ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

ADDING DATA SCIENCE AND MORE INTELLIGENCE TO OUR ACCELERATOR TOOLBOX

S.G. Biedron*, Element Aero, Chicago, IL 60643 USA

Abstract

In this paper, we briefly review the generalities of data science, give examples of uses in the physical sciences and engineering, and discuss recent applications to accelerator science and engineering. Specifically, we give a couple of accelerator-centered examples and then point the readers to several recent documents and workshops that summarize data science for science as well as data science applied to accelerators. We also point to high-performance computing and robust simulation tools as part of the data science "tool" for accelerators.

INTRODUCTION

Today, phrases such as "Artificial Intelligence (AI)" and "Machine Learning (ML)" are used repeatedly and often interchangeably in the popular press to describe a sort-of magic bullet that can solve everything. Even business advertisements push the terms AI and ML for solving everything from control of self-driving cars to predicting financial fraud before it happens. Although these phrases/terms are abundant now in all forms of media, what they actually mean in general terms is *data science*. Data science is in fact not a new field, but computing technologies and advanced algorithms have immensely expanded the applications over the last decade. The particle accelerator community is now adapting these concepts to their problems and systems. This includes the expanded use of high-performance computing (HPC) resources coupled with robust codes based on first principles.

Here, we provide a very general overview of data science including intelligent methods of examining and using data. We then describe how high-performance computing and advanced simulations tools formulated on basic principles are part of the data science approaches in accelerator design and operation. Then, we mention a few examples of several applications using intelligence and data in particle accelerators including in controls. Next, we point the audience to reference materials for recent workshops, resources of past efforts in related fields, and recent initiatives for ethical design of intelligent systems by international professional technical societies. Finally, we remind ourselves we cannot be decoupled from several items, such as computing advances and systems engineering practices, when applying data science to accelerators.

GENERAL DESCRIPTION OF DATA SCIENCE

We understand the reader would enjoy having a simplistic definition of data science. In fact, however, much like everything in the universe, humans attempt to "bin" nature into segments or disciplines. Nature is in fact not segmented. Nature is nature. With this in mind, let's try to define data science.

Looking to a source that is updated by the masses, Wikipedia states that "Data science is a multi-disciplinary field that uses scientific methods, processes, algorithms and systems to extract knowledge and insights from structured and unstructured data" [1]. Data science to the author means the appropriate combination/integration of the gathering of data, knowledge of where the data originates, simulations from first principles, simple exploratory data analysis, mathematics, statistics, machine learning, optimization, visualization, high performance computing, deep learning, and using (if needed) specialized computer hardware [e.g., field-programmable gate arrays (FPGAs); application-specific integrated circuits (ASICs), including tensor processing units (TPUs), neural engines; etc.] all for a specific application. The point being that data science is understood differently by different people, and the concepts are far from new. Further, terms and concepts such as Neural Networks (NNs) [2] and machine learning (ML) have been around for more than a half century. For instance, in 1943, neurophysiologist Warren McCulloch and mathematician Walter Pitts wrote a paper on how the neurons of an artificial neutral network might work. They illustrated their concept through a simple neural network of electrical circuits [2]. Moreover, Arthur Samuel, in his 1959 article suggested that "[Machine Learning is the] field of study that gives computers the ability to learn without being explicitly programmed" [3].

Here, we describe one view that the mathematics and statistics of data science is a collision of numerical analysis and AI, as shown in Figure 1. While this does not reflect the source of data, simulations, or computing platforms/tools, etc., we hope this helps in understanding some of the components of data science.

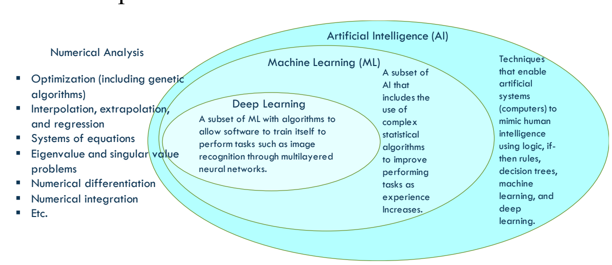

Figure 1: One concept of selected elements of the mathematics and statistics of data science.

Numerical analysis can contain mathematical optimization (including genetic algorithms); interpolation, extrapolation, and regression; systems of equations; eigenvalue and singular value problems; numerical differentiation; and numerical integration. *Artificial intelligence* (AI) can

*sgbiedron@elementaero.com

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

be thought of as including techniques that enable artificial systems (computers) to mimic human intelligence using logic, if-then rules, decision trees, machine learning, and deep learning. *Machine learning (*ML) can be viewed as a subset of AI that includes the use of complex statistical algorithms to improve task performance as experience increases. *Deep learning* (DL) can be described as a subset of ML with algorithms that allow software to train itself to perform tasks such as image and speech recognition through multi-layered neural networks.

