

# Response to NITRD, NCO, NSF Request for Information on Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan

J. Amundson<sup>1</sup>, J. Annis<sup>1</sup>, C. Avestruz<sup>2,3</sup>, C. R. Bom<sup>4</sup>, D. Bowring<sup>1</sup>, J. Caldeira<sup>2</sup>, G. Cerati<sup>1</sup>,  
C. Chang<sup>3,5</sup>, S. Dodelson<sup>6</sup>, J. Einstein-Curtis<sup>1</sup>, D. Elvira<sup>1</sup>, A. Farahi<sup>6</sup>, P. J. Fox<sup>1</sup>, M. J. Franklin<sup>7</sup>,  
K. Genser<sup>1</sup>, L. Gray<sup>1</sup>, O. Gutsche<sup>1</sup>, P. Harris<sup>8</sup>, J. Kinney<sup>9</sup>, R. Kondor<sup>7,10</sup>, J. B. Kowalkowski<sup>1</sup>,  
R. Kutschke<sup>1</sup>, R. Mandelbaum<sup>6</sup>, S. Mrenna<sup>1</sup>, B. Nord<sup>1,3,5</sup>, A. Para<sup>1</sup>, K. Pedro<sup>1</sup>, G. N. Perdue<sup>1</sup>,  
H. B. Runesha<sup>11</sup>, A. Scheinker<sup>12</sup>, P. Spentzouris<sup>\*1</sup>, J. St. John<sup>1</sup>, N. Tran<sup>1</sup>, S. Trivedi<sup>13</sup>,  
L. Trouille<sup>14,15,16</sup>, L. M. Walkowicz<sup>17</sup>, and W. L. K. Wu<sup>3</sup>

<sup>1</sup>Fermi National Accelerator Laboratory, Batavia, IL 60510, USA

<sup>2</sup>Enrico Fermi Institute at the University of Chicago, Chicago, IL 60637, USA

<sup>3</sup>Kavli Institute for Cosmological Physics at the University of Chicago, Chicago, IL 60637, USA

<sup>4</sup>Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, Brazil

<sup>5</sup>Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

<sup>6</sup>McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University,  
Pittsburgh, PA 15213, USA

<sup>7</sup>Department of Computer Science, The University of Chicago, Chicago, IL 60637, USA

<sup>8</sup>Massachusetts Institute of Technology, Cambridge, MA 02139, USA

<sup>9</sup>Google, Inc.

<sup>10</sup>Department of Statistics, The University of Chicago, Chicago, IL 60637, USA

<sup>11</sup>Research Computing Center, The University of Chicago, Chicago, IL 60637, USA

<sup>12</sup>Los Alamos National Laboratory, Los Alamos, NM 87545, USA

<sup>13</sup>Institute for Computational and Experimental Research in Mathematics, Brown University,  
Providence, RI 02903, USA

<sup>14</sup>Department of Citizen Science, The Adler Planetarium, Chicago, IL 60605, USA

<sup>15</sup>Zooniverse

<sup>16</sup>Center for Interdisciplinary Exploration and Research in Astrophysics, Northwestern University,  
Evanston, IL 60208, USA

<sup>17</sup>Department of Astronomy, The Adler Planetarium, Chicago, IL 60605, USA

October 26, 2018

## 1 Introduction

This document presents a response to the question of whether and how the National Artificial Intelligence Research and Development Strategic Plan (NAIRDSP) should be updated from the perspective of Fermilab, America's premier national laboratory for High Energy Physics (HEP). We believe the NAIRDSP should be extended in light of the rapid pace of development and innovation in the field of Artificial Intelligence (AI) since 2016, and present our recommendations below. AI has profoundly impacted many areas of human life, promising to dramatically reshape

\*corresponding author: [spentz@fnal.gov](mailto:spentz@fnal.gov)

society — e.g., economy, education, science — in the coming years. We are still early in this process. It is critical to invest now in this technology to ensure it is deployed efficiently, safely and ethically. Science and society both have a strong need for accuracy, efficiency, transparency, and accountability in algorithms, making investments in scientific AI particularly valuable. Thus far the US has been a leader in AI technologies, and we believe as a national laboratory it is crucial to help maintain and extend this leadership. Moreover, investments in AI will be important for maintaining US leadership in the physical sciences.

