

Opportunities in Machine Learning for Particle Accelerators

Editors

A. Edelen and C. Mayes

SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

D. Bowring

Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

Contributors

D. Ratner

SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

A. Adelman, R. Ischebeck, and J. Snuverink

Paul Scherrer Institut, Villigen, Switzerland

I. Agapov and R. Kammering

Deutsches Elektronen-Synchrotron, Hamburg, Germany

J. Edelen

Radiasoft, LLC, Boulder, Colorado 80301, USA

I. Bazarov

Cornell University, Ithaca, NY 14853, USA

G. Valentino

University of Malta, Msida, Malta

J. Wenninger

CERN, Geneva, Switzerland

Abstract

“Machine learning” (ML) is a subfield of artificial intelligence. The term applies broadly to a collection of computational algorithms and techniques that train systems from raw data rather than *a priori* models. ML techniques are now technologically mature enough to be applied to particle accelerators, and we expect that ML will become an increasingly valuable tool to meet new demands for beam energy, brightness, and stability. The intent of this white paper is to provide a high-level introduction to problems in accelerator science and operation where incorporating ML-based approaches may provide significant benefit. We review ML techniques currently being investigated at particle accelerator facilities, and we place specific emphasis on active research efforts and promising exploratory results. We also identify new applications and discuss their feasibility, along with the required data and infrastructure strategies. We conclude with a set of guidelines and recommendations for laboratory managers and administrators, emphasizing the logistical and technological requirements for successfully adopting this technology. This white paper also serves as a summary of the discussion from a recent workshop held at SLAC on ML for particle accelerators [1].

I. INTRODUCTION

Machine learning (ML) is a subfield of artificial intelligence centered on algorithms that can learn to complete tasks using data. ML is currently experiencing a renaissance thanks to recent computational, theoretical, and practical advances in the field. Recently, various ML approaches have been successfully demonstrated in real-world tasks for computer vision [2], anomaly detection [3], and language translation [4]. Those wishing to use ML are able to access an increasingly wide variety of high performance computing resources. For example, computing clusters at universities and other research institutions are now widely accessible, along with a variety of web-based services. The greater availability of large data sets and the variety of new open source high-level software libraries has also helped to improve the reach of existing ML techniques and to facilitate rapid prototyping of new algorithms. These technical advances go hand-in-hand with new real-world applications that motivate ML research and funding in the public and private sectors, as well as provide feedback that informs further algorithmic development.

Particle accelerators in particular are some of the largest, most data-intensive, and most complex scientific systems in existence. The interrelations between machine subsystems are complicated and often nonlinear, the system dynamics involve large parameter spaces that evolve over multiple relevant time scales, and accelerator systems can be difficult to model *a priori*. Relevant problems for accelerators include, for example, analysis of large quantities of archived data, accurate and fast modeling of accelerator systems, detection of aberrant machine behavior, optimization of accelerator design, and active tuning and control. At present, ML-based approaches are technologically mature enough to be brought to bear on a wide variety of problems within these domains.

We expect that ML will become an increasingly valuable tool to meet new demands for beam energy, brightness, reliability, and stability. In this white paper we review ML techniques currently being investigated at particle accelerator facilities, including several promising experimental demonstrations. We also identify new applications and discuss their feasibility, along with the required data and infrastructure strategies.

Much of the content here has been guided by presentations and discussion arising from the inaugural International Committee for Future Accelerators (ICFA) Workshop on Machine Learning for Particle Accelerators, held from February 27 through March 2, 2018 at SLAC

National Accelerator Laboratory [1]. A similar workshop was held at Daresbury Lab on January 30th and 31st, 2018 [5]. For some additional introductory material on the topic of ML for particle accelerators, see [6].

II. GENERAL APPLICATIONS OF MACHINE LEARNING

ML has been successfully applied to a variety of other scientific/engineering problems. Here we review some applications beyond accelerator science for historical context.

The modern wave of ML algorithms (e.g. deep neural networks), first began to see success in computer vision applications. For a high-level overview of some recent successes in classic computer vision domains, e.g. object recognition in natural scenes, see Refs. [2, 7]. In astronomy, neural networks have been used extensively for automated sorting and analysis of sky survey images and rejection of artifacts [8–12]. In particle detectors, ML methods are used to identify and classify particle tracks [13]. More recently, this has included deep neural networks [14–17], and a review of these developments can be found in [18]. Notably, a 2014 “Kaggle” code competition also drew numerous solutions from the public to identify ATLAS detector signals related to the Higgs boson. Contestants used neural networks, decision trees, and ensemble methods [19].

In the medical community, ML has been used for automated diagnosis of medical conditions. ML algorithms have recently shown improvements over traditional screening techniques in identifying/classifying skin and cervical cancers[20, 21]. Other efforts showed that a ML algorithm was able to beat cardiologists at detecting heart arrhythmia viewed on an electrocardiogram [22]

As an industrial example, robotic arms have been controlled using ML algorithms to perform sorting and packaging operations. These algorithms have been used reliably in operation for years [23]. Additionally, ML has been in use by the U.S. post office for the interpretation of hand-written addresses since the 1990s with an accuracy of above 98%, saving 100 million dollars during the first year of deployment [24]

With regard to anomaly detection and machine protection, the Joint European Torus has used ML for instability detection [25, 26] and fault prediction [26, 27] in its fusion research facility. ML was also recently used to optimize tokamak settings for plasma confinement, resulting in over 50 percent improvement in key performance metrics such as energy loss rate and total plasma energy [28].

