



Advances in AI for Science: DOE's National Laboratories

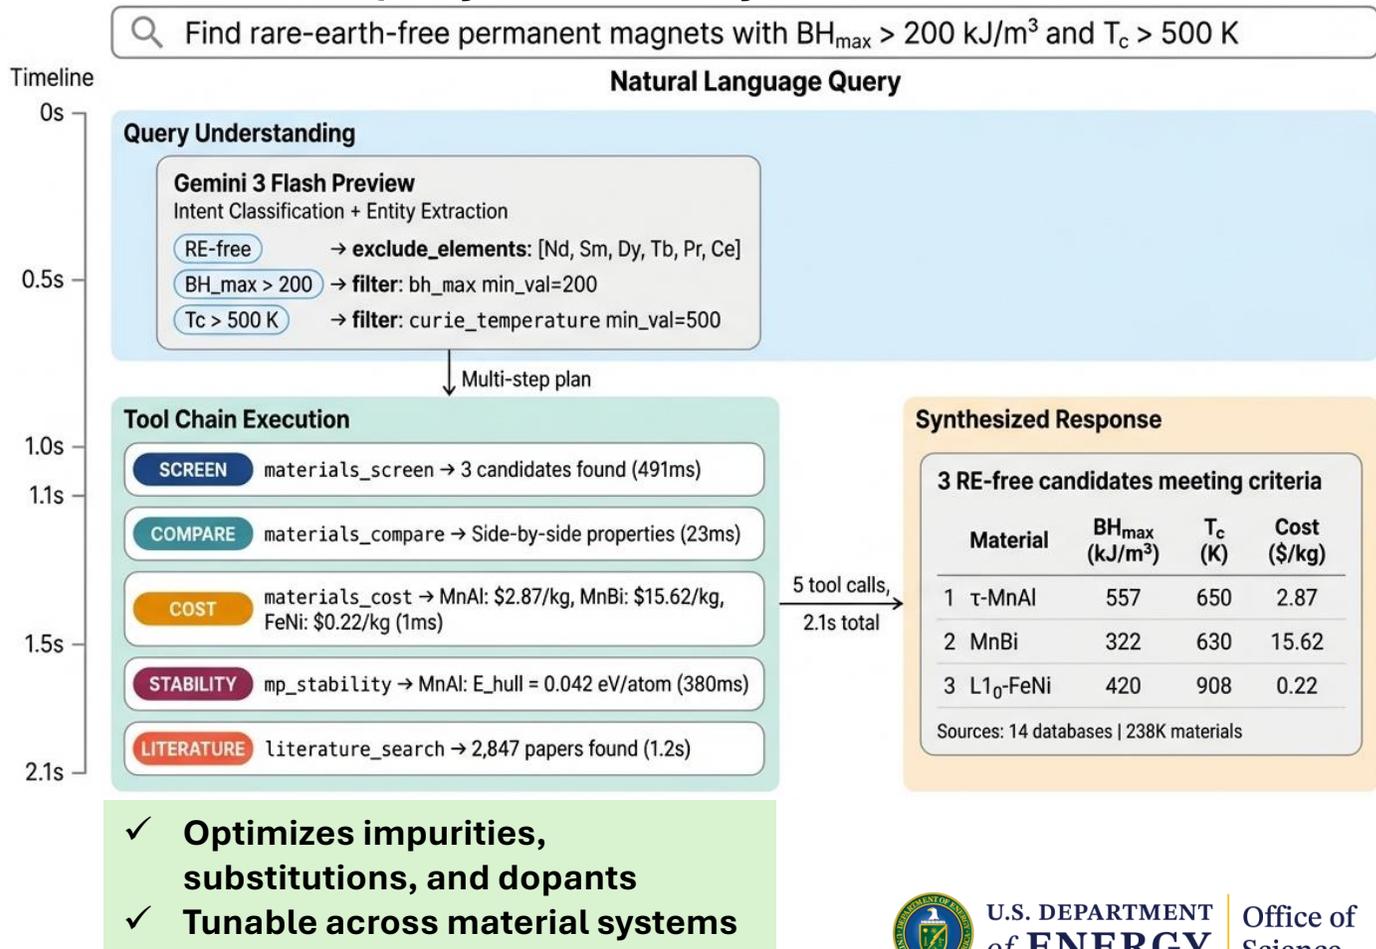


U.S. DEPARTMENT
of **ENERGY**

Office of
Science

MAG-GPT: A RAG-Augmented Foundation Model for Magnetic Materials Discovery

“Query to Discovery in Seconds”



Motivation: High-performance magnets rely on REEs, creating supply chain vulnerabilities. Domestic production is costly and complex, while new materials must deliver predictable performance with reduced REEs.

Approach

- Build curated data retrieval and contextual grounding framework.
- Integrate RAG-augmented LLMs to accelerate magnet design while integrating supply chain intelligence (*Retrieval-Augmented Generation*).

Impact

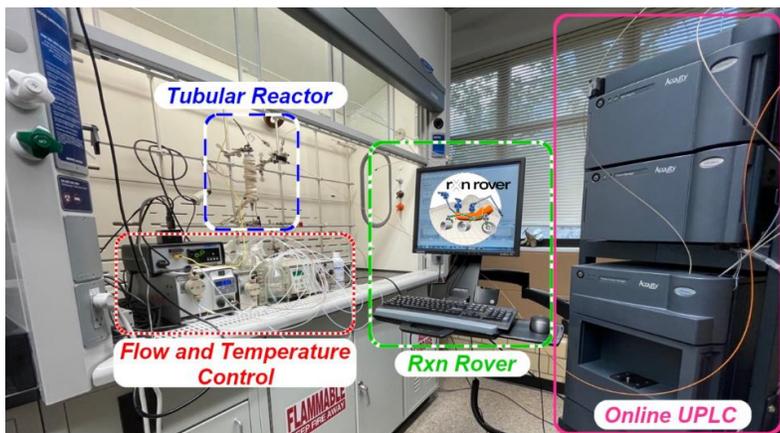
- >100× speed-up in targeted material search.
- Strengthens U.S. leadership in magnet design

S&T Challenge

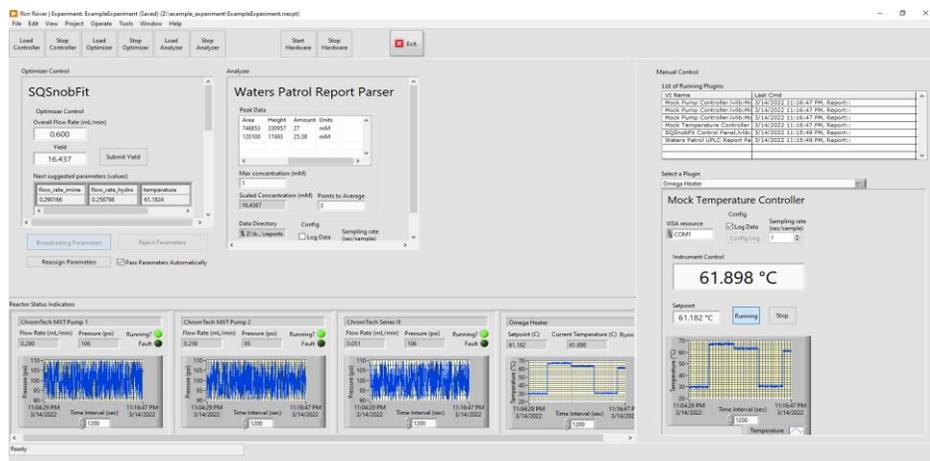
- Aligns with “RE-Free Magnet Innovation”, and “Supply Chain Resilience and Scalability”.



Rxn Rover: User-Friendly AI-Guided Laboratory Automation



Single reactor accelerated by Rxn Rover



Rxn Rover user interface



Motivation

Self-driving laboratories (SDLs) combine AI intelligence with automation technologies to create fully closed-loop processes. These platforms can provide on-demand synthesis, expedite process development, and improve data reproducibility. However, SDL platforms are tied to a specific hardware platform or require significant programming experience to set up.

Approach

- Modular, open-source automation software, **Rxn Rover**, designed for flexible, AI-augmented experiment capabilities
- Plugin-based design for custom laboratory flexibility and simple no-code setup

Impact

- Turns multi-day experiments into single-day continuous experiments with 10-fold increase in time efficiency
- Improves safety and reproducibility by minimizing human interaction
- Enables AI/ML supervision for rapid reaction optimization

S&T Challenge:

Aligns with "Achieving AI-Driven Autonomous Laboratories" and "Securing America's Critical Materials Supply."

Self-Driving Laboratories for Critical Materials Separation

Motivation

Critical materials separation depends on advanced ligands, but conventional ligand development and process evaluation are time-consuming and labor-intensive. Rapid progress requires both automated synthesis and fast screening of separation performance.

Approach

- Digitized flow reactor, controlled by Rxn Rover, enables automated synthesis with real-time monitoring
- AI algorithm supervision to minimize the steps and time for process optimization
- High-throughput evaluation of rare-earth separation on a robotic platform

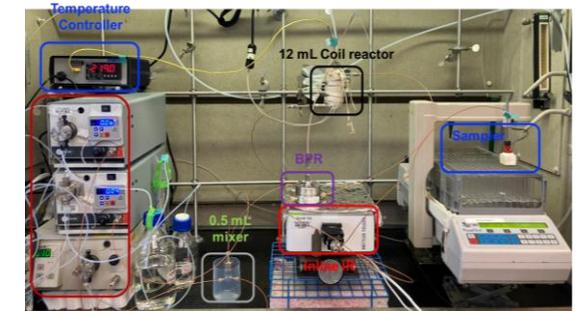
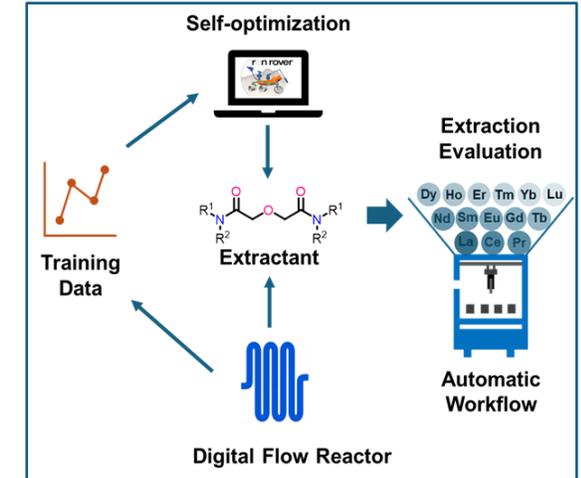
Impact

- Reduced experimental effort by over 40%
- Enabled 180 automated separation screenings in 48 hours
- Established the foundation for fully self-driving labs for critical materials

S&T Challenge:

Aligns with “Securing America’s Critical Minerals Supply” and “Achieving AI-Driven Autonomous Laboratories.”

