Artificial Intelligence for Nuclear Deterrence Strategy 2023

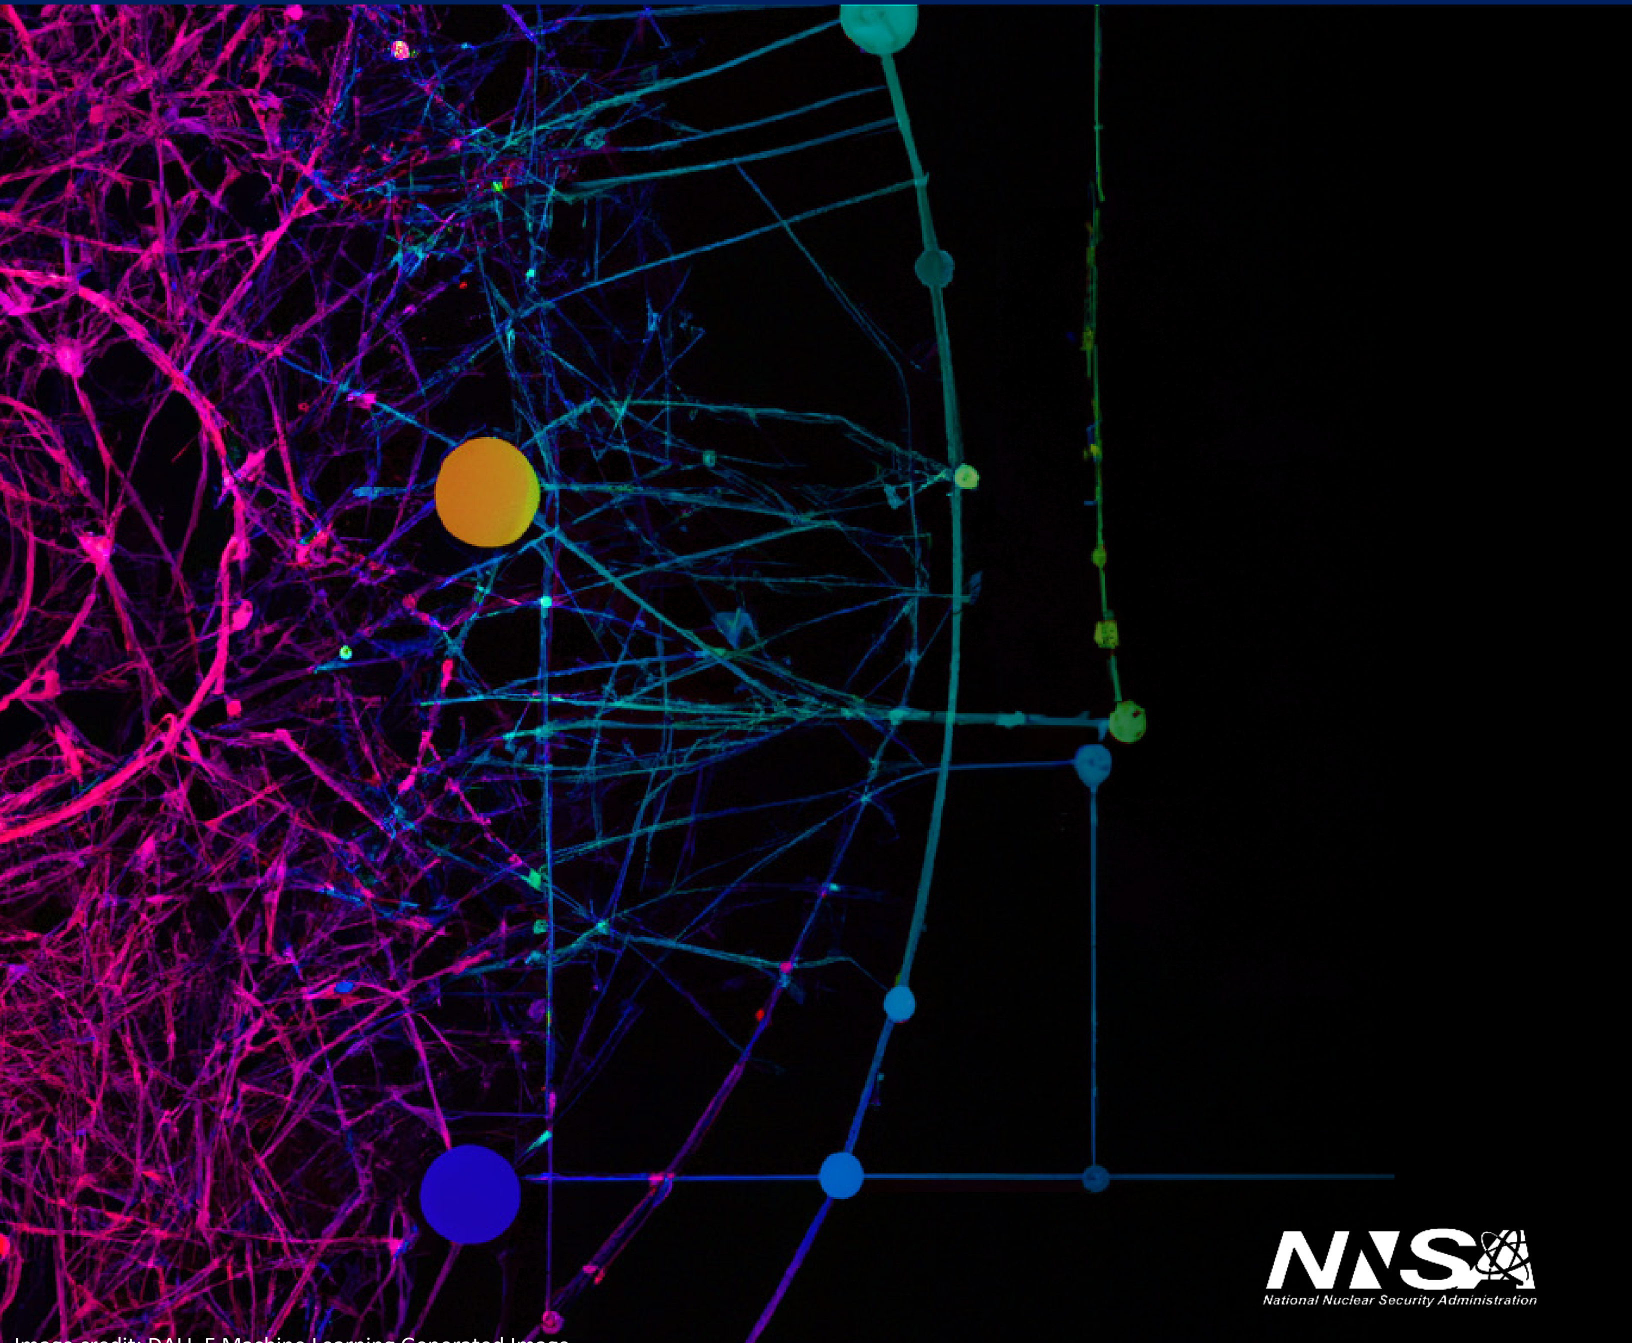

Image credit: DALL-E Machine Learning Generated Image

FOREWORD

For almost three decades, the Advanced Simulation and Computing (ASC) program has pioneered the use of high-performance computing technologies to create a unique modeling and simulation capability to address the challenges arising from the 1992 ban on underground nuclear testing. To achieve this feat, the program has balanced investments between fundamental and applied research and development, covering an extensive tools suite of hardware, software, algorithms, mathematics, and physics.

The incredible opportunities which are emerging from artificial intelligence have truly revolutionary possibilities. We believe this technology can be applied to every aspect of our nuclear deterrent mission, accelerating the time needed to solve some of the nation's toughest science challenges, but only if investments are made to re-purpose the technology to apply to the unique challenges posed by the NNSA mission.

In this strategy, we describe how the ASC program plans to deploy artificial intelligence methods in the unique high-security and high-consequence environment in which the NNSA must operate.

As we look to the third decade of the ASC program, our dedicated and talented laboratory staff will diligently work to blend the very best of our high-performance modeling and simulation capabilities with novel artificial intelligence tools, ensuring that ASC stays at the forefront of scientific delivery in support of the U.S. Stockpile Stewardship Program.

