AI-ENABLED DIGITAL TWINS AND OPTIMIZATION WORKFLOWS FOR ACCELERATOR CONTROL[*]

J. Bird[1,4], T. Griffin[2,4], K. Makino[1,4], K. R. Ahmed[1,4], S. Hossain[1,4],
J. R. Delayen[1,4], S. U. De Silva[1,4], F. Liu[2], A. Seryi[1], M. Sosonkina[3],
T. Satogata[4,5], Q. Su[3,4], B. Terzić[1,4], L. Van Riesen-Haupt[1,4], M. Yadav[†1,2,4]
[1]Department of Physics, Old Dominion University, Norfolk, VA, USA
[2]School of Data Science, Old Dominion University, Norfolk, VA, USA
[3]Department of Electrical and Computer Engineering, Old Dominion University, Norfolk, VA, USA
[4]Center for Accelerator Science, Old Dominion University, Norfolk, VA, USA
[5]Thomas Jefferson National Accelerator Facility, Newport News, VA, USA

Abstract

We propose to develop advanced machine learning (ML) models, such as physics informed neural network (PINN) based surrogate models, to accurately represent accelerator phase space transport. These surrogate models will enable precise diagnosis and prediction of beam phase space evolution along the beamline, facilitating real-time control and optimization. The developed models will be tested using the Upgraded Injector Test Facility (UITF) at Thomas Jefferson National Accelerator Facility (JLab), providing a pathway toward ML-driven enhanced diagnostics and beamline control in operational accelerator environments. The primary aim will be to facilitate this by developing machine learning models that outperform traditional simulations in speed and precision. We will build a virtual beamline, train a reinforcement learning (RL) controller across varied calibration scenarios, and then transfer it to the real machine. Beyond operation, fast and accurate models are also essential for design optimization workflows using machine learning methods that iterate through design parameters. A long-term goal of this work will be to establish such workflows and apply them to the design of a compact accelerator at Old Dominion University (ODU).

INTRODUCTION

Particle accelerators are essential tools in modern science and technology, supporting applications ranging from high-energy physics and materials science to medical and industrial imaging. The development of next-generation accelerator systems requires improved beam control, enhanced efficiency, and reduced operational costs. In alignment with the Genesis Mission, these challenges motivate the exploration of advanced computational, artificial-intelligence, and data-driven approaches for accelerator modeling, diagnostics, and control.

Recent progress in machine learning (ML) and high-performance computing (HPC) have enabled new strategies for accelerator modeling, design, and control. In particular, ML-based surrogate models and data-driven optimization techniques offer the potential to significantly expedite simulation workflows and enable real-time diagnostics in complex beamline environments. Digital twins are virtual replicas of accelerators that continuously link modeled behavior with observed performance to facilitate real-time optimization and minimize experimental overhead [1–4].

To realize this vision, we propose a comprehensive research framework that leverages interdisciplinary expertise from accelerator physicists, plasma physicists, and data scientists. The foundation of our approach relies on robust infrastructure, utilizing ODU's HPC resources and industry-standard codes (such as Elegant [5] and QuickPIC [6]) for simulations and ML operations. Experimental data and validation are supported by ODU's Center for Accelerator Science (CAS) and measured beam profiles from Jefferson Lab accelerators.

The framework is intended to support three related goals: AI-enhanced diagnostics and control, development of an on-campus compact accelerator testbed at ODU, and validation of digital-twin methods using operational accelerator facilities.

NEAR-TERM RESEARCH STRATEGY AND UITF SYNERGY

The long-term goal of this effort is to support an advanced accelerator testbed at ODU's Center for Accelerator Science, including a conventional linac and possible plasma wakefield acceleration stages. As a near-term step, the project focuses on conventional linac modeling and control. The design parameters of the planned ODU linac are closely related to the UITF at Jefferson Lab, making UITF an ideal proxy for model development, data collection, and validation. The schematic layout of the UITF is illustrated in Fig. 1. Our near-term workflow has three components: collecting experimental data from UITF, refining simulation models of the beamline, and developing ML tools trained on both simulation outputs and measured data.