AI and automation accelerate rare-earth separation (top) via automated synthesis (middle) and high-throughput screening (bottom)



Gridmind: A Multi-Agent AI Framework for Grid Decision Support

Motivation

Analyzing and **managing the electric power grid are complex tasks**. They involve many different types of simulations and each simulation requires specialized tools, programming skills and deep technical knowledge. These **disconnected workflows** can slow down decision-making when fast, clear actions are most needed.

Approach

- Developed GridMind, an **agentic AI orchestration framework** for the power grid that fuses grid data, forecasts, domain tools, and advanced analytics into a unified decision workflow
- Coordinates specialized agents and physics-based tools to rapidly produce **prioritized, traceable, decision-ready recommendations**

Impact

- Allows operators to run and interpret complex power system analyses using **natural language**
- **Orchestrates specialized AI agents** across weather, outage risk, grid analysis, and validation

Capability Leap

- Combines AI reasoning with physics-based solvers and fast surrogate models to **screen many possible future scenarios quickly**

S&T Challenge

Aligns with “Scaling the Grid to Power the American Economy”

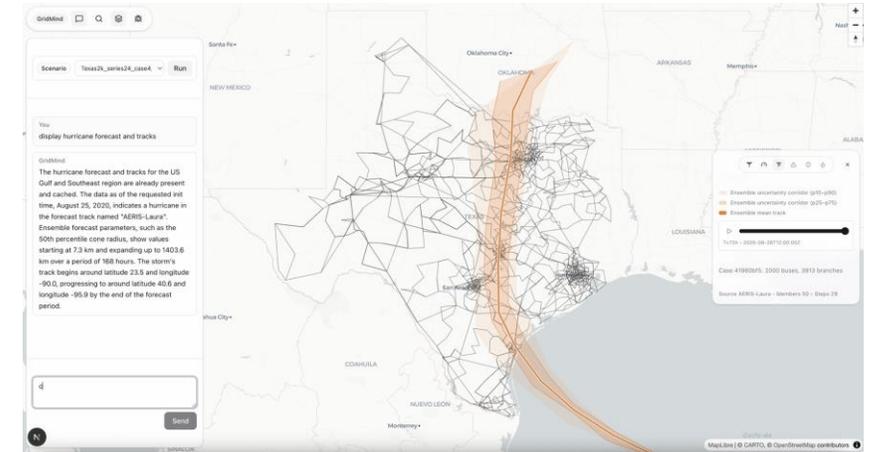


Figure top: GridMind showing hurricane-track uncertainty
Figure bottom: GridMind showing cascading failure hotspots

LUMINA: A Grid Foundation Model for Evaluating Risks to the Grid

Motivation

When the grid is under stress, **grid operators have minutes to make the right decision**. Today, analysis is slow, which **limits how many scenarios can be tested** before a decision is made. As a result, the rare combinations of factors that drive the most severe failures are often excluded from both the analysis and the decisions that follow.

Approach

- Developed LUMINA, a foundational model that screens **thousands of cases in seconds** for violations and load-shed risk
- The highest risk cases are then verified using simulation-grade physics, producing a shortlist of results that are both fast and defensible

Impact

- Foundation-style pretraining achieves equivalent feasibility with 83.6% fewer training steps on large grids – a **key step toward generalizable grid foundation models**
- Open-sourced as the LUMINA suite to accelerate reproducible, physics-consistent AI for power grid operations

Capability Leap

- Evaluates 43,000 scenarios in 5 minutes, accounting for **99% of outage risk** probability
- Traditional solver evaluates 30 scenarios in same time (**0.5% risk probability**)

S&T Challenge

Aligns with “Scaling the Grid to Power the American Economy”

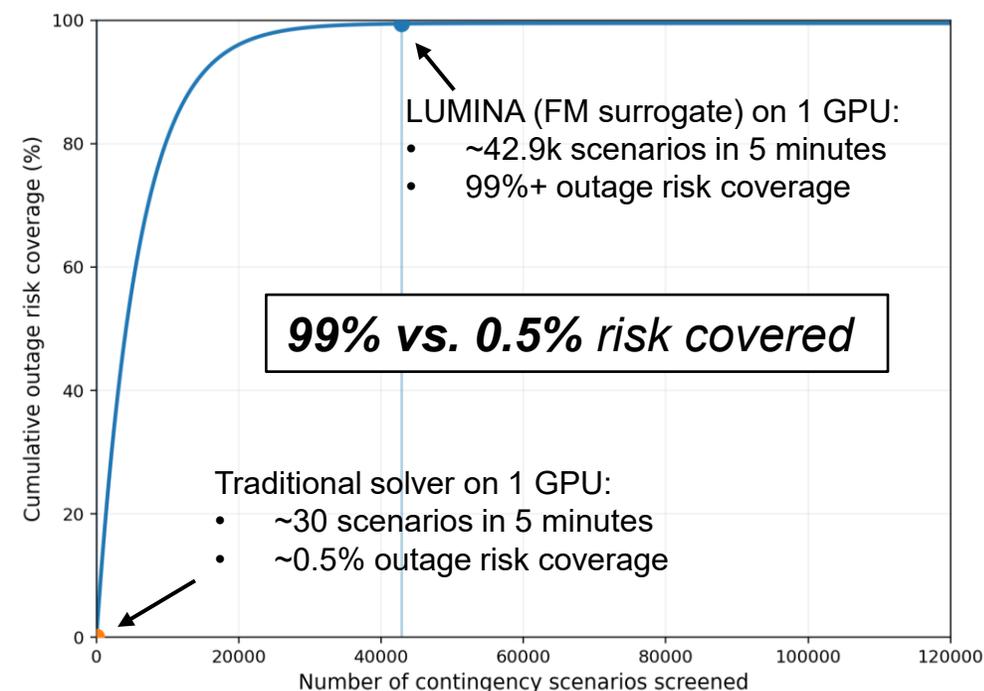


Figure: Risk coverage and scenarios screening matrix

COMB-FLOW-AI: A Frontier AI Platform for Transforming Combustion and Energy Technologies

Motivation

Fuel combustion technologies constitute 84% of total primary energy production in the U.S. However, clean slate design cycles remain extremely prolonged (~5-15 years) and costly. Traditional simulations are too expensive to navigate high-dimensional spaces effectively, and typical ML surrogates lack generalizability.

Approach

- Developed COMB-FLOW, an integrated frontier AI platform with two pillars
- Multimodal AI foundation models enable **fast and predictive digital twins** for multiscale-multiphysics modeling of combustion systems
- Multi-agent AI pipelines for orchestration of end-to-end workflows **speed up design** conceptualization and optimization of novel fuel-combustion systems

Impact

- **Near-real-time forecasting** and control of combustion device operation
- Prediction and causal understanding of **safety-critical extreme/rare combustion events**
- Rapid co-design of fuels and combustion systems

Capability Leap

- Reduces fuel-engine **design cycles by ~10-100x**
- Reduces maintenance, repair, and overhaul (MRO) **costs by ~2-3x**

Aligns with S&T Challenges

- Reenvisioning Advanced Manufacturing and Industrial Productivity
- Securing U.S. Leadership in Data Centers
- Achieving AI-Driven Autonomous Laboratories

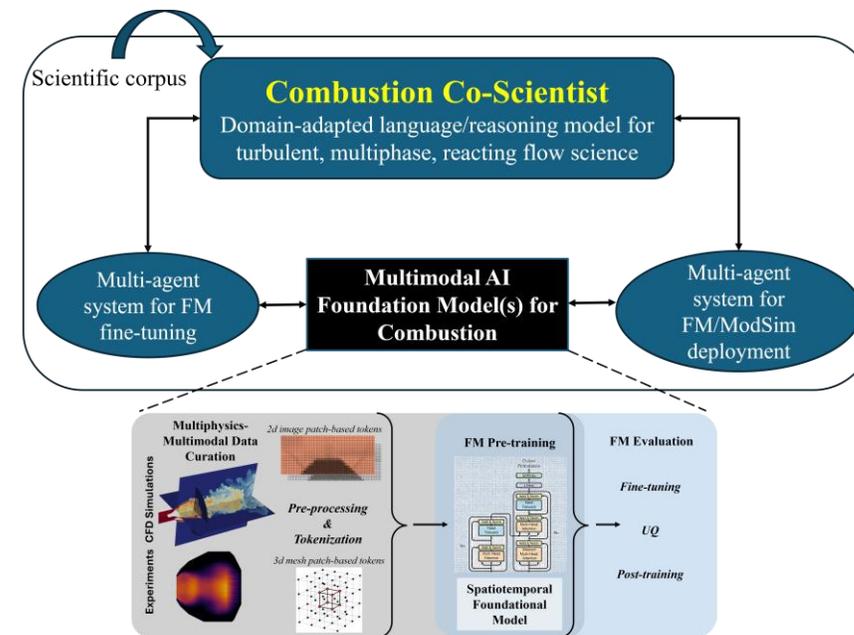


Figure: A schematic of the **Genesis Mission Seed Model Team's COMB-FLOW-AI platform**

AI-Boosted Particle Accelerator Operations & Design

Motivation

Particle accelerators are strategic infrastructure to science, industry and medicine. AI is poised to transform accelerator operations, design and research at DOE (including the largest accelerators, among the most complex machines ever built), academia, and industry.