Since ML is the center of many questions posed to the author, we will say a few more words about this area of AI. ML algorithms are essentially learning methods using, for instance, neural networks (NNs). Although NNs are algorithms that are discussed widely in the media, they do not necessarily have to be NNs, as one can use many other algorithms such as Gaussian processes or Viterbi algorithms. One must choose a learning (training) method that is appropriate from the many that exist and is already suited to your data (e.g., image, text, numbers, etc.).

Again, nature is not segmented, but ML can be thought of as having four variations. *Supervised learning* typically deals with categorical target variables or continuous target variables. Applications such as medical imaging can be addressed through the technique of classification for categorical target variables. Techniques such as regression can help with continuous target variables (an example being stock market predictions). *Unsupervised learning* deals with cases where the target variable is not available. For instance, strategies for predicting what a consumer might also purchase if they purchase one type of item (known as "market basket analysis") are based on association techniques. Clustering techniques are used, for example, to determine how to segment customers into sub-populations for a given product or service. *Semi-supervised learning* concentrates on categorical target variables and uses clustering techniques. Applications can include lane-finding with GPS data and using classification techniques for text classification. *Reinforcement learning* concentrates on classification methods for categorical target variables for cases of optimizing marketing and on control for cases where the target variable is not available, such as in control of driverless cars and certain particle accelerator control schemes.

The general steps to a training method include: 1) get data from a test or simulation; 2) prepare the data (may also be the step for analysis if you are using more standard numerical analysis techniques), which might include data reduction or so-called "cleaning" and "splitting" the data into training and model test data; 3) choose the learning method; 4) allow the machine (e.g., NN) to learn; 5) evaluate the model; 6) refine, refine, and refine the model; and then 7) execute the predict/use step! Of course, this is highly over-simplified, and many details have been left out regarding the details of the algorithms themselves, but this presents a general picture of the steps in data science.

Data science is a complex gathering of one or more techniques; there is no one magic solution, and it is certainly not a black box. People, for instance, can spend their entire careers in concentrated areas of data science for genres of specific data sets. Also, keep in mind that the individual(s) that developed a certain algorithm were perhaps concerned about making a better algorithm and not about our exact data set in accelerators. The algorithm might have been developed on random data. We accelerator people, on the other hand, have "more interesting" data sets, including errors in data and embedded physics. This can also be said for the data that is generated from accelerator-enabled applications. For our field, a warning: there might not be just one answer to the data science approach for a particular challenge, so be wise in your approach(es). To reinforce, this is not a black box. Many algorithms exist and there are infinite combinations. Also remember what George E. P. Box said: "All models are wrong, but some are useful."

Besides the algorithms representing/providing the model, our colleague Reza Pirayeshshirazinezhad gently reminds us (repeatedly) that the complexity of the data should be enough to make the algorithm efficient for modeling the system. The complexity of the data should represent all the variables with sufficient accuracy. In an autonomous system, for instance, there are many components working together, and the data taken from the system might not sufficiently represent the finer aspects of the system. This makes the (for instance) ML algorithm inefficient in modeling the system with enough accuracy. In particle accelerators, for example, the data should represent all the variables of the system(s) under consideration. Looking toward one example—the low-level radio frequency (LLRF) sub-systems for superconducting particle accelerators—the data should represent the range of the controller parameters, the noise and errors in the system(s), all the cryomodules (if in a superconducting system), all the different types of the controllers, etc. If, for example, PID controllers are used, and we would like to provide an efficient controller, we need to experiment using a large enough number of parameters over an accurate and sufficiently wide range of the parameters. If the range is not accurate or wide enough, an ML (or other) model can only represent that given range of the controller parameters. If we only use a PD controller, the ML (or other) model can only represent the controllers with PD controllers, not PID controllers. As a result, the complexity of the data should meet the requirement(s) of the system.

AI can be interpreted as an estimation of the system. This estimation can include a variety of parameters, including the dynamics of the system and the feedback of the controller. These dynamics, in particle accelerators, represent the equation of the system. These dynamics can help in effectively estimating the system. This estimation is primarily based on the dynamics of the system, the sensors and the feedback associated with the sensors, the controllers, and the states of the system. The Kalman filter is one of the approaches that works well with linear systems. There are many approaches to optimally estimate the system. Other famous approaches are the extended Kalman filter, the ensemble Kalman filter, batch state estimation, and particle filtering. These filters can help reduce the noise and make the estimation more optimal.