CURRENT EFFORTS AND RESULTS

ODU provides a strong environment for this work through its Accelerator Cluster, which combines expertise from the Center for Accelerator Science, the Departments of Physics

[*] Work supported by ODU STaR grant
[†] myadav@odu.edu

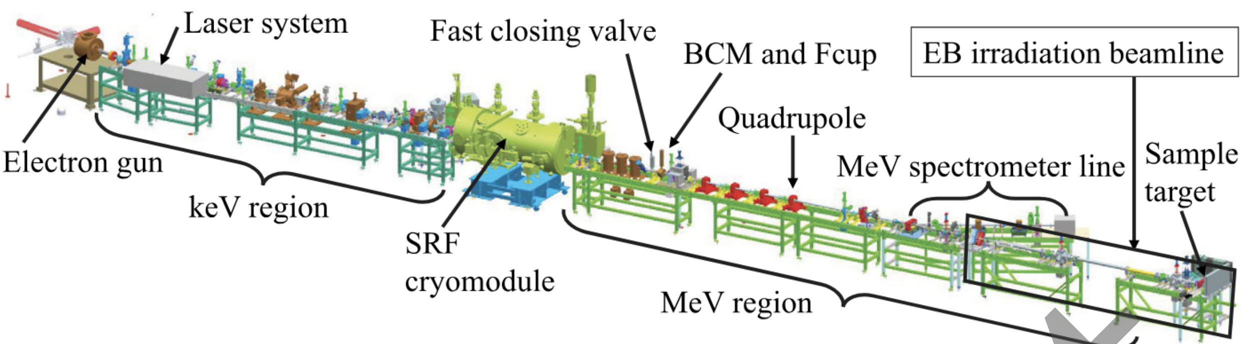

Figure 1: UITF at Jefferson Lab for testing new algorithms and framework including the ability of these models to transfer between facilities.

and Electrical & Computer Engineering, and the School of Data Science. In collaboration with Jefferson Lab, the project uses simulation tools such as Elegant (Fig. 2) and QuickPIC together with experimental validation from CEBAF and UITF.

Current efforts focus on developing a digital twin for accelerator beamlines. Beam dynamics models are being constructed and validated against measured behavior, while ML tools such as surrogate models and phase-space reconstruction methods are being developed to enable rapid beam prediction and future optimization [7, 8].

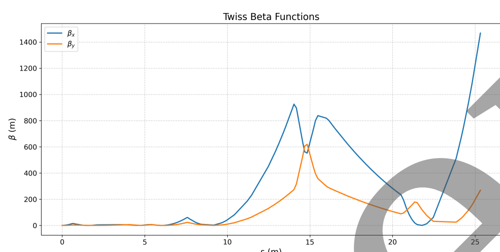

Figure 2: Twiss parameters generated using elegant.

Physics Simulations

The digital twin effort begins with a reliable simulation framework for the UITF beamline. Beam dynamics simulations are being performed using Elegant, with lattice configurations and operating parameters chosen to reflect the physical machine. Initial studies focus on validating lattice models and benchmarking simulated observables, such as beam position monitor (BPM) responses.

Interfaces between Elegant and Python-based analysis tools are being developed to automate parameter scans, process simulation outputs, and connect simulations to ML workflows. Simulations are also being deployed on HPC resources to support large-scale parameter studies and future measurement-informed calibration.

Machine Learning

Within the digital-twin framework, we are developing surrogate modeling and optimization methods for accelerator systems. A key direction is phase-space reconstruction, where ML models infer difficult-to-measure beam

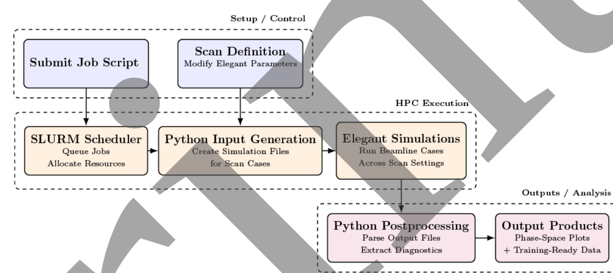

Figure 3: ODU Campus HPC workflow for the elegant simulation code for dataset generation.

properties from limited diagnostics. This is important because full phase-space measurements are often invasive, and simulation-to-experiment discrepancies complicate model calibration. To address these challenges, we plan to investigate physics-informed models [9] trained on simulation data generated using accelerator simulation programs such as Elegant. As preparatory work, we have developed a workflow (Fig. 3) on the ODU campus HPC platform that can programmatically modify Elegant parameters and automate large-scale simulation studies.

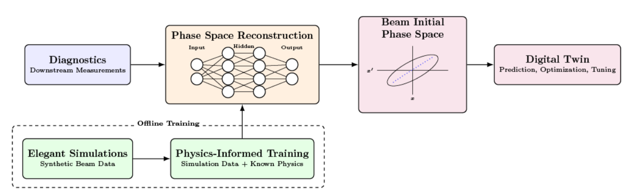

Figure 4: Workflow for phase-space reconstruction using PINN.

These models are planned to learn mappings between diagnostic signals and underlying beam dynamics while incorporating the physical constraints of the accelerator system. One example direction is beam phase-space reconstruction (Fig. 4) in situations where the full phase space cannot be directly measured [10]. Experimental measurements are planned to be used to validate and calibrate the model outputs for real-world accelerator applications. Within the broader digital twin framework, the trained models are also planned to serve as fast surrogate models for beam dynamics prediction and accelerator optimization, helping maintain consistency between simulated and observed system behavior despite limited direct measurement data.

Data Collection and Generation

The UITF serves as the primary experimental platform and validation source for this effort. As a compact linear accelerator capable of delivering spin-polarized electron beams at energies up to 10 MeV, its reconfigurable design supports specialized low-energy beam physics studies. These characteristics make the UITF uniquely suited for benchmarking *Elegant* beam dynamics simulations against measured machine behavior, providing the necessary empirical foundation for digital twin development.

Data collection for model calibration and machine learning training relies on the integration of several key sources. The JLab Archiver provides a high-fidelity, continuous history of machine Process Variables (PVs), such as magnet currents, vacuum levels, and RF gradients. These data will be used to calibrate the simulation model and provide training inputs for ML-based digital-twin development. A central goal is to establish a direct connection between the numerical model and archived machine settings so that simulations can reproduce specific operational states.

In addition, planned UITF studies will investigate false-positive RF cavity triggers, potentially linked to ion trapping and release in the beam pipe. Understanding these effects will help refine the lattice model, account for non-ideal machine behavior, and improve the data pipeline needed for real-time optimization and diagnostic support.

Plasma Wake Field Acceleration

The digital-twin framework will also be extended toward plasma wakefield acceleration (PWFA) as a candidate approach for compact accelerator sources. PWFAs can sustain accelerating gradients on the order of GV/m, allowing compact acceleration over centimeter-scale distances. A key challenge is to model, diagnose, and tune these systems with sufficient accuracy for stable beam production.

Large-scale plasma simulation tools such as QuickPIC and OSIRIS [11] will be used to generate datasets for surrogate modeling, parameter optimization, and non-invasive diagnostics. An optimization framework has been developed in [12] that optimize the beam current profile from QuickPIC simulations (Fig. 5). In particular, betatron radiation emitted by electrons in the wakefield can provide information about beam divergence, oscillation amplitude, and stability. These signals can be used to train ML models for beam characterization and digital-twin development [13–17].

Expected outcomes include identification of plasma and driver regimes for efficient acceleration, improved understanding of beam-quality preservation, and quantitative prediction of betatron-radiation signatures for ML-driven [18] diagnostics.

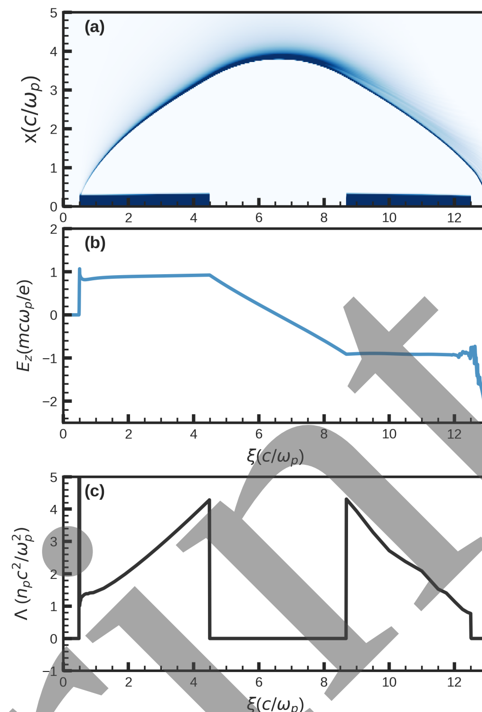

Figure 5: (a) Charge density of plasma electrons excited by the optimized drive beam and the optimized trailing beam. (b) On-axis accelerating field given the optimized current profile. (c) The optimized beam density per unit length for the drive and trailing beams.

OUTLOOK AND CONCLUSION

We have outlined an AI-enabled digital-twin framework that connects accelerator simulations, experimental data, and machine-learning models for beam prediction, diagnostics, and future control. Initial work focuses on UITF as a validation platform, with longer-term extensions toward compact accelerator design and plasma-based acceleration. This effort provides a foundation for scalable digital-twin workflows that can support both accelerator operations and advanced accelerator R&D.

ACKNOWLEDGEMENTS

This work is supported by the Old Dominion University Division of Research and Economic Development Stimulating Transdisciplinary Research (STaR) Collaboration Program. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Contract No. DE-AC05-06OR23177. This work is also supported by NSF CSSI Grant No. 2513760. Computational and storage resources were provided by the Wahab and Turing clusters at ODU. Additional computational resources were provided by NERSC under Grant No. ERCAP0035820.

REFERENCES

[1] J. Qiang, “Differentiable self-consistent space-charge simulation for accelerator design”, *Phys. Rev. Accel. Beams*, vol.

26, no. 2, p. 024601, 2023.
doi:10.1103/PhysRevAccelBeams.26.024601

[2] Y. Wang, Y. Yin *et al.*, “Semi-supervised learning of dynamical systems with neural ordinary differential equations: A teacher-student model approach”, in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 14, 2024, pp. 15698–15705.
doi:10.1609/aaai.v38i14.29498

[3] J. Kaiser *et al.*, “Bridging the gap between machine learning and particle accelerator physics with high-speed, differentiable simulations”, *Phys. Rev. Accel. Beams*, vol. 27, no. 5, p. 054601, 2024.
doi:10.1103/PhysRevAccelBeams.27.054601

[4] A. Edelen *et al.*, “Opportunities in machine learning for particle accelerators”, 2018.
doi:10.48550/arXiv.1811.03172

[5] M. Borland, “ELEGANT: A flexible SDDS-compliant code for accelerator simulation”, Argonne National Laboratory, IL, USA, Rep. LS-287, 2000. doi:10.2172/761286

[6] W. An *et al.*, “An improved iteration loop for the three dimensional quasi-static particle-in-cell algorithm: QuickPIC”, *J. Comput. Phys.*, vol. 250, pp. 165–177, 2013.
doi:10.1016/j.jcp.2013.05.020

[7] J. Gonzalez-Aguilera *et al.*, “Towards fully differentiable accelerator modeling”, in *Proc. IPAC’23*, Venice, Italy, May 2023, pp. 2797–2800.
doi:10.18429/JACoW-IPAC2023-WEPA065

[8] T. Hellert *et al.*, “Agentic artificial intelligence for multistage physics experiments at a large-scale user facility particle accelerator”, *Phys. Rev. Res.*, vol. 8, no. 1, p. L012017, 2026.
doi:10.1103/jtqy-9jz1

[9] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”, *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019. doi:10.1016/j.jcp.2018.10.045

[10] C. Emma *et al.*, “Machine learning-based longitudinal phase space prediction of particle accelerators”, *Phys. Rev. Accel. Beams*, vol. 21, no. 11, p. 112802, 2018.
doi:10.1103/PhysRevAccelBeams.21.112802

[11] R. A. Fonseca *et al.*, “OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators”, in *Comput. Sci.*, Springer, Berlin, Heidelberg, 2002, pp. 342–351, doi:10.1007/3-540-47789-6_36

[12] Q. Su *et al.*, “Optimization of transformer ratio and beam loading in a plasma wakefield accelerator with a structure-exploiting algorithm”, *Phys. Plasma*, vol. 30, no. 5, p. 053108, 2023, doi:10.1063/5.0142940

[13] M. Yadav *et al.*, “Reconstruction of beam parameters and betatron radiation spectra measured with a Compton spectrometer”, *Phys. Rev. Accel. Beams*, vol. 28, no. 4, p. 042802, 2025. doi:10.1103/PhysRevAccelBeams.28.042802

[14] M. Schram *et al.*, “Uncertainty aware machine-learning-based surrogate models for particle accelerators: Study at the Fermilab Booster Accelerator Complex”, *Phys. Rev. Accel. Beams*, vol. 26, no. 4, p. 044602, 2023.
doi:10.1103/PhysRevAccelBeams.26.044602

[15] M. Yadav *et al.*, “Reconstructing Gamma-ray Energy Distributions from PEDRO Pair Spectrometer Data”, 2024.
doi:10.48550/arXiv.2409.02113

[16] A. Hanuka *et al.*, “Physics model-informed Gaussian process for online optimization of particle accelerators”, *Phys. Rev. Accel. Beams*, vol. 24, no. 7, p. 072802, 2021.
doi:10.1103/PhysRevAccelBeams.24.072802

[17] L. P. Kaelbling, M. L. Littman, and A. W. Moore, “Reinforcement learning: A survey”, *J. Artif. Intell.*, vol. 4, p. 237, 1996.
doi:10.48550/arXiv.cs/9605103

[18] J. Strube *et al.*, “Artificial Intelligence for improved facilities operation in the FNAL LINAC”, in *Proc. IPAC’23*, Venice, Italy, May 2023, pp. 4491–4494.
doi:10.18429/JACoW-IPAC2023-THPL033