Approach

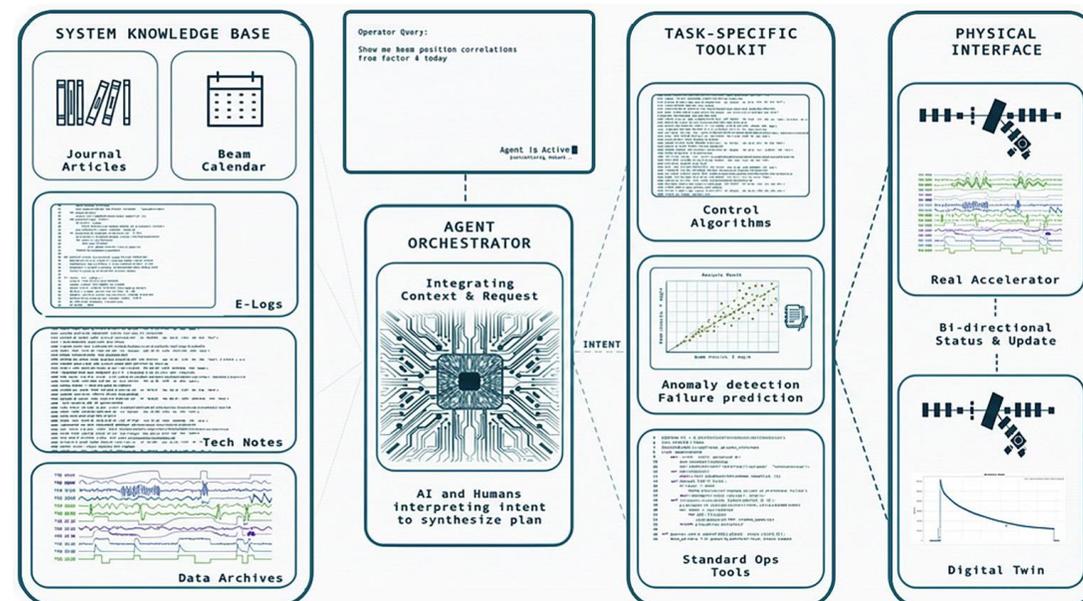
- Develop and implement language interfaces, AI assistants, and information synthesis at unprecedented scale
- Deploy AI-assisted-to-fully-autonomous design and operations with preventive fault detection and auto-fault-recovery
- Treat nation's accelerator ecosystem as a single, intelligent system

Impact

Accelerate discovery, industrial and medical applications with enhanced, intelligent, automated operations and novel accelerator designs that deliver better beams (e.g., higher energy and precision)

S&T Challenge

Aligns with “Enhancing Particle Accelerators for Discovery”



Agentic AI Osprey for accelerator operations

- Installed at 8 facilities in first phase
- Streamlines access to machine parameters for control (millions of parameters when fully deployed)
- >100x faster operator queries & response for key activities

One codebase, eight sites, Genesis Mission's infrastructure, a shared platform for transforming particle accelerators in the age of AI.

Seeing and Designing Materials at the Level of Atoms

Motivation

Nanoscale atomic arrangements and electronic details drive material functionalities. Both designing new materials and achieving autonomous science-driven AI labs require the ability to observe and predict atomic arrangements and their dynamics.

Approach

Tight integration of physics-driven simulations with AI enable unprecedented atomic structure visualization and discovery acceleration.

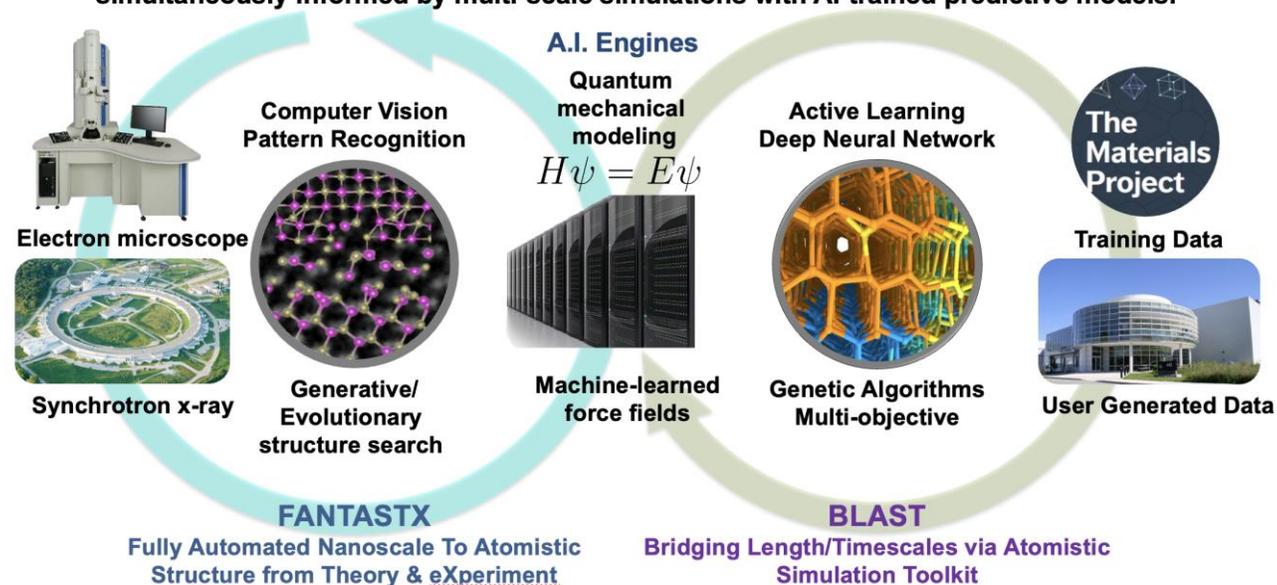
Impact

- Discover new phases of 2D materials
- Visualize single defects at interfaces
- Accelerate microelectronics and energy storage research

Aligns with S&T Challenges

- Achieving AI-Driven Autonomous Laboratories
- Designing Materials with Predictable Functionality

Vision: Automated AI-guided interpretation of experimental data and driving of experimental probes, simultaneously informed by multi-scale simulations with AI-trained predictive models.



Catalysis Foundry: AI-Powered Catalysis Discovery

Motivation

Tackle three major challenges in heterogeneous catalysis:

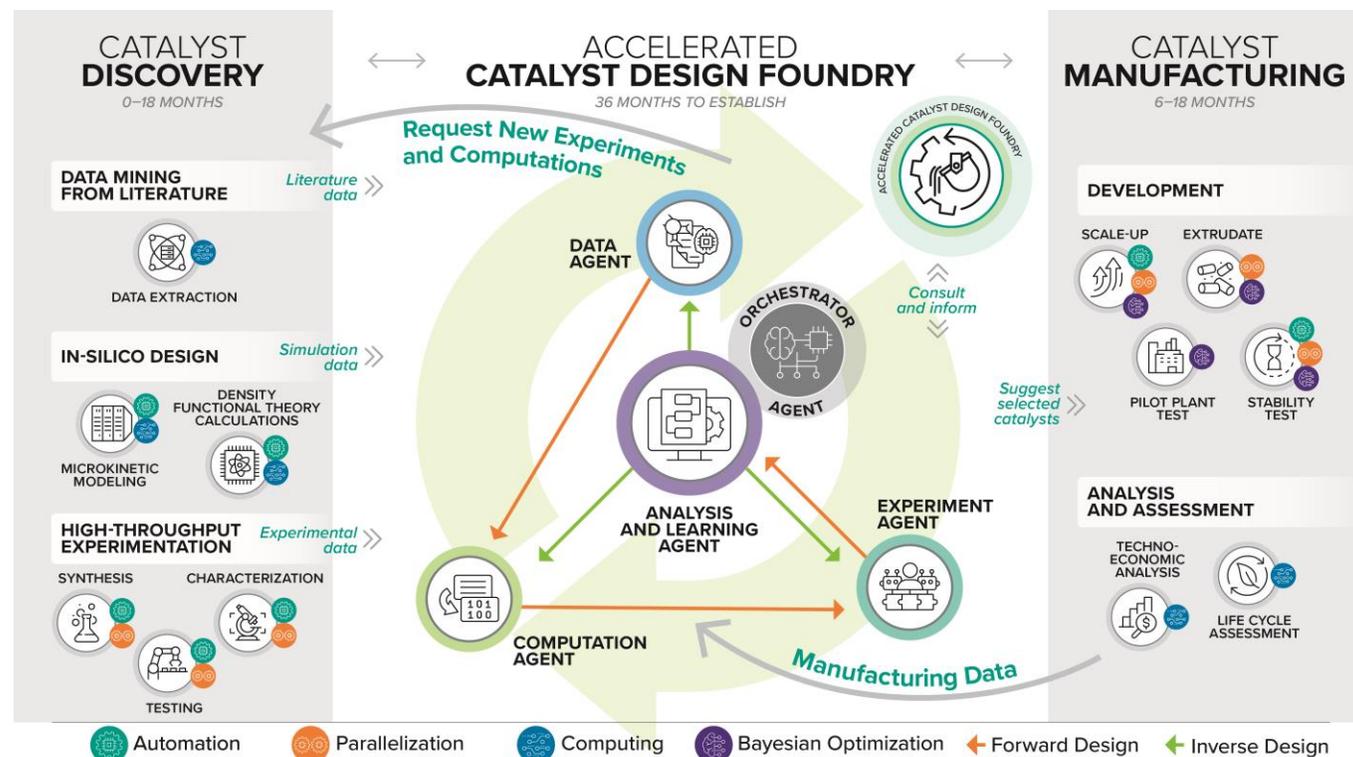
1. Identify the true active site under reaction condition
2. Solve complex structure-property relationships
3. Optimize synthesis and reaction conditions

Approach

1. Accelerated computational predictions, autonomous catalysis synthesis and testing, and automated data collection and learning under the umbrella of an agentic framework.
2. A multimodal foundation model that combines a generative AI and GNN models for forward and inverse catalyst design.