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

As was already mentioned, there are many algorithms in data science; if using an AI approach, one must choose them based on the system and what is required from the system. The algorithm is also highly affected by the complexity of the data, and the algorithm can be hybrid and hierarchical. For example, an adaptive controller with different hybrid optimization can be hierarchically mounted on a particle accelerator by a reinforcement algorithm. This adds more complexity to the system, and so the system works more efficiently if the algorithm(s) are implemented correctly. This is the beauty of AI, the beauty that helps solve advanced and complex problems and can be applied to systems like particle accelerators.

DATA SCIENCE IN OUR ACCELERATOR TOOLBOX

Our accelerator community is composed of many disciplines and sub-disciplines. We are typically trained formally in physics, electrical engineering, and/or mechnical engineering. The skill set, "our toolbox," however, is filled with many items already, including special relativity, electromagnetism, collective effects, the photoelectric effect, laser physics, superconductivity, materials science, vacuum science, vacuum electronics/rf sources, instrumentation, controls engineering, computational techniques, radiation effects, health physics, dosimetry, pulsed power, survey and alignment, etc. In the last few years, our field has been reluctant to accept high-performance computing and data science (including mathematical optimization and AI techniques). Why? Perhaps since we as a community already have such a comprehensive skill set (a full "toolbox"), this may have led to a slow adaptation of data science, as we thought we could solve problems by other methods already in our toolbox. Adding data science means that we need more training and more people to do the job. We assert that we must accept this, as we want our systems to perform at a higher reliability and in unique ways. In short, this added complexity as well as demands on overall performance necessitate more tools in our toolbox, including data science.

ADDING SYSTEMS ENGINEERING TO OUR ACCELERATOR TOOLBOX IS IMPORTANT FOR DATA SCIENCE

Building data science and intelligence into future accelerators begins on Day 1 (as should controls!). We as a community have not consistently applied systems engineering principles to particle accelerators. When planning a new or upgrading an aging machine, we need to remind ourselves of the need to use more modern techniques when "architecting" the machine. What we mean is that we should think of using computing and data science more globally around the machine and "fusing" the entire system. Systems engineering is already practiced on many defense, space, and industrial software and systems. If we want to think globally about our systems, we need to employ systems engineering practices. These are not easy concepts to carry on throughout the lifetime of a project. Systems engineering is usually formally taught at the graduate level, but many resources can be found through systems engineering societies, such as the International Council on Systems Engineering [4].

If we want to use data science tools, then we need to architect this goal into the system (or system upgrades)—real-time data logging, local computing (e.g., GPUs, FPGAs), application-specific integrated circuits, co-location to a powerful cluster or HPC, better diagnostics, better electronics, etc. From our experience, for instance, powerful and smart controls tools cannot be mounted onto systems that are not architected to handle such data science applications.

A FEW EXAMPLES AND WHERE TO FIND EVEN MORE

Here we provide only a few examples of data science in accelerators and related fields. We also provide resources to find other related activities within our own and other fields. There are millions of examples of data science in many disciplines from which to learn.

The first example is one use of high-performance computing in the design of machines. Care must be taken when setting up the simulations and dealing with the enormous amounts of simulated data; the codes must be capable of parallelization, for instance. In recent discussions with Michael Borland of Argonne National Laboratory [5], he expressed that the Advanced Photon Source Upgrade (APS-U) beam physics effort was provided with ~60M core hours per year on Argonne internal computing clusters with excellent support. Having expert support on the data science side is imperative. He noted that they also had computing resources at the Argonne Leadership Computing Facility and that similar computing resources were being utilized by the Advanced Light Source's Upgrade Team at Lawrence Berkeley National Laboratory (LBNL). For the work on the APS-U, the high-performance computing resources at Argonne were used to run *elegant*, a general-purpose Argonne-developed code for linacs, transport lines, and rings. It was used to generate the frequency map analysis for the APS-U final design lattice [6] using a massively parallel multi-objective genetic optimizer.

He stressed that to reduce design costs, HPC is essential as it: provides high-fidelity models with fewer compromises, leading to fewer (often no) iterations with hardware prototypes; makes best use of expensive engineering and physics staff; and provides robust statistical data on expected performance. He also offered his opinion on the common reasons HPC is underutilized. From his point of view, he felt that often the staff is not trained to use HPC systems and software; that the software was chosen for familiarity and convenience, not HPC capabilities; and that management undervalues computation and fails to provide access to HPC systems and support.

Our colleague at LBNL, Jean-Luc Vay [7], offers similar views that extreme computing enables advances by offering predictive, detailed, beam and accelerator science. Using Cori at the National Energy Research Scientific

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

Computing Center (NERSC), they provided two examples of improvements in reducing the time to solution using data science—HPC + optimizers + novel algorithms + analysis. The first was for the Parallel Luminosity Optimization in an rf accelerator. Using the parallelized code BeamBeam3D [8,9] plus an optimizer on Cori, the time to solution was about 24 hours. If this had been run serially with the optimizer it would have taken 35 years. The second example studied the pulse front tilt in laser-plasma using WarpX [10,11]. In this case, the time to solution comparison was that the parallel code in the boosted frame ran in ~4 hours and the serial code in the lab frame would have taken ~500 years.