III. EARLY HISTORY OF USAGE FOR PARTICLE ACCELERATORS

ML techniques have been applied to particle accelerators since the late 1980s. Much early discussion during the late 1980s and early 1990s focused on applying rule-based systems to accelerator control and tuning [29–33]. In the early 1990s, scientists at Los Alamos National Lab had some experimental success with neural-network-based ion source control [34–36]. Other early studies at the University of New Mexico focused on orbit/trajectory control [37–41], fault detection and management [42, 43], and root-cause analysis of errors (e.g. distinguishing between bending magnet misalignment and field defects [44]). General AI/ML platform for beamline tuning were also planned [45–48]. None of these systems were eventually used routinely as part of an accelerator’s main control system.

The lack of clear success in bringing ML to regular use in accelerator systems was partly due to limitations in the then-available hardware, algorithms, and software packages, as well as the limited accessibility of good data sets and simulation tools. Similar situations were encountered in other scientific fields where ML approaches were tried before they had reached sufficient technological maturity relative to the challenges of the particular application.

IV. AREAS OF APPLICABILITY FOR PARTICLE ACCELERATORS

ML-based solutions to challenges encountered in particle accelerators are now under development and yielding promising results. As ML becomes a routine tool in this field, other applications are likely to emerge. Here we review several use cases.

A. Anomaly detection and machine protection

The operation of today's particle accelerators is increasingly complex and data-intensive, and it typically involves a large number of process variables, nonlinear behavior, and many interacting subsystems. The identification of conditions that could negatively impact machine performance requires monitoring of many interacting subsystems, and conventional analysis methods are often insufficient. Using machine data to identify poor conditions and the root cause of errors encountered during operation is therefore a growing challenge. In addition, improving the reliability of machine protection systems can help pave the way to producing beams of higher intensity/stored energy.

Anomaly and breakout detection algorithms may provide added capabilities beyond those of existing machine protection systems by detecting subtle behaviors of key variables prior to negative events. This approach has been used for detection of quench precursors in both superconducting magnets [49] and superconducting RF cavities [50].

In the context of system control, anomaly detection algorithms can also be used to identify and throw away bad signals. The LHC has used such algorithms to identify bad readings from beam position monitors and remove them prior to a standard correction procedure [51], as well as to assist in automated collimator alignment [52].

B. System Modeling

Another challenge in accelerator science is the need to obtain machine models that (a) are sufficiently accurate and (b) can execute quickly enough to be useful during operations, possibly with input from machine instrumentation (i.e. **online modeling**), and in large-scale offline optimization.

Many simulation tools are too slow in practice to be used directly in control systems or to provide guidance to operators during machine operation. In recent years, some hardware-

accelerated modeling codes have been developed that only require modest computing resources, for example, GPU-accelerated `elegant` [53] and `HPSim`, a GPU-accelerated code for online modeling of ion linacs [54, 55]. However, most online models at present rely on simplified representations of the relevant physics: accuracy is traded for speed. Furthermore, ensuring these models accurately represent the behavior of a particular machine still requires substantial work in most cases and is not always possible due to practical constraints (e.g. lack of diagnostics, lack of personnel, lack of time for systematic machine studies).

System modeling with ML enables one to learn representations that combine information from physics-based simulations with measured data, which can be useful when measured data only covers a limited portion of the system behavior that needs to be included for prediction or control. This has been demonstrated for the injector and low-energy beamline at Fermilab’s FAST facility [56]. Furthermore, the execution speed of such ML models can be faster than conventional accelerator simulations, enabling fast-executing representations of slowly-executing, high-fidelity simulations to be created. This technique is known as **surrogate modeling**. For example, in one early study a neural network surrogate model of the RF gun and low-energy beamline at FAST sped up the simulation time from around 20 minutes to under a millisecond [57]. The resulting simulation tools can be useful for providing real-time results in a control room environment or for planning experiments and doing offline controller design. For some examples, see [57–59].

Neural networks in particular are also well-suited to tasks which deal with image data, and this enables image-based diagnostics to be used directly in accelerator models, both as outputs and inputs. For example, to take into account variability in a drive laser for a photocathode gun, a virtual cathode image of the laser spot has been used as one input to a model to predict downstream beam parameters at FAST [57].

In terms of offline optimization and design, current approaches take advantage of heavy parallelization when a high-fidelity physics model is required (for example see [60, 61]). However, gaining access to the necessary computing resources for this is not always possible, and in practice simulation speed often limits the extent of parameter space exploration during the optimization and design process. In this case, ML models that are sufficiently well-validated could be used instead to interpolate across optimum fronts of accelerator performance from a set of sparsely-sampled training examples. This ultimately could lead to an increase in the throughput of optimization for challenging accelerator design problems

(such as dynamic aperture maximization for diffraction-limited storage ring light source [62]) or allow faster exploration of various trade-offs between competing parameters.

C. Virtual Instrumentation / Virtual Diagnostics

Accelerators are outfitted with a variety of instruments to provide measurements of both the beam itself and other key variables that impact the beam (e.g. RF signals). These instruments are often referred to as "diagnostics," in the sense that they provide some indication of the performance and operating state of the accelerator. A **virtual instrument (also referred to as a "virtual diagnostic")** provides an estimate of what an instrument would read when such a reading is unavailable, thus allowing non-invasive "measurements" in locations where it may be impossible to put an instrument, where only destructive instruments are available and thus cannot be used during normal operations, where the update rate of an existing instrument is lower than needed for experimental analysis or for control, or where the instrument is not sensitive across the entire operating range.