Impact

- Accelerate catalyst discovery through a physics-aware foundation model, autonomous lab and AI agents
- Create AI-ready catalysis datasets
- Establish unprecedented fundamental understanding in heterogeneous catalysis



Aligns with S&T Challenges

- Reenvisioning Advanced Manufacturing and Industrial Productivity
- Achieving AI-Driven Autonomous Laboratories
- Designing Materials with Predictable Functionality

AI Transforms Semiconductor X-ray Metrology

Motivation

Non-destructive 3D X-ray nano imaging with large fields of view provides a unique tool for pre-competitive R&D and scale-up phase

Approach

Training Vision AI models to learn the mapping between scattering and object spaces accelerates reconstruction by **100x**

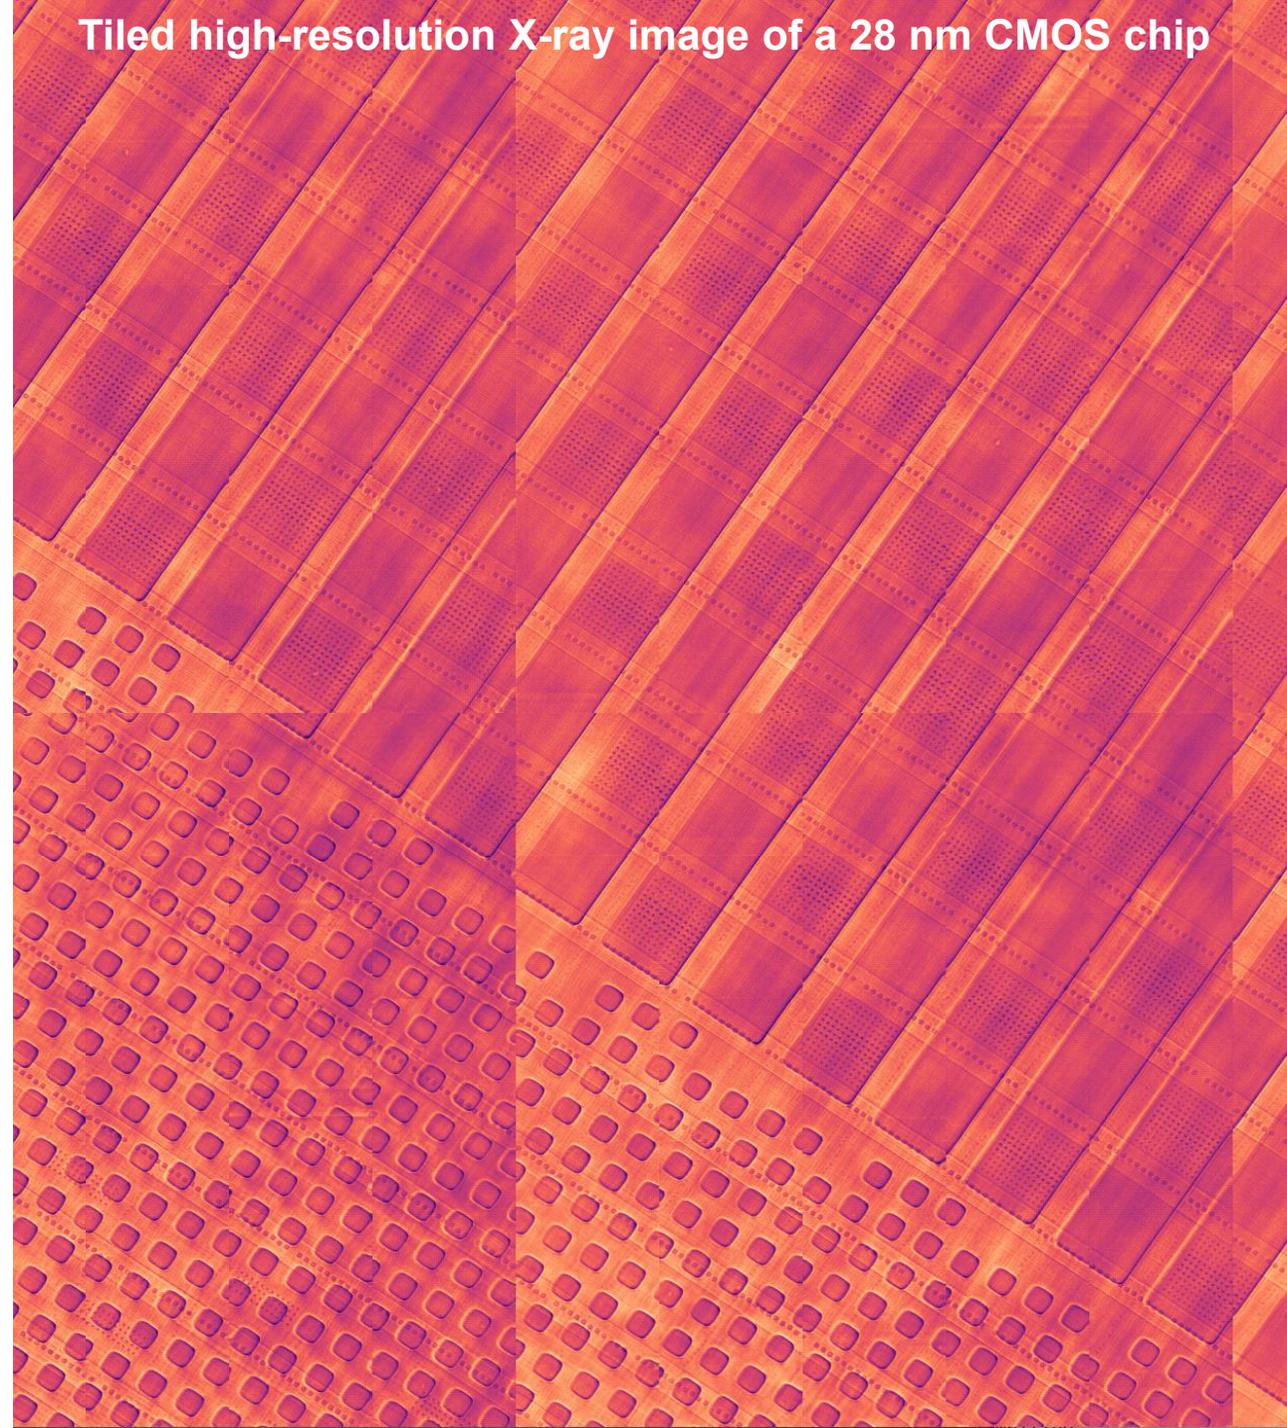
Impact

Real time inspection of microelectronics opening the door for high throughput characterization

Aligns with S&T Challenges

- Enhancing Particle Accelerators for Discovery
- Recentering Microelectronics in America

Done as part of SYNAPS-I, a collaboration between 5 light sources and neutron sources to revolutionize X-ray and neutron imaging with AI.



Programming Matter for Biotech

Motivation

Researchers led by Nobel Laureate David Baker create enzymes with native-like catalytic efficiencies from AI diffusion-flow models that are studied experimentally using X-ray crystallography at light sources

Approach

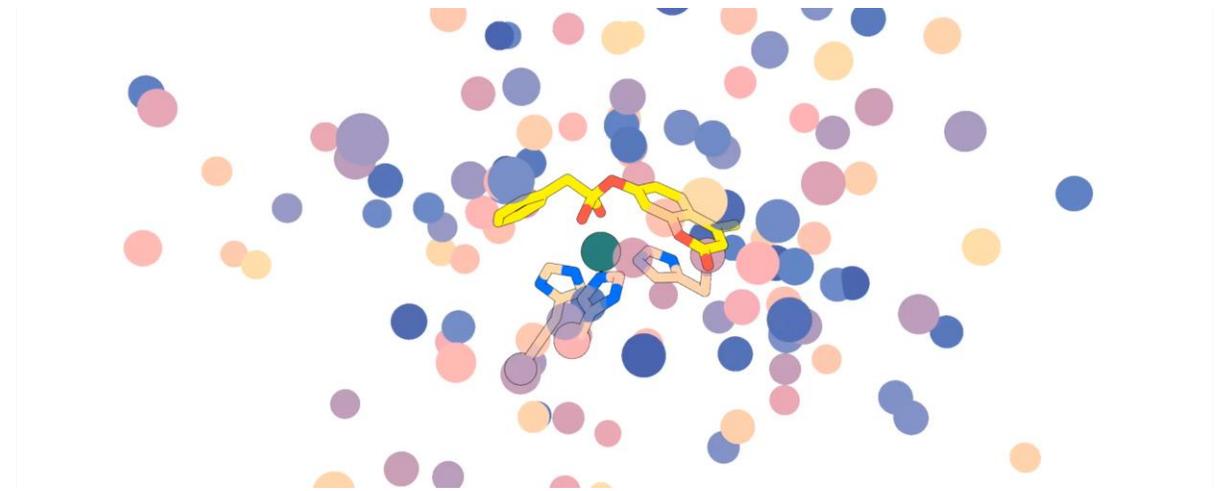
Diffusion flow methods (similar to diffusion models for image generation) are used to build artificial scaffolds that can stabilize a desired catalytic function

Impact

Matter and function are becoming fully “programmable” and tuned to the biotech needs on demand

Aligns with S&T Challenges

- Enhancing Particle Accelerators for Discovery
- Scaling the Biotechnology Revolution



Kim, D., Woodbury, S.M., Ahern, W. et al. Computational design of metallohydrolases. *Nature* 649, 246–253 (2026).