Our colleague Jochem Snuverink [12] at the Paul Scherrer Institute (PSI) has recently launched a collaboration called "PACMAN: Particle Accelerators & Machine Learning." They hope to use ML to operate particle accelerators such as the Large Hardron Collider (LHC) as safely as possible by minimizing beam losses, having better control of accelerator parameters, preventing unnecessary machine interruptions, and improving beam dynamics modeling. One example is the recent work with L. Coyle (EPFL and CERN) where, using the LHC physics fills from 2018, they built a surrogate model of the LHC at injection energy. With that model, they determined the best working point. Then, they performed experimental scanning of the various parameters they had in the model and measured the lifetimes. They found that the measured lifetime is the highest and also reflects the prediction of the surrogate model. They believe that larger emittances in some cases, most likely due to collective instabilities, resulted in errors in the trained model that in the future can be further improved. Their work is continuing.

Colleagues such as Anna Solopova are working on "SRF Cavity Fault Classification Using Machine Learning at CEBAF" [13], and Gabriella Azzopardi is working on "Operational Results of the LHC Collimator Alignment using Machine Learning" [14].

Our own work began as a quest to have human-centered control of a near-autonomous accelerator operation as well as rapid switching between free-electron laser (FEL) frequencies while managing the associated instabilities. The main focus has been exploring control of accelerators and lasers [15].

At the Australian Synchrotron, the Linac Coherent Light Source (LCLS), and the Fermi@Elettra at Sincrotrone Trieste [16-21], we experimentally stabilized the electron beam energy and energy spread in the presence of klystron phase and amplitude jitter. We also demonstrated a novel controller that maintained the maximum transmission through the machine.

Further, we have explored long, variable time delays due to large thermal time constants in water transport, fast-switching in FELs, virtual diagnostics, in laser control, and control of medical accelerators, where our work continues [22-33].

In other related fields, we can immediately find data science examples that can be mapped over to accelerators. Employing AI algorithms has demonstrably sped up calculations (e.g., one million times for analysis of strong lenses) [34] and produced clear cost savings (30% increase in effective detector volume for neutrino flavor tagging with neural networks) [35,36]. A recent talk by Gabriel Perdue at Fermilab clearly visually illustrates the similarity of the predictions and DL reconstruction in neutrino physics [37]. At Los Alamos National Laboratory some of this research—including work in shock physics using AI-based data science techniques—has resulted in a number of publications [38-43].

Another example is in the field of Optics and Imaging. There is a special issue of *IEEE Access* entitled "Advanced Optical Imaging for Extreme Environments" [44]. We want to bring the reader's attention to this special edition as optical imaging systems have many similarities to particle accelerators, including that "…these optical systems and optical processing methods adhere to age-old architectures that drive practical solutions towards unsustainable complexity under ever-increasing performance requirements", and that "Researchers need to reinvent optical devices, systems, architectures, and methods for extreme optical imaging. On the other hand, artificial intelligence has become a topic of increasing interest for researchers. It is foreseeable that artificial intelligence will be the main solution of the next generation of extreme optical imaging." We invite you to explore this special edition.

One final example is applying machine learning to formation control of multiple spacecraft. Spacecraft are required to be autonomous, since they are far from Earth and any sort of intervention. At launch, there should be an onboard agent dictating different operational scenarios of the system in order that the spacecraft can be fully autonomous and efficient, and follow the requirements of the mission. High-precision formation control of multiple spacecraft, just like the high-precision control in particle accelerators, requires advanced controls and AI approaches. An ML approach can help the control system provide more energy-efficient inputs in the actuators and more precise formation of the spacecraft. This has been illustrated in research for a NASA program entitled, "Virtual Telescope for X-Ray Observation" [45,46]. These techniques provide us with highly precise control and an energy-efficient spacecraft that can reveal the nature of space, including black holes. This example shows the importance of ML for autonomous systems.

After summarizing these few examples and before pointing the reader to the many available resources for data science, we want to mention the work of individuals from several technical professional societies who are thinking in terms of systems engineering of intelligent systems. There are important ethical discussions underway regarding integrating intelligent data science into designs. The Institute of Electrical and Electronics Engineers (IEEE) is active in investigating these topics. One significant activity was started in 2016 and is referred to as the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. It is an open, global, and inclusive (regionally, culturally, gender, etc.) community of experts and interested persons

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

from technology and human science. Its mission is to ensure that every stakeholder involved in the design and development of autonomous and intelligent systems is educated, trained, and empowered to prioritize ethical considerations so that these technologies are advanced for the benefit of humanity. One outcome from this initiative is a standard [47]. In addition, the IEEE P7000 series of standardization projects is an ongoing process that seeks to gather more concrete guidelines through consensus in a meaningful time frame [48].