For example, a virtual instrument for the beam phase space based on high-fidelity, GPU-accelerated physics simulations has been demonstrated at LANSCE. The simulation reads inputs from the machine and within seconds provides predictions about the resulting beam, including live plots of the phase space [55]. In a slightly different approach demonstrated at FACET, free parameters in an online model with only simple physics included were automatically tuned to match the output of one instrument reading (in this case the output of a transverse deflecting cavity) by using the prediction accuracy on a separate instrument as a guide (in this case a spectrometer) [63]. This approach can provide estimates of what the additional instrument would show non-invasively when only the more readily available diagnostics can be used, and it also only requires a simple physics model.

New kinds of virtual instrumentation can also be created using ML. For example, artificial neural networks and support vector regression techniques have been used at LCLS to identify correlations between information obtained from slowly-updating instruments (e.g. photon energy, spectral shape of X-ray pulses) and more abundant, quickly-updating outputs (e.g. readings from beam position monitors). These correlations can be exploited to extract information that is of interest to the users about each X-ray pulse quickly and with high fidelity [64].

Taking this concept a step further, the output of image-based diagnostics can also be predicted directly (thus creating a virtual version of the instrument). This has been used, for example, to predict the images produced by a destructive multi-slit transverse phase space measurement [1, 56, 65] and to predict the longitudinal phase space as displayed by a transverse deflecting cavity [66, 67].

D. Tuning, Control, and Rapid Switching Between Operating Conditions

Particle accelerator facilities have a wide range of operational needs when it comes to tuning, optimization, and control. For example, free electron laser (FEL) user facilities have a strong incentive to reduce time spent switching between user-requested operating conditions. For other machines, users may require long-term machine stability or long-term optimization of specific beam parameters. Facilities focused on neutrino science, for example, may seek to maximize the number of protons per hour on a target. In order to meet these demands, particle accelerators rely on databases of previous settings, on fine-tuning of machine settings by operators or online optimization routines, and dedicated feedback and control systems to maintain stable operation. In each of these cases, ML can supplement existing techniques and procedures.

ML-based methods can reduce the need to restore old settings directly from a database. For example, ML models mapping a desired system output (such as particular beam parameters) to the required settings could cut down on the time it takes to switch between different operating conditions. The model could also be re-trained over time, as new data accumulates and as operations change. This would shorten the time required to determine settings for some previously unseen operating state. For an exploratory example in simulation, see [59], in which a neural network was trained to switch between settings needed in the injector and beamline of a compact FEL for different electron beam energies while maintaining a good match into the FEL's undulator.

ML can also be used to learn specific update rules for moving closer to an optimum. By building up representations of machine behavior, these algorithms implicitly use examples of past machine behavior to help guide the search toward optimal settings for a given operating condition. One approach, reinforcement learning using neural networks, has been used at LCLS for tuning the undulator taper, resulting in a doubling of the power in self-seeding

mode [68].

Another approach, Bayesian optimization using Gaussian process models, has been used to tune quadrupole settings in the LCLS beamline [69]. Bayesian optimization can also explicitly learn to avoid disruptive combinations of settings during tuning [70, 71].

ML models and control policies can also be useful for systems that need to be controlled at a rate that is faster than the system settling time. For example, ML-based models have been trained on measured data to predict the temperature response of a photoinjector [6, 72] and the resonant frequency response of a radio frequency quadrupole at Fermilab [73], both of which exhibit long-period thermal behavior but benefit from rapid control adjustments. Predictive control over the photoinjector using the model resulted in a substantially faster settling time, making it faster to switch between RF power settings while staying on-resonance. Similarly, researchers at the Korea Atomic Energy Research Institute have created a steady-state ion source model from measured data (mapping input parameters to their final steady state output values), with the eventual aim of using it in model predictive control [74]. Ion sources are particularly difficult to model from physics principles alone and typically exhibit highly nonlinear behavior and long response times, making them appealing targets for ML-based modeling and predictive control.

E. Advanced Data Analysis

ML methods can aid data analysis and discovery of new physics, as already seen in other fields. Section II describes computer vision algorithms that are used for high-throughput galaxy classification in astronomical surveys, or particle track identification in high energy physics detectors. At accelerator laboratories, computer vision can also be useful for processing data from image-based beam diagnostics. For example, one could automatically incorporate image data into machine models and controllers, use image data to automatically flag aberrant behavior, or automatically extract useful information for users and accelerator physicists.

Modern computational methods and ML can also be exploited to handle reconstruction tasks, such as recovering additional information about the beam from limited measurements. For example, kernel density estimation (KDE), is an unsupervised method to describe a density distribution. KDEs were used at the Muon Ionization Cooling Experiment (MICE) to

measure the phase space density of muon beams. This method enhances the precision of beam emittance measurements [75, 76]. As another example, the field of compressed sensing incorporates prior knowledge from physical models, data, or simulations to reconstruct under-sampled targets. Researchers at UCLA and SLAC have used compressed sensing and priors trained on simulations to speed-up convergence of electron ghost imaging at the Pegasus beam line [77].