AI advantage:

AI not only transforms enzyme design but also accelerates serial crystallography at light sources

Seeing the Invisible in Additive Manufacturing

Motivation

To understand materials properties manufactured using additive methods, X-ray imaging at ultra-fast speeds are used during printing to understand defect formations

Approach

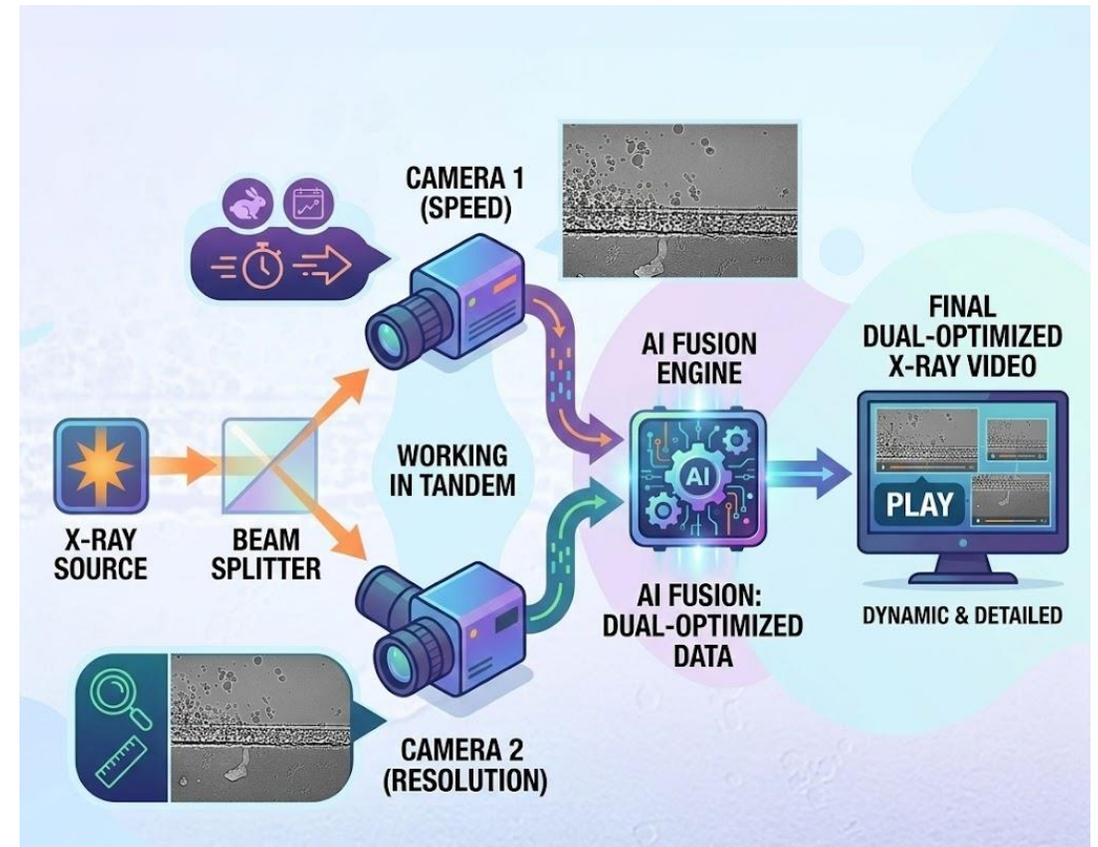
Training AI models on multiple modalities enables fusing high temporal and high spatial resolution at high fidelity images, allowing scientists to access information not seen by either method individually

Impact

- Extract quantitative information about printing process at relevant time scales and resolutions
- Build models that accelerate materials printing and validation of components

S&T Challenge:

Aligns with “Enhancing Particle Accelerators for Discovery”

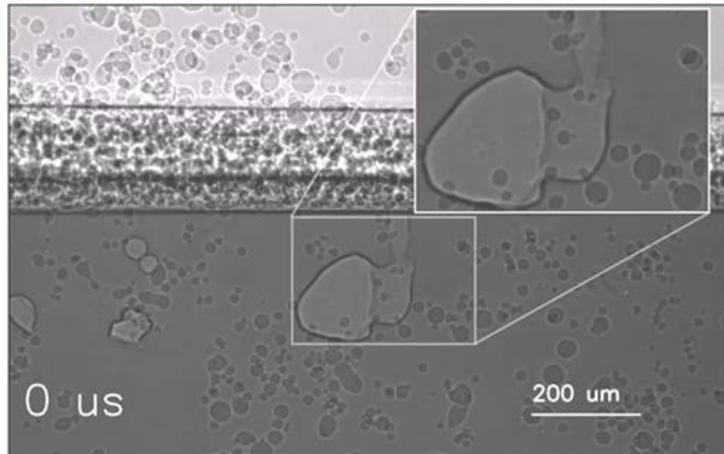


AI speed-up:
14x in imaging speed (1MHz) at high resolution

Seeing the Invisible in Additive Manufacturing

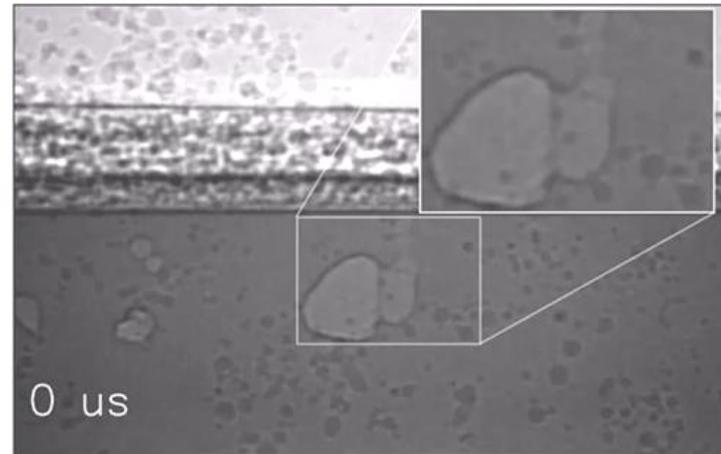
Deployed at APS 32-ID beamline and used for experiments

Slow – High Res.



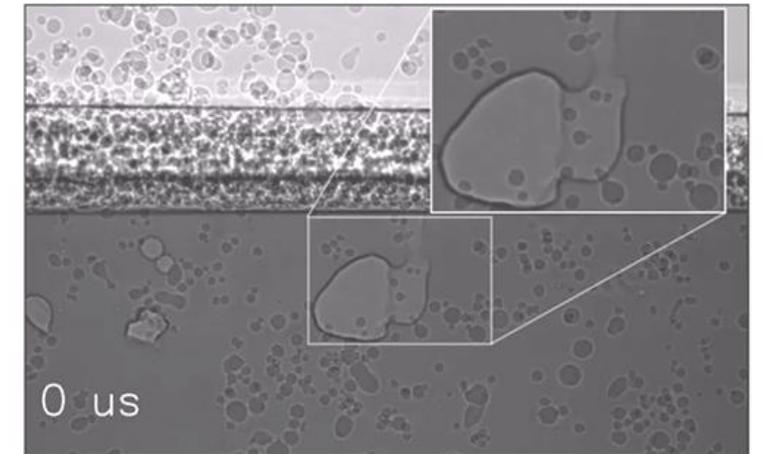
Phantom camera (1.31 $\mu\text{m}/\text{px}$, 71.43 kHz)

Fast – Low Res.



Shimadzu camera (3.82 $\mu\text{m}/\text{px}$, 1 MHz)

AI: Fast + High Res.



Reconstruction (1.31 $\mu\text{m}/\text{px}$, 1 MHz)

CM2US: Critical Minerals and Materials to Unlock Supply

Motivation

Securing the critical minerals supply chain is one of the defining economic and national security challenges of this decade. Today, decisions about mineral resources, materials innovation, and supply chain investment are made using fragmented data and disconnected models, siloing efforts and expertise.

Approach

Using AI and advanced analytics, CM2US will be able to connect critical areas of the ecosystem: Resource Discovery – include identifying and characterizing mineral resources; Material Discovery – include accelerating the development of critical materials; Supply Chain Intelligence – include understanding economic viability, risks, and dependencies, and identifying strategic actions. By linking these domains, CM2US can uncover hidden relationships, create feedback loops between science, technology, and markets, and generate insights that help guide smarter investments, faster innovation, and more secure supply chains.

Impact

Improved capability to anticipate risks, identify opportunities, and accelerate technologies critical for energy security and industrial competitiveness

S&T Challenge

Aligns with “Securing America’s Critical Minerals Supply”



Imaging the Proton's Interior Quark-Gluon Landscape

Motivation

Facilities like Jefferson Lab and the forthcoming Electron-Ion Collider produce Petabytes of data that contains information on the quarks and gluons inside protons and nuclei. This information is not direct, and a grand challenge is to infer from this data images of the interior quark-gluon structure of visible matter.

