Artificial Intelligence for Nuclear Physics

Tanja Horn, report to NSAC

THE CATHOLIC UNIVERSITY OF AMERICA

Jefferson Lab
Thomas Jefferson National Accelerator Facility

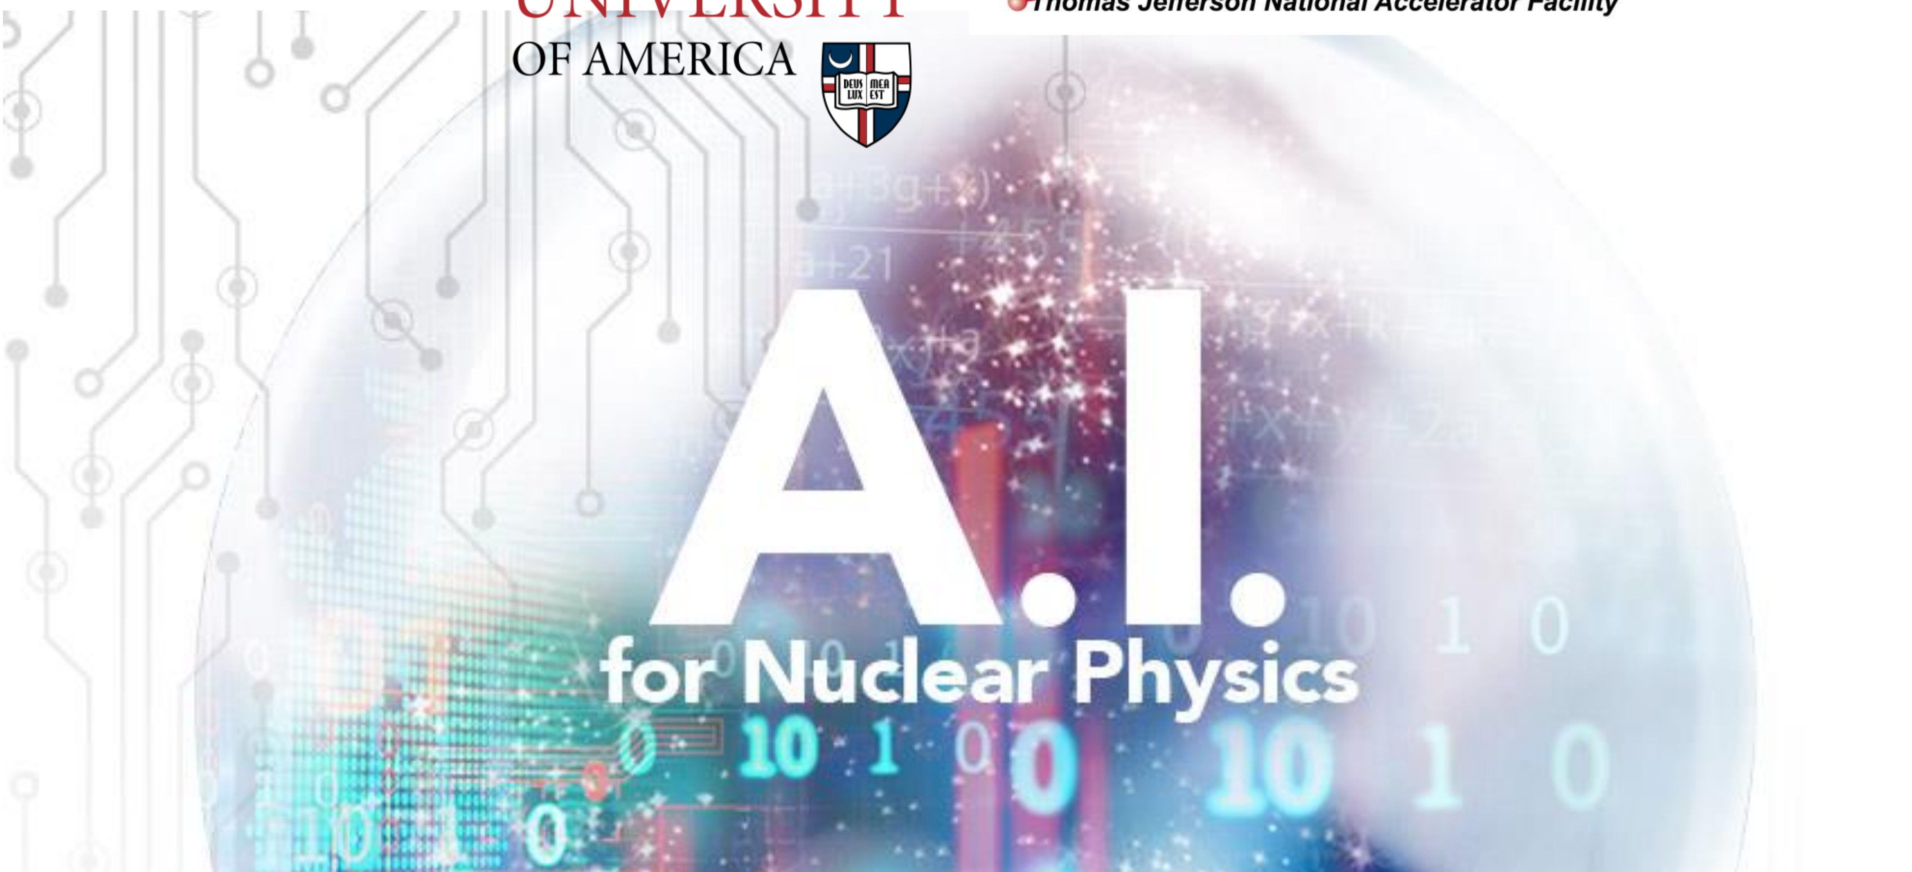

2020 ASCAC Subcommittee on: 'AI/ML, Data-intensive Science and High-Performance Computing'
Use shorthand: Subcommittee on 'AI for Science'

The Charge Letter:

❑The letter sets the context of the challenge to the subcommittee:

- Artificial Intelligence and Machine Learning (AI/ML) have the potential for providing new insights and even new discoveries from this data, including the correlation of experimental and computational data.
- However, the technical aspects of "AI/ML for Science" may be more challenging than currently envisioned. Over the last few years, several workshops and subcommittee reports have identified and enumerated the scientific opportunities and some challenges from the intersection of AI/ML with data-intensive science and high performance computing.

❑The letter requires the sub-committee deliver a report that:

- Assesses the opportunities and challenges from Artificial Intelligence and Machine Learning for the advancement of science, technology, and Office of Science missions.
- Identifies strategies that ASCR can use, in coordination with the other SC programs, to address the challenges and deliver on the opportunities.
- Notes that, due to the cross-cutting nature of this effort, in assembling this subcommittee, we need to include members of, and recommendations from the other Office of Science Federal Advisory Committees, as well as Industry and other Federal experts.

Membership of the AI4Science Subcommittee

ASCAC Members – Tony Hey, Jack Dongarra

AITO – Fred Streitz

HPC, AI and Big Data academic members:

- Geoffrey Fox (Indiana)
- Ewa Deelman (USC)
- Joel Saltz (Stony Brook)
- Dan Stanzione (TACC)
- Rebecca Willett (Chicago)

Representatives from other Office of Science Advisory Committees:

- HEPAP – Mike Hildreth (Notre Dame)
- BERAC – Kerstin Kleese van Dam (BNL)
- BESAC – Abbas Ourmazd (Wisconsin)
- **NSAC – Tanja Horn (The Catholic University of America)**
- FESAC – Phil Snyder (General Atomics)

NSB – Anneila Sargent (Caltech)

Summary of the Subcommittee's activities

- February 2020: One 2-day briefing meeting in Washington DC
- March/April 2020: Remaining briefing meetings via Zoom
- April/May 2020: Subcommittee discussion and report writing task allocation
- July 2020: First draft of report produced in June and circulated to all SC Advisory Committees asking for comments, corrections and suggestions for improvement
- August 2020: Subcommittee email discussion of suggestions and consequent reworking of several sections of the draft report
- 1 September 2020: Subcommittee discussion and agreement on final draft of the report
- 24-25 September 20: Presentation to ASCAC and approval of the report

The AI for Science Report - Structure

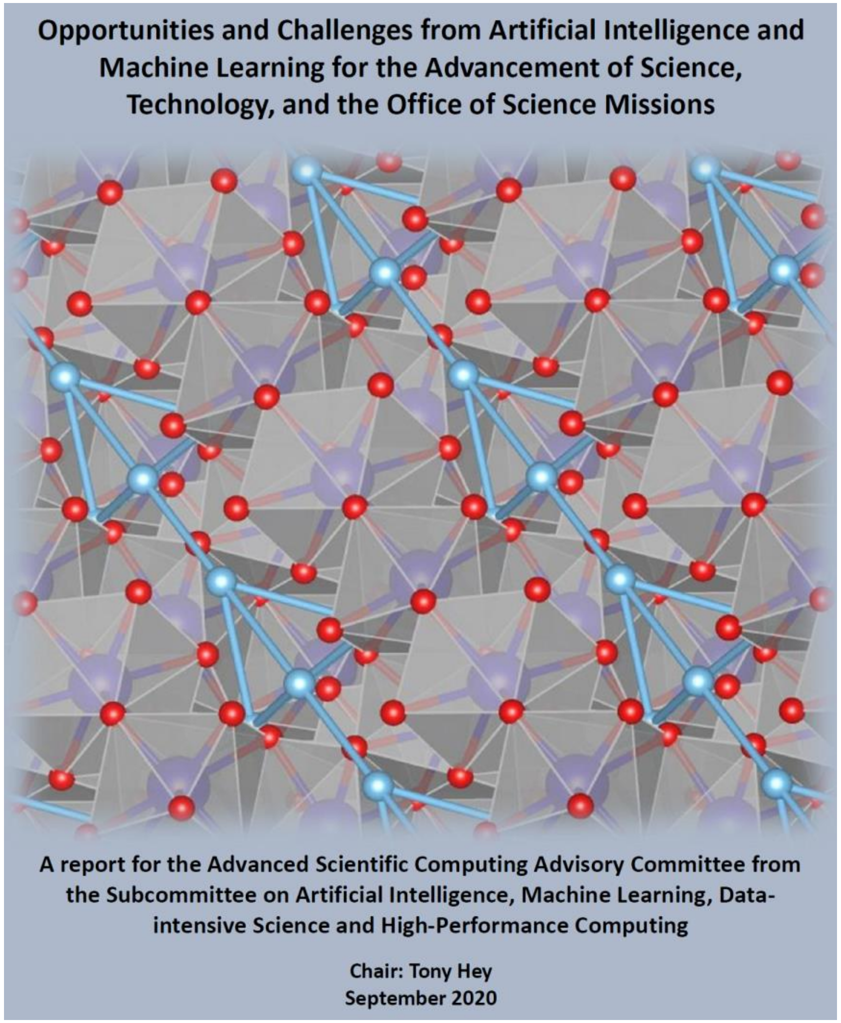

https://science.osti.gov/ascr/ascac/Meetings/202009

Executive Summary

Introduction

Context

Key Findings

Recommendations

Report

1. Introduction and Background
2. Charge Letter to ASCR
3. Subcommittee Information Gathering Activities
4. DOE as lead agency for AI/ML applied to Facilities Science
5. Opportunities and challenges from AI/ML for the advancement of science, technology, and the Office of Science missions
6. Strategies for the DOE Office of Science to address the challenges and deliver on the opportunities.
7. Summary of conclusions