HEP is concerned with the investigation of the most fundamental building blocks of matter and the forces that govern the universe. Like all the scientific disciplines, HEP will be transformed by the capabilities of AI for large-scale data analysis, pattern recognition, and anomaly detection. The HEP community contributes to the execution of a national strategy by advancing new technologies and by utilizing AI to solve some of the most difficult problems in science. We are early adopters of powerful tools, like AI, because we must be — to complete our mission we must push the best tools available to their limits and beyond. Some kinds of AI algorithms, like machine learning and deep learning, have quickly become important tools of the trade for analysis of big data sets. Given anticipated future data sizes and the complexity of problems, we have barely scratched the surface.

AI is more than algorithms. AI comprises data, sensors, computing platforms, and analysis techniques. HEP is well positioned to make unique contributions to each of these facets. As a community, we work with data at enormous scale and develop novel, state-of-the-art detector systems. We readily embrace new computing platforms, and our focus on understanding bias and quantifying uncertainty will help with issues, such as algorithm explainability, cost-benefit analysis, and risk estimation in arenas outside of the sciences. It is not enough to just design a new AI algorithm — it must be connected to the real world, and we have the technical skills and use cases to be effective in this role. Moreover, the data-driven algorithms of AI critically lack a statistical theoretical basis. It's crucial that we identify areas of AI theory where the physical sciences (data and methods) can inform AI tools and capabilities; this will create a positive feedback mechanism to the benefit of all parties.

The applications should focus on accelerating science and on adding discovery methods for new science avenues. This should include the analysis of data, as well as the control of systems and instruments. There should also be feedback of science domain knowledge and methods into developing the theory and tools of AI. In order to facilitate the engagement of our community and the development of the necessary science applications, we believe it is essential to build partnerships across industries and academia in order to provide access to AI tools, develop the necessary infrastructure, and to share appropriate data sets for advancing these fields. Discussed below are potential areas of research, development, application, and collaboration in AI that, from our perspective, will enable technological advancements to address scientific needs for HEP, as well as unique capabilities within the HEP community that can contribute to advancing AI.

## **2 The Intersection of HEP and AI**

### **2.1 Advanced Analysis for Discovery with Big Data**

High Energy Physics is deeply invested in the analysis of very large and complex data sets that are based on consistent underlying systems of meaning. Indeed, the scale and complexity of modern data sets are exceeding the capacity of conventional data analysis algorithms and models. AI is evolving into an indispensable analysis tool for these data. Moreover, we believe that HEP data has the potential to enable key insights into the theory and algorithms of AI itself.

#### **2.1.1 AI accelerates HEP science**

Modern implementations of AI algorithms are accelerating discovery for many critical tasks, like classification, measurement, and simulation [1–3]. Employing these algorithms has demonstrably sped up calculations (e.g., one million times for analysis of strong lenses [4]), and produced clear cost savings (30% increase in effective detector volume for neutrino flavor tagging with neural networks; [5, 6]). These faster algorithms accelerate progress in science, and in some cases they lead to paradigm shifts: in situations where it once required one day for a human to analyze a single object, it may now take just a few seconds, pushing the frontier of discovery.

While AI is driving new analysis techniques, we find a suite of challenges in carrying these successes into the future and building on them. Most AI systems currently used for science must be trained on data sets assembled by humans; these systems therefore incorporate human biases into calculations. Nevertheless, high-precision hypothesis testing and modeling require accurate results that are free of bias. Even more, we must be able to interpret and explain the results coming from AI at least as well as we can from standard algorithms. But large models with millions of parameters, as is typical of AI algorithms, tend to lack interpretability. This is because the individual parameters (e.g.,

the weights in a neural network) lack clear physical meaning, and the measurements of error and uncertainty on these parameters differ in their statistical meaning from more conventional algorithms. Modern AI algorithms are presently accelerating measurements, but they will be too slow for even many near-future data sets; we need to continue to reduce training and inference time. HEP experiments require large amounts of simulated data for calibration of instruments and mathematical models, but these simulations remain prohibitively slow. Finally, there are demonstrated successes in AI producing models of phenomena that are already known to us, but, how do we leverage the speed and flexibility of AI to make genuinely new discoveries — outside the realm of current knowledge and beyond current training sets — to produce new understanding of the Universe?