Approach

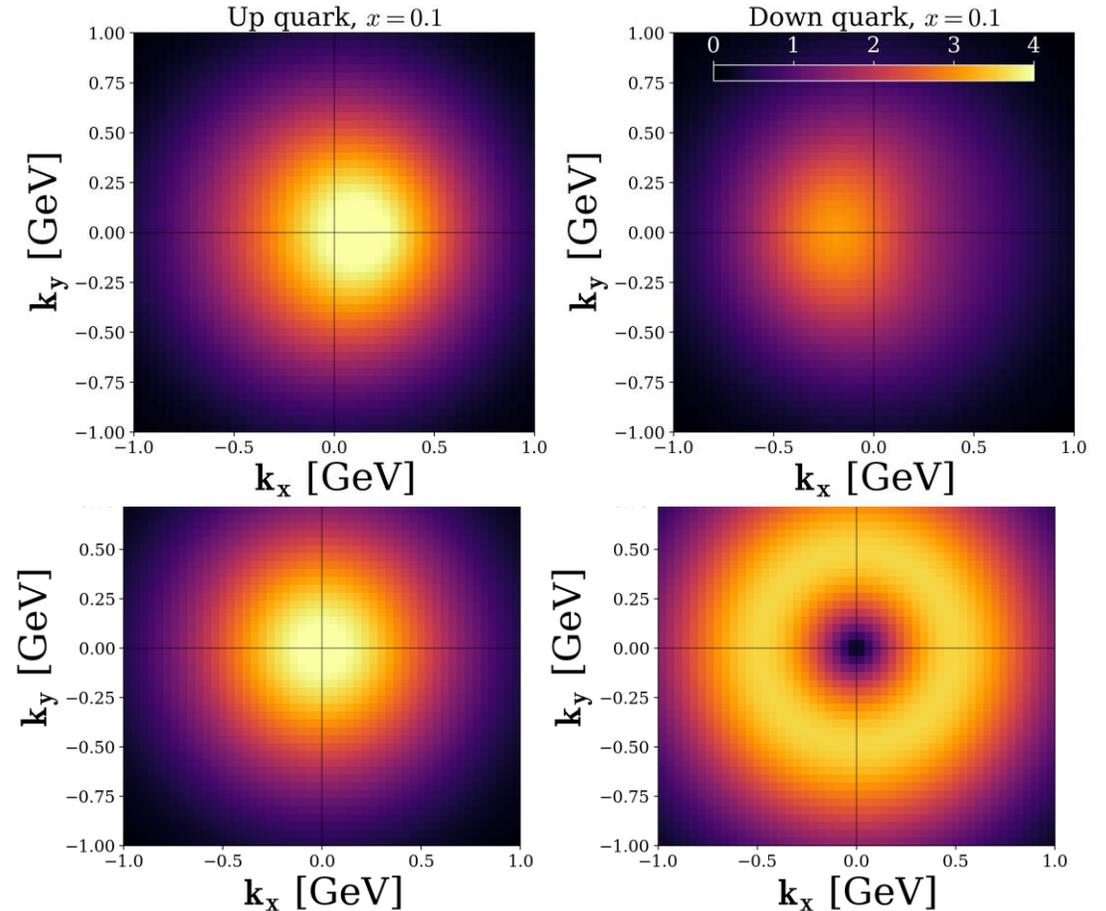
Build an agentic AI system to accelerate this inference process to real-time in a way that integrates literature-grounded reasoning, statistical inference, data collection, visualization, and anomaly detection into a unified workflow.

Impact

Greatly accelerate time between data taking and discovery. The AI system will be able to make assertions on the data, develop new physics theories, and test those theories through data analysis.

S&T Challenge

Aligns with “Unifying Physics from Quarks to the Cosmos”



Momentum Space Images of the Quarks in a Proton obtained from data using our AI Agentic System

Exploring the Frontier-e Simulation with Agents

Motivation

Large-scale, many PB sized cosmological simulations hold a wealth of new information about structure formation in the universe. Exploring these simulations is a complex task and often, even the questions are not well defined.

Approach

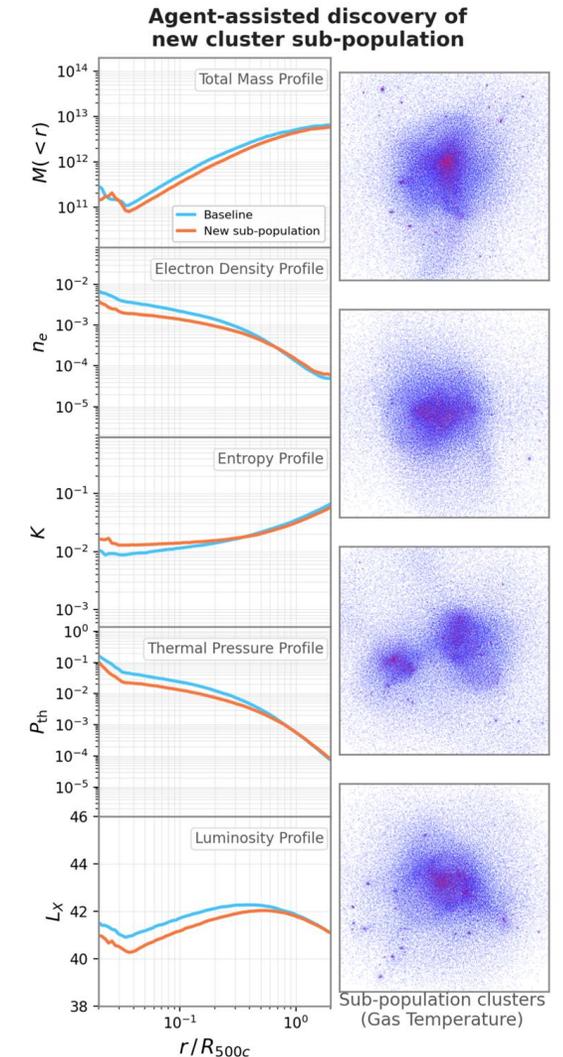
We have built an agentic system to explore such simulations. The system integrates cloud-based reasoning models, domain-specific foundation models and services such as the AmSC Intelligent Interfaces as well as a set of reliable, scalable analysis tools.

Impact

The agentic system not only speeds up scientific analysis by an order of magnitude but has also already delivered new scientific insights.

S&T Challenge

Aligns with “Unifying Physics from Quarks to the Cosmos”



Preparing QCD Data for Foundation Models

Motivation

Next-generation nuclear physics experiments generate uniquely rich, high-dimensional data, but the lack of standardized, AI-ready data structures limits their use in foundation models. Common data pipelines and metadata are essential to enable scalable AI across experiments and let foundation models learn physics from a broad body of experimental data.

Approach

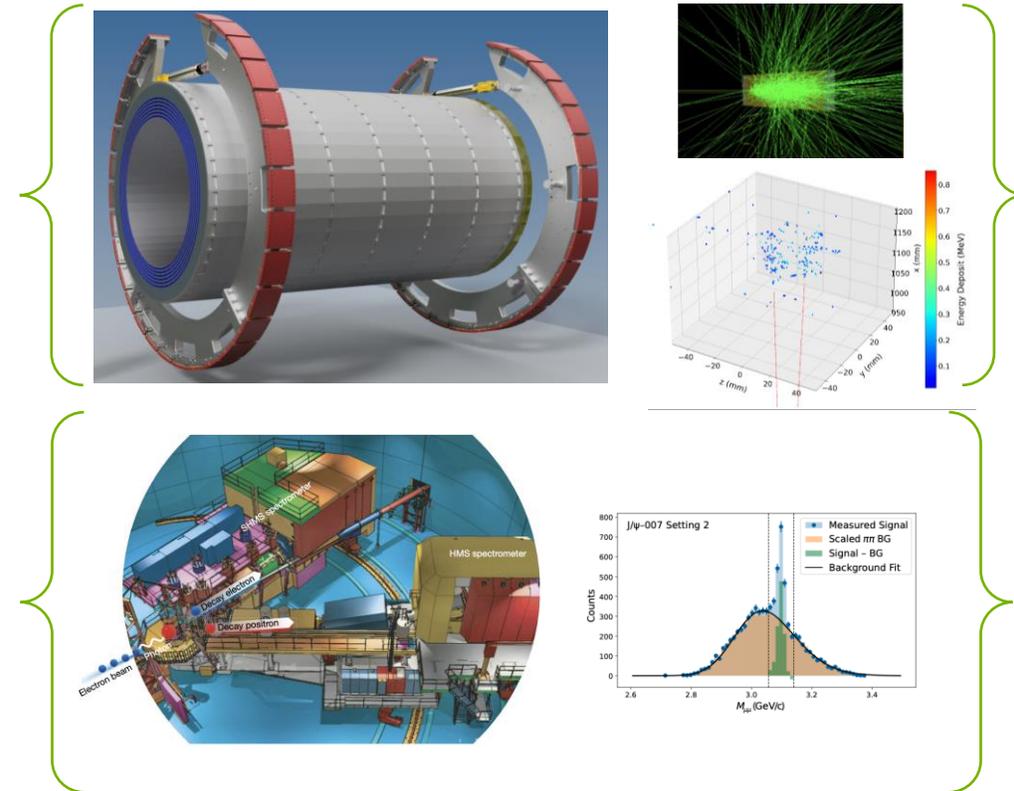
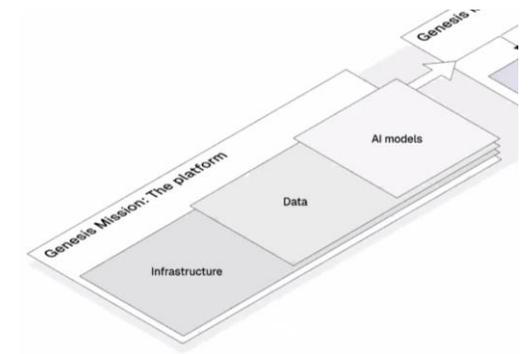
- Use the Argonne-designed ePIC Barrel Imaging Calorimeter as a testbed for developing AI-ready data standards that combine time-resolved detector signals with spatial and energy-resolved particle-shower images.
- Build and release a realistic simulated EIC dataset with rich metadata, including detector geometry, expert knowledge, and validation tools.
- Curate and release a real dataset from the Hall C J/ψ -007 experiment, demonstrating the standard scales from complex to simpler detectors and providing benchmark real J/ψ measurements.