Thuc T. Hoang

Director, Office of Advanced Simulation and Computing
& Institutional Research and Development Programs
Office of Defense Programs
National Nuclear Security Administration
Department of Energy

Authors and Contributors

Los Alamos National Laboratory

James Ahrens
Nathan Debardeleben
Scott Doebling
Aimee Hungerford
Michael K. Lang
Ying Wai Li
Earl Lawrence
Jason Pruet
Timothy Randles

Sandia National Laboratories

Shawn Burns
Eric Cyr
John Feddema
Jennifer Gaudioso
Robert Hoekstra
Justin Newcomer
Ron Oldfield
Siva Rajamanickam
Jaideep Ray

Lawrence Livermore National Laboratory

Scott Futral
Katie Lewis
John R. Neely
David Richards
Brian Spears
Brian Van Essen

National Nuclear Security Administration

Kevin Elzie
David Etim
Simon Hammond
Thuc Hoang
K. Michael Lang
Anthony Lewis
James Peltz
Michelle Quirk
Emily Simpson

We would also like to acknowledge inputs and comments from the 2023 JASON AI Letter Study committee members who had an opportunity to review the ASC *AI4ND Strategy*.

EXECUTIVE SUMMARY

The NNSA Advanced Simulation and Computing (ASC) program is responsible for developing validated modeling and simulation capabilities to design, certify, and qualify the nuclear weapons stockpile of the United States. For almost three decades, ASC has successfully paired these capabilities with the most powerful high-performance computing systems in the world to underpin accurate and timely analysis for a safe, secure, and reliable nuclear deterrent.

As part of ASC's work to meet the ongoing and future needs of the Nuclear Security Enterprise, the program has invested in several projects to explore the use of artificial intelligence (AI) and machine learning (ML) methods to enhance and extend these capabilities for the NNSA mission space. This ASC Artificial Intelligence for Nuclear Deterrence (AI4ND) strategy has been developed based on lessons learned from these initial exploratory projects over the past few years.

The ASC vision for the next decade is a future where the most complex modeling and simulation capabilities are delivered at speed to NNSA, transforming how the U.S. ensures an effective and agile deterrent to meet current and future national security needs. To achieve this vision, ASC must rapidly augment our world-leading simulation capabilities with scalable, trusted, and efficient data-driven tools, including ML methods, and greater use of complementary AI technologies.

The program will adapt technologies from the external community, embedding additional capabilities and approaches relevant to the high-security, high-consequence mission of the Nuclear Security Enterprise. Objectives will be to increase the agility and responsiveness of the weapons complex in designing future warheads, improve the efficiency of the NNSA production sites involved in manufacturing complex components for the stockpile, and deepen the insights obtained from the simulation and experimental data.

To achieve this vision, the AI4ND strategy defines four areas in which the ASC program will focus its efforts and investment to ensure that AI methods can be appropriately applied to the unique and demanding needs of the Nuclear Security Enterprise:

- **Application of Artificial Intelligence Methods and Technologies to Nuclear Security Mission Areas:** the demonstration and application of AI to the Nuclear Security Enterprise and high-consequence applications will be accomplished by partnering with key stakeholders in the weapons design, production, and analysis community.

- **Foundational R&D in Machine Learning Methods and Technologies**: the development of ML tools and techniques that enable successful application in sparse or limited data environments where model accuracy constraints are likely to be much tighter than in industry or academia. In addition, the methods that will be developed will need to scale to the substantial data environment associated with the simulation of complex nuclear physics phenomena.

- **Scalable and Performant Data Infrastructure:** the availability of rich, curated data sets will be critical to the use of ML within ASC. Investment will be required to create a secure hardware and software infrastructure that connects users across the design and production agencies of the Nuclear Security Enterprise. Ensuring the

environment is scalable into the future and provides sufficient performance to prevent model training and inference from becoming a bottleneck will be an essential component to a successful execution of this strategy.

- **Enabling the Data-Driven Workforce of the Future**: ASC's most important asset is its unique workforce of laboratory technical staff who provide expertise in a wide variety of technical areas, including physics, engineering, mathematics, and advanced computing. ASC will invest in training and developing a pipeline of additional staff to engage across projects and activities, with the goal of providing data analytics and complex data-driven modeling. Attracting and retaining the best workforce will likewise mean demonstrating that ASC is performing cutting-edge research in AI methods and applying them to the nation's most challenging problems. ASC will collaborate with industry, academia, and other U.S. agencies to leverage existing knowledge, experienced staff, and best practices. Additionally, ASC will invest and innovate in areas that are more narrowly restricted to the nuclear deterrence (ND) mission space, in addition to open science research with academia, to develop staff pipeline needed to build the laboratory workforce for the future.