A number of promising avenues can pave a path to the future of scientific discovery with AI. First, we must integrate statistical models with AI algorithms (e.g., measuring clear error bars on AI outputs). This will permit us to estimate and remove bias, and to produce statistically interpretable measures of uncertainty. Next, we should incorporate elements of conventional physical models (whose parameters have physical meaning) into AI algorithms to improve interpretability. This also has the potential to improve algorithm speed: by leveraging well-understood physical laws through these physics parameters, it can limit the search space during optimization. AI generative models can quickly produce simulated data sets, and we must learn how to deploy algorithms like these to more quickly produce simulations for calibrations and analyses [3]. When data becomes sufficiently large, these algorithms will need to run on distributed systems [7].

Beyond merely speeding up our analysis frameworks, AI has the potential to point us toward new theories of particle physics and cosmology. First, we discuss the connection between mathematical symmetries in AI algorithms and symmetries in nature. HEP science is governed by symmetries of space, time, and energy. AI algorithms typically exploit only translational symmetry to optimize the search for patterns in data composed of flat images. Recently, it has been shown that they can be generalized and implemented in a spherical context [8]. This is more appropriate, for example, for data that has rotational and spherical symmetries, like the night sky, which is on a spherical surface and which we see in upcoming large-scale cosmological experiments like CMB-S4 (Cosmic Microwave Background - Stage 4) [9] and LSST (Large Synoptic Survey Telescope) [10]. Also, consider the cases of force fields for molecular dynamics and properties of chemical compounds [11, 12]. The symmetries here are yet more complex. By matching the symmetry of the neural network architecture to the symmetry of the physical problem, we have a new opportunity for AI to learn the underlying physics in nature from data.

We also note the discovery potential that could be unleashed with ‘unsupervised’ AI algorithms. In these algorithms, there is no longer a training set from which the algorithm can learn patterns in the data before it sees the new target data; it must learn as it sees the data for the first time. For example, clustering algorithms measure how similar each datum is from all others [13]. Modern cosmological experiments, like LSST, will produce sufficiently large and complex data sets that searching for indications of new phenomena is prohibitively time-consuming for conventional non-AI algorithms. In these data sets, unsupervised algorithms could point us to outliers in the data, which could contain noise artifacts or new unexpected physical phenomena. Finally, to effectively and efficiently tackle enormous data rates and volumes, while also preserving the capacity for serendipitous discovery, it is important to leverage the complementary strengths of humans and machines. By keeping humans in the loop, these scenarios can embrace the speed of the algorithms while making them more accurate and interpretable. This has been demonstrated in citizen science efforts in galaxy classification [14].

### 2.1.2 HEP brings insight to AI

AI algorithms have been very successful on a number of standard data sets. For example, natural images of everyday objects (ImageNet [15, 16]) and handwritten digits (MNIST [17]) are canonical benchmarks for computer vision. Also, large-scale object detection, segmentation, and captioning data sets (COCO [18]) are available, as are standard environments for studying AI with games (OpenAI gym [19]).

HEP has unique capabilities in data acquisition and simulation. HEP data is highly structured, self-consistent, and emerges from first principles. We can produce very large, richly-labeled simulated datasets (that are easy to share), and as a field we specialize in understanding how they differ from nature. These data sets have meaning and “explainability” well beyond current natural image data sets, because they are developed from physical principles. HEP data sets can provide rich, new benchmarks for all kinds of AI algorithms, from computer vision to reinforcement learning.

It is crucial to mitigate concentration of expertise into pools surrounding large private datasets in order to have broad participation across the economy. The 2016 national strategy highlighted the need to develop shared public datasets and environments, and we should strongly extend those efforts. This supports workforce development and many other HEP initiatives that can result in opening public datasets for algorithm development — almost any project can include

an open data component, but we can also build projects specifically around constructing platforms for dataset hosting and challenges. The democratization of frameworks and tools [20–22], and the broad availability of commercial cloud computing services have opened the field to many participants — the difficult missing ingredient for initializing an AI research effort is no longer code, or access to modern compute infrastructure, but access to large, richly structured datasets. HEP simulations can play an important role in solving this problem.