Not all improvements need to be focused around truly intelligent data science. There are many controls algorithms, for instance, that can be adapted from other fields first, if addressing controls challenges. The most challenging problems should be left to the most "intelligent" approaches. We can benefit from looking to publications such as those available through the IEEE on intelligent systems [49]. There are also many other related IEEE activities including special sections of journals that address intelligent systems [50].

In the last fifteen months, four workshops have been held, demonstrating the interest in artificial intelligence techniques as applied to particle accelerators as well as physics in general [51-54]. More than fifteen have been held in high-energy physics in the last several years [see "Promise and Challenges of Machine Learning in Particle Physics, Astrophysics, and Cosmology" by Kyle Cranmer in ref. 53]. The first, "Intelligent Controls for Particle Accelerators," was held on 30-31 January, 2018, at Daresbury Laboratory [51]. The second, "ICFA Beam Dynamics Mini-Workshop: Machine Learning Applications for Particle Accelerators," was held on 27 February-2 March, 2018, at SLAC [52]; this also has a follow-up report written by many of those who could attend [55]. The third, "Physics Next: Machine Learning," was held by the American Physical Society Physical Review publications group in New York on 8-10 October, 2018 [53]. This was an exploratory meeting to help determine if a machine learning journal would be needed in the physics realm due to the recent uptick in AI-centered data science for physics. The conclusion was that one is not needed since there already exist so many other pathways for publication, and machine learning and data science in general are tools of many fields. The fourth workshop was the "2nd ICFA Workshop on Machine Learning for Charged Particle Accelerators," that was held from 26 February-1 March, 2019, at the Paul Scherrer Institute (PSI) [54]. Finally, a recent Department of Energy workshop, "Basic Research Needs Workshop on Compact Accelerators for Security and Medicine [56]," was held in early May, 2019. The workshop report, due to be released in August, 2019, will help guide research priorities for accelerators for several applications in security and medicine.

CONCLUSIONS

This paper and its accompanying presentation provide a very general overview of data science and its role in our accelerator science and engineering toolbox. We encourage the reader to explore the references, including the recent workshops on concepts and applications of data science for accelerators.

Our accelerator user colleagues at Fermilab, including James Amundson, remind us that "[AI/Intelligence is one way of looking at things] and it is more than just algorithms, it is data, sensors, computing platforms and analysis techniques!" Our accelerator systems are complex, and we need all the right tools to understand and make them better.

A few things to keep in mind with data science are: 1) to effectively apply data science to accelerators, we must have knowledge of the system and the physics, engineering, materials, etc. behind it; 2) we need to choose the right algorithm(s) to explore the system, not just about the data interpretation we already know (in other words what might be hidden and how do you find it?); 3) we need to find the right data set for the question(s) we have (e.g., do we need more or better diagnostics, are we sure we are measuring the correct things, etc.); and 4) we need to realize that we ourselves have to improve our own methods as new or better data science methods are discovered for a given application (e.g., we the authors are taking something we examined with genetic algorithms and are now using a deep belief network algorithm to formulate our control of laser systems).

A few areas of research we believe:

- Include using ML to find glitches in our software (simulation software as well as controls software);
- Include the use of specialized computing hardware (e.g., FPGAs, application-specific integrated circuits, systems-on-a chip) that are in wide use by our cousins in HEP and are natural stepping-stones toward dedicated AI hardware;
- Think like industry by not jumping to a product. In the case of the quest for compact accelerators, the more data we have on a "prototype" the better and cheaper we can make the next versions to be produced and sold. Intelligent methods will be important here;
- Develop new intelligent techniques and algorithms specifically to address the underlying physics in an accelerator and peripheral systems (e.g., lasers);
- Use data science intensively to further improve controls.

Further, our field is behind in data science, and we need to convince the managers as well as the funding agencies that data science is low-hanging fruit in the field of particle accelerators.

Finally, keep in mind that data science is only human—the only working model of a universe is a universe.

ACKNOWLEDGMENTS

The author expresses special thanks to Jim Amundson, Aasma Aslam, Gabriella Azzopardi, Trudy Bolin, Michael Borland, John Cary, Mark Curtin, Massimo Dal Forno, Jorge Alberto Diaz Cruz, Catherine Eyberger, Richard Farnsworth, Steve Lidia, Evelyne Meier, Stephen Milton, Michael Papka, John Petillo, Reza Pirayeshshirazinezhad, Robert Rosner, Quentin Saulter, Anna Solopova, Salvador

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

Sosa, Jochem Snuverink, Josh Stein, Catherine Sweeney, Jean-Luc Vay, and Ji Qiang for many discussions in preparation for this paper and its accompanying talk. I also acknowledge all of the persons with whom I have collaborated on the subject of data science and with whom I have co-authored papers and reports.