TABLE OF CONTENTS

INTRODUCTION

The last decade has seen an enormous growth in the use of artificial intelligence (AI). Vast improvements in computational performance afforded by general-purpose graphics processors, combined with large and accessible datasets from the internet, have enabled industry and academia to develop progressively larger and more sophisticated machine learning methods. While these have been successfully used in numerous commercial areas, the application of these methods to complex and demanding scientific environments is still in its relative infancy, and, while showing great promise, will remain a longer-term community activity.

Recently, the Advanced Simulation and Computing (ASC) program funded several projects to explore the application of machine learning (ML) methods to the mission spaces of the Nuclear Security Enterprise. These missions often encounter limited or sparse data sets, tight constraints associated with the model correctness driven by the high consequence of error, and the complex parameterization of the problems being explored. Despite these challenges, there are immense opportunities to the Nuclear Security Enterprise that can be realized with a larger and sustained research program.

There are several areas where AI/ML have the potential to reduce cost or reduce schedule in the discovery, design optimization, manufacturing and certification, and maintenance and surveillance phases of a nuclear warhead system (Figure 1). The phases are shown in the figure below and discussed in the next section.

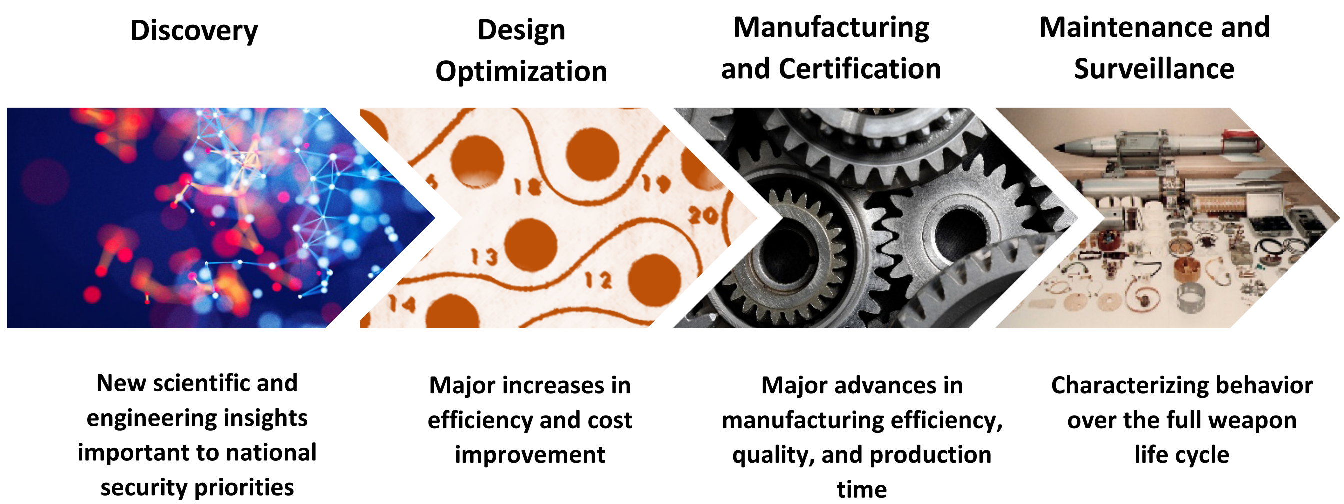

Figure 1 - Life Cycle of a Typical Warhead System.

The AI/ML technologies have potential to make significant advances for the full weapons life cycle. Expected outcomes from a sustained program include outpacing competitor-nation programs, increased analysis efficiency, and a broadening of the range of technology options that can be provided to national defense mission areas. DOE and the NNSA are prepared for this challenge due to qualified multi-discipline teams, robust expertise with high-performance and exascale computing resources, and extensive experience in delivering on high-security and high-consequence mission scope. Achieving these goals with new ML-enabled capabilities requires mission-focused strategic investments in laboratory staff, foundational methods, and infrastructure across the Nuclear Security Enterprise.