Impact

- Establish a common, open framework for high-dimensional nuclear physics data, making future experiments ready for foundation-model workflows.
- Produce curated simulated and real datasets with metadata, enabling reproducible AI benchmarks and opening legacy data to the community.
- Accelerate discovery by allowing AI models to reason across experiments and identify new phenomena.

S&T Challenge

Aligns with “Unifying Physics from Quarks to the Cosmos”



“Preparing QCD Data for Foundation Models” is a multifaceted joint effort between BNL, Argonne, LBNL, ORNL, and Jefferson Lab as part of the American Science Cloud (AmSC) Data Providers Program (DaPP).

Reusable, Domain-Agnostic AI Co-Scientist Platform

Motivation

Multi-agent AI systems have shown great promise for complex reasoning in scientific workflows, particularly those that have general-purpose capability that emphasizes orchestration of tools and integration of retrieval-augmented generation in scientific workflows.

Approach

- Developed AISAC (AI Scientific Assistant Core), a hierarchical multi-agent architecture enabling **complex task orchestration** and long-context structured reasoning
- Integration of RAG, external data sources, and computational tools with Model Context Protocol (MCP) to connect to external agents
- Designed to retain faithful trace for **full human audit** to build trust

Impact

- Adopted across **3+ projects** as a common capability including Genesis Mission's COMB-FLOW-AI, AI for Energy Process Safety with USDOT, and FM4R2P
- Supports consistent reasoning and tool orchestration across workflows, reducing human effort and **increasing throughput ~10-50x**
- Reusable AI Co-scientist foundation **adaptable to multiple scientific domains**

Capability Leap

- Bridges experimental agentic AI and production scientific practice — giving domain scientists governed, auditable multi-agent workflows without having to develop them

S&T Challenge

Aligns with “Achieving AI-Driven Autonomous Laboratories”

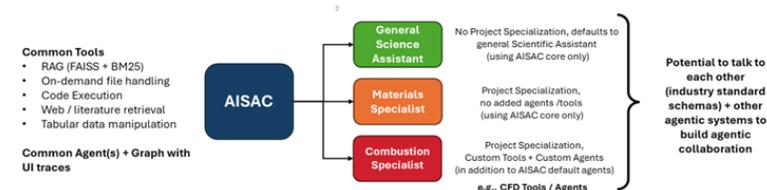
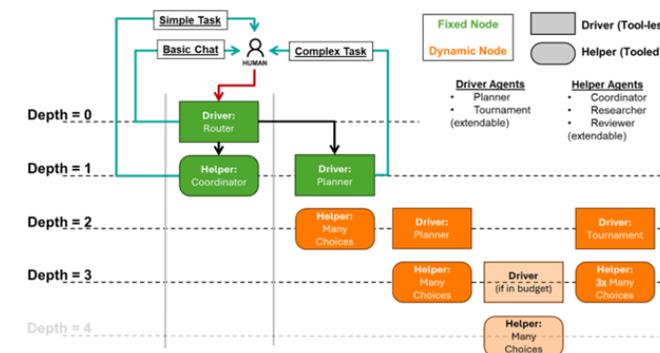
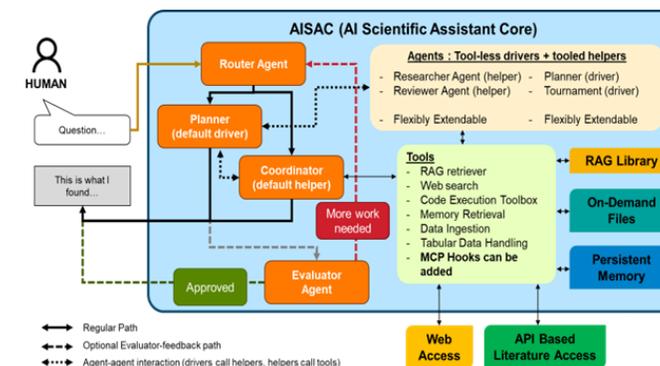


Figure: A schematic of the AISAC platform

AI for Transportation: Accelerating U.S. Economic Growth through Enhanced Mobility

Motivation

Transportation drives **30% of U.S. energy** demand and costs the economy over **\$800B annually** in congestion and inefficiency. Today's planning tools cannot jointly address **energy, freight, and supply chain** challenges at the speed decision-makers require.

Approach

- Developed MOVE-FM (Mobility Optimization & Vehicle-Energy Foundation Model), a **hybrid AI model** built on Argonne's multi-scale transportation simulation ecosystem
- Connect vehicle, fleet, freight, and system levels
- Enable **real-time scenario analysis** without rerunning costly simulations
- Leverage **Argonne's unique models and datasets** unavailable elsewhere

Impact

- A **first-of-its-kind foundation model** enabling real-time, cross-scale scenario analysis of mobility, freight, and energy impacts without rerunning costly simulations
- Enables **MPOs and state DOTs** to evaluate energy, congestion, and infrastructure trade-offs in real time
- Strengthens **domestic supply chain resilience** through freight and logistics optimization

Capability Leap

- **>100x faster scenario analysis**, planning goes from months to minutes

Aligns with S&T Challenges

- Securing America's Critical Minerals Supply
- Scaling the Grid to Power the American Economy

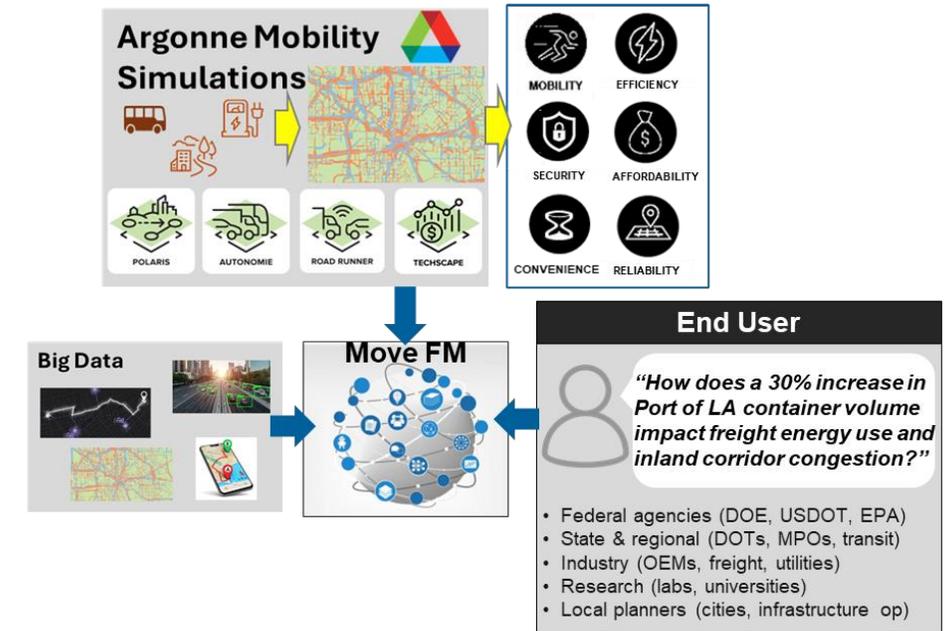


Figure: A schematic of the MOVE-FM platform

Enhancing Electrocatalyst Activity

Motivation

Proton-exchange membrane fuel cells (PEMFCs) efficiently convert the chemical energy in fuel to electricity, but require platinum electrocatalysts. Replacement of platinum, a critical material, with Earth-abundant metals will enable wide-spread use of this technology

Approach

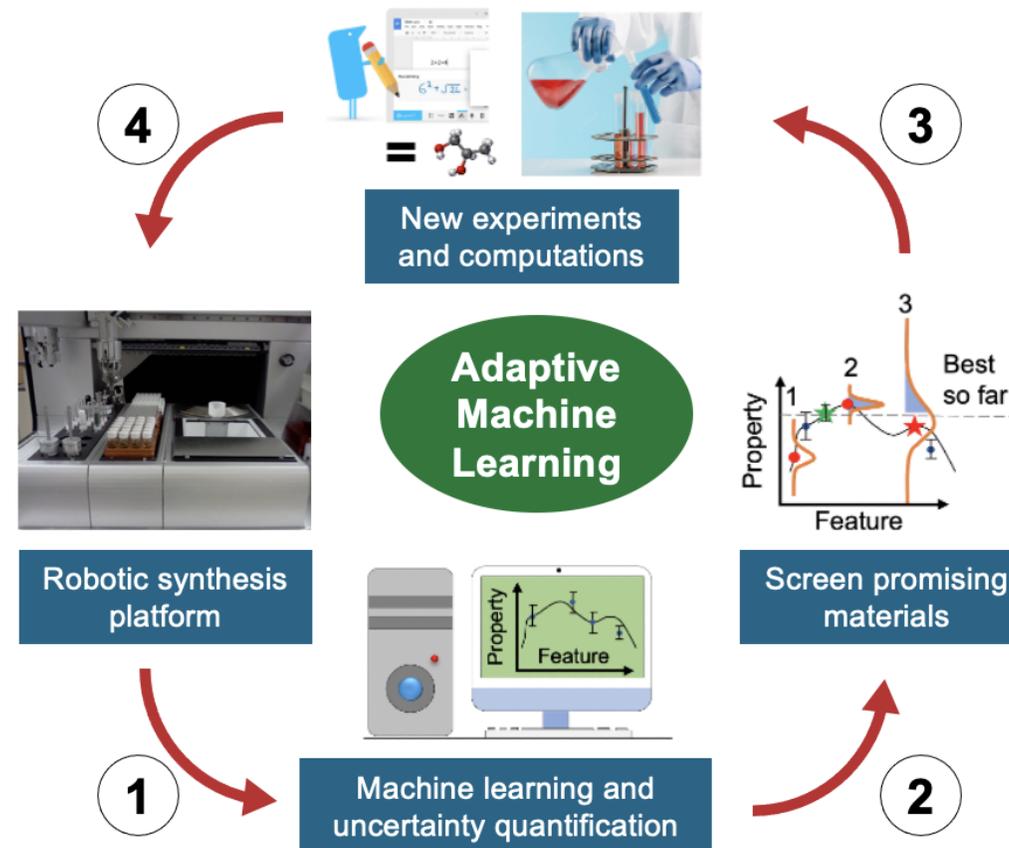
Using AI models trained on electrocatalytic activity and X-ray spectroscopy data for FeNC catalysts synthesized using high-throughput approaches to guide synthesis parameters.