1. Introduction and Background

The National Artificial Intelligence R&D Strategic Plan

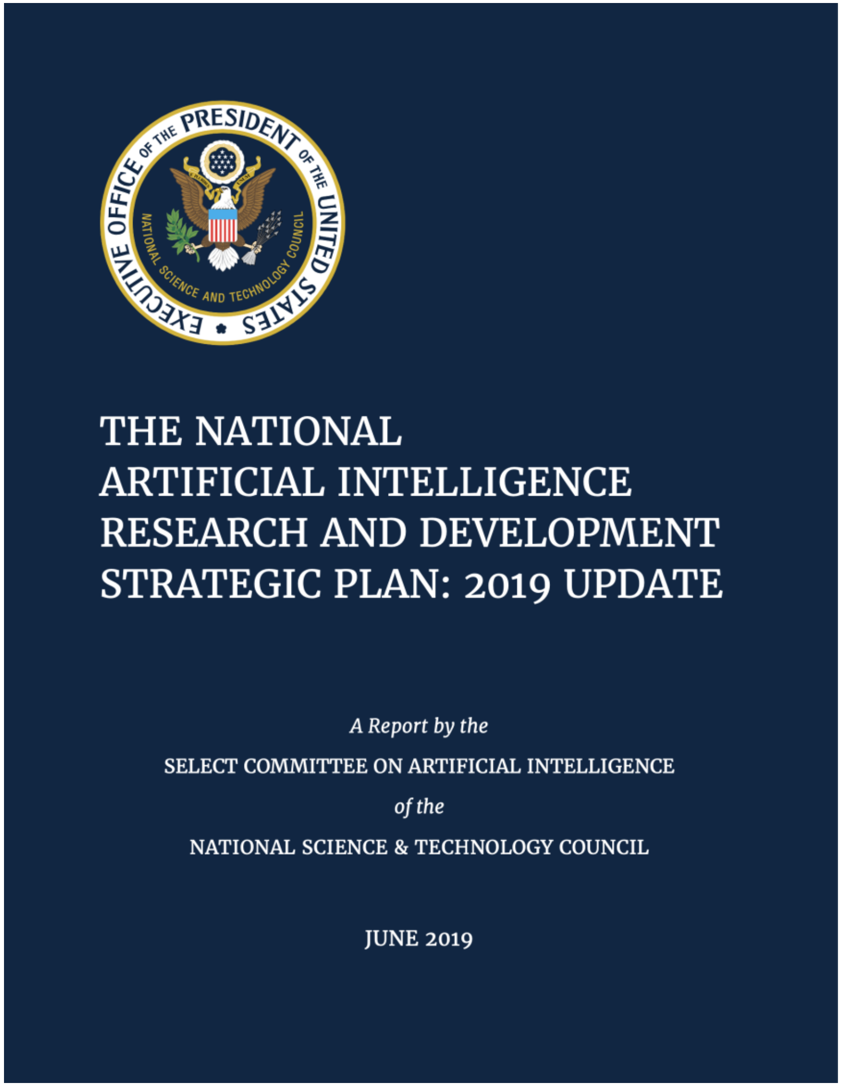
THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN: 2019 UPDATE

A Report by the

SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE

of the

NATIONAL SCIENCE & TECHNOLOGY COUNCIL

JUNE 2019

On February 11, 2019, the President signed Executive Order 13859, *Maintaining American Leadership in Artificial Intelligence*.[1]

This order launched the American AI Initiative, a **concerted effort to promote and protect AI technology and innovation in the United State**s.

The Initiative implements a **whole-of-government strategy** in collaboration and engagement with the private sector, academia, the public, and **like-minded international partners**.

Among other actions, key directives in the Initiative **call for Federal agencies to prioritize AI research and development (R&D) investments**, enhance access to high-quality cyberinfrastructure and data, ensure that the Nation leads in the development of technical standards for AI, and provide education and training opportunities to prepare the American workforce for the new era of AI.

AI for Science – What's Next After Exascale

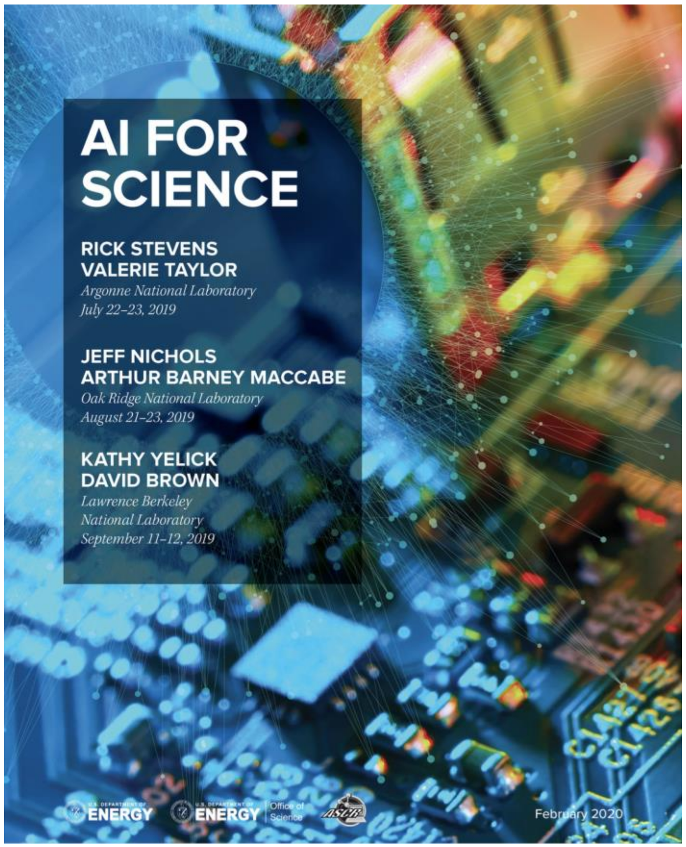

https://www.anl.gov/ai-for-science-report

- Over 1,000 scientists participated in four town halls during the summer of 2019
- Research Opportunities in AI
  - Biology, Chemistry, Materials,
  - Climate, Physics, Energy, Cosmology
  - Mathematics and Foundations
  - Data Life Cycle
  - Software Infrastructure
  - Hardware for AI
  - Integration with Scientific Facilities
- Modeled after the Exascale Series in 2007

5. Opportunities and Challenges from AI and ML for Office of Science Missions

Advances expected from use of AI/ML in Science

- Accelerate the design, discovery, and evaluation of new materials
- Advance the development of new hardware and software systems, instruments and simulation data streams
- Identify new science and theories revealed as a result of increasingly high-bandwidth instrument data streams
- Improve experiments by inserting inference capabilities in control and analysis loops
- Enable the design, evaluation, autonomous operation, and optimization of complex systems from light sources and accelerators to instrumented detectors and HPC data centers
- Advance the development of self-driving laboratories and scientific workflows
- Increase the capabilities of exascale and future supercomputers by capitalizing on AI surrogates
- Automate the large-scale creation of “FAIR” (Findable, Accessible, Interoperable, Reusable) data

6. Strategies for the Office of Science to deliver on opportunities of *AI for Science*

Four components needed for successful AI in Science Initiative

- Application-specific solutions based on hardware/software/algorithm co-design

- Research in AI algorithms and foundations

- Development of AI software infrastructure

- AI-specific computing architectures and hardware

Successful integration of these four components will require

- A full partnership between all areas of the Office of Science
- Engagement of the National Laboratories and their user facilities
- Involvement of the university and private industry research community
- Mechanisms for collaborative projects with agencies such as the NSF, NIH, NIST and DOD
- Collaboration with expert organizations from similarly minded countries
- An organized process for dissemination to the scientific community

6.2 AI Applications

Priority research directions identified for each of the five Office of Science R&D programs:

- Basic Energy Sciences
- Biological and Environmental Sciences
  - Climate and Environmental Sciences
  - Biological Systems Science
- Fusion Energy Sciences
- High Energy Physics
- **Nuclear Physics**

Note that several generic priorities in these submissions have not been separated from those specific to each research area

Unique Features of Nuclear Physics

❑ Multi-scale, highly correlated, and high dimensionality nature of NP

- Approximate symmetries found in NP are thought to have origins in the underlying **strong interaction** and also in the **many-body physics** of the problems
- Data-rich: billions of events with processes that are not yet first-principles calculable
- Surrogate models, e.g., for computation-intensive theory

❑ Diversity of data sets in NP

- Experiments **range from small- and intermediate scale to very large detector programs**
- Nuclear theory is concerned with how quarks and gluons interact to form hadrons and how these interact to form atomic nuclei

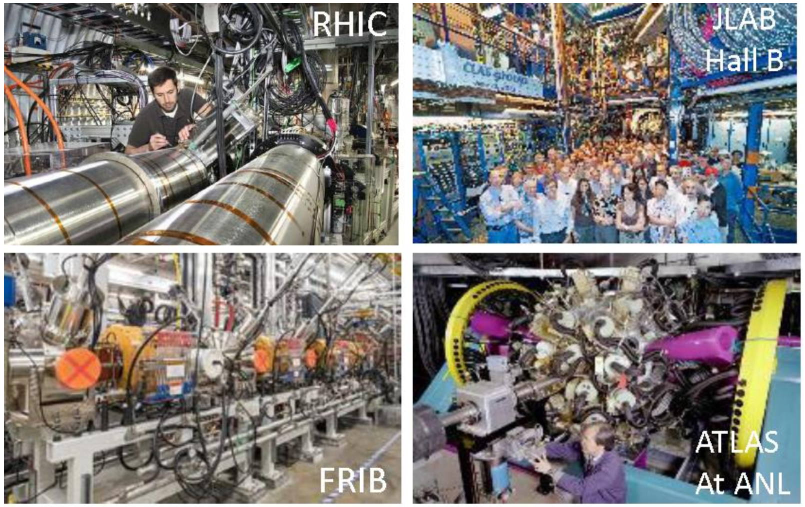

❑ Availability of NP data on short timescales

- **Transfer of knowledge** to other disciplines and sectors are themes expected to benefit society at large and that will play central roles.
- NP has the opportunity to build in cutting edge tools from the start and as **experiments are relatively short and have many different configurations**

AI in Nuclear Physics – Grand Challenges

- **Harness the physics program of the Electron-Ion Collider (EIC)**
  - AI/ML will help guarantee maximum science output from the EIC
- **Realize the science potential of FRIB**
  - A variety of AI/ML tools will be developed to address specific needs including beam generation, event characterization, detector response, experiment optimization and data analysis
- **Event Reconstruction in Nuclear Physics**
  - AI techniques for reconstruction of tracks in time projection chambers at FRIB, and for heavy ion collisions
- **Improve Tracking Algorithms**
  - AI/ML to significantly improve tracking at all NP accelerator facilities
- **Particle Identification**
  - AI/ML to complement existing Monte Carlo methods for PID
  - Gamma-Ray Energy Tracking Array (GRETA): AI/ML to reconstruct the path of multiple gamma rays from measured interaction positions and deposited energies

Conclusions and Executive Summary

Context

Artificial Intelligence (AI)
How can computational technologies be developed & used to assist, augment or automate human skills?

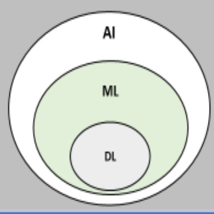

Machine Learning (ML) - Foundational basis for AI.

ML algorithms make predictions, decisions, & estimates from building a mathematical model or "learning" based on probabilities, samples, or training data.

Deep Learning (DL) - Includes neural network-trained approaches for tasks such as spam filtering, fraud/anomaly detection, image analysis.

Cognitive Skills (AI)	Specific Tasks (ML)	Neural Networks (DL)
Vision & Perception Natural Language Processing Search & Planning Problem solving Knowledge reasoning	Classification Clustering & regression Simplified or surrogate models Feature extraction Pattern recognition	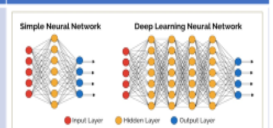

Workshop Report on Basic Research Needs for Scientific Machine Learning: Core Technologies for Artificial Intelligence (DOE/ASCR, 2019)

Data Engineer 	**People who are expert at** • Operating at low levels close to the data, write code that manipulates • They may have some machine learning background. • Large companies may have teams of them in-house or they may look to third party specialists to do the work.
Data Analyst 	**People who explore data through statistical and analytical methods** • They may know programming; May be an spreadsheet wizard. • Either way, they can build models based on low-level data. • They eat and drink numbers; They know which questions to ask of the data. Every company will have lots of these.
Data Steward 	**People who think to managing, curating, and preserving data.** • They are information specialists, archivists, librarians and compliance officers. • This is an important role: if data has value, you want someone to manage it, make it discoverable, look after it and make sure it remains usable.

What is a data scientist? Microsoft UK Enterprise Insights Blog, Kenji Takeda
http://blogs.msdn.com/b/microsoftenterpriseinsight/archive/2013/01/31/what-is-a-data-scientist.aspx

Key Findings (1)

- **Finding A**

The growing convergence of AI, Data, and HPC provides a once in a generation opportunity to profoundly accelerate scientific discovery, create synergies across scientific areas, and improve international competitiveness.

- **Finding B**

Science can greatly benefit from AI methods and tools. However, commercial solutions and existing algorithms are not sufficient to address the needs of science automation and science knowledge extraction from current and future DOE facilities and data.

- **Finding C**

Adopting ***AI for Science*** technologies throughout the Office of Science will enable US scientists to take advantage of the tremendous new advances in the DOE's scientific user facilities.

- **Finding D**

Realizing the potential for a generational shift in scientific experimentation at the DOE Laboratories due to science-driven AI/ML technologies requires far more than simply compute power and encompasses the full spectrum of computing infrastructures, ranging from ubiquitous sensors and interconnectivity across devices to real-time monitoring and data analytics, and will require a concerted and coordinated R&D effort on AI/ML algorithms, tools, and software infrastructure.

Key Findings (2)

- **Finding E**

The DOE Labs are uniquely positioned to integrate AI/ML technologies across a host of scientific challenges thanks to the enviable culture of co-design teams consisting of scientific users, instrument providers, theoretical scientists, mathematicians and computer scientists that has proven so successful in the Exascale Computing Project.

- **Finding F**

The impact of a DOE-driven AI/ML strategy for science will have national implications far beyond the Office of Science and will drive new industrial investments, including accelerating engineering designs, synthesizing materials, and optimizing energy devices, as well as advancing hardware and software computing capabilities.

- **Finding G**

A workforce trained in advanced AI/ML technologies would play a pivotal role in enhancing US competitiveness.

- **Finding H**

Partnering with other Agencies and with international efforts will be important to deliver on the ambitious goals of an ***AI for Science*** initiative.

Recommendations (1)

- **#1 Creation of a 10-year AI for Science Initiative**

In order to create the world-leading AI systems and applications needed to drive scientific productivity and discovery in science and technology dramatically beyond that achievable with traditional scientific supercomputing, we recommend that the DOE Office of Science start a ten-year program to develop an ambitious AI for Science initiative, as recommended in the recent PCAST report [9].

- **#2 Structure of an SC AI for Science Initiative**

It is recommended that this AI for Science initiative be structured around four major AI R&D themes:

  - AI-enabled applications
  - AI algorithms and foundational research
  - AI software infrastructure
  - New hardware technologies for AI

Recommendations (2)

- **#3 An Instrument-to-Edge Initiative**

The subcommittee believes that ASCR, in close cooperation with BES and with the other science programs in the Office of Science, should work with scientists, users, and the broad academic community to define requirements, conduct research, competitive procurement and design a highly integrated end-to-end system and software stack that connects instruments at the edge to the needed AI computing resources. Integrating national and global data sources (large scale experimental facilities, observational networks terrestrial & space-based, etc.) poses unique opportunities and challenges that require addressing foundational research in the context of leading-edge scientific experiments

- **#4 Training, focusing, and retention of AI/ML workforce**

A strong research program will crucially rely on a complementary education and skills component, which is as important as providing adequate infrastructure support. As emphasized in the ASCR ECP Transition report [10], this is also a timely and important opportunity to focus SC efforts to create a more diverse and inclusive workforce. A continuing supply of high-quality computational and data scientists available for work at DOE laboratories is of vital importance.

Recommendations (3)

- **#5 Inter-Agency Collaboration**

The subcommittee, therefore, recommends that the SC explore new opportunities to work with both NSF and NIH in areas where there would be a clear benefit for scientific progress under a DOE-led AI for Science initiative. There may also be opportunities to work with other US funding agencies, such as NIST and DOD, in areas of mutual interest.

- **#6 International Collaboration**

There is a need for broad-based, coordinated action by like-minded international partners to harness the global scientific software community to address the tremendous opportunities in data-intensive science stemming from the huge increase in scientific data collection rates. To realize this vision for an international cooperative effort, the Office of Science needs to:

  - Provide a framework for organizing the software research community
  - Create a thorough assessment of needs, issues, and strategies
  - Initiate development of a coordinated software roadmap
  - Encourage and facilitate collaboration in education and training
  - Engage and coordinate the vendor community in cross-cutting efforts

Activities in Nuclear Physics: A.I. for NP Workshop

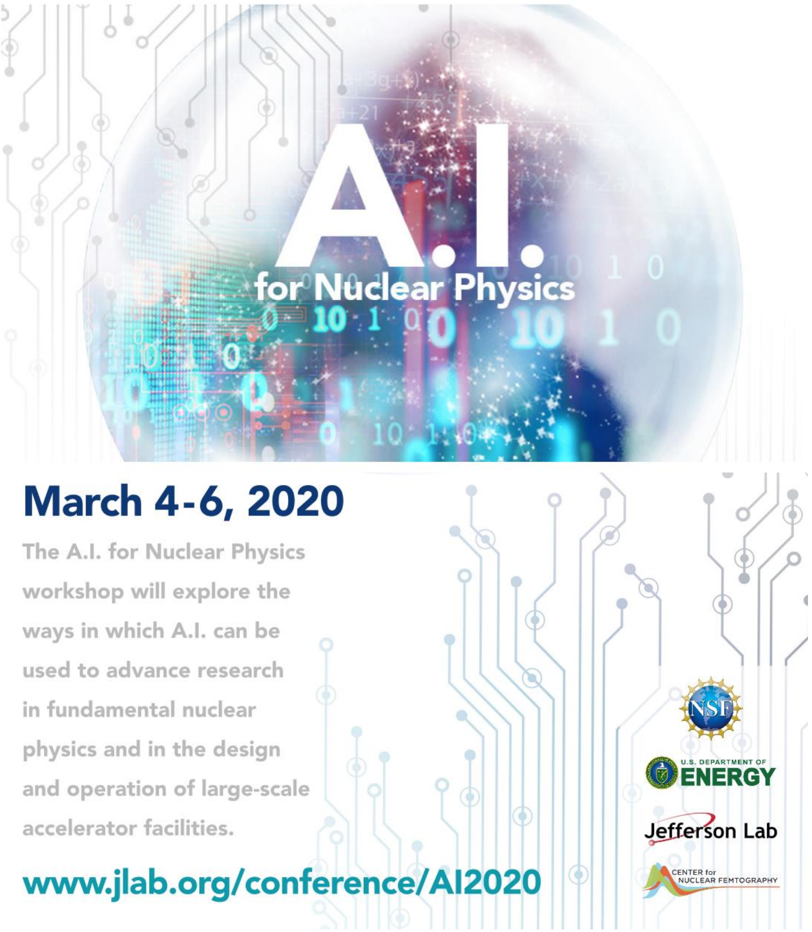

- ❑ 3 day workshop with 180 attendees with plenary breakout format
- ❑ 1 day pre-meeting 'hackathon'
- ❑ 8 Break out groups to address aspects of Theory, Experiment, and Accelerator
- ❑ The results were summarized in a report containing an assessment of ongoing efforts

arXiv:2006.05422v2 and EPJA in press

- Also summary of the breakout session and additional topical areas (relativistic heavy Ions, Project 8, Next and Wanda) not present at the workshop
- Identified need for workforce development and education and a need for cross disciplinary collaborations.

Tremendous interest in AI/ML in the Nuclear Physics Community

AI/ML in Nuclear Theory and Lattice QCD

- **Development of a spectroscopic-quality nuclear energy density functional**
  - Crucial for understanding rare isotopes
- **Discovering nucleonic correlations and emergent phenomena**
  - Discover correlations in calculations of nuclear wave functions that use underlying forces
- **Neutron star and dense matter equation of state**
  - Deduce nuclear matter equation of state from intermediate-energy heavy-ion collisions data
- **Ensemble generation in lattice QCD**
  - Scalability, compact variables, sign problem
- **Inverse problem: spectral function, Parton Distributions**
- **Phase transitions and estimators for correlation functions**

Nuclear Theory Examples

- **Example 1: Bayesian Model Averaging to Quantify Limits of the nuclear landscape**

  Constrained density functional theory calculations in multidimensional collective spaces with Bayesian model averaging

  L. Neufcourt et al., Phys. Rev. C 101 (2020) 044307

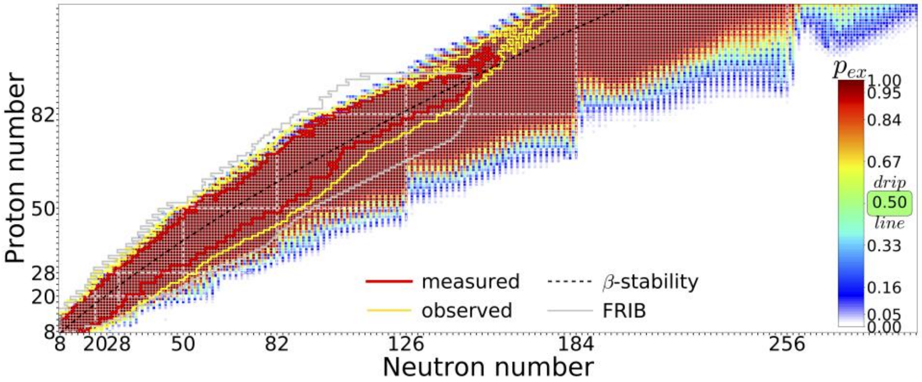

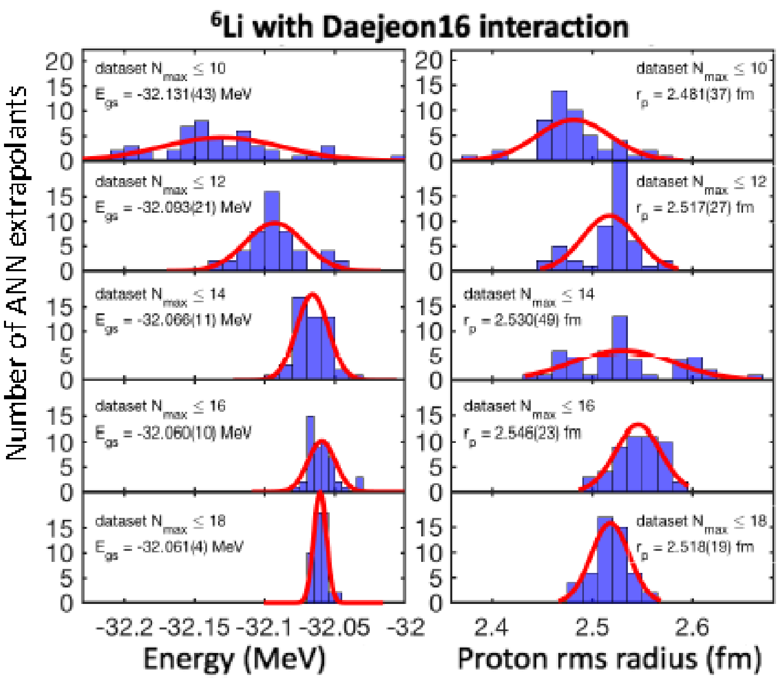

- **Example 2: Deep Learning for Nuclear Binding Energy and Radius**

  Demonstrated predictive power of ANNs for converged solutions of weakly converging simulations of the nuclear radius

  G. A. Negoita et al., Phys. Rev. C 99 (2019) 054308

Lattice QCD Examples

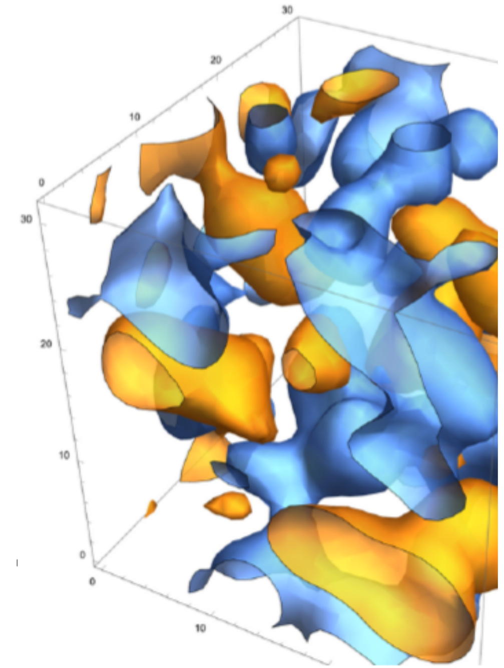

❑ Example 1: Ensemble generation

Multi-scale algorithms: parallels with image recognition

P. E.Shanahan et al., Phys. Rev. D 97 (2018) 094506

Generative models to replace hybrid Monte-Carlo: parallels with image generation

M.S. Albergo et al., Phys. Rev. D 100 (2019) no.3 034515

❑ Example 2: Speed up Hadron Correlator Computation

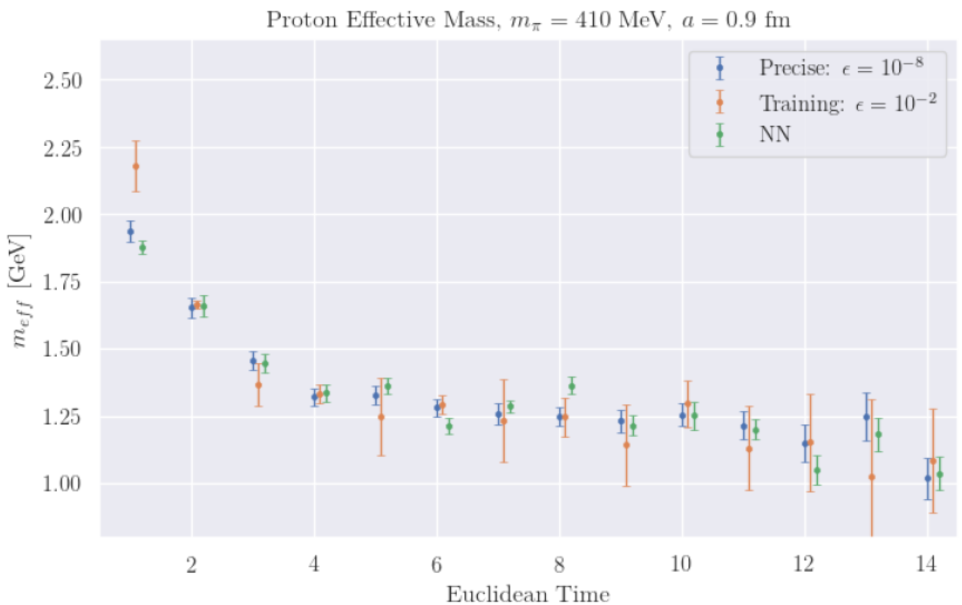

Boosted Decision Trees and ANNs to reduce the cost of iterative solvers for quark propagator by relating solutions to the system computed at different precision

G. Pederiva et al., “Machine Learning Algorithms for Hadron Correlators from Lattice QCD”, in preparation

Experimental Methods: AI to Detect Complex Signatures

- **Near Term: Improved Analysis**
  - Improved sensitivity
  - Faster Analysis → faster scientific output

- **Long Term:**
  - Holistic approach to experimentation
  - Standardized data formats
  - Experiment design not limited by computation

Experimental Methods Examples

❑ Example 1: Convolutional Neural Networks search for neutrino events

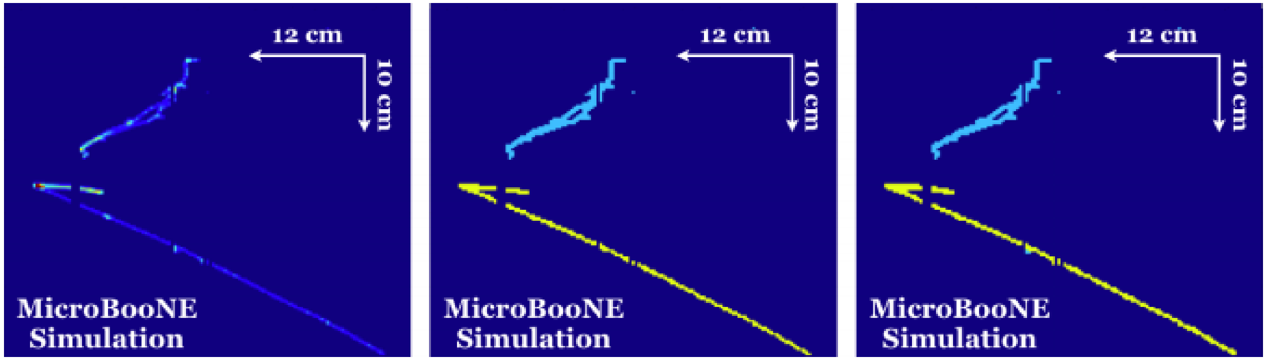

MicroBooNE Collaboration, Phys. Rev. D 99 (2019) no.9, 092001

CNNs capable of semantic segmentation (3D-Mask R-CNN) to sift through the thousands of interactions that occur every day and create 3-D images of the most interesting ones.

❑ Example 2: Boosted Decision Trees to Search for Exotic Mesons in GlueX

- Isolate events of interest from a disproportionately large background using boosted decision trees.
- Preliminary studies show that these algorithms have the potential to offer vast improvements in both signal selection efficiency and purity over more traditional techniques.