MOTIVATIONS FOR THE USE OF ARTIFICIAL INTELLIGENCE IN THE NUCLEAR DETERRENCE MISSION

AI and ML will be evaluated for potentially accelerating the design and production of systems that support the U.S. nuclear stockpile, and ultimately, nuclear deterrence (ND). AI/ML are technologies that hold the promise of reducing turn-around time for material discovery, accelerating model development and maturation, automating activities in manufacturing and certification, and simplifying the maintenance and surveillance of ND systems. In addition, AI/ML will assist in improving the NNSA's ability to assess current and new environments for weapon systems, and reduce design margins in environmental specifications prior to experimental data being available.

Discovery

Scientific discovery is the bedrock of the ND mission because ASC can produce capabilities and solutions to solve challenges. For example, routinely discovering and manufacturing-sustainable materials that meet the stringent requirements of weapons systems would be transformative. ML algorithms that leverage high-fidelity experimental characterization data, nano- and meso-scale computational simulation, and existing analyst workflows could be used to provide accurate multiscale methods for the prediction of material properties. AI/ML-enabled multiscale materials modeling have the potential to improve the resolution and reduce the cost of computational simulation-based predictions and impact the ND mission to quickly identify new materials with designed properties, such as high explosives with improved safety performance.

Design Exploration and Optimization

ML has the potential to accelerate ASC modeling and simulation capabilities; hence designers can interactively explore a parameter space larger than what is possible with human-controlled simulation alone. ML can realize major efficiencies by reducing complex design parameter spaces and optimizing weapons parts and system designs for metrics such as manufacturability, longevity, reliability, and cost. AI/ML-generated surrogate models have also shown solutions that are up to three orders of magnitude faster than direct numerical simulation, with only limited degradation in accuracy. These types of acceleration in computation would allow weapons designers to explore design options that were previously unavailable because of high computational cost.

Manufacturing and Certification

The production and certification of weapons components and systems is one of the most challenging and lengthy tasks in the Nuclear Security Enterprise. AI/ML has the potential to accelerate and automate these steps when used in a tight collaboration loop involving designers and production experts. The NSE frequently requires small quantities of unique parts where the impact of any individual component failure can be catastrophic. ML can accelerate the collection of inspection and testing data, and its comparison to computational simulations of the desired functionality. Similar to AI-assisted medical imaging, NNSA could employ real time AI-assisted defect screening for material and part certification. Automated inspection and testing processes will bring significant savings in cost and time over the human effort that is currently performed manually. Developing an AI solution to this screening and failure prediction problem will have a major impact on multiple-component structural certification bottlenecks and enable rapid and credible deployment of new materials and designs.

Maintenance and Surveillance

The NNSA labs are developing new technologies that will enable in-situ surveillance of stockpile components and systems in the future. ML can enable fast comparisons between computational digital twins of components and systems using captured surveillance data, supporting the identification of defects, anomalies, and out-of-specification variations. New AI foundation models extended with a corpus of classified nuclear weapons data should reduce the burden of generating significant findings reports by helping to rapidly identify the root cause of defects. The ability to automatically digest technical documents, compare surveillance data to computational simulation data of its digital twin, create summaries and perform root cause analysis would be transformative for the NNSA mission.

Enabling Capabilities and Investment Areas

Foundational R&D in Machine Learning Methods and Technologies

To support the high-security, high-consequence character of the NNSA mission space, ASC must invest in methodologies that diverge or exceed those currently being pursued by industry or academia. Although ASC will leverage industry technology wherever feasible, the program recognizes that the scientific rigor and consequence of inaccuracy associated with its mission place greater emphasis on the development of trustworthy AI/ML solutions that will exceed most, if not all, demands from external communities. The program will have to ensure that data related to national security remains secured and uncompromised, challenging the use of entirely open or commercial solutions. Investment in fundamental research areas performed in close collaboration with subject matter experts across ASC, the rest of the Nuclear Security Enterprise, academia, and industry will ensure leading edge solutions for the NNSA mission. Three areas identified for investment are:

Physics-Informed Machine Learning (PIML)

Data-driven models will be used to represent physical processes in nuclear weapon engineering problems, with key research targeting how physical constraints can be embedded in the model and its training. Learning from limited data with models that capture the governing physics will help lead to trustworthy AI/ML. PIML models can provide results orders of magnitude faster than conventional engineering simulators and thus allow the automated discovery of optimal, non-intuitive, component designs. They will also allow for the construction of digital twins of full weapon systems, to accelerate manufacturing via digital engineering practices and to forecast aging effects, when calibrated to weapon surveillance data.