## 2.2 Precise Control of Complex Systems

Automation in our society is accelerating, and promises to provide new levels of efficiency, cost savings, and improved standards of living. Key examples are automation in transportation [23] and industrial control [24]. Self-driving cars and AI-augmented factory workforces are projected to improve safety and energy efficiency. Enhanced facility management is already showing great energy-saving benefits: Consider Google’s success with data center management. Over the past few years, Google has developed algorithms that leverage AI to recommend efficient ways to cool their data centers [24–26]. The resulting cost savings were significant enough that Google has automated the implementation of these recommendations, reducing power consumption by 40% and freeing up engineers to focus on new challenges.

In HEP, we manage some of the most complex systems in the world. For example, particle accelerators have thousands of elements to control and precisely tune within very tight spatial and timing tolerances. In cosmology, telescope arrays scan the sky for highly diverse target sets (especially in surveys like LSST), and optimizing movements of the telescope is critical. Nearly every experiment at a national laboratory is dependent on quality operation of the accelerator complex. If automated control based on deep reinforcement learning [27, 28] and other strategies is able to reduce downtime of machines, and free operators from mundane burdens to focus on more demanding and innovative work, the entire program benefits. It is important for long-term efficiency and cost-effectiveness to begin automating the operation and control of these systems. We must use extremely fast systems for this, necessitating the development of real-time AI-based controls.

AI research on accelerator optimization and usage is increasingly important, because it can provide significantly better performance with relatively small cost. This is an increasingly competitive field internationally with many research avenues open for development.

Within HEP, accelerator control proposals to date have considered only the optimization and automation of a single small system, e.g., [29], but the real power of controls will materialize in the operation of large networks. Modern reinforcement learning research is focused on hierarchical problem solving – moving from global strategies all the way down into highly localized execution spaces [30]. Linking together multiple control systems across an accelerator complex into a larger network with greater responsibility and more opportunity for optimization is the end game, but execution of such a complex vision is a many-year project. We need to do a great deal of research to make such a system possible, and even more to make it safe and understandable. Furthermore, the physical complexity of accelerator systems and the huge number of difficult-to-learn constraints (for an exploratory algorithm) mean that it will be critical to combine deep domain expertise with algorithmic research in order to build a functioning system.

One crucial challenge to the automatic control and optimization of large, complex systems like particle accelerators becomes apparent when we consider the scenario when both the control components and the observed phenomena vary in time. This limits the accuracy and capability of model-based methods: even a deep neural network (trained with extremely large amounts of data to precisely track the desired beam characteristics) will only be accurate up until both the components and the beam drift from their original values. On the other hand, there are many model-independent approaches to optimization and control of uncertain and time-varying systems — e.g., Extremum Seeking (ES) algorithms [31]. While ES is able to simultaneously tune many parameters of highly dynamic systems in a completely model-independent way, it is susceptible to getting stuck in a local minimum, as are all adaptive, model-independent feedback-based approaches. Therefore, it is critical to combine the two fields — AI and adaptive feedback. With the adaptive AI approaches, models like neural networks provide instant, global approximations for parameter settings, based on a knowledge of the general features of large systems. Then local, model-independent, feedback-based approaches can zoom in on and track the optimal time-varying parameters.

When both our machines and the data they’re observing produce moving targets, we require novel approaches that engage both adaptive AI and model-independent approaches for just-in-time feedback to control mechanisms.

## 2.3 AI at the Edge: Fast, Intelligent Sensors

Many real-world problems require direct interaction between the computational resource and the physical environment. Sensor application complexity and integration are rising, and real-time collection and reaction times are declining. For example, in industrial settings, stringent requirements on sensitivity, safety, proximity, speed, and accuracy necessitate

automation. Sensors in these settings are not only numerous, but also diverse and sophisticated, ranging from high-resolution CCD detectors to high-speed particle detectors. This multiplicity of sensors can easily exceed what a person or a conventional computing system can accomplish in achieving the overall mission of such a facility.

AI promises to address computing problems in this setting through a number of avenues: fast conversion of sensor signals to actionable information (e.g. classification or measurement of physical quantities); the intelligent synthesizing of information; and the routing of these data to larger global decision support and data analysis systems. For AI systems to perform nearly any of these interactive functions, it will require real-time interaction with the environment — i.e., to be embedded at the point of sensing. Systems that include large-scale sensor read-out and processing naturally have extremely heterogeneous compute environments — utilizing the best computing elements for the job, accounting for factors like energy use and robustness against temperature, dirt, radiation, interference, and shock. AI promises to perform well (speed and accuracy) under these conditions, embedded directly in the problem domain.