Experimental Methods Examples

- **Example 3: automated (ML driven) design of observables – looking for generic solutions/experiment independent**

Discovery of new observables by NN

Discovery of theoretical models via automated analysis

- Previous: Finding most sensitive observables to a model parameters → *Lai; 1810.00835*

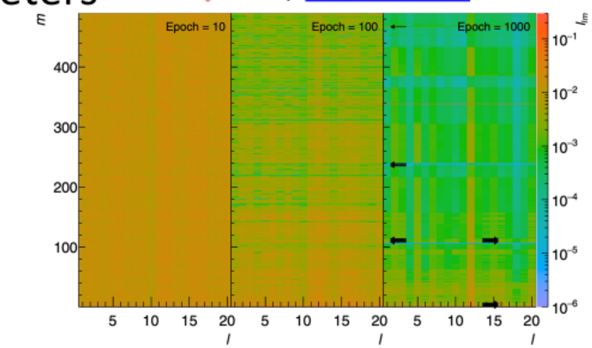

- Current focus (from a longer list): How much information is contained in high-energy particle collisions and jets?

Extract knowledge on complex processes (e.g. jet quenching) directly from data - human understandable result (!)
→ new guidance to experiment(s)
→ critical input for theory

Lai, Mulligan, Ploskon, Ringer - in preparation

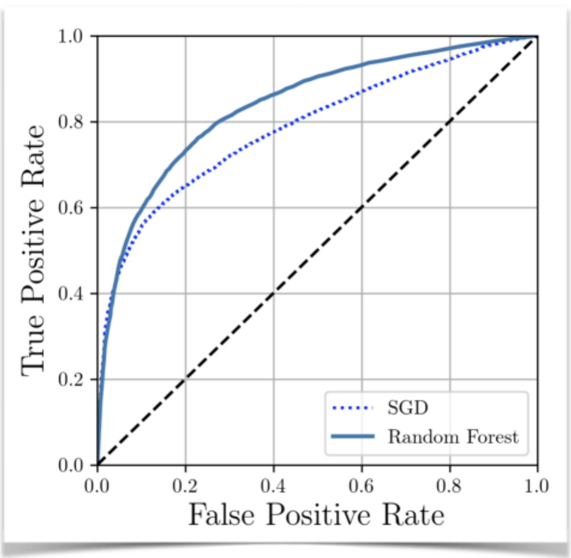

Next challenge?

→ Hadronization

- Long standing problem
- Impact in both NP and HEP
- Guidance for EIC experiments

Simulations: AI/ML for detector design and to extract physics

- **Accelerating Simulations**
  - Calorimeter, PID
- **AI driven detector design**
  - Detector Design Optimization for Electron-Ion Collider
- **HPC Utilization**
- **ML for event generators**

Example 1: Detector Design Optimization

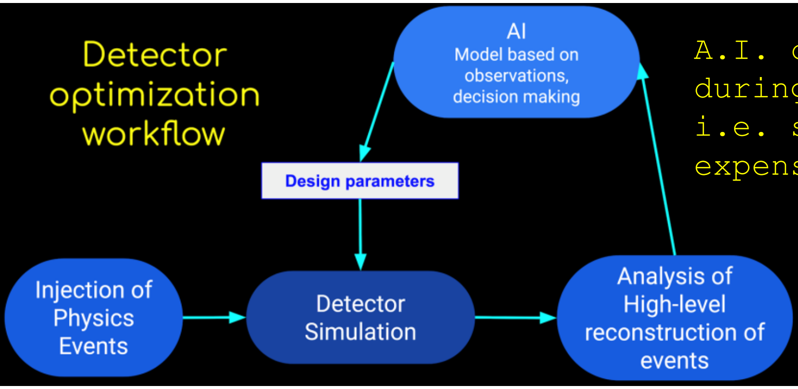

A.I. can optimize the design of complex during the R&D of large-scale detectors, i.e. simulating noisy and computationally expensive black-box functions

- **Bayesian Optimization (BO), Evolutionary Algorithms (EA), etc**
- **Multi-objective optimization (MOO) in multi-dimensional design space**

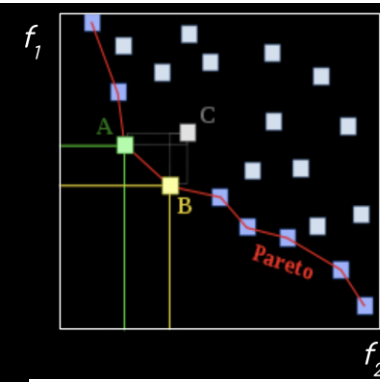

Dual-RICH @EIC: **First EIC paper using AI**, an automated, highly parallelized, self-consistent framework based on BO+ML to optimize the Geant simulation of the dual-RICH. O(10)parameters optimized.

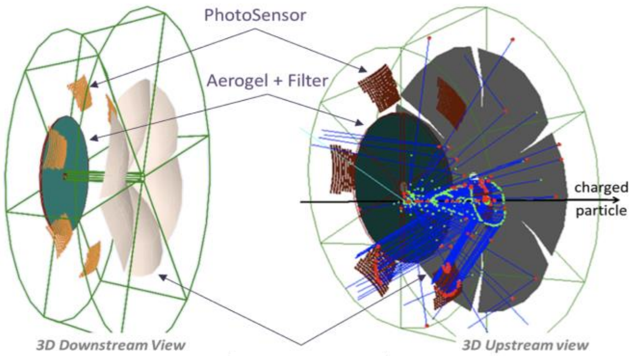

E. Cisbani, A. Del Dotto, C. Fanelli, M. Williams et al
2020 JINST 15 P05009

R&D of novel composite aerogel+fibers: Aerogels with low refractive index are very fragile. We are designing with the AI optimizing **mechanical strength** and **resolution** using evolutionary MOO. Geant4 + Autodesk (gmsh+elmer)

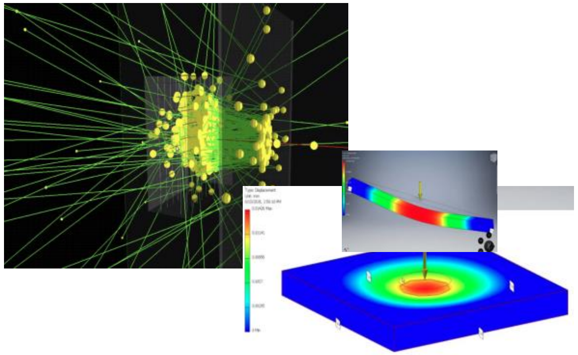

The team: V. Berdnikov, J. Crafts, E. Cisbani, C. Fanelli, T.H. , R.. Trotta

EIC Electron Endcap EM Calorimeter: Optimization of glass/crystal material selection in shared rapidity regions including mechanical constraints from EIC detector. MOO to make decision on EEemCAL resolution (how it affects physics of interest), and crystal/glass cost optimization. EIC SW: ESCalate.

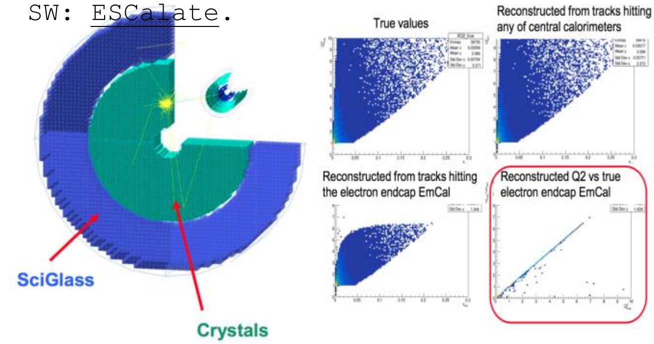

The team: V. Berdnikov, M. Bondi', C. Fanelli, Y. Furletova, T.H., I. Larin, D. Romanov, R. Trotta

Example 2: Event Classification and Tracking

Motivation:

The largest CPU resource driver for event reconstruction is charged particle tracking

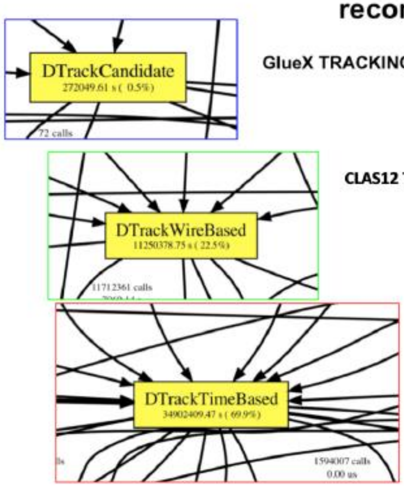

GlueX TRACKING CPU Usage:

DTrackCandidate: ~2.0%
DTrackWireBased: 22.5%
DTrackTimeBased: 69.9%

Tracking Total: 94.4%

CLAS12 TRACKING CPU Usage:

DCPatternRec ~2%
HitBasedTracking 58%
TimeBasedTracking 37%

Tracking Total: 97%

Targeted Areas of Improvement for CLAS12

- Processing **speed**
  - More efficient noise rejection
  - Combinatorics (ghost tracks)

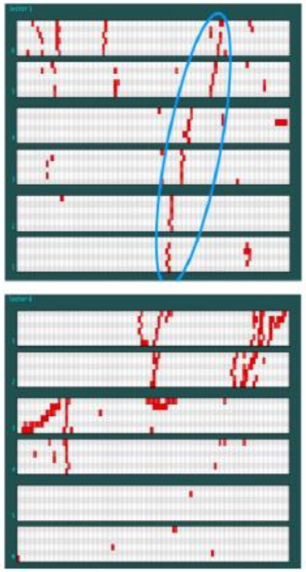

Pilot project
We can start with pre-trained **VGG16** architecture to identify tracks in our drift chambers; reduce the data sample that tracking has to work with.
Using Adversarial Neural Network we can clean up the hits that belong to the tracks: reducing number of combinatorics.
Extension:
Use regression network on the top to calculate track parameters and pass it to tracking code to minimize Kalman-Filter iterations.

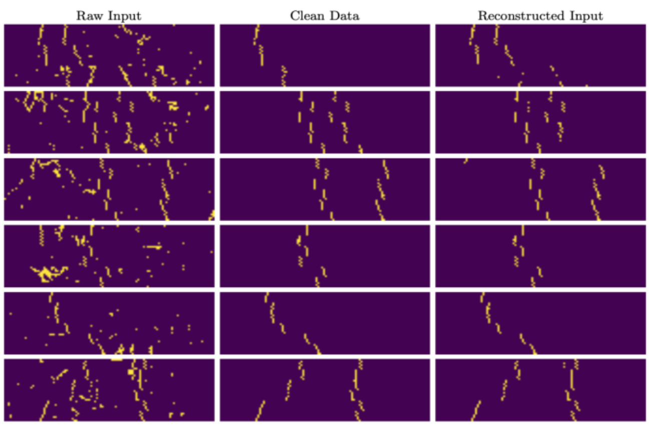

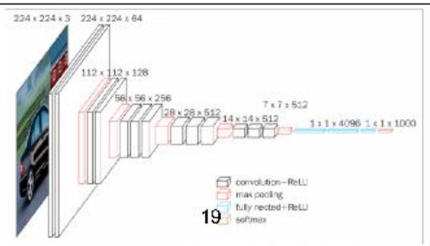

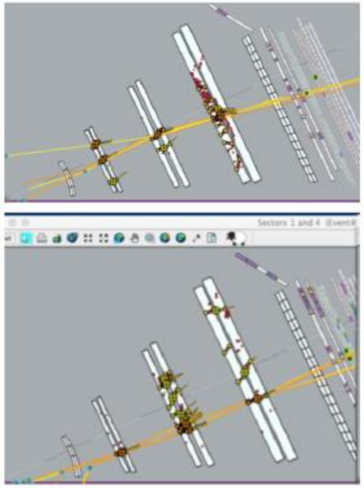

Jefferson Lab

Neural network algorithms perform comparably to traditional methods, but with 6x speedup, allowing for faster analysis

Control and Optimization of Complex Accelerators: A primary focus for developing the enormous potential of AI/ML in NP

- **Accelerator Science**
  - Optics and lattice design
  - Beam instrumentation design and optimization
  - Reinforcement learning for controls

- **Accelerator Operations**
  - Optics and lattice optimization
  - Target, charge stripper, collimation system
  - Anomaly detection and mitigation

Example 1: Superconducting RF Cavity Fault Classification

Anomaly detection and machine protection: ML-based solutions to challenges encountered in particle accelerators are yielding promising results.

C. Tennant et al., Phys. Rev. Accel. Beams **23**, 114601 (2020)

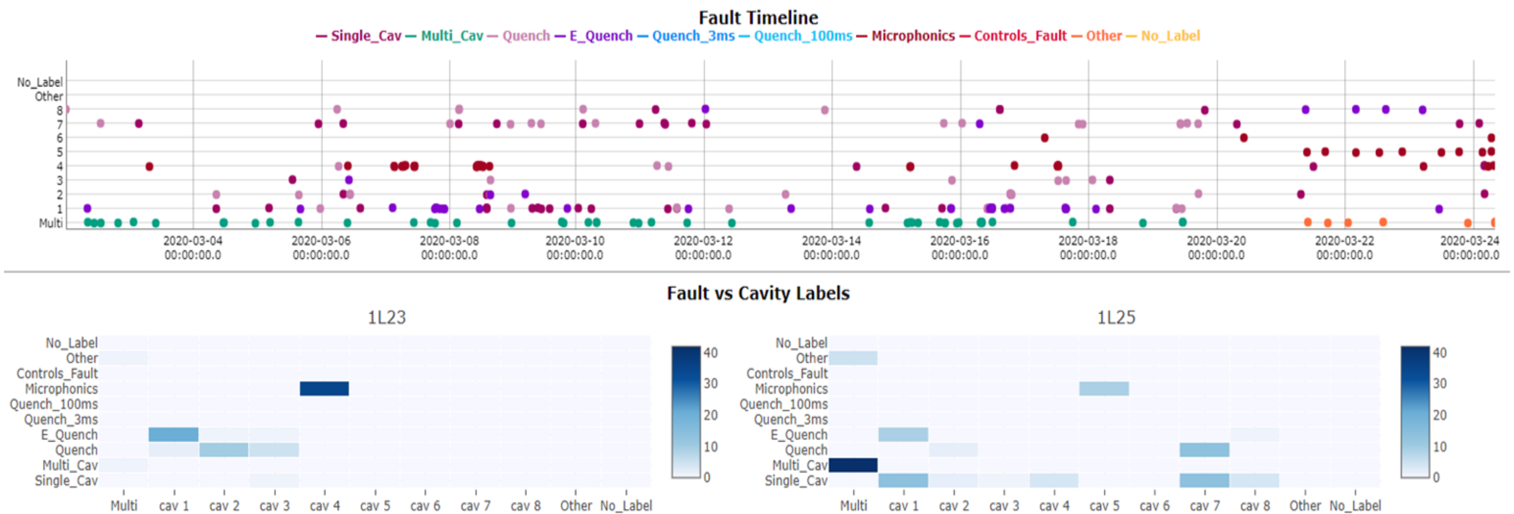

Example 2: Data quality monitoring

Image analysis significantly speeds up online monitoring and reduces data mislabeling

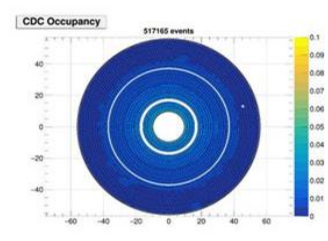

chunk 0092: Good @ 0.96700119972229

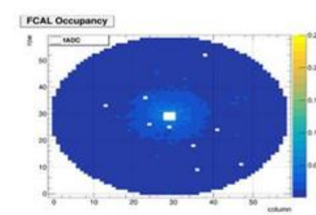

chunk 0291: Good @ 0.9995253086090088

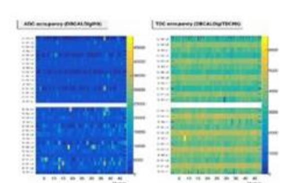

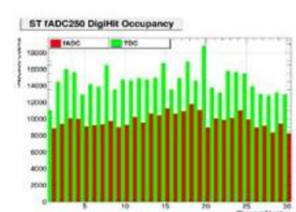

A.I. for NP: Priority Research Directions

- Game Changers in Nuclear Theory
  - LQCD: sign problem, extraction of physical observables, propagator inversion
  - Global QCD analysis
  - Identifying rare events
  - Microscopic description of nuclear fission, origin of the elements, quantified computation of heavy nuclei, correlations and emergent phenomena, spectroscopic quality nuclear density functional, neutron start and dense matter equation of state
- Holistic approach to experimentation – expert systems to increase scientific output
  - Intelligently combine disparate data sources
  - Real time analysis and feedback
- Experiment Design not limited by computation
  - Data compactification, sophisticated triggers, and fast online analysis
- Improving simulation and Analysis
  - Use AI/ML to improve the sensitivity of current instruments and accuracy of data
  - Decrease simulation and analysis time
- Accelerator Design and operations

A.I. for NP: Community Identified Needs and Communalities

- ❑ Need for workforce development
  - ➢ Educational activities
  - ➢ Need for broader community
  - ➢ Need for collaboration

- ❑ Need for problem-specific tools
  - ➢ NP applications are unique in that they are often aimed at accelerating calculation, e.g.,
    - Evaluation of models where one can use AI techniques to identify the most promising calculative pathways
    - Simulations where AI-determined parameterizations can be used to circumvent performance limiting elements

- ❑ Enabling infrastructure for AI in NP
  - ➢ Need for standardized frameworks
  - ➢ Need for comprehensive data management
  - ➢ Need for adequate computing resources

- ❑ Need for uncertainty quantification

AI in Nuclear Data

- Large potential of AI/ML algorithms to address critical Nuclear Data problems – already used for many tasks in the "nuclear data pipeline"

Workshop on Applied Nuclear Data Activities (WANDA) in March 2020: https://conferences.lbl.gov/event/292/

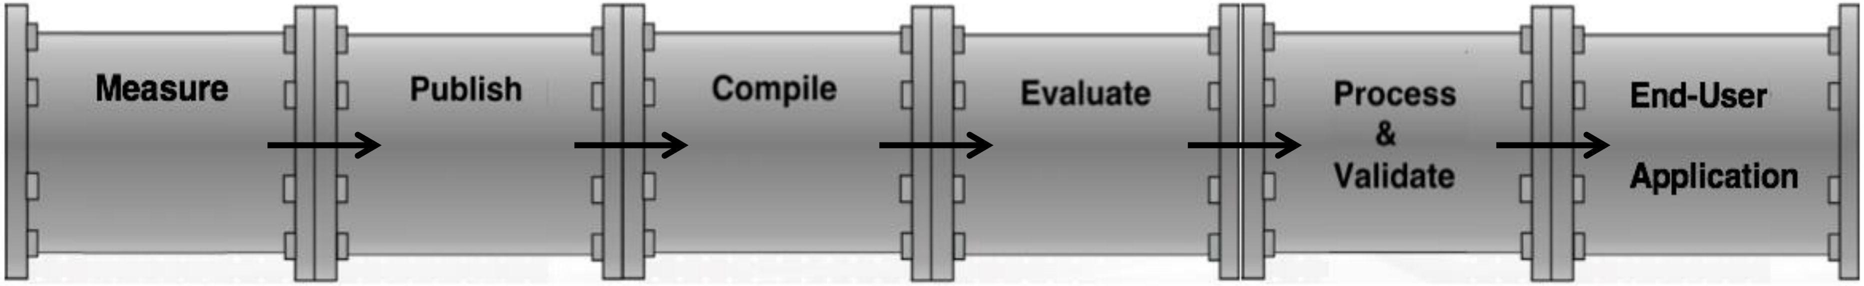

- ML use anticipated to grow exponentially in nuclear data
  - offers new approaches to longstanding problems
  - TensorFlow and Pytorch libraries speed up ML utilization
  - young researchers eager to use ML
- new trends include
  - **transforming workflows** with ML-based approaches
  - **"physics-aware" ML** models
  - **using ML to guide** experiments, theory, and evaluations

Examples AI in Nuclear Data

L. Neufcourt et al, PRL 122 (2019) 062502

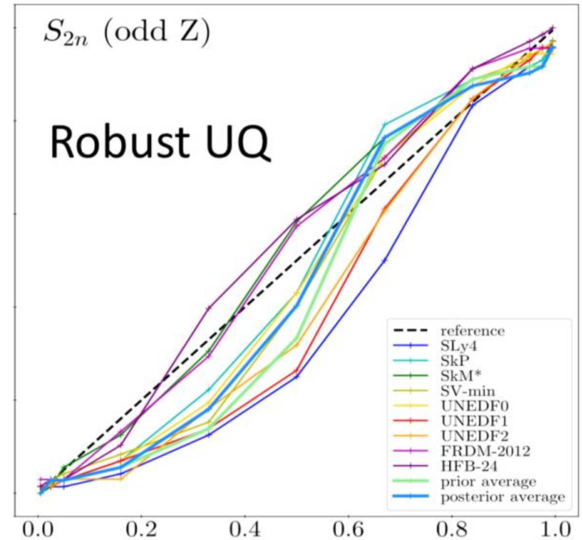

❑ Example 1: Physics aware ML models

Predict ground- and excited state energies from theory model; better predictions than traditional evaluation tools

❑ Example 2: ML-guided search

Random forests were used to augment expert knowledge in pinpointing errors in nuclear data and benchmark experiments leading to bias in simulating criticality benchmarks, e.g., ML found $^{19}F(n,inl)$ missed by experts

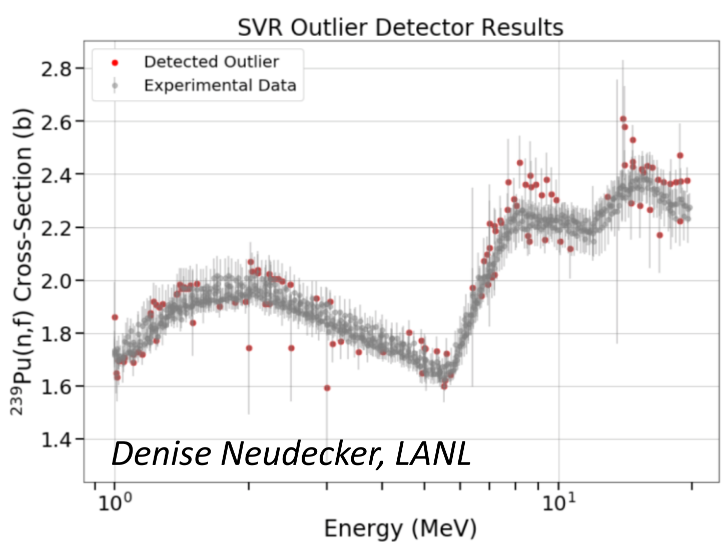

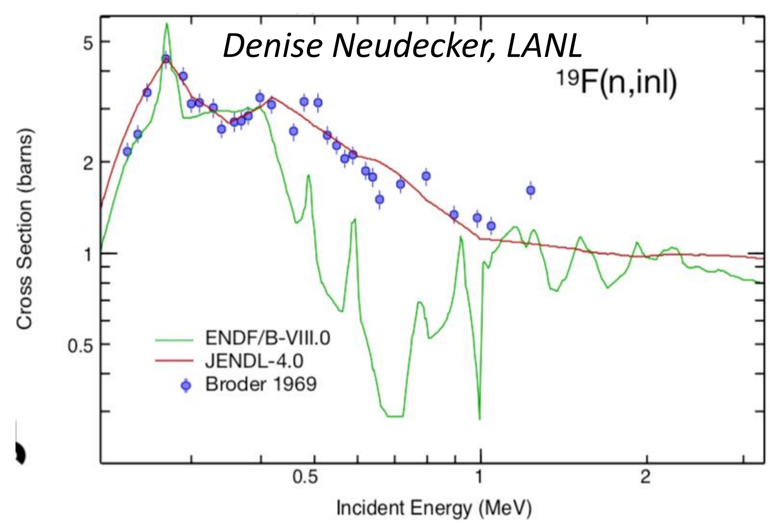

Activities in NP: 2021 AI4NP Winter School

https://indico.jlab.org/event/409/overview

AI4NP WINTER SCHOOL

11-15 January 2021
Virtual

11-15 January 2021
Virtual
US/Eastern timezone

- Overview
- Timetable
- Registration
- Participant List

Artificial Intelligence (AI) is a rapidly developing field focused on computational technologies that can be trained, with data, to augment or automate human skill. A subset of AI is machine learning (ML), which is usually grouped into supervised, unsupervised and reinforcement learning. Nuclear Physics is big data: the gigantic data volumes produced in modern experiments now and over the next decade are reaching scales and complexities that require computational methods for tasks such as big data analytics, design of new detectors, controls, and calibration systems. AI has the potential to provide the methodologies to optimize operating parameters and perform theoretical calculations of nuclear many-body systems.

The AI4NP Winter School will give the participants a deeper understanding on what Artificial Intelligence and Machine Learning are and how they can be used to analyze nuclear physics data, design new detectors, controls, and calibration systems for nuclear physics experiments and perform theoretical calculations of nuclear many-body systems. The AI4NP lecture topics will emphasize active Nuclear Physics research, both experiment and theory, that relies on AI/ML techniques, as well as synergies between the computer science and the NP communities and inspire areas for possible collaboration in order to foster vital contributions to urgent and long-term challenges for nuclear physics.

Organizers:
Paulo Bedaque (UMD), Amber Boehnlein (JLab), and Tanja Horn (CUA)
Sponsored by Department of Energy, Office of Science, Office of Nuclear Physics

THE CATHOLIC UNIVERSITY OF AMERICA

Jefferson Lab

U.S. DEPARTMENT OF ENERGY

- 361 registered participants
  - Daily attendance ~100-200
  - Experience level ranging from absolute beginner to expert
- Four major lecture topics
  - Neural Networks and DL
  - Variational Monte Carlo and ML
  - Detector Design Optimizations
  - Data set feature extractions

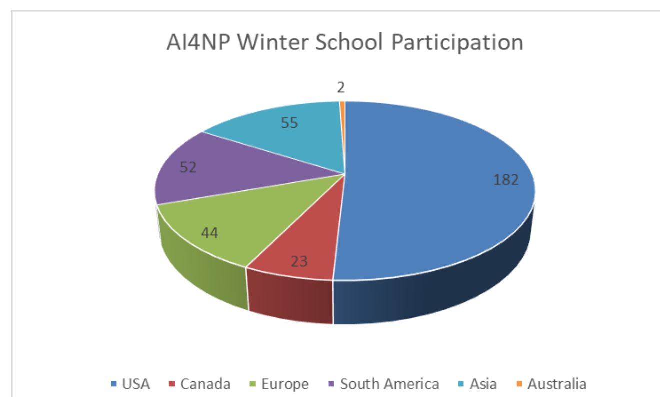

Outlook

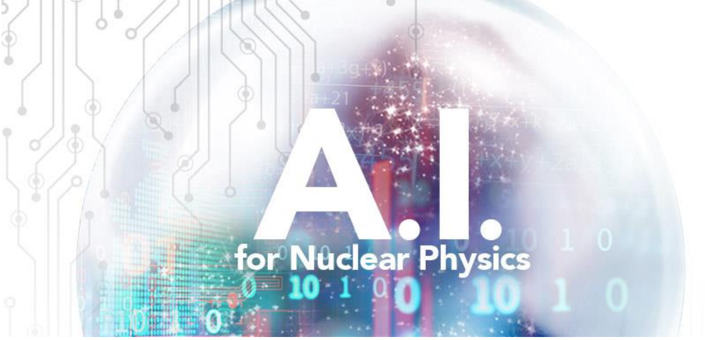

- The areas where NP research can benefit from AI/ML are ubiquitous
- NP researchers already have the talent and many of the tools required for this revolution – much progress in nuclear theory
- NP addresses challenges that are not addressed in current technologies
- NP presents data sets that expose limitations of cutting edge methods
- Strong collaborations between NP, AI/ML/data science, and industry would be beneficial for all parties – cross-disciplinary funding programs can help facilitate connections to computer science
- Education is key to increase the level of AI-literacy – research programs and curricula in AI/ML can help to attract students

Tremendous interest in AI/ML in the Nuclear Physics Community

The AI for Science Report: Recommendation #1

1. Creation of a 10-year *AI for Science* Initiative

- In order to create the world-leading AI systems and applications needed to drive scientific productivity and discovery in science and technology dramatically beyond that achievable with traditional scientific supercomputing, we recommend that the DOE Office of Science start a ten-year program to develop an ambitious ***AI for Science*** initiative, as recommended in the recent PCAST report [9].
- This program should encompass foundational research into new, science-aware AI methodologies, specifically designed for DOE mission-critical challenges, and AI solutions that can be deployed in operational settings at leading DOE research facilities.
- The initiative should provide a clear, guided roadmap from research to deployment. The DOE laboratories can play a key role here, offering leading-edge exascale supercomputers and large experimental facilities generating increasingly large scientific datasets, as well as providing critical expertise in mathematics, computer science, and experience with DOE mission-specific applications.
- No other agency has the breadth, critical mass, or recent large project management experience to undertake this cross-disciplinary ***AI for Science*** challenge. However, there is a clear case for the benefits of collaboration with other agencies and other countries, to leverage existing expertise to maximum advantage. Partnerships with other funding agencies and other countries are therefore strongly encouraged.

The AI for Science Report: Recommendation #2

2. Structure of an SC *AI for Science* Initiative

- It is recommended that this ***AI for Science*** initiative be structured around four major AI R&D themes:
  - AI-enabled applications
  - AI algorithms and foundational research
  - AI software infrastructure
  - New hardware technologies for AI
- The subcommittee believes that this ten-year ***AI for Science*** initiative should be funded at the same scale as the successful Exascale Computing Initiative (ECI) and Exascale Computing Project (ECP). Essential for the success of such an initiative is that the work of these four themes must be closely-coupled in a manner similar to that used in the ECP, as the advances and improvements in one area can inform advances and improvements in other areas.
- Figure 3 illustrates an overview of a possible roadmap for such an ***AI for Science*** initiative. As for the ECI and ECP, the roadmap for this proposed AI for Science initiative envisages an initial ‘incubation’ research phase of coordinated projects with co-design centers connecting the four major themes. Partnerships across all Office of Science domains, with participation from universities and private industry, would be initiated early in the program.
- The goal of this research phase is to specify the application grand challenges and AI/ML tools and services required as deliverables in the more focused project R&D and Deployment phases, where broad engagement of the DOE research community becomes critical. Since these applied R&D and Deployment phases will inevitably generate new questions and challenges, having the research phase continuing and overlapping with the R&D and Deployment phases will significantly increase the chances of success for the ***AI for Science*** Project.

Figure 3: Structure of SC AI for Science 10-year Initiative

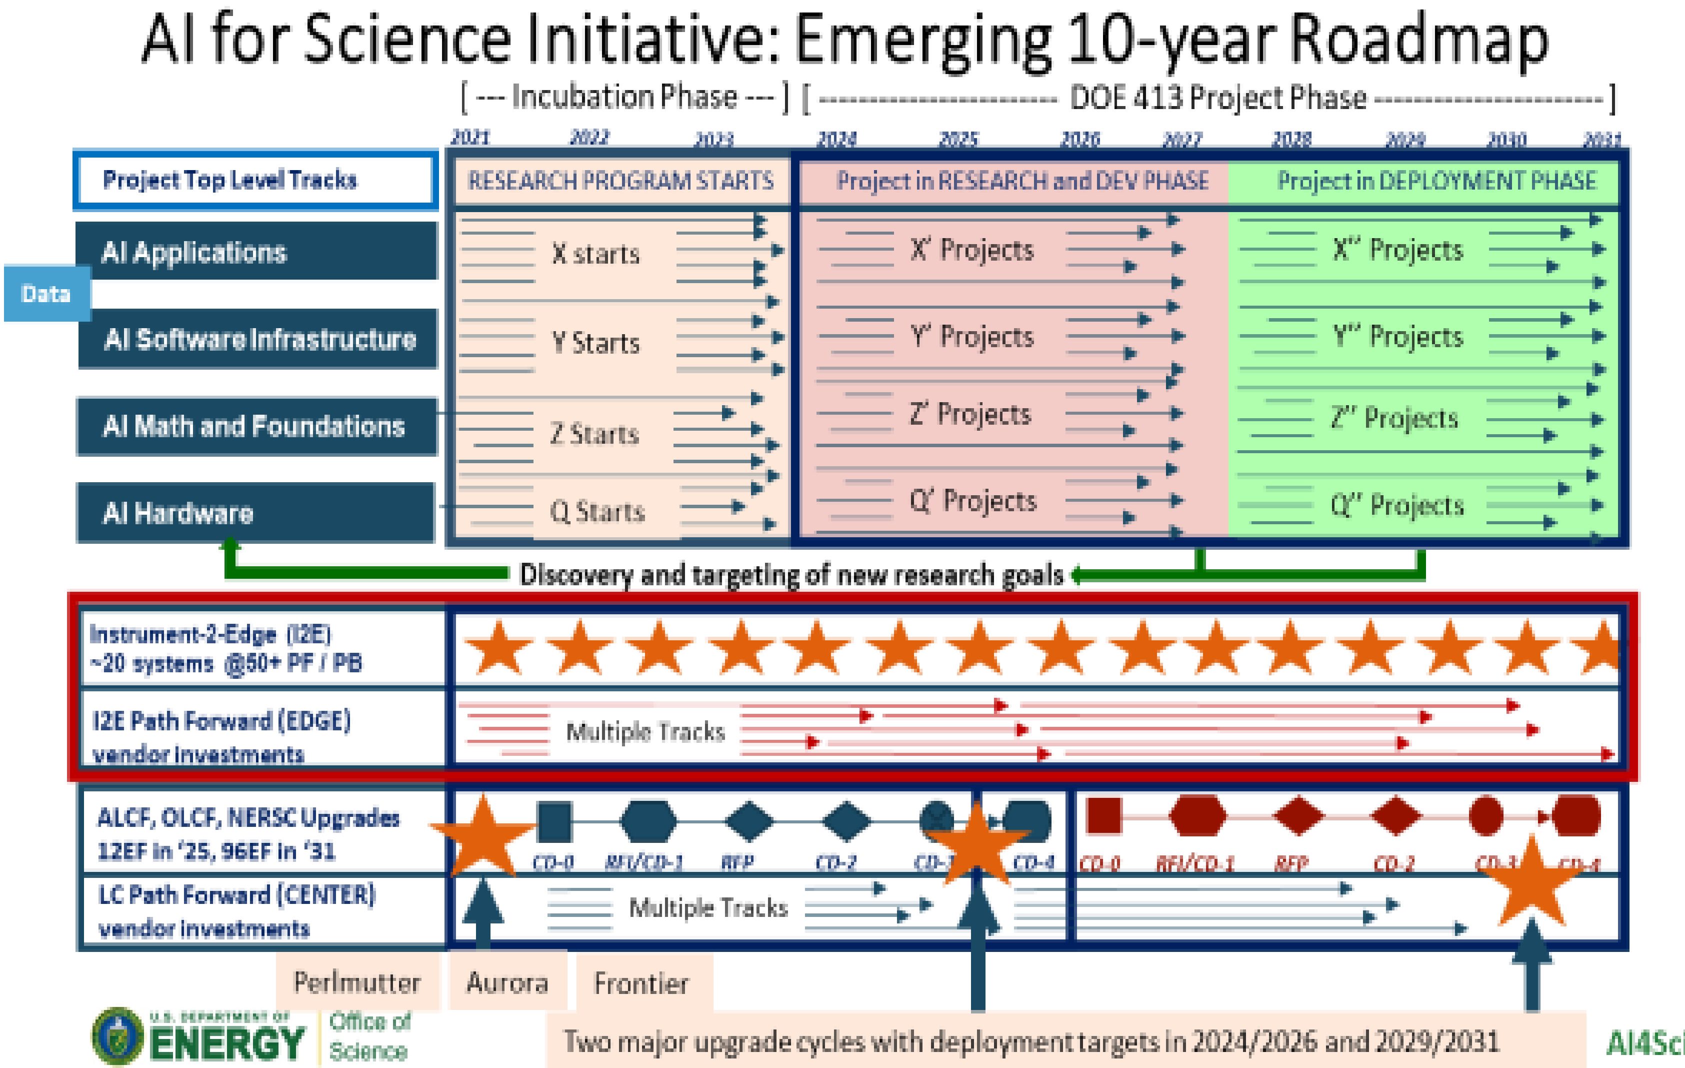

The AI for Science Report: Recommendation #3

3. An Instrument-to-Edge Initiative

- The subcommittee believes that ASCR, in close cooperation with BES and with the other science programs in the Office of Science, should work with scientists, users, and the broad academic community to define requirements, conduct research, competitive procurement and design a highly integrated end-to-end system and software stack that connects instruments at the edge to the needed AI computing resources. Integrating national and global data sources (large scale experimental facilities, observational networks terrestrial & space-based, etc.) poses unique opportunities and challenges that require addressing foundational research in the context of leading-edge scientific experiments.

- Integrated systems for acquiring, analyzing, transforming, storing, and maintaining scientific results, capturing provenance, and contributing broadly accessed analytical workflows within DOE supported computational infrastructure could be transformative. There are, however, severe challenges that will need to be confronted in terms of privacy, security, commercial licensing of data, and integrated data services.

- Building on ASCR's co-design experience in ECP, application users, software infrastructure developers, AI/ML researchers, and Lab and industry hardware specialists should be encouraged to define, develop, and contribute to a common software stack for AI/ML Edge computing resources across the different facilities. The software infrastructure should support some generic services at the facilities but also allow the easy creation of specialized AI-based software pipelines specific to the facility and capable of supporting coupling to particular instruments in some cases.

The AI for Science Report: Recommendation #4

4. Training, focusing, and retention of AI/ML workforce

- Industry, national laboratories, government, and broad areas of academic research are making more use than ever before of AI, ML, and simulation-based decision-making. This trend is apparent across many domains such as energy, manufacturing, finance, and transportation. These are all areas in which AI is playing an increasingly significant role, with many more examples across science, engineering, business, and government. Research and innovation, both in academia and in the private sector, are increasingly driven by large-scale computational approaches using AI and ML technologies.

- With this significant and increased use comes a demand for a workforce versed in technologies necessary for effective and efficient AI/ML-based computational modeling and simulation and big data analytics, as well as the fundamentals of AI/ML algorithms. Graduates with the interdisciplinary expertise needed to develop and/or utilize AI techniques and methods in order to advance the understanding of physical phenomena in a particular scientific, engineering, or business field and also to support better decision-making are in high demand.

- A strong research program will crucially rely on a complementary education and skills component, which is as important as providing adequate infrastructure support. As emphasized in the ASCR ECP Transition report [10], this is also a timely and important opportunity to focus SC efforts to create a more diverse and inclusive workforce. A continuing supply of high-quality computational and data scientists available for work at DOE laboratories is of vital importance.

- In high performance modeling and simulation, for example, the DOE Computational Science Graduate Fellowship (CSGF) program has successfully provided support and guidance to some of the nation's best scientific graduate students, and many of these students are now employed in DOE laboratories, private industry, and educational institutions. We need a similar fellowship program to meet the increasing requirement for computational and data scientists trained to tackle exascale and data-intensive computing challenges. In addition, the DOE SC should explore the possibilities for collaboration with the NSF about the provision of relevant training programs in AI/ML technologies and their application to science.

The AI for Science Report: Recommendation #5

5. Inter-Agency Collaboration

- Although the NSF has long been regarded as the lead agency for fundamental AI research, DOE is clearly the lead agency for research involving the intersection of 'Big Science, Big Data, and Big Computing.' DOE has not only established national and international leadership in HPC and supercomputing but is also a leader in the application of AI/ML technologies to the very large scientific datasets generated by their large-scale experimental facilities.

- With the NIH, the DOE SC has a successful collaboration with the National Cancer Institute (NCI) in the CANDLE project [11]. DOE is now developing an MOU with both the NSF and NIH on a program of collaborative research in Computational Neuroscience. The subcommittee, therefore, recommends that the SC explore new opportunities to work with both NSF and NIH in areas where there would be a clear benefit for scientific progress under a DOE-led ***AI for Science*** initiative. There may also be opportunities to work with other US funding agencies, such as NIST and DOD, in areas of mutual interest.

The AI for Science Report: Recommendation #6

6. International collaboration

- There is a need for broad-based, coordinated action by like-minded international partners to harness the global scientific software community to address the tremendous opportunities in data-intensive science stemming from the huge increase in scientific data collection rates. Computational and data analytical methods driven by AI/ML are now universally accepted as indispensable for future progress in science and engineering.
- International leadership in ***AI for Science*** over the coming decade will hinge on the realization of an integrated set of programs spanning the four interdependent areas noted above – AI-enabled applications, AI algorithms and foundational research, AI software infrastructure, and new hardware technologies for AI. Scientists in nearly every research field in every country will now depend on the development of such software infrastructure for high-end computing and big data analytics to open up new research fields and to dramatically increase their research productivity.
- Such AI/ML software infrastructure and algorithms capable of scaling up to exascale systems will underpin the work of global scientific communities working together on problems of global significance and enable them to leverage distributed resources in transnational configurations. In terms of feasibility, the dimensions of the task – totally re-thinking, re-imagining, and expanding, in the period of just a few years, the massive software foundation of computational and data science to meet the new realities of ***AI for science*** – are simply too many and too large for any one country to undertake on its own.
- To realize this vision for an international cooperative effort, the Office of Science needs to:
  - Provide a framework for organizing the software research community
  - Create a thorough assessment of needs, issues, and strategies
  - Initiate development of a coordinated software roadmap
  - Encourage and facilitate collaboration in education and training
  - Engage and coordinate the vendor community in cross-cutting efforts
- In DOE's Office of Science, ASCR is well suited to lead on bringing the international community together to work on these challenges.