Limited and Sparse Data Sets

The explosion in the use of machine learning over the past decade, especially with highly accurate deep learning technologies, is due to progressively more powerful hardware and the accessibility of large amounts of data through the internet. For NNSA mission space, there are relatively few full-scale experiments, and limited underground weapon tests. Data from ASC's state-of-the art modeling and simulation will be required to augment existing data sets. Investments are needed in the development of novel methods that combine multiple data sources and improve training accuracy with the limited experimental data available. In other cases where data is similar to open science results (e.g., materials), open foundation models trained to support the Nuclear Security Enterprise mission offer a potential solution.

Verification, Validation, Uncertainty Qualification, and AI Trustworthiness

NNSA currently employs best-practiced verification, validation, and uncertainty qualification to ensure trust in applying its modeling and simulation capabilities to its high-consequence mission space. While much of industry is not implementing this same level of rigor in their solutions, NNSA must assess the credibility of a solution, and, in specific circumstances, interpret model predictions to gain knowledge of why or how a prediction was made. In adding AI/ML methods to the portfolio, the trustworthiness of the tools must be maintained. Trustworthy AI/ML is a rapidly evolving technical consideration where industry, academia, and government agencies are developing essential building blocks for the community. To support that work, ASC will explore how to integrate existing verification/validation and uncertainty quantification with emerging trustworthy AI/ML concepts that include bias, interpretability, and explainability. Until these issues are addressed, more conservative approaches that restrict AI systems to validated solutions or to a human-in-the-loop workflow will be required to support using these methods when making decisions in high-consequence applications.

Nuclear Security Enterprise-wide Data Infrastructure

The construction of complex ML models is typically performed using large volumes of well characterized data. While the definition of large volume is unique to each problem being explored, typically tens of thousands to millions of training samples may be required to construct models with a sufficient level of fidelity. A well designed and rich data infrastructure is needed to develop AI/ML solutions for the Nuclear Security Enterprise.

Performant Data Storage Technologies and Systems

The collection, storage, and retrieval of large volumes of data remains a significant challenge in terms of cost, aggregate system performance, and scale for all practitioners whether they are from industry, academia, or government agencies. This can be particularly challenging in scientific environments where data may need to be captured at exceptionally high speeds from experimental facilities or large-scale HPC platforms. ASC will invest in R&D to support the development of future data storage hardware and software systems that meet our program requirements for efficiently storing complex scientific simulation and experimental data.

Federated Data Environment

A diverse set of data sources across the Nuclear Security Enterprise will be required to construct a future family of mission-relevant machine learning models. A federated infrastructure, with storage resources decentralized across the Nuclear Security Enterprise, but managed under a common framework or set of application programming interfaces (APIs), will better enable sharing and use of data that might otherwise be inaccessible. Likewise, the ability to move compute to where data lives could help overcome limitations due to the high latencies and limited bandwidth involved in accessing data collections over fully remote connections. Extension and upgrade of the classified and unclassified inter-site networks will allow data to be exchanged or copied locally for more efficient processing. Federation of resources is inherently more scalable than a single data storage location and will support the integration of existing and new data stores throughout the complex with local sites having autonomy over data formats and security and retention policies.

Data Access Interfaces

Integrating existing and new data infrastructure across programs, projects, and sites requires great flexibility. The use of a single API to all datasets is likely to be unachievable across the complex and diverse needs of the Nuclear Security Enterprise. Moreover, no single data layout will support the optimization of data formats to match the upstream creation methods (for instance, utilizing a format which maps to an experimental instrument), resulting in lower performance. Rather than mandate a single approach to data formats or data access, the ASC program will promote the description and publishing of APIs by dataset owners to allow data virtualization across a collection federated data sources. Such an approach will allow users from across the Nuclear Security Enterprise to access data, regardless of its location or format, while still benefiting from its optimized and curated structure.

