Sandra Biedron (University of New Mexico and Element Aero), **Kevin Brown** (BNL), **Aric Hagberg** (LANL), **Manel Martinez-Ramon** (UNM), **Stephen Milton** (LANL) **Malachi Schram** (PNNL)

Increasingly complex, large-scale infrastructures, such as particle accelerators and reactors, can benefit from AI algorithms that factor in real-time analysis, and these advanced computational tools may prove indispensable for optimizing factors such as performance and power consumption while at the same time maintaining optimal operations with little human expert intervention. Despite the importance of these complex, large-scale infrastructures, there only exist incomplete, limited AI models of them that are in vague resemblance to the real systems they are modeling [1]. This stems somewhat from a lack of intensive models-based systems engineering incorporating AI concepts from the birth of a project/concept, but more significantly from the abyss lying between the various communities, e.g. applied mathematics, designers/builders of complex systems (e.g. particle accelerator engineers, nuclear engineers, systems engineers), experts in artificial intelligence, experts in controls, and experts in high-performance computing [2].

Specifically, successful operation of particle accelerators requires as system that can successfully manage many interleaved control problems over a variety of operating modes, different sampling rates, and with multitude of constraints.

Use of AI, specifically machine learning (ML) methodologies, is particularly well suited to help here. A few examples of this whose solution can be attempted by reinforcement learning (ML) techniques specifically include:

- Bringing complex accelerator systems into am optimal state as fast as possible (changing modes fast across a large number of parameters);
- Maintaining optimal conditions over long periods and specifically compensating for drift (changing many potentially couple and nonlinear variables by small increments, slowly, based on a trained ML);
- Optimizing a machine performance condition or important parameter (e.g., maximizing colliding beam luminosity, or to optimize for cooling)

The third one is an example where neither measurements nor model are sufficient to find an optimal state for a given figure of merit. The situation is just too complex, so a heuristic is typically used as a starting point, modeling is done, and the machine experts do the final tuning by hand. This is kind of a *holy grail* but there are subsystems that follow the same paradigm (e.g., finding the optimal positions of 28 degrees of motional freedom in the SPS slow extraction septum positions to minimize beam loss, or finding the 250 or so correct collimator angles and positions in the Large Hardon Collider collimators for minimized backgrounds at both the ATLAS and CMS detectors).

All of these "features" make particle accelerators highly complex systems that can greatly benefit from the application of Data Science/Artificial Intelligence/Machine Learning/High-Performance Computing (DS.AI/ML/HPC) with areas such as design optimization (e.g. through deep reinforcement learning [3]), computational decision support and to help with system resilience. Although there are activities being derived from another recent workshop and report, *Workshop*

Report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence [4], these activities only scratch the surface of what is required for highly complex systems [5,6] such as particle accelerators. Moreover, the Advanced Scientific Computing Research (ASCR) Priority Research Directions (PRDs) and their recent calls for proposals and forthcoming research are focused on solving specific applied mathematical and HPC challenges.

Several communities need a dedicated exploration space where they can experimentally test using DS/AI/ML for a variety of complex systems. Even reinforcement learning (RL) based controllers, for instance, are challenging to apply to accelerators as the entire environment cannot be fully modeled. During the last five years a variety of extensions and modifications to RL have been introduced that need to be explored in the field of particle accelerators. For example, in [7] the concept of Safe RL has been introduced. Extensive exploration into incorporating Long Short Term Memory networks (LSTMs) to incorporate "ideas" already learned as well as transfer learning from other similar systems is also required. This is an unexplored example of Deep Reinforcement Learning [8] combined with transfer learning [9]. In [10], an approach to use the framework of stochastic optimal control with path integrals is promising.

Further yet, we need to also compare many networks and find ways of quantifying the uncertainty as applied to many complex systems, as the ones used, for example, in Bayesian frameworks combined with variational inference. Such additional research will help us understand many things, including how to best model complex systems, train and create emulators, and provide insight how to better monitor and control complex devices.

There is still much research and development to do.

[1] "AI for Science Report" https://www.anl.gov/ai-for-science-report
[2] Basic Research Needs Workshop on Compact Accelerators for Security and Medicine, Tools for the 21 Century May 6-8, 2019, https://science.osti.gov/-
/media/hep/pdf/Reports/2020/CASM_WorkshopReport.pdf?la=en&hash=AEB0B318ED0436B1C5FF4E
E0FDD6DEB84C2F15B2
[3] Yonekura, K., Hattori, H. "Framework for design optimization using deep reinforcement learning," Structural and Multidisciplinary Optimization, 60, 1709–1713 (2019). https://doi.org/10.1007/s00158-019-02276-w
[4] "Workshop Report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence," http://www.osti.gov/biblio/1478744
[5] "Assessing the Reliability of Complex Models – Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification," National Research Council 2012. Washington, D.C.: The National Academies Press. http://www.nap.edu/catalog.php?record_id=13395
[6] "A Multifaceted Mathematical Approach for Complex Systems," Report of the DOE Workshop on Mathematics for the Analysis, Simulation, and Optimization of Complex Systems. https://www.osti.gov/servlets/purl/1093585

[7] Garcia, J . Fernandez, F. A comprehensive survey on safe reinforcement learning. Journal of Machine Learning Research, 16, 1437–1480 (2015).
[8] Mnih, V., Kavukcuoglu, K., Silver , D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidje-land, A. K., Ostrovski, G., et al.Human-level control through deep reinforcement learning. Nature,518,529–533 (2015).
[9] Taylor, Matthew E., and Peter Stone. "Transfer learning for reinforcement learning domains: A survey." Journal of Machine Learning Research 10, 1633-1685 (2009).
[10] Theodorou, E. A., Buchli, J., and Schaal, S..A generalized path integral control approach to reinforcement learning. Journal of Machine Learning Research, 11, 3137–3181 (2010).