "Physics-constrained 3D convolutional neural networks for electrodynamics", *APL Machine Learning*, vol. 1, no. 2, 2023. doi:10.1063/5.0132433

[15] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics", in *International Conference on Machine Learning*, pp. 2256–2265, 2015. doi:10.48550/arXiv.1503.03585

[16] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models", *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020. doi:10.48550/arXiv.2006.11239

[17] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models", 2020. doi:10.48550/arXiv.2010.02502

[18] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models", in *International Conference on Machine Learning*, pp. 8162–8171, 2021. https://proceedings.mlr.press/v139/nichol21a.html.

[19] T. Karras, M. Aittala, T. Aila, and S. Laine, "Elucidating the design space of diffusion-based generative models", *Advances in Neural Information Processing Systems*, vol. 35, pp. 26 565–26 577, 2022. doi:10.48550/arXiv.2206.00364

[20] L. Yang *et al.*, "Diffusion models: A comprehensive survey of methods and applications", *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–39, 2023. doi:10.1145/3626235

[21] S. Yuan and S. Dordevic, "Diffusion models for conditional generation of hypothetical new families of superconductors", *Scientific Reports*, vol. 14, no. 1, p. 10 275, 2024. doi:10.1038/s41598-024-61040-3

[22] W. H. Pinaya *et al.*, "Brain imaging generation with latent diffusion models", in *MICCAI Workshop on Deep Generative Models*, pp. 117–126, 2022. doi:10.1007/978-3-031-18576-2_12

ISBN: 978-3-95450-219-6 ISSN: 2226-0366 doi: 10.18429/JACoW-LINAC2024-WEXA005

[23] Z. Guo *et al.*, "Diffusion models in bioinformatics and computational biology", *Nature Reviews Bioengineering*, vol. 2, no. 2, pp. 136–154, 2024. doi:10.1038/s44222-023-00114-9

[24] K. E. Wu *et al.*, "Protein structure generation via folding diffusion", *Nature Communications*, vol. 15, no. 1, p. 1059, 2024. doi:10.1038/s41467-024-45051-2

[25] A. Scheinker, "Conditional guided generative diffusion for particle accelerator beam diagnostics", *Scientific Reports*, vol. 14, p. 19 210, 2024. doi:10.1038/s41598-024-70302-z

[26] A. Scheinker, "cDVAE: Multimodal generative conditional diffusion guided by variational autoencoder latent embedding for virtual 6d phase space diagnostics", 2024. doi:10.48550/arXiv.2407.20218