Machine Learning Architectures and Systems

Alongside federated data storage systems, ASC anticipates an increased demand for high-performance machine learning training systems as machine learning models are used more broadly in the weapons design community. The increased demand will be driven by not only the development of many models but also the need to frequently perform retraining as additional simulation or experimental data becomes available. Model training and inference resources will need to be incorporated into traditional CPU and GPU systems allowing data to seamlessly move between memory spaces during, and adjacent, to HPC simulations. These workflows necessitate new APIs for managing data across accelerators and will require co-design to ensure performance. Edge computing will also be used to couple lightweight compute for data analysis (potentially in mission areas such as warhead surveillance) and data collection and model inference for manufacturing and facility operations.

Workforce, Collaborations and Partnerships

The scale of investment in ML methods by the broader community far exceeds the resources that can be afforded by the ASC program alone. Benefiting from the immense technical innovation in the community necessitates partnering with technology leaders with the aim of augmenting existing solutions to make them usable within the NNSA. For these collaborations to have impact, ASC must develop open challenge problems and proxy applications that closely relate to NNSA's mission space to enable focused research partnerships. This section outlines relationships the ASC program will forge to benefit from the broader ecosystem investment in AI/ML and to develop a capable workforce pipeline into the future.

Partnerships with Industry

By committing to a high degree of collaboration with industry, ASC can benefit from the R&D already taking place and fund exploration of areas unique to its mission. Many areas ripe for collaboration, such as digital twins for manufacturing and natural language foundation models, would accelerate the uptake of these tools even if they need to be retrained or redeveloped for NNSA-specific use cases. NNSA benefits by augmenting lab expertise with the industry methods and investment.

Partnerships with Academia

By partnering with universities, the NNSA can take advantage of increasing student interest in the field of AI/ML. The ASC program can influence the curriculum and research efforts to expose students to AI/ML in scientific applications, in contrast to commercial applications. This model has succeeded for the ASC Predictive Science Academic Alliance Program (PSAAP) which requires multidisciplinary teams to work on annual overarching simulation goals and serves as an excellent laboratory staffing pipeline.

Partnerships across the U.S. Government

Meeting its mission requires ASC to have a close partnership with its existing and new partners. The DOE Office of Science programs are positioned to demonstrate ML technologies and methods that will serve as a foundation for innovative advances to strengthen NNSA capabilities. Computing resources and human talent developed by ASC have long been critical for supporting broader global security applications, as well as the technical knowledge gained in working on the stockpile stewardship mission providing a frame of reference for evaluating foreign programs. Similarly, contributions from global security missions are essential for guiding future efforts by ASC. Assessments by the intelligence community, along with actions by non-proliferation and counter-proliferation programs, provide context for the future evolution of U.S. military capabilities.

There are numerous AI/ML efforts underway across federal space. ASC is looking to leverage ongoing work and partnerships with the DOE Office of Science, Department of Defense, other U.S. agencies, and internally with other NNSA program offices. As opportunities are presented with other government agencies, ASC will look to identify key stakeholders and participants to reduce stove piping and strengthen collaboration in advancing nuclear security.

International Partnerships

ASC is already collaborating internationally with programs such as France's Alternative Energies and Atomic Energy Commission (CEA), the United Kingdom's Atomic Weapons Establishment (AWE), and Japan's RIKEN Center for Computational Science, as opportunities allow and where the benefits of doing so are clear to the broader ASC goals. Similarly, the program is carefully tracking other international programs, such as Australia's HPC and ML investments, and the European Exascale (EuroHPC[1]) effort that will deploy multiple exascale systems.

ASC Workforce of the Future

For AI/ML to have mission impacts, ASC needs to develop and foster interdisciplinary teams with AI/ML practitioners, data librarians, data scientists and laboratory subject matter scientists. An FY 2024 exemplar activity is to tune foundation models with classified data resulting in a knowledge assistant for nuclear weapons designers, especially to allow new hires an easier path to understand the prior art.

Through leveraging our collaborations, we can create new staff pipelines and expose current staff to best practices. Adopting standard AI/ML software and methods developed externally will ensure that ASC is using best practices that open the opportunity to growing the NNSA workforce.

CONCLUSIONS

Successful execution of this initiative will position the ASC program to effectively meet its strategic objectives of increasing confidence in its tools and capabilities, enabling efficiencies across the nuclear security enterprise and providing a pathway to attract and retain highly specialized talent needed for NNSA missions. Accomplishments through this AI4ND strategy, in collaboration with academia and industry and other U.S. government agencies, will enable the ASC program to realize its vision of being responsive and agile to an evolving stockpile, via a more modern and powerful computational approach to high-performance modeling and simulation.