Clustering and correlation analysis can also aid exploration of complicated data sets found in accelerator simulations and measurements. Researchers at Karlsruhe Institute of Technology used K-means clustering to identify dominant modes of the microbunching instability at the Karlsruhe Research Accelerator (KARA) ring. The results led to identification of new features in the longitudinal phase space density [78]. Similarly, clustering has been used to identify different types of quenches in superconducting RF cavities at DESY [79]. Such analysis can also provide guidance on promising avenues of research during early stages of larger projects. Finally, at the Paul Scherrer Institut's High Intensity Proton Accelerator (HIPA), researchers used correlation analysis to investigate which machine parameters contributed most to beam losses at different points along the accelerator [1].

V. TECHNICAL CONSIDERATIONS

Researchers should be aware of the practical requirements that must be met in order to apply ML to accelerator systems successfully. This includes, for example, the proper retention and formatting of machine data and improving the availability and scale of computing resources (both for training and deployment).

A. Data Requirements and Availability

The success of ML algorithms depends on the amount of available training data; specifically, the quality of a given ML solution can improve substantially with a larger numbers of training examples. Fortunately, particle accelerators are typically instrumented to accommodate this requirement. Control at many modern facilities is driven by either digital signal processors (DSPs) or field programmable gate arrays (FPGAs), and signal monitoring can occur at MHz rates or faster. Some modern facilities can produce terabytes of data per day, and this typically covers only the most relevant diagnostics information.

At most facilities, data is archived but typically only used for post-mortem analysis after failures, or for transient studies on a specific topic of interest. In many cases there is already a sufficient amount of data in the archive to be put to use more routinely for ML applications. Much effort, however, will need to be spent in data preprocessing in order to make use of this data. Some ML techniques could also enable researchers to analyze a greater proportion of this data (for example, clustering), thus enabling it to be put to better use for improving operations and for understanding machine behavior.

B. Availability and scale of computing resources

ML-enabled data analysis is computationally intensive and often requires high-performance computing (HPC) platforms. Fortunately, these resources are becoming more accessible and can be made available for accelerator data processing. In addition, newer computing architectures, such as GPUs designed specifically for ML [80], have enabled faster training of ML algorithms. GPUs can currently be purchased for just a few hundred dollars (e.g. NVIDIA GeForce GTX 1080) to a few thousand dollars (e.g. NVIDIA Tesla P100) and can easily support a small group of users.

Furthermore, there are now many options for doing ML online through cloud-based services. For instance, Google Colaboratory [81] allows one to use a virtual GPU through a web browser at no charge. Amazon Web Services [82] and Google Cloud [83] both offer dedicated services for ML. These resources provide a convenient way of scaling up from small initial efforts without needing to purchase hardware or pay for its ongoing maintenance.

C. Data formatting considerations

Common ML scenarios in industry assume either structured data (hashed data relating, e.g., customer name and visits to a website) or homogeneous data sets such as text entered into search engines. The relevant data for accelerator operations is mostly unstructured and heterogeneous: floating point data and images. Furthermore, in some cases the data may not all be logged in a uniform way. Different instruments may log data to different storage systems, or use different timestamps.

We can divide our data into two major domains: accelerator data and experiment-related data (e.g. data generated in collider detectors or light sources). The needs of these domains, as well as the corresponding data flow and infrastructure, can greatly differ. For many years, the HEP and light source user community has been using carefully designed data infrastructures to allow processing and analysis of data sets. By contrast, the accelerator domain mostly stores and handles data in a very heterogeneous and ad hoc manner. These subsystems may store data on local disks, store data at a low rate, or not store data at all.

In order to extract unknown correlations or error signatures using ML, more comprehensive data processing, faster data transport, and increased data storage capacity will be needed. To make accelerator data available for ML algorithms, laboratories must standardize and classify this data as early as possible in the acquisition chain. To facilitate the use of ML frameworks, the accelerator community needs to consider classification, labeling, and storage requirements in the control system architecture, ensuring synchronicity while limiting the data rates to appropriate levels.

Much of the software requirements for making use of this data are already available and not accelerator-specific. Free platforms and software libraries already exist for manipulating large datasets. Furthermore, data analysis, ML, numerical, and visualization libraries in modern programming languages (e.g. Python and its associated scientific libraries, such

as scikit-learn) can help to facilitate rapid development and deployment of ML algorithms. Person-hours at accelerator laboratories can be focused on the application of these tools to specific problems. In addition, examples of code using ML for different use cases is widely available.

VI. RECOMMENDATIONS

A. Institutional support

The successful completion of ML projects in accelerator facilities depends on laboratory support beyond individual researchers. This includes support for instrumentation, control, operations, and computing staff, and their coordination. If possible, we recommend that accelerator laboratories form a dedicated ML group.

Facilities may also need to provide additional funding and personnel for infrastructure improvements to make full use of ML. This might include actions such as investing in or facilitating access to additional computing resources, or making improvements to data archiving systems.

Managers should not underestimate the effort needed to curate data to be used for ML. In addition to making use of archived data, development and testing of some algorithms (e.g. for control) will require dedicated machine time and studies.

B. Education and Intersection with Other Fields

Mutually-beneficial collaboration between laboratories, universities, and industrial partners is essential to bring the latest developments in computer science to particle accelerators. Laboratories should actively seek these collaborations.