DOE National Labs, such as Fermilab, are well-situated to utilize AI because of their experience with the construction and operation of large-scale data acquisition and distributed control systems. In addition, all necessary features for such interactions are present within their experimental facilities (the detection apparatus, as well as the equipment and buildings), which will benefit immensely from the inclusion of these new AI at-the-edge capabilities. To utilize these new technologies, new techniques and tools will be necessary to handle increases in data rates. The larger computing facilities where sensor data is collected and processed will also require real-time analysis to form results and produce accurate reactions. For a large science experiment, this means ensuring that data is understood and correct. For a science collaboration, this means incorporation of newly collected data to form scientific results. National Labs are in an excellent position for contributions because of their knowledge of “big data” science and algorithms for data analysis that can benefit AI.

Specialized compute hardware — such as FPGAs, ASICs, and systems-on-a-chip — comprise the base layers of the systems described above. These components can (and will) incorporate AI inference to particle accelerator processing, performing the time-critical functions — doing so at low power consumption rates and near to the collection and data transfer points. In the future, the higher layers of these edge-enabled systems will contain GPUs and specialized AI chips, such as TPUs [32] and Neutral Engines [33]. FPGAs are being worked into the higher-level components because of their ability to be dynamically reconfigured and handle large degrees of parallelism.

Because of the experience developed over many decades with reconfigurable and special-purpose hardware, Fermilab and other national labs are ideally suited to contribute to the construction of software (and firmware) algorithms to increase their utility. An example of the need for very fast real-time data acquisition and handling is the LHC trigger systems at the front of the detectors, which must react in hundreds of nanoseconds. Continuous read-out detectors, such as those present in neutrino experiments, need data acquisition completed in milliseconds. Many applications of future AI systems will involve analysis of real-time image data collected by cameras or other detectors, and in many cases these data may involve time evolution (“movies”) requiring processing at a scale far beyond the capabilities of humans. The 50-picosecond time resolution of the dedicated cameras in modern smartphones, for example, enables robust and reliable face recognition through 3-D true depth imaging. High-resolution and -depth-sensing cameras may allow future AI-enabled systems to perform complicated tasks exceeding that of human capabilities, and can greatly simplify the computational complexity to improve the reliability of AI systems. Construction of large cameras with high spatial resolution and picosecond time resolution is well within current technical capabilities. The biggest challenge is in developing the integrated electronics systems to interface such cameras to AI.

Exploration into using industry tools for AI have already begun for some data analysis problems to address real-time aspects of data processing. Google worked with the NASA to determine if machine learning could accelerate the search for exoplanets. This project, using Google’s AutoML service and cloud TPUs, reduced single retrieval times from several days at 94 percent accuracy down to seconds, while increasing accuracy to 96 percent [34]. These efficiency gains will allow NASA to now convert a serialized exploration into one that can run hundreds of retrievals simultaneously on the cloud. LSST is considering using Jupyterlab, coupled with Kubernetes, to aid in real-time interactive analyses [35].

## 2.4 Quantum Technology

Investing in new hardware platforms and paradigms is an important part of long-term support for AI. In particular, investments in quantum information science (QIS) will play an important role in the long-term impact of AI. Quantum technology is a key component of a comprehensive response to the post-Moore’s law era of computing [36]. Furthermore, quantum sensors will enable entirely new programs for fundamental science. It is critical to assume and maintain leadership in this technology and its applications.

Fermilab is seeking to leverage AI to help advance quantum science. We are poised to become a leader in QIS over the next five years. AI algorithms provide a powerful and efficient mechanism for tuning complex quantum systems in regimes where very fast times to solution are critical [37]. The current state of the art for qubit coherence is tens of microseconds, so the time available to work with these systems is very short.

Machine learning can help control and optimize a quantum computer. The applications of quantum information science rely on optimal control of quantum devices to realize fast and high-fidelity state preparation, gate operations and readout [38]. Designing these optimal control protocols involves the optimization of highly complex systems and hence is extremely challenging. Artificial intelligence is a powerful tool for the protocol design process to achieve performance beyond human capacity.

In terms of AI algorithms themselves, classical computers will dominate AI in the near and medium term, but over the longer time horizon of five to ten years from now, quantum advantages will greatly expand the reach and power of AI for certain problems. Recent algorithmic advances in quantum machine learning, e.g. [39], and successful proof-of-principle algorithmic tests, e.g. [40], have shown that quantum machine learning will eventually be able to leverage neural network connectivity beyond what is available to classical computers. It is critical to establish leadership in a technology that may come to dominate AI applications.

We need to develop a quantum information- and AI-aware workforce. The convergence of quantum science and AI at Fermilab make this an ideal training ground for integrating these technologies with difficult, real-world applications supported by the right combination of expertise and infrastructure.

## 2.5 Workforce Development

Workforce development is one of the most crucial challenges posed by AI technologies. AI has the potential to be a powerful engine of job creation and economic growth, as well as to improve the efficiency and safety for careers across almost all job sectors, including scientists. Properly educating and training America's AI-enabled workforce will require academia, industry, and government working in tandem as an effective team.

We need to train not only the next generation of professional scientists to use these tools for research, but all citizens so they can become familiar and adept in the application of these tools in their careers. Google has established programs in this area [41, 42], and a number of universities have established new majors [43] or institutes of research and education [44]. National laboratories have significant experience in developing workshops and schools for training scientists on the job, for example well-established theory and accelerators schools [45]. National labs can play a key role in that effort as a natural collaborative facilitator and nexus for these groups to partner on education and training ventures. Fermilab in particular, with its open access policy and partnerships with the tech industry, is well positioned to facilitate such activities.

## 3 Organizational and Assessment Considerations

Organization and assessments are important components of initiatives and projects in general. Evaluating the degree to which Federal investments are successful is crucial for continuing support, as well as defining future R&D directions for these programs. For the development of AI algorithms and applications for HEP, it is important to take into account that the field is in its early days. AI has the potential for transformative impact on discovery science, and science is poised to provide insight and unique data sets for AI algorithm development. It will be crucial to maintain US leadership in the development and deployment of AI, and that will require bringing all of our technical resources into proper alignment and fully utilizing the individual strengths of each field.

The broad applicability, uniqueness of scientific data, and need for algorithm theory development all necessitate partnerships between industry, national laboratories and universities that have such capabilities. Given that algorithm experts are spread out among labs, academia and industry, a mix of such institutions is needed to carry out AI algorithm R&D. The national laboratories are an ideal platform for cooperation between academia, industry, and government. HEP has much to offer to this enterprise in terms of large and richly detailed public data sets, a deep understanding of how to use simulation in real-world systems, cutting edge requirements for low-latency applications, and an important role in the educational and workforce development infrastructure of the country. National laboratories feature the organization and infrastructure to manage large projects and have strong network effects throughout science in terms of drawing researchers together and serving as a locus for a wide variety of efforts. Given the broad, multi-disciplinary approach needed to successfully make progress in AI, it is natural to consider them as hubs for such investment.

DOE national laboratories are also well-situated to play a role as clearinghouses for data and algorithms. In particular, HEP science requires the generation, storage, and curation of massive, complex data sets, and laboratories

such as Fermilab have cross-cutting resources, programs, and infrastructure that could be leveraged to share these data sets with the broader research community.

National laboratory campuses are designed to provide facilities for experimentation and computing. These facilities could be coupled to leadership-class compute facilities which provide massively large-scale GPU-enabled machines with thousands of nodes for production-scale work. We envision a powerful ecosystem where utilization of small-scale resources (laboratories, academia) is coordinated with utilization of large-scale resources of similar systems (laboratories with leadership-class facilities). This would allow personnel at laboratories and nearby institutional partners to access both the hardware and software environments of this computing infrastructure, giving them time to experiment and prepare for deployment of algorithms on the production-scale machines. This facilitates forward-thinking about hardware capabilities and the opportunity to get familiar with the devices for when they are deployed more broadly. This also provides complementary access and spreads out expertise, experience, and insight for next-generation devices. Yet broader access to these AI-enabled computers may be facilitated with cloud infrastructure that is built to present a common portal for submission and management of workflows to the heterogeneous computing resources. Cloud capabilities will be especially important for fast development and testing of algorithms, and for unimpeded access to compute resources, which we believe will be essential for HEP science. Fermilab has already developed and deployed infrastructure that allows access to cloud, leadership-class, and on-campus resources through a common interface (e.g., HEPCloud [46]).

To summarize, the frontier of scientific discovery is algorithmic in nature, and national labs can provide key resources and expertise as a nexus for research, collaboration, innovation, and workforce development to realize this future.

## References

- [1] D. Guest, K. Cranmer, and D. Whiteson, “Deep Learning and its Application to LHC Physics,” 2018.
- [2] 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,” *JINST*, vol. 11, no. 09, p. P09001, 2016.
- [3] M. Paganini, L. de Oliveira, and B. Nachman, “Calogan: Simulating 3d high energy particle showers in multilayer electromagnetic calorimeters with generative adversarial networks,” *Phys. Rev. D*, vol. 97, p. 014021, Jan 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.97.014021>
- [4] 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.
- [5] P. Adamson *et al.*, “Constraints on oscillation parameters from  $\nu_e$  appearance and  $\nu_\mu$  disappearance in nova,” *Phys. Rev. Lett.*, vol. 118, p. 231801, Jun 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.118.231801>
- [6] M. A. Acero *et al.*, “New constraints on oscillation parameters from  $\nu_e$  appearance and  $\nu_\mu$  disappearance in the nova experiment,” *Phys. Rev. D*, vol. 98, p. 032012, Aug 2018. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.98.032012>
- [7] A. Mathuriya, D. Bard *et al.*, “CosmoFlow: Using Deep Learning to Learn the Universe at Scale,” *ArXiv e-prints*, Aug. 2018.
- [8] R. Kondor, Z. Lin, and S. Trivedi, “Clebsch-Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network,” *ArXiv e-prints*, Jun. 2018.
- [9] K. N. Abazajian *et al.*, “CMB-S4 Science Book, First Edition,” *ArXiv e-prints*, Oct. 2016.
- [10] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and *et al.*, “LSST Science Book, Version 2.0,” *ArXiv e-prints*, Dec. 2009.
- [11] A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, “Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons,” *Physical Review Letters*, vol. 104, no. 13, p. 136403, Apr. 2010.
- [12] A. P. Bartók, R. Kondor, and G. Csányi, “On representing chemical environments,” *Phys. Rev. B*, vol. 87, p. 184115, May 2013. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.87.184115>
- [13] Y.-C. Hsu and Z. Kira, “Neural network-based clustering using pairwise constraints,” *ArXiv e-prints*, Nov. 2015.
- [14] M. R. Beck, C. Scarlata, L. F. Fortson, C. J. Lintott, B. D. Simmons, M. A. Galloway, K. W. Willett, H. Dickinson, K. L. Masters, P. J. Marshall, and D. Wright, “Integrating human and machine intelligence in galaxy morphology classification tasks,” *MNRAS*, vol. 476, pp. 5516–5534, Jun. 2018.
- [15] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.
- [16] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” in *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS’12. USA: Curran Associates Inc., 2012, pp. 1097–1105, <http://dl.acm.org/citation.cfm?id=2999134.2999257>.
- [17] Y. LeCun and C. Cortes, “MNIST handwritten digit database,” 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [18] COCO: Common Objects in Context. (2018). [Online]. Available: <http://cocodataset.org/#home>
- [19] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” 2016.



- [20] M. Abadi *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [21] F. Chollet *et al.*, “Keras,” <https://keras.io>, 2015.
- [22] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in pytorch,” in *NIPS-W*, 2017.
- [23] Wired. (2018) Self-Driving Cars: The Complete Guide. [Online]. Available: <https://www.wired.com/story/guide-self-driving-cars/>
- [24] DeepMind. (2018) Safety-first AI for autonomous data centre cooling and industrial control. [Online]. Available: <https://deepmind.com/blog/safety-first-ai-autonomous-data-centre-cooling-and-industrial-control/>
- [25] J. Gao. Machine Learning Applications for Data Center Optimization. [Online]. Available: <https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/42542.pdf>
- [26] W. Knight. (2018) Google just gave control over data center cooling to an AI. [Online]. Available: <https://www.technologyreview.com/s/611902/google-just-gave-control-over-data-center-cooling-to-an-ai/>
- [27] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. London, England: The MIT Press, 2018.
- [28] S. S. Mousavi, M. Schukat, and E. Howley, “Deep reinforcement learning: An overview,” *CoRR*, vol. abs/1806.08894, 2018.
- [29] 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): Busan, Korea, May 8-13, 2016*, 2016, p. THPOY020. [Online]. Available: <http://lss.fnal.gov/archive/2016/conf/fermilab-conf-16-603-ad.pdf>
- [30] K. Shao, Y. Zhu, and D. Zhao, “Starcraft micromanagement with reinforcement learning and curriculum transfer learning,” *CoRR*, vol. abs/1804.00810, 2018. [Online]. Available: <http://arxiv.org/abs/1804.00810>
- [31] A. Scheinker and D. Scheinker, “Constrained extremum seeking stabilization of systems not affine in control,” *Int J Robust Nonlinear Control*, vol. 28, pp. 568–581, Jul 2017. [Online]. Available: <https://doi.org/10.1002/rnc.3886>
- [32] Google. (2017) An in-depth look at Google’s first Tensor Processing Unit (TPU). [Online]. Available: <https://cloud.google.com/blog/products/gcp/an-in-depth-look-at-googles-first-tensor-processing-unit-tpu>
- [33] Wired. (2017) Apple’s ‘Neural Engine’ Infuses the iPhone With AI Smarts. [Online]. Available: <https://www.wired.com/story/apples-neural-engine-infuses-the-iphone-with-ai-smarts/>
- [34] Google. (2017) To find intelligent life beyond earth, NASA FDL turns to artificial intelligence. [Online]. Available: <https://cloud.google.com/nasa-fdl/>
- [35] LSST. (2018) SQR-018: Investigations into JupyterLab as a basis for the LSST Science Platform. [Online]. Available: <https://sqr-018.lsst.io/>
- [36] A. Aspuru-Guzik *et al.* (2015-02) ASCR report on quantum computing for science. [Online]. Available: <https://science.energy.gov/~media/ascr/pdf/programdocuments/docs/ASCRQuantumReport-final.pdf>
- [37] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “Implementing a universal gate set on a logical qubit encoded in an oscillator,” *Nature Communications*, vol. 8, no. 1, p. 94, 2017. [Online]. Available: <https://doi.org/10.1038/s41467-017-00045-1>
- [38] J. Zhang, Y. xi Liu, R.-B. Wu, K. Jacobs, and F. Nori, “Quantum feedback: Theory, experiments, and applications,” *Physics Reports*, vol. 679, pp. 1 – 60, 2017, quantum feedback: theory, experiments, and applications. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0370157317300479>

- [39] G. Verdon, J. Pye, and M. Broughton, “A Universal Training Algorithm for Quantum Deep Learning,” *ArXiv e-prints*, Jun. 2018.
- [40] A. Mott, J. Job, J.-R. Vlimant, D. Lidar, and M. Spiropulu, “Solving a Higgs optimization problem with quantum annealing for machine learning,” *Nature*, vol. 550, pp. 375 EP –, Oct 2017. [Online]. Available: <http://dx.doi.org/10.1038/nature24047>
- [41] Google. (2018) Learn with Google AI: Making ML education available to everyone. [Online]. Available: <https://www.blog.google/technology/ai/learn-google-ai-making-ml-education-available-everyone/>
- [42] ——. (2018) Learn from ML experts at Google. [Online]. Available: <https://ai.google/education/>
- [43] Forbes. (2018) Carnegie Mellon Establishes The First Undergrad Degree In AI. [Online]. Available: <https://www.forbes.com/sites/peterhigh/2018/05/21/carnegie-mellon-establishes-the-first-undergrad-degree-in-ai/#3597f42c8caf>
- [44] T. N. Y. Times. (2018) M.I.T. Plans College for Artificial Intelligence, Backed by \$1 Billion . [Online]. Available: <https://www.nytimes.com/2018/10/15/technology/mit-college-artificial-intelligence.html>
- [45] U. P. A. School. (2018) U.S. Particle Accelerator School: Education in Beam Physics and Accelerator Technology. [Online]. Available: <http://uspas.fnal.gov/>
- [46] B. Holzman *et al.*, “HEPCloud, a New Paradigm for HEP Facilities: CMS Amazon Web Services Investigation,” *ArXiv e-prints*, Sep. 2017.