[1] https://www.etp4hpc.eu/eurohpc.html

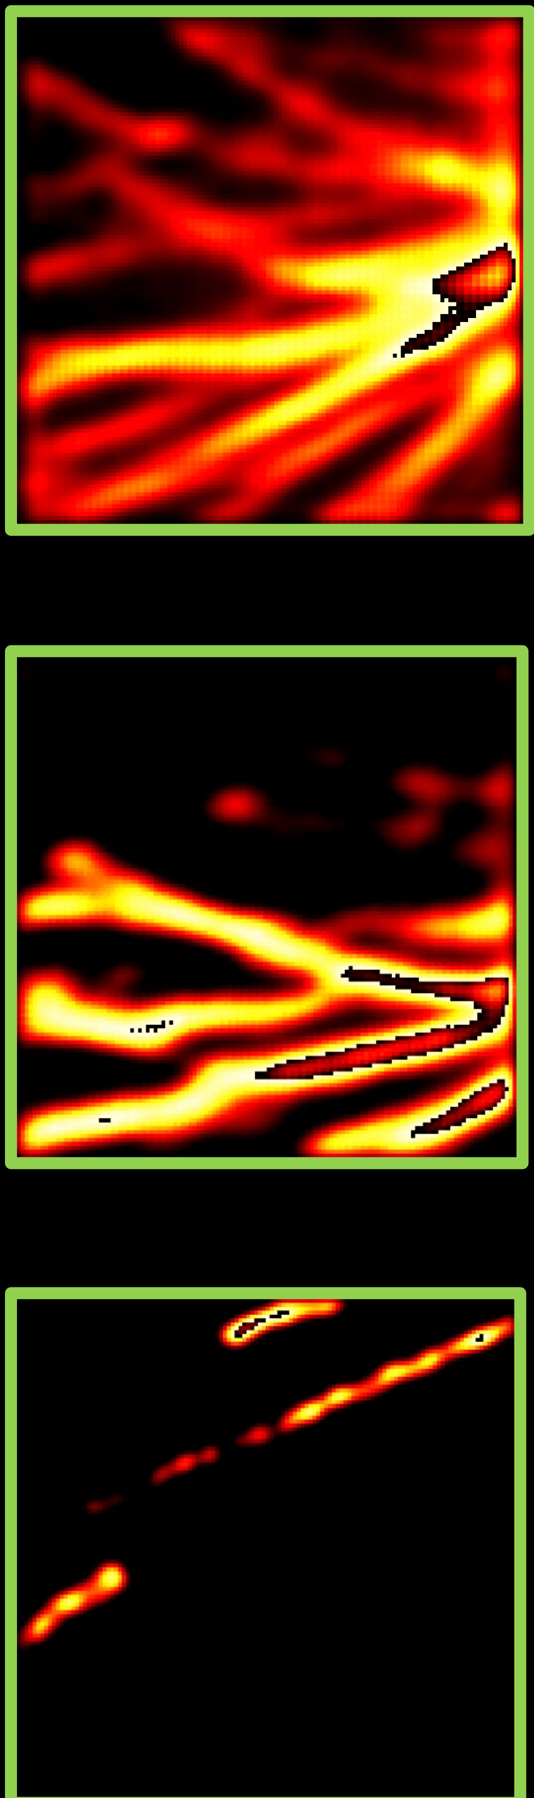

Image credit: machine learning predictions for displaced molten material (spatter) in metal alloy laser powder bed fusion (LPBF). This semi-supervised learning technique allows pixel-by-pixel classification with high-speed thermal camera imagery. Sean Tronsen, Colt Montgomery and Nathan DeBardeleben, Los Alamos National Laboratory.

Established by Congress in 2000, NNSA is a semi-autonomous agency within the U.S. Department of Energy responsible for enhancing national security through the military application of nuclear science. NNSA maintains and enhances the safety, security, and effectiveness of the U.S. nuclear weapons stockpile without nuclear explosive testing; works to reduce the global danger from weapons of mass destruction; provides the U.S. Navy with safe and effective nuclear propulsion; and responds to nuclear and radiological emergencies in the U.S. and abroad.

DOE/NNSA-0145

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration