ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL010

SUMMARY OF THE 3rd ICFA BEAM DYNAMICS MINI-WORKSHOP ON MACHINE LEARNING APPLICATIONS FOR PARTICLE ACCELERATORS*

K. A. Brown†[1], C-A Dept., Brookhaven National Laboratory, Upton, NY, USA
S. G. Biedron[2,3], Element Aero, Chicago, IL, USA
[1] also at ECE Dept., Stony Brook University, Stony Brook, NY, USA
[2] also at the Center for Bright Beams, Cornell University, NY, USA
[3] also at ECE and ME Depts., University of New Mexico, NM, USA

Abstract

The 3rd ICFA Beam Dynamics Mini-Workshop on Machine Learning (ML) Applications for Particle Accelerators was held in Chicago, Il, USA, on November 1-4, 2022. This was an in-person workshop focused on ML techniques as applied to accelerator operations, design, and simulations. There were 76 attendees representing 26 institutions from around the world. A total of 59 abstracts were submitted allowing us to build a diverse program with both oral and poster presentations. The workshop was sponsored by the Center for Bright Beams (CBB), with support from the National Science Foundation and by RadiaSoft, an industry leader in high-level research and design and scientific consulting for beamline physics and machine learning. CBB supported eight graduate students for this meeting. The workshop was approved as a mini workshop by the International Committee for Future Accelerators (ICFA) Beam Dynamics Panel. In this report we provide a summary of the workshop and directions of future efforts.

BRIEF HISTORY

At the time of the 1st ICFA Beam Dynamics Mini-Workshop on Machine Learning Applications for Particle Accelerators, in early 2018 [1, 2], there was growing interest in the topic but only a few groups seriously researching methods. There was some notable work prior to 2018, including the use of ML to improve beam collimation at the LHC [3,4], the use of Bayesian Optimization (BO) at SLAC [5], use of Genetic Algorithms for chromaticity correction [6], optimizing nonlinear dynamics with particle swarm and genetic algorithms [7], and use of genetic algorithms for maximizing dynamic aperture [8]. Although there were discussions at conferences and at various laboratories on the topic prior to 2010 [9–14], general tools and computing resources were still lagging the demands of particle accelerator systems. What is encouraging is many people were exploring the use of neural networks to model accelerator problems and some were having significant success. Between 2010 and today there have been hundreds of papers presented at conferences and dozens of papers published in peer reviewed journals. The 2nd ICFA workshop was held a little over one year later, in 2019 [15], demonstrating how, over just one-year, impressive progress was made at many labs. The 3rd workshop was delayed due to the Covid19 pandemic and was finally able to take place in Chicago in November of 2022. With over two and one-half years after the second workshop, this third workshop showed that the pandemic did not slow us down and the progress over that long period has shown the technologies for using ML in accelerators are now becoming much more sophisticated and mature. There are now two open-source repositories of tools people can use, XOpt/Badger based at SLAC [16], and COI/GeOFF based at CERN [17].

The 3rd ICFA Beam Dynamics Mini-Workshop on Machine Learning Applications for Particle Accelerators was held from 1 to 4 November 2022 in Chicago, IL, USA, as an in-person event [18]. The workshop had representation from all regions. As is the tradition for this workshop series, the first day was dedicated to tutorials and demonstrations. The goal of this workshop was to continue to work on building a world-wide community of researchers and engineers interested in applying artificial intelligence and machine learning technologies to particle accelerators. All the presentations, posters, and session summaries are available on the workshop website as a record of the presentations and discussions.

The community of accelerator scientists and engineers focusing on machine learning applications is strong and supportive. This was clear during the first day of tutorials and collaboration discussions. The tutorials were well prepared, with help from RadiaSoft, who provided a cloud-based platform for attendees to work through examples. ML tools are starting to be developed and shared. XOpt/Badger and the COI/GeOFF ML toolsets are being adopted by other laboratories and collaborations are growing.

The types of problems in accelerator systems can be split into four topics: Tuning/optimization/control, Prognostics/alarm handling/anomaly-breakout detection, Data analysis, and Simulations/modeling.

This workshop series aims to collect and unify the community's understanding of the relevant state-of-the-art ML techniques, provide tutorials on machine learning for accelerator physicists and engineers, and seed collaborations between laboratories, academia, and industry.

* Work supported Under Contract No. DE-Sc0012704 with The Auspices of the U.S. Department of Energy
† brownk@bnl.gov

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL010

3rd MINI-WORKSHOP ON ML FOR ACCELERATORS

The first day of the workshop was dedicated to tutorials and the day ended with presentations and discussions on building collaboration, with a focus on the XOpt/Badger and COI/GeOFF toolsets. The next two and a half days were divided into nine sessions: BO, Reinforcement Learning (RL), Prognostics, Modeling, Facilities, Optimization, Analysis, a poster session, and summaries from session chairs. There was an equal number of talks and posters covering a wide range of topics. The poster session was a vital part of the workshop, and presenters were encouraged to embrace a better poster design process to feature a core message and maximize the insights transmitted from their work. All presentations, whether oral or poster represented very high-quality work.

Tutorials and Collaboration

For this instance of the workshop, we decided that the community was more advanced and so we focused the Tutorials on a few specific topics.

Bayesian optimization is one of the most widely used methods that applies to many accelerator problems. Ryan Roussel, SLAC, provided an excellent overview on BO methods and demonstrated how to apply the methods to accelerator problems. An excellent paper published in PRAB by Ryan and his collaborators covers the topic of multi-objective BO [19], which he discussed in the tutorial. There are many recent examples of applying BO for accelerator problems; BO to optimize trajectory alignment for electron cooling at RHIC [20], BO for a recoil mass separator [21], BO for tuning with safety constraints [22], and BO for the beam injection process [23].

Antonin Sulc, DESY, went through anomaly detection (AD) methods. Anomaly detection is a popular topic since if offers the promise of detecting problems more quickly. There are different classes of anomalies; point anomalies, group anomalies, contextual anomalies, hidden anomalies, and more. Good examples of applications of anomaly detection in accelerators are Uncertainty aware AD to predict errant beam pulses [24] and AD at the European XFEL [25].

Kishansingh Rajput, TJNAF, gave an excellent hands-on tutorial on Reinforcement Learning. This ML method is extremely important for optimization problems where the system has some state that is influenced by its environment. A classic example is autonomous cars, where the state of the car must follow its awareness of the environment (e.g., perhaps via a camera). For accelerators there are many problems that fit well into such a model. At CERN they are studying sample-efficient reinforcement learning for the accelerator controls [26]. At FNAL they are looking at real-time correction in the Fermilab Booster [27]. For superconducting rf guns, a study of deep reinforcement learning for the SeaLab project shows how it can be used to optimize multiple parameters quickly and efficiently [28]. And for the medium energy beam transport section of the China Accelerator Facility for Superheavy Elements they have used reinforcement learning for orbit correction [29].

Natalie Isenberg, BNL, presented an introduction to the important topic of uncertainty quantification (UQ) for ML applications. The aim of UQ is to place statistics on the models. This includes spatiotemporal statistics, sensitivity analysis, uncertainty propagation, parameter estimation, data assimilation, as well as other inverse problems. This topic is beginning to get more attention in the accelerator community as people gain experience using ML applications. At SLAC they studied how to quantify uncertainty in virtual diagnostics [30]. They also looked more generally at UQ for deep learning applications [31].

Corey Adams of Argonne National Laboratory presented an overview of the Argonne Leadership Computing Facility, a DOE user facility, that has many resources for AI/ML research and activities enabled by the available hardware and hardware coming online [32].

Bayesian Optimization

BO has many applications for accelerator systems, as it is for global optimization of black box functions. We quite often don't have physics or even detailed engineering models for the many accelerator subsystems used to control our machines. We typically don't need such models, as we have operators who learn how to get the machines to perform as needed to meet experimental needs. These statistics-based models allow strategies to be developed to help optimized subsystems quickly and dependably. Three excellent examples of BO were presented in the first session. SLAC is using BO to tune particle accelerator emittance with partial measurements (results shown in Fig. 1). At Argonne they are using BO for the ATLAS ion LINAC to improve performance. And at BNL, they used BO to align beam trajectories for the low energy electron cooler at RHIC.

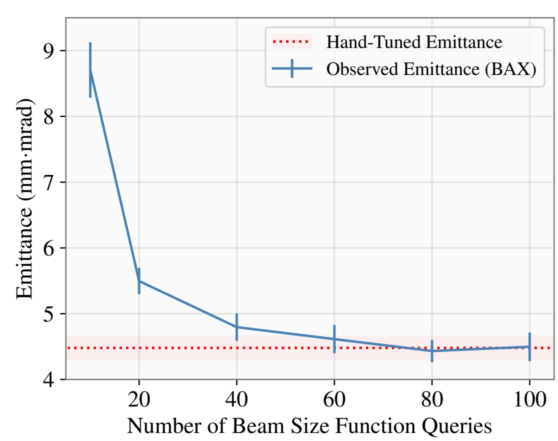

Figure 1: Live optimization of the emittance at FACET-II (courtesy of Sara Mishovich.)

Reinforcement Learning

There were three talks on RL. At FNAL they are using ML to regulate the extracted beam spill in the Fermilab De-

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL010

livery ring. Their approach is to construct a differentiable simulator that trains a neural network to regulate the spill extraction rate instantaneously. Preliminary results show such an approach can outperform a PID controller. At CERN, they are using ML to optimize rf manipulations to split and merge bunches. The parameters to adjust are the rf phase and amplitude. At the University of Malta, they are looking generally at ways to apply RL in beam-based feedback systems.

Prognostics

Fault detection and prediction is notoriously challenging, as we are hoping that some algorithm or autonomous entity can monitor signals for us and alert us when something doesn't look correct. Such an entity can tirelessly monitor that signal and watch for anomalies. It may even detect subtle changes that seriously deviate from past behavior. In this session four talks covered different methods for fault detection and discussed the challenges. At TJNAL ML is used for improved SRF operation at CEBAF. These systems are complex, with thousands of signals to monitor, and must operate under strict stability guidelines. At SLAC, beam-based rf station fault identification is used at LCLS. At DESY, they are working on beam trajectory anomaly detection. Fermilab is disentangling beam losses using real-time edge computing. And at TJNAF, they are developing uncertainty aware anomaly detection to predict errant beam pulses.

Modeling

Modeling covers Data models as well as accelerator modeling. BNL is using ML for nonlinear dynamics for storage ring design. At the University of Chicago ML methods are being developed for end-to-end differentiable accelerator modeling. SLAC described their experience with integrating online physics models, adaptive ML models, and model-based controls. And BNL presented how ML tools are being developed for modeling luminosity tuning for the future electron ion collider.

Facility

Facility talks were a little more diverse and tended to focus on operations. By integrating natural language processing into electronic logbooks BNL has greatly improved the power of logbooks. This goes well beyond simple search but allows the logbook to learn from each end user the things they most need to know as well as learn what the overall operations team needs at any given moment. At CERN they are looking at how ML can lead towards more autonomous accelerators. Also, at CERN they are working on getting ML into the control room, with ML frameworks under development and in use. At BNL data analysis and control of an MeV ultrafast electron diffraction system is being improved using ML.

Optimization

In this session four talks discussed more complex optimization problems in accelerators. Using RL at the Karlsruhe Research Accelerator (KARA) and the Ferninfrarot Linac- und Test Experiment (FLUTE), along with surrogate models and parallel BO, they are tackling many different problems to gain high stability beams, pulse optimization, virtual diagnostics, and improved operations. Fermilab is working on optimizing the performance of the LINAC RF systems using ML. At PETRA III (DESY), they are using ML for insertion device gap compensation. At the APS, ANL, they are using ML to help resolve operations problems in optimization and anomaly prediction.

Analysis

The analysis session tended to focus on applying ML to different kinds of measurements. At SLAC they are using Neural Networks and differential simulations to reconstruct the phase space from beam measurements. CERN is developing physics informed neural networks for improving neural network predictions. At FRIB, they are developing ML tools for transverse 2D phase-space tomography using beam position measurements. Cornell is collaborating with BNL to use ML in simulation studies and orbit response measurements to improve the AGS models.

THE NEXT WORKSHOP

We are pleased to announce that the 4$^{\text{th}}$ Workshop on Machine Learning Applications for Particle Accelerators will be held 5-8 March, 2024 in Gyeongju, Republic of Korea. More information can be found at www.indico.kr/e/ml2024/.

ACKNOWLEDGEMENTS

This event was partially supported by the U.S. National Science Foundation under Award PHY-1549132, the Center for Bright Beams. We wish to thank Anna Petway (BNL), Christina Blas (BNL), Bryan Callaghan (BNL), and Linh Nguyen (BNL) in supporting the workshop organization. We also acknowledge support from RadiaSoft.

REFERENCES

[1] A. Edelen *et al.*, "Opportunities in Machine Learning for Particle Accelerators", `doi:10.48550/arXiv.1811.03172`

[2] ICFA Beam Dynamics Mini-Workshop: Machine Learning Applications for Particle Accelerators, `https://conf.slac.stanford.edu/icfa-ml-2018/`.

[3] G. Valentino, R. Aßmann, S. Redaelli, and N. Sammut, "Simulator for beam-based LHC collimator alignment", *Phys. Rev. ST Accel. Beams*, vol. 17, p. 021003, Feb. 2014. `doi:10.1103/PhysRevSTAB.17.021003`

[4] G. Valentino, R. Aßmann, R. Bruce, S. Redaelli, A. Rossi, N. Sammut, and D. Wollmann, "Semiautomatic beam-based LHC collimator alignment", *Phys. Rev. ST Accel. Beams*, vol. 15, p. 051002, May 2012. `doi:10.1103/PhysRevSTAB.15.051002`

[5] M. McIntire, D. Ratner, and S. Ermon, "Sparse Gaussian Processes for Bayesian Optimization", in *Proc. UAI 2016*, New York City, NY, USA, Jun. 2016. `http://auai.org/uai2016/proceedings/papers/269.pdf`

[6] M. P. Ehrlichman, "Genetic algorithm for chromaticity correction in diffraction limited storage rings", *Phys. Rev. Accel. Beams*, vol. 19 p. 044001, 2016. doi:10.1103/PhysRevAccelBeams.19.044001

[7] X. Huang and J. Safranek, "Nonlinear dynamics optimization with particle swarm and genetic algorithms for SPEAR3 emittance upgrade", *Nucl. Instrum. Meth. A*, 757, pp. 48-53, 2014. doi:10.1016/j.nima.2014.04.078

[8] Y. Li, W. Cheng, L.H. Yu, and R. Rainer, "Genetic algorithm enhanced by machine learning in dynamic aperture optimization", *Phys. Rev. Accel. Beams*, vol. 21, p. 054601, 2018. doi:10.1103/PhysRevAccelBeams.21.054601

[9] J. E. Spencer, "Accelerator Diagnosis and Control by Neural Net", in *Proc. PAC'89*, Chicago, IL, USA, Mar. 1989, pp. 1642–1645.

[10] D. Nguyen, M. Lee, R. Sass, and H. Shoaee, "Accelerator and Feedback Control Simulation Using Neural Networks", in *Proc. PAC'91*, San Francisco, CA, USA, May 1991, pp. 1437–1440. doi:10.1109/PAC.1991.164660

[11] Y. Kijima, M. Mizota, K. Suzuki, and K. Yoshida, "A Beam Diagnostic System for Accelerator using Neural Networks", in *Proc. EPAC'92*, Berlin, Germany, Mar. 1992, pp. 1155–1158.

[12] E. Bozoki and A. Friedman, "Neural Networks and Orbit Control in Accelerators", in *Proc. EPAC'94*, London, UK, Jun.-Jul. 1994, pp. 1589–1592.

[13] D. Schirmer, T. Buening, P. Hartmann, and D. Mueller, "Electron Transport Line Optimization using Neural Networks and Genetic Algorithms", in *Proc. EPAC'06*, Edinburgh, UK, Jun. 2006, paper WEPCH013.

[14] M. Pieck, "Artificial Intelligence Research in Particle Accelerator Control Systems for Beam Line Tuning", in *Proc. LINAC'08*, Victoria, Canada, Sep.-Oct. 2008, paper MOP103, pp. 314–316.

[15] 2nd ICFA Mini-Workshop on Machine Learning for Charged Particle Accelerators, https://indico.psi.ch/event/6698/.

[16] XOpt/Badger, https://github.com/slaclab/Badger

[17] CERNML-COI, https://gitlab.cern.ch/geoff/cernml-coi

[18] 3rd ICFA Beam Dynamics Mini-Workshop on Machine Learning Applications for Particle Accelerators, https://www.bnl.gov/mlaworkshop2022/.

[19] R. Roussel, A. Hanuka, and A. Edelen, "Multiobjective Bayesian optimization for online accelerator tuning", *Phys. Rev. Accel. Beams*, vol. 24, p. 062801, June 2021. doi:10.1103/PhysRevAccelBeams.24.062801

[20] Y. Gao, W. Lin, K. A. Brown, X. Gu, G. H. Hoffstaetter, J. Morris, and S. Seletskiy, "Bayesian optimization experiment for trajectory alignment at the low energy RHIC electron cooling system", *Phys. Rev. Accel. Beams*, vol. 25, p. 014601, Jan. 2022. doi:10.1103/PhysRevAccelBeams.25.014601

[21] S.A. Miskovich, et al., "Online Bayesian optimization for a recoil mass separator", *Phys. Rev. Accel. Beams*, vol. 25, p. 044601, Apr. 2022. doi:10.1103/PhysRevAccelBeams.25.044601

[22] J. Kirschner, M. Mutný, A. Krause, J. Coello de Portugal, N. Hiller, and J. Snuverink, "Tuning particle accelerators with safety constraints using Bayesian optimization", *Phys. Rev. Accel. Beams*, vol. 25, p. 062802, Jun. 2022. doi:10.1103/PhysRevAccelBeams.25.062802

[23] C. Xu, T. Boltz, A. Mochihashi, A. Santamaria Garcia, M. Schuh, and A. Müller, "Bayesian optimization of the beam injection process into a storage ring", *Phys. Rev. Accel. Beams*, vol. 26, p. 034601, Mar. 2023. doi:10.1103/PhysRevAccelBeams.25.062802

[24] W. Blokland, et al., "Uncertainty aware anomaly detection to predict errant beam pulses in the Oak Ridge Spallation Neutron Source accelerator", *Phys. Rev. Accel. Beams*, vol. 25, p. 122802, Dec. 2022. doi:10.1103/PhysRevAccelBeams.25.122802

[25] A. Eichler, J. Branlard, and J. H. K. Timm, "Anomaly detection at the European X-ray Free Electron Laser using a parity-space-based method", *Phys. Rev. Accel. Beams*, vol. 26, p. 012801, Jan. 2023. doi:10.1103/PhysRevAccelBeams.25.012801

[26] V. Kain, S. Hirlander, B. Goddard, F. Maria Velotti, G. Zevi Della Porta, N. Bruchon, and G. Valentino "Sample-efficient reinforcement learning for CERN accelerator control", *Phys. Rev. Accel. Beams*, vol. 23, p. 124801, Dec. 2020. doi:10.1103/PhysRevAccelBeams.23.124801

[27] J. St. John, et al., "Real-time artificial intelligence for accelerator control: A study at the Fermilab Booster", *Phys. Rev. Accel. Beams*, vol. 24, p. 104601, Oct. 2021. doi:10.1103/PhysRevAccelBeams.24.104601

[28] D. Meier, L. Vera Ramirez, J. Völker, J. Viefhaus, B. Sick, and G. Hartmann, "Optimizing a superconducting radiofrequency gun using deep reinforcement learning", *Phys. Rev. Accel. Beams*, vol. 25, p. 104604, Oct. 2022. doi:10.1103/PhysRevAccelBeams.25.104604

[29] X. Chen, Y. Jia, X. Qi, Z. Wang, and Y. He, "Orbit correction based on improved reinforcement learning algorithm", *Phys. Rev. Accel. Beams*, vol. 26, p. 044601, Apr. 2023. doi:10.1103/PhysRevAccelBeams.26.044601

[30] O. Convery, L. Smith, Y. Gal, and A. Hanuka, "Uncertainty quantification for virtual diagnostic of particle accelerators", *Phys. Rev. Accel. Beams*, vol. 24, p. 074602, Jul. 2021. doi:10.1103/PhysRevAccelBeams.26.044601

[31] A. Ananda Mishra, A. Edelen, A. Hanuka, and C. Mayes, "Uncertainty quantification for deep learning in particle accelerator applications", *Phys. Rev. Accel. Beams*, vol. 24, p. 114601, Nov. 2021. doi:10.1103/PhysRevAccelBeams.24.114601

[32] https://www.alcf.anl.gov/.