Physicists tend to already have a good grounding in probability, statistics, and data analysis. Therefore, they already have the fundamentals in place to appreciate how ML algorithms function. At the university level, physics students could be offered optional courses in ML, which build upon what they would have covered in probability and statistics courses. Lectures in ML could be introduced in the various particle accelerator schools, such as the US Particle Accelerator School or the CERN Accelerator School. Existing education/research networks, such as the Center for Bright Beams in the US, could be leveraged to recruit graduate students to work on a common problem. In addition, accelerator facilities could host a local seminar series aimed at exposing accelerator physicists to current applied ML research.

C. Outreach and Citizen Science

Encouraging non-physicists to collaborate on ML problems could allow laboratories to outsource some of this work. Examples of fora in which this could take place are online competitions such as Kaggle, in which participants are asked to provide a solution to a specific problem, with the winner being awarded a prize. As an example, a successful competition was organized in 2014 by the high energy physics community [19], and it attracted over 1700 participants. This can also serve as a public relations and outreach exercise, as well as a way to attract students and outside researchers to work with accelerator facilities. Hackathons are an alternative way of getting accelerator physicists and ML experts together in the same room for a couple of days to work intensely on domain specific ML related problems.

-
- [1] SLAC National Accelerator Laboratory, “ICFA beam dynamics mini-workshop: Machine learning applications for particle accelerators,” (2018).
 - [2] Y. LeCun, Y. Bengio, and G.E. Hinton, “Deep learning,” *Nature* **521**, 436–444 (2015).
 - [3] J. O. Awoyemi, A. O. Adetunmbi, and S. A. Oluwadare, “Credit card fraud detection using machine learning techniques: A comparative analysis,” in *2017 International Conference on Computing Networking and Informatics (ICCNi)* (2017) pp. 1–9.
 - [4] J. Hirschberg and C.D. Manning, “Advances in natural language processing,” *Science* **349**, 261–266 (2015).
 - [5] Daresbury Laboratory, “Intelligent controls for particle accelerators workshop,” (2018).
 - [6] A. L. Edelen, S. G. Biedron, B. E. Chase, D. Edstrom, S. V. Milton, and P. Stabile, “Neural Networks for Modeling and Control of Particle Accelerators,” *IEEE Trans. Nucl. Sci.* **63**, 878–897 (2016), arXiv:1610.06151 [physics.acc-ph].
 - [7] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural Networks* **61**, 85 – 117 (2015).

- [8] N.M. Ball and R.J. Brunner, “Data mining and machine learning in astronomy,” *International Journal of Modern Physics D* **19**, 1049–1106 (2010), <https://www.worldscientific.com/doi/pdf/10.1142/S0218271810017160>.
- [9] M. Way, J. Scargle, K. Ali, and A. Srivastava, eds., *Advances in Machine Learning and Data Mining for Astronomy* (New York: Chapman and Hall/CRC, 2012).
- [10] P. Graff, F. Feroz, M. P. Hobson, and A. Lasenby, “Neural networks for astronomical data analysis and bayesian inference,” in *2013 IEEE 13th International Conference on Data Mining Workshops* (2013) pp. 16–23.
- [11] P. Graff, F. Feroz, M.P. Hobson, and A. Lasenby, “Skynet: an efficient and robust neural network training tool for machine learning in astronomy,” *Monthly Notices of the Royal Astronomical Society* **441**, 1741–1759 (2014).
- [12] R. Biswas, L. Blackburn, J. Cao, R. Essick, K.A. Hodge, E. Katsavounidis, K. Kim, Y.-M. Kim, E.-O. Le Bigot, C.-H. Lee, J.J. Oh, *et al.*, “Application of machine learning algorithms to the study of noise artifacts in gravitational-wave data,” *Phys. Rev. D* **88**, 062003 (2013).
- [13] P. Bhat, “Multivariate analysis methods in particle physics,” *Annual Review of Nuclear and Particle Science* **121**, 281–309 (2011).
- [14] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M.D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, “A convolutional neural network neutrino event classifier,” *Journal of Instrumentation* **11** (2016), <https://doi.org/10.1088/1748-0221/11/09/P09001>.
- [15] P. Baldi, Paul D. Sadowski, and Daniel Whiteson, “Searching for exotic particles in high-energy physics with deep learning.” *Nature communications* **5**, 4308 (2014).
- [16] P.C. Bhat, “Multivariate analysis methods in particle physics,” *Annual Review of Nuclear and Particle Science* **61**, 281–309 (2011), <https://doi.org/10.1146/annurev.nucl.012809.104427>.
- [17] S. Whiteson and D. Whiteson, “Machine learning for event selection in high energy physics,” *Engineering Applications of Artificial Intelligence* **22**, 1203 – 1217 (2009).
- [18] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel, A. Aurisano, K. Terao, and T. Wongjirad, “Machine learning at the energy and intensity frontiers of particle physics,” *Nature* **560**, 41–48 (2018).
- [19] C. Adam-Bourdarios, G. Cowan, C. Germain-Renaud, I. Guyon, B. Kégl, and D. Rousseau, “The Higgs Machine Learning Challenge,” *Proceedings, 21st International Conference on*

Computing in High Energy and Nuclear Physics (CHEP 2015): Okinawa, Japan, April 13-17, 2015, J. Phys. Conf. Ser. **664**, 072015 (2015).

- [20] A. Esteva, B. Kuprel, R.A. Novoa, J. Ko, S.M. Swetter, H.M. Blau, and S. Thrun, “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature* **542**, 115–118 (2017).
- [21] J. Zhang and Y. Liu, “Cervical cancer detection using svm based feature screening,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention* (Springer, 2004) pp. 873–880.
- [22] P. Rajpurkar, A. Y. Hannun, M. Haghpahani, C. Bourn, and A. Y. Ng, “Cardiologist-Level Arrhythmia Detection with Convolutional Neural Networks,” *ArXiv e-prints* (2017), arXiv:1707.01836 [cs.CV].
- [23] M.P. Deisenroth, G. Neumann, and J. Peters, “A survey on policy search for robotics,” *Foundations and Trends in Robotics* **2**, 1–142 (2013).
- [24] S. N. Srihari and E. J. Kuebert, “Integration of hand-written address interpretation technology into the united states postal service remote computer reader system,” in *Proceedings of the Fourth International Conference on Document Analysis and Recognition*, Vol. 2 (1997) pp. 892–896 vol.2.
- [25] A. Murari, D. Mazon, N. Martin, G. Vagliasindi, and M. Gelfusa, “Exploratory data analysis techniques to determine the dimensionality of complex nonlinear phenomena: The l-to-h transition at jet as a case study,” *IEEE Transactions on Plasma Science* **40**, 1386–1394 (2012).
- [26] A. Murari, A. Buscarino, L. Fortuna, M. Frasca, M. Iachello, and G. Mazzitelli, “Identifying jet instabilities with neural networks,” 2012 16th IEEE Mediterranean Electrotechnical Conference , 932–935 (2012).
- [27] A. Murari, P. Arena, A. Buscarino, L. Fortuna, and M. Iachello, “On the identification of instabilities with neural networks on jet,” *Nucl. Instrum. Methods Phys. Res. A* **720**, 2 – 6 (2013), selected papers from the 2nd International Conference Frontiers in Diagnostic Technologies (ICFDT2).
- [28] E. A. Baltz, E. Trask, M. Binderbauer, M. Dikovskiy, H. Gota, R. Mendoza, J. C. Platt, and P. F. Riley, “Achievement of sustained net plasma heating in a fusion experiment with the optometrist algorithm,” *Scientific Reports* **7**, 6425 (2017).

- [29] T. Higo, H. Shoaee, and J Spencer, “Some applications of ai to the problems of accelerator physics,” in *Conf. Proc.*, Vol. 870316 (1986) p. 701.
- [30] D.P. Weygand, *Artificial intelligence and accelerator control*, Tech. Rep. (Brookhaven National Lab., Upton, NY (USA), 1987).
- [31] P. Skarek and L. Varga, “Multi-agent cooperation for particle accelerator control,” *Expert Systems with applications* **11**, 481–487 (1996).
- [32] D.E. Schultz and P.A. Brown, “The development of an expert system to tune a beam line,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **293**, 486–490 (1990).
- [33] E. Fiesler and S.R. Campbell, “Hybrid neural networks and their application to particle accelerator control,” in *Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation II*, Vol. 3812 (International Society for Optics and Photonics, 1999) pp. 132–143.
- [34] J.A. Howell, C.W. Barnes, S.K. Brown, G.W. Flake, R.D. Jones, Y.C. Lee, S. Qian, and R.M. Wright, “Control of a negative-ion accelerator source using neural networks,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **293**, 517–522 (1990).
- [35] W.C. Mead, P.S. Bowling, S.K. Brown, R.D. Jones, C.W. Barnes, H.E. Gibson, J.R. Goulding, and Y.C. Lee, “Optimization and control of a small-angle negative ion source using an on-line adaptive controller based on the connectionist normalized local spline neural network,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **72**, 271–289 (1992).
- [36] W.C. Mead, S.K. Brown, R.D. Jones, P.S. Bowling, and C.W. Barnes, “Adaptive optimization and control using neural networks,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **352**, 309–315 (1994).
- [37] E. Bozoki and A. Friedman, “Neural network technique for orbit correction in accelerators/storage rings,” in *AIP Conference Proceedings*, Vol. 315 (AIP, 1994) pp. 103–110.
- [38] D. Nguyen, M. Lee, R. Sass, and H. Shoaee, *Accelerator and feedback control simulation using neural networks*, Tech. Rep. (Stanford Linear Accelerator Center, Menlo Park, CA (USA), 1991).

- [39] Y. Hitaka, M. Shirakata, H. Sato, M. Yokomichi, and M. Kono, “Numerical methods for the orbit control at the kek 12 gev-ps,” in *Proceedings of the European Particle Accelerator Conference* (2004) pp. 2664–2666.
- [40] K.H. Kim, K. Shim, J. Choi, M.H. Cho, W. Namkung, and I.S. Ko, “Simulation of the global orbit feedback system for pohang light source,” in *Proceedings of the European Particle Accelerator Conference* (2000) pp. 1906–1908.
- [41] D. Schirmer, P. Hartmann, T. Büning, and D. Müller, “Electron transport line optimization using neural networks and genetic algorithms,” in *Proceedings of EPAC* (2006).
- [42] N.R. Jennings, E.H. Mamdani, J.M. Corera, I. Laresgoiti, F. Perriolat, P. Skarek, and L. Varga, “Using archon to develop real-world dai applications. 1,” *IEEE expert* **11**, 64–70 (1996).
- [43] F. Perriolat, P. Skarek, L.Z. Varga, and N.R. Jennings, “Using archon-3. particle acceleration control,” *IEEE Expert* **11**, 80–86 (1996).
- [44] Y. Kijima, K. Yoshida, M. Mizota, and K. Suzuki, “A beam diagnostic system for accelerators using neural networks,” in *Proceedings of the European Particle Accelerator Conference* (1992) pp. 1155–1157.
- [45] C.R. Stern, W.B. Klein, G.F. Luger, M. Kroupa, and R.T. Westervelt, “Tuning and optimization at brookhaven and argonne: Results of recent experiments using a portable intelligent control system,” in *APS Meeting Abstracts* (1997).
- [46] W.B. Klein, C.R. Stern, G.F. Luger, and E.T. Olsson, “An intelligent control architecture for accelerator beamline tuning,” in *Proceedings of the fourteenth national conference on artificial intelligence and ninth conference on Innovative applications of artificial intelligence* (AAAI Press, 1997) pp. 1019–1025.
- [47] W. Klein, C. Stern, G. Luger, and E. Olsson, “Designing a portable architecture for intelligent particle accelerator control,” in *Particle Accelerator Conference, 1997. Proceedings of the 1997*, Vol. 2 (IEEE, 1997) pp. 2422–2424.
- [48] W. Klein, R. Westervelt, and G. Luger, “Developing a general purpose intelligent control system for particle accelerators,” *Journal of Intelligent & Fuzzy Systems* **7**, 1–12 (1999).
- [49] M. Wielgosz, A. Skoczeń, and M. Mertik, “Using LSTM recurrent neural networks for monitoring the LHC superconducting magnets,” *Nucl. Instrum. Meth.* **A867**, 40–50 (2017), arXiv:1611.06241 [physics.ins-det].

- [50] A. S. Nawaz, S. Pfeiffer, G. Lichtenberg, and H. Schlarb, “Self-organized critical control for the european xfel using black box parameter identification for the quench detection system,” in *2016 3rd Conference on Control and Fault-Tolerant Systems (SysTol)* (2016) pp. 196–201.
- [51] E. Fol, R. Tomas Garcia, and P. Henning, “Evaluation of Machine Learning Methods for LHC Optics Measurements and Corrections Software,” (2017), presented 23 Oct 2017.
- [52] G. Valentino, R. Bruce, S. Redaelli, R. Rossi, P. Theodoropoulos, and S. Jaster-Merz, “Anomaly Detection for Beam Loss Maps in the Large Hadron Collider,” *Proceedings, 8th International Particle Accelerator Conference (IPAC 2017): Copenhagen, Denmark, May 14-19, 2017*, J. Phys. Conf. Ser. **874**, 012002 (2017), [MOPAB010(2017)].
- [53] J. R. King, I. V. Pogorelov, K. M. Amyx, M. Borland, and R. Soliday, “GPU acceleration and performance of the particle-beam-dynamics code Elegant,” (2017), arXiv:1710.07350 [physics.comp-ph].
- [54] X. Pang and L. Rybarcyk, “GPU accelerated online multi-particle beam dynamics simulator for ion linear particle accelerators,” *Comput. Phys. Commun.* **185**, 744–753 (2014).
- [55] X. Pang, “Advances in Proton Linac Online Modeling,” in *Proceedings, 6th International Particle Accelerator Conference (IPAC 2015)* (2015) p. WEXC2.
- [56] A. L. Edelen, J.P. Edelen, D. Bowring, B.E. Chase, D.R. Edstrom, J. Steimel, P.J.M. van der Slot, and S.G. Biedron, “Recent Applications of Neural Network-Based Approaches to the Modeling and Control of Particle Accelerators,” in *9th International Particle Accelerator Conference (IPAC 2018): Vancouver, BC, Canada, April 29-May 4, 2018* (2018) p. THYGBE2.
- [57] Auralee Edelen, Sandra Biedron, Jonathan Edelen, and Stephen Milton, “First Steps Toward Incorporating Image Based Diagnostics into Particle Accelerator Control Systems Using Convolutional Neural Networks,” in *Proceedings, 2nd North American Particle Accelerator Conference (NAPAC2016)* (2017) p. TUPOA51, arXiv:1612.05662 [physics.acc-ph].
- [58] A. Adelman, “On Uncertainty Quantification in Particle Accelerators Modelling,” (2015), arXiv:1509.08130 [physics.acc-ph].
- [59] A.L. Edelen, S.G. Biedron, J.P. Edelen, S.V. Milton, and P.J.M. van der Slot, “Using a neural network control policy for rapid switching between beam parameters in an fel,” in *Proceedings*

- of the 38th International Free Electron Laser Conference* (2017) pp. 406–409.
- [60] Y. Wang, M. Borland, H. Shang, R. Soliday, A. Xiao, “Recent Progress on Parallel ELEGANT,” in *Proceeding ICAP 2009: San Francisco, CA (2009)*.
- [61] Y. Ineichen, A. Adelman, C. Bekas, A. Curioni, P. Arbenz, “A fast and scalable low dimensional solver for charged particle dynamics in large particle accelerators,” *Comput. Sci. Res. Dev.*, **28**, 185–192 (2013).
- [62] X. Huang and J. Safranek, “Online optimization of storage ring nonlinear beam dynamics,” *Phys. Rev. ST Accel. Beams* **18**, 084001 (2015), arXiv:1502.07799 [physics.acc-ph].
- [63] A. Scheinker and S. Gessner, “Adaptive method for electron bunch profile prediction,” *Phys. Rev. ST Accel. Beams* **18**, 102801 (2015).
- [64] A. Sanchez-Gonzalez, P. Micaelli, C. Olivier, T. R. Barillot, M. Ilchen, A. A. Lutman, A. Marinelli, T. Maxwell, A. Achner, M. Agåker, *et al.*, “Accurate prediction of x-ray pulse properties from a free-electron laser using machine learning,” *Nature Communications* **8**, 15461 EP – (2017).
- [65] A. L. Edelen, J.P. Edelen, D.R. Edstrom, A. Halavanau, P. Piot, and S.G. Biedron, “Neural Network Virtual Diagnostic for the FAST Low Energy Beamline,” in *9th International Particle Accelerator Conference (IPAC 2018): Vancouver, BC, Canada, April 29-May 4, 2018* (2018) p. WEPAF040.
- [66] C. Emma, A.L. Edelen, G. White, and M. Hogan, “Virtual diagnostic for phase space prediction at FACET-II,” in *Advanced Accelerator Concepts Workshop: Breckenridge, CO, USA, August 12 - 17, 2018* (2018).
- [67] C. Emma, A. L. Edelen, G. White, M. Hogan, B. O’Shea, V. Yakimenko, “Machine Learning-based Longitudinal Phase Space Prediction of Particle Accelerators,” accepted in *Physical Review Accelerators and Beams*, (Oct. 2018).
- [68] J. Wu, X. Huang, T. Raubenheimer, and A. Scheinker, “Recent On-Line Taper Optimization on LCLS,” in *Proceedings, 38th International Free Electron Laser Conference, FEL2017* (2018) p. TUB04.
- [69] M. McIntire, T. Cope, S. Ermon, and D. Ratner, “Bayesian Optimization of FEL Performance at LCLS,” in *Proceedings, 7th International Particle Accelerator Conference (IPAC 2016)* (2016) p. WEPOW055.

- [70] F. Berkenkamp, A. P. Schoellig, and A. Krause, “Safe Controller Optimization for Quadrotors with Gaussian Processes,” Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA) (2016), <https://doi.org/10.1088/1748-0221/11/09/P09001>.
- [71] F. Berkenkamp, A. Krause, and A. P. Schoellig, “Bayesian Optimization with Safety Constraints: Safe and Automatic Parameter Tuning in Robotics,” arXiv (2016), arXiv:1602.04450.
- [72] A. L. Edelen *et al.*, “Initial experimental results of a machine learning-based temperature control system for an RF gun,” in *Proceedings, 6th International Particle Accelerator Conference (IPAC 2015)* (2015) p. MOPWI028, arXiv:1511.01883 [physics.acc-ph].
- [73] A. Edelen, S. Biedron, D. Bowring, B. Chase, J. Edelen, S. Milton, and J. Steimel, “Neural Network Model Of The PXIE RFQ Cooling System and Resonant Frequency Response,” in *Proceedings, 7th International Particle Accelerator Conference (IPAC 2016)* (2016) p. TH-POY020, arXiv:1612.07237 [physics.acc-ph].
- [74] Y.B. Kong, M.G. Hur, E.J. Lee, J.H. Park, Y.D. Park, and S.D. Yang, “Predictive ion source control using artificial neural network for rft-30 cyclotron,” Nucl. Instrum. Methods Phys. Res. A **806**, 55 – 60 (2016).
- [75] T.A. Mohayai, P. Snopok, and D. Neuffer (MICE), “A Non-Parametric Density Estimation Approach to Measuring Beam Cooling in MICE,” in *9th International Particle Accelerator Conference (IPAC 2018) Vancouver, BC Canada, April 29-May 4, 2018* (2018) p. TUPML063, arXiv:1806.01834 [physics.acc-ph].
- [76] T.A. Mohayai (MICE), “First Demonstration of Ionization Cooling in MICE,” in *9th International Particle Accelerator Conference (IPAC 2018) Vancouver, BC Canada, April 29-May 4, 2018* (2018) p. FRXGBE3, arXiv:1806.01807 [physics.acc-ph].
- [77] S. Li, F. Cropp, K. Kabra, T. J. Lane, G. Wetzstein, P. Musumeci, and D. Ratner, “Electron ghost imaging,” Phys. Rev. Lett. **121**, 114801 (2018).
- [78] T. Boltz *et al.*, “Studies of longitudinal dynamics in the micro-bunching instability using machine learning,” in *9th International Particle Accelerator Conference (IPAC 2018) Vancouver, BC Canada, April 29-May 4, 2018* (2018) p. THPAK030.
- [79] A.S. Nawaz, G. Lichtenberg, S. Pfeiffer, and P. Rostalski, “Anomaly Detection for Cavity Signals - Results from the European XFEL,” in *9th International Particle Accelerator Conference*

(IPAC 2018) Vancouver, BC Canada, April 29-May 4, 2018 (2018).

[80] <https://www.nvidia.com/en-us/data-center/tesla/>.

[81] <https://colab.research.google.com>.

[82] <https://aws.amazon.com/machine-learning/>.

[83] <https://cloud.google.com/ml-engine/>.