Special thoughts to the author's deceased father (Emil Biedron) who piqued the author's interest in controls by building a control system with her for her model train set, and by teaching her about automobile and aircraft control, and analysis of related intelligent systems. Discussions included using intelligence to understand complex systems such as patterns and predictions of transportation in his work at the Chicago Area Transportation Study and the Northeast Illinois Planning Commission.

Special and admiring thoughts go to our recently deceased colleague Greg LeBlanc with whom we worked on the subject of data science for accelerators since "the old days."

REFERENCES

[1] Wikipedia definition of data science accessed 30 May, 2019, `https://en.wikipedia.org/wiki/Data_science`

[2] W. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *B. Math. Biophys.*, vol. 5, p. 115, 1943.

[3] A.L. Samuel, "Some studies in machine learning using the game of checkers," *IBM J. Res. Dev.*, vol. 3, no. 3, p. 210, 1959.

[4] `https://www.incose.org`

[5] Michael Borland, private communication, May, 2019.

[6] `https://www.aps.anl.gov/Accelerator-Operations-Physics/Software`

[7] Jean-Luc Vay, private discussions.

[8] `http://blast.lbl.gov/blast-codes-beambeam3d/`

[9] J. Qiang, S. Lidia, R.D. Ryne, and C. Limborg-Deprey, "Three-dimensional quasistatic model for high brightness beam dynamics simulation," *Phys. Rev. ST Accel. Beams*, vol. 9, p. 044204, 2006, `https://doi.org/10.1103/PhysRevSTAB.9.044204`

[10] https://github.com/ECP-WarpX

[11] J.-L. Vay et al., "Warp-X: A new exascale computing platform for beam–plasma simulations," Nucl. Instrum. Methods A vol. 909, pp. 486-479, 2018, `https://doi.org/10.1016/j.nima.2018.01.035`

[12] Jochem Snuverink, private discussions.

[13] A.D. Solopova et al., "SRF cavity fault classification using machine learning at CEBAF," presented at IPAC'19, Melbourne, Australia, May 2019, paper TUXXPLM2, this conference.

[14] G. Azzopardi, A. Muscat, G. Valentino, S. Redaelli, B. Salvachua, "Operational results of the LHC collimator alignment using machine learning," presented at IPAC'19, Melbourne, Australia, May 2019, paper TUZZPLM1, this conference.

[15] S.G. Biedron, "Free-electron laser activities," presented at U.S. Navy's MW-Panel Meeting, Arlington, Virginia, October 2004.

[16] E. Meier, S.G. Biedron, G. LeBlanc et al., "Electron beam energy and bunch length feed forward control studies using an artificial neural network at the Linac Coherent Light Source," Nucl. Instrum. Methods A, vol. 610, no. 3, p. 629, 2009.

[17] E. Meier, M. Morgan, S.G. Biedron et al., "Development of a combined feed forward-feedback system for an electron linac," Nucl. Instrum. Methods A, vol. 609, nos. 2-3, p. 79, 2009.

[18] E. Meier, S.G. Biedron, G. LeBlanc, M.J. Morgan, "Development of a novel optimization tool for electron linacs inspired by artificial intelligence techniques in video games," Nucl. Instrum. Methods A, vol. 632, no. 1, p. 1, 2011.

[19] E. Meier, S.G. Biedron, G. LeBlanc, M.J. Morgan, "Artificial intelligence systems for electron beam parameters optimisation at the Australian Synchrotron," in Proc. IPAC'10, Kyoto, Japan, May 2010, paper TUPEC025, pp. 1770-1772.

[20] E. Meier, S.G. Biedron, G. LeBlanc et al., "Electron beam energy stabilization using a neural network hybrid controller at the Australian Synchrotron Linac," in Proc. PAC'09, Vancouver, BC, Canada, May 2009, paper TU5RFP050, pp. 1201-1203.

[21] E. Meier, S.G. Biedron, G. LeBlanc et al., "Electron beam energy stabilisation test results using a neural network hybrid controller at the Australian Synchrotron and LINAC Coherent Light Source Projects," in Proc. FEL2009, Liverpool, UK, August 2009, paper FROA02, pp. 766-771.

[22] AL. Edelen, S.G. Biedron, S.V. Milton et al., "Initial experimental results of a machine learning-based temperature control system for an RF gun," in Proc. IPAC'15, Richmond, VA, USA, May 2015, pp. 1217-1219, `doi:10.18429/JACoW-IPAC2015-MOPWI028`

[23] A.L. Edelen, S.G. Biedron, B.E. Chase, D. Edstrom, S.V. Milton, P. Stabile, "Neural networks for modeling and control of particlie accelerators," IEEE Trans. Nucl. Sci., vol. 63. no. 2, pp. 878-897, 2016, `doi:10.1109/TNS.2016.2543203.` (Invited paper)

[24] A. Edelen, S. Biedron, S. Milton et al., "Resonant frequency control for the PIP-II injector test RFQ: control framework and initial results," in Proc. NA-PAC2016, Chicago, IL, USA, October 2016, pp. 109-112, `doi:10.18429/JACoW-NAPAC2016-MOPOB17`

[25] D. Bowring, B. Chase, J. Czajkowski et al., "Resonance control for Fermilab's PXIE RFQ," in Proc. IPAC'16, Busan, Korea, May 2016, pp. 447-450, `doi:10.18429/JACoW-IPAC2016-MOPMW026`

[26] A.L. Edelen, S.G. Biedron, S.V. Milton et al., "Neural network model of the PXIE RFQ cooling system and resonant frequency response," in Proc. IPAC'16, Busan, Korea, May 2016, pp. 4131-4133, `doi:10.18429/JACoW-IPAC2016-THPOY020`

[27] A. Morin, S.G. Biedron, S.V. Milton et al., "Trajectory response studies at the Thomas Jefferson National Accelerator Facility energy recovery linac and free electron laser," in Proc. 16th Annual Directed Energy Symposium, March 2014, Huntsville, Alabama, Directed Energy Professional Society, `https://protected.networkshosting.com/depsor/store/merchandise/TOCs/DEsymp14-TOC.html`

[28] A. Morin, C. Tennant, S. Biedron, D. Douglas, S. Milton, "Control systems development for the Thomas Jefferson National Accelerator facility free electron laser: preliminary results for trajectory response data," in Proc. 15th Annual Directed Energy Symposium, November 2012, Albuquerque, New Mexico, Directed Energy Professional Society.

[29] A. Edelen, S. Biedron, J.P. Edelen, S.V. Milton, P.J.M van der Slot, "Using a neural network control policy for rapid switching between beam parameters in an FEL," in Proc.

ISBN: 978-3-95450-208-0 doi:10.18429/JACoW-IPAC2019-TUZPLM1

FEL2017, Santa Fe, NM, USA, August 2017, pp. 480-483, doi:10.18429/JACoW-FEL2017-WEP031

[30] A. Edelen, S. Biedron, "First steps toward incorporating image based diagnostics into particle accelerator control systems using convolutional neural networks," in Proc. NAPAC2016, Chicago, IL, USA, October 2016, pp. 390-393, doi:10.18429/JACoW-NAPAC2016-TUPOA51

[31] A. Edelen et al., "Recent application of neural network-based approaches to modeling and control of particle accelerators," presented at IPAC'18, Vancouver B.C., Canada, April-May 2018, slides THYGBE2, http://www.JACoW.org/ipac2018/talks/thygbe2_talk.pdf

[32] J.A.D. Cruz et al., RF control with data science, to be submitted.

[33] Work ongoing with A. Aslam, T. Bolin, J. Cary, M. Curtin, R. Pirayeshshirazinezhad, and S. Sosa on simulations and control techniques for accelerators and lasers for accelerators.

[34] Y.D. Hezaveh, L.P. Levasseur, and P.J. Marshall, "Fast automated analysis of strong gravitational lenses with convolutional neural networks," Nature, vol. 548, pp. 555-557, Aug. 2017.

[35] P. Adamson et al., "Constraints on oscillation parameters from νe appearance and νμ disappearance in NOνA," Phys. Rev. Lett., vol. 118, p. 231801, June 2017, https://doi.org/10.1103/PhysRevLett.118.231801

[36] M.A. Acero et al., "New constraints on oscillation parameters from νe appearance and νμ disappearance in the NOvA experiment," Phys. Rev. D, vol. 98, p. 032012, Aug. 2018, https://doi.org/10.1103/PhysRevD.98.032012

[37] Gabriel Perdue, "DL in neutrino experiments," ML in Science and Engineering, CMU, Pittsburgh, June 8, 2019.

[38] D.J. Walters et al., "Bayesian calibration of strength parameters using hydrocode simulations of symmetric impact shock experiments of Al-5083," J. Appl. Phys., vol. 124, no. 20, p. 205105, Nov. 2018, doi:10.1063/1.5051442 (also published as LA-UR-18-20884).

[39] D.T. Orban, D. Keefe, A. Biswas, J.P. Ahrens, and D.H. Rogers, "Drag and track: a direct manipulation interface for contextualizing data instances within a continuous parameter space: application to shock physics," presented at IEEE VIS 2018, Berlin, Germany, October 2018, http://ivlab.cs.umn.edu/papers/Orban-2018-Drag AndTrack.pdf (also published as LA-UR-18-22844).

[40] A. Biswas, J.P. Ahrens, and E.C. Lawrence, "High-dimensional data analysis and visualization using high fidelity emulators with uncertainty," presented at IEEE VIS 2018, Berlin, Germany, October 2018, http://ayanbiswas.net/vis18m-sub1022-cam-i7.pdf (also published as LA-UR-18-22810).

[41] D. Orban et al., "Continuous workflow and exploration of disparate data types using Cinema:Bandit: a toolset for beamline science demonstrated on XFEL shock physics experiments," LANL report LA-UR-19-22551, 2019. In journal review.

[42] C.M. Biwer, S.C. Vogel, M.M. McKerns, and J.P. Ahrens, "Spotlight: automation of Rietveld analyses using an ensemble of local optimizers," LANL report LA-UR-18-30288, 2018.

[43] S.C. Vogel et al., "Interactive visualization of multi-data-set Rietveld analyses using Cinema: Debye-Scherrer," J. Appl. Crystallogr., vol. 51, no. 3, pp. 943-951, June 2018, https://doi.org/10.1107/S1600576718003989

[44] https://ieeeaccess.ieee.org/special-sections/advanced-optical-imaging-for-extreme-environments/

[45] R. Pirayesh, A. Naseri, F. Moreu, S. Stochaj, N. Shah, J. Krizmanic, "Attitude Control Optimization of a Two-CubeSat Virtual Telescope in a Highly Elliptical Orbit," in Space Operations: Inspiring Humankind's Future, H. Pasquier, C.A. Cruzen, M. Schmidhuber, Y.H. Lee, Eds., Cham, Switzerland: Springer, 2019, pp. 233-258.

[46] A. Naseri, R. Pirayesh, R.K. Adcock, S.J. Stochaj, N. Shah, and J. Krizmanic, "Formation Flying of a Two-CubeSat Virtual Telescope in a Highly Elliptical Orbit," in Proc. 2018 Space Ops Conference, Marseille, France, 2018, https://doi.org/10.2514/6.2018-2633

[47] Private discussions Alpesh Shah, alpesh.shah@ieee.org, and Sam Sciacca, s.sciacca@ieee.org

[48] https://ethicsinaction.ieee.org/standards/p7000/

[49] See, for example, IEEE Transactions on Neural Networks and Learning Systems; IEEE Intelligent Systems; IEEE Transactions on Fuzzy Systems; IEEE Transactions on Nuclear Science; IEEE Transactions on Human-Machine Systems; IEEE Transactions on Cybernetics; IEEE Transactions on Control Systems Technology; IEEE Transactions on Automation Science and Engineering; IEEE Transactions on Industrial Informatics; IEEE Transactions on Systems, Man, and Cybernetics: Systems; IEEE Transactions on Circuits and Systems; IEEE Transactions on Pattern Analysis and Machine Intelligence; IEEE Transactions on Cognitive and Developmental Systems; IEEE Transactions on Image Processing; IEEE Robotics and Automation Letters; and IEEE Transactions on Emerging Topics in Computational Intelligence.

[50] S.G. Biedron, "Adding data science and more intelligence to our accelerator toolbox," presented at IPAC2019, Melbourne, Australia, May 2019, slides TUZPLM1, this conference.

[51] Intelligent Controls for Particle Accelerators Workshop, held 30-31 January 2018, Daresbury Laboratory, UK, www.cockcroft.ac.uk/events/ICPA/

[52] ICFA Beam Dynamics Mini-Workshop: Machine Learning Applications for Particle Accelerators, held February 27-March 2, 2018, SLAC National Accelerator Laboratory, Menlo Park, California, USA, conf.slac.stanford.edu/icfa-ml-2018/

[53] Physics Next: Machine Learning, held by the American Physical Society Physical Review publications group on October 8-10, 2018, Riverhead, Long Island, New York, USA, https://journals.aps.org/physics-next/2018/machine-learning

[54] 2nd ICFA Workshop on Machine Learning for Charged Particle Accelerators, held 26 February, 2019 to 1 March, 2019 Paul Scherrer Institut, Villigen, Switzerland, https://indico.psi.ch/event/6698/

[55] A. Edelen et al., "Opportunities in Machine Learning for Particle Accelerators," arXiv:1811.03172, November 2018, arxiv.org/abs/1811.03172

[56] Basic Research Needs Workshop on Compact Accelerators for Security and Medicine, May 6-8, 2019, Tysons, VA, https://www.orau.gov/brncasm2019/default.htm