Impact

- Accelerate the development of more active and durable Pt-free PEMFC catalysts
- Identify the spectroscopic signatures correlated with catalyst activity, thus identifying active sites

Aligns with S&T Challenges

- Designing Materials with Predictable Functionality
- Recentring Microelectronics in America



1. M. R. Karim *et al.*, *ACS Appl. Energy Mater.* 3, 9083 (2020).
2. W. J. M. Kort-Kamp *et al.*, *J. Power Sources* 559, 232583 (2023).

AI-Driven Co-design for Microelectronics

Motivation

End-to-end co-design can transform the way we conduct microelectronics research to unlock new pathways for materials research, computer architecture and systems design. Co-design of the full design stack is challenging because of the need to develop direct relationships between nonadjacent levels of stack – materials, devices, architectures, systems, system software, algorithms, applications, and experimental science workflows. Developing such relationships can require years of data collection and analysis.

Approach

- Use frontier AI within an agentic orchestrated framework to model multi-level relationships between the different levels of the microelectronics design stack
- Start with the initial challenge of software refactoring and hardware synthesis co-design, and then move to other levels
- Use the data generated from the agentic framework to develop a digital twin of the stack and fine tune a foundation model specific to microelectronics

Impact

- Accelerate the co-design iterations to provide deeper and broader explorations of the computing design space
- Significantly reduce the overall design cycle time by a least two orders of magnitude
- Change the way we conduct microelectronics research to make it possible to easily explore relations among any levels of the design stack, thereby significantly increasing productivity
- Significantly reduce the time to discover the materials needed for the devices to address the needs of DOE applications by at least of factor of two

S&T Challenge

Aligns with “Recentring Microelectronics in America”

Aim: AI-Driven Co-Design for Microelectronics



DOI: <https://doi.org/10.2172/1545772>

Initial Challenge: Software Refactoring and Hardware Synthesis Co-Design



AlphaFold for Microelectronics

Motivation

Just as AlphaFold revolutionized biology by predicting protein structures, our project aims to predict device performance directly from atomistic structures.

Approach

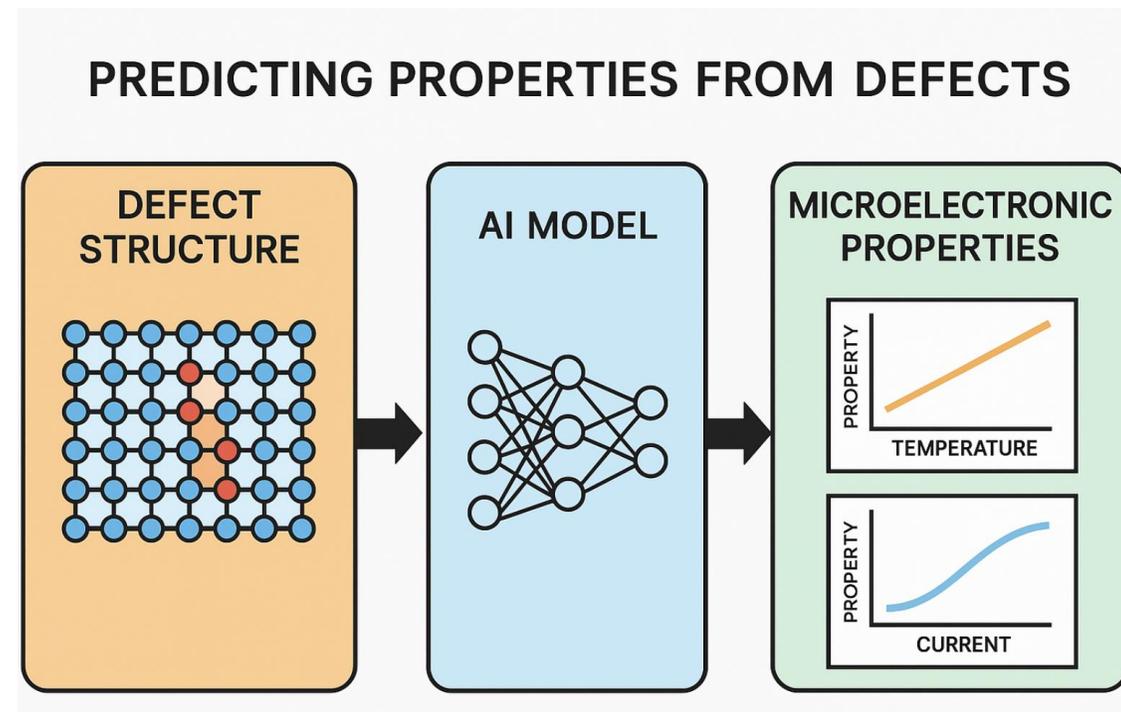
We are building a high-performance AI ecosystem called “The **Materials Discovery Cloud (MDC)**” that unifies multimodal data streams - integrating high-fidelity simulations, experimental data from facilities like APS and CNM, and scientific literature to architect next-generation microelectronics

Impact

Enables rapid prediction of defect-driven device performance, accelerating the design of reliable, efficient next-generation microelectronics.

Aligns with S&T Challenges

- Recentring Microelectronics in America
- Achieving AI-Driven Autonomous Laboratories
- Designing Materials with Predictable Functionality



AERIS: Argonne's Earth Systems Model for Reliable and Skillful Predictions

Motivation

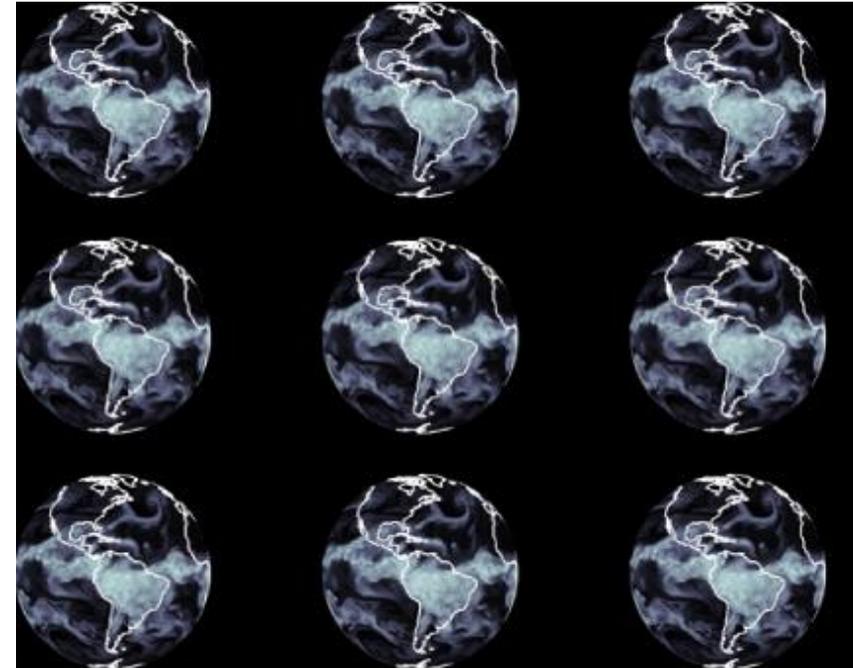
- Scalable AI/ML models present an immediate opportunity to advance weather forecasting and decision-making
- Argonne domain expertise in atmospheric science enabled the lab to accelerate AI/ML model development and deployment at scale on high-performance computing
- Immediate applications focus on subseasonal-to-seasonal (S2S) weather forecasting (e.g., 90 days) using large, probabilistic ensembles approach

Approach

1. Develop state-of-the-art shifting-windows diffusion transformer machine learning model
2. Forecast skillfully to 14 days, stable out to 90 days, delivering potential for subseasonal-to-seasonal (S2S) forecasts
3. Produce largest AI for science model at time of development, scaling to 80B parameters
4. Train on 120,960 GPUs on Argonne's Aurora supercomputer

Impact

- Generate large ensemble, probabilistic global forecasts (current global resolution of 30km)
- Leverage novel forecasting capability to enhance energy sector operational planning
- Account for seasonal temperature and precipitation variations affecting energy demand, load forecasting, hydropower generation, thermal cooling water availability



AERIS 1B Swin
Initialized: 2020-01-01T0z

Forecast: +90 days h
Variable: Specific Humidity (Q)

AERIS was named a finalist for the Association for Computing Machinery's (ACM) Gordon Bell Prize in high-performance computing

S&T Challenge

Aligns with "Scaling the Grid to Power the American Economy"

Argonne Helps Nuclear Industry Embrace AI and Reduce Delays

Motivation

- Use AI to improve safety, maintenance planning, and plant efficiency
- Make AI deployable in regulated nuclear environments
- Cut licensing delays without weakening oversight or trust

Approach

- Validate AI in a test facility with full safety review
- Build the Regulatory Context Protocol (RCP) to automate applicant-regulator exchanges
- Deploy PRO-AID digital twins for early fault diagnosis

Impact

- Better regulatory information quality and fewer workflow delays
- Faster advanced reactor licensing and time-to-grid
- Earlier detection of sensor bias and cooling failures

S&T Challenge

Aligns with “Delivering Nuclear Energy that is Faster, Safer, Cheaper”



The Mechanisms Engineering Test Loop (METL) facility where Parameter-Free Reasoning Operator for Automated Identification and Diagnosis (PRO-AID) was successfully tested.

AI validation

RCP automation

PRO-AID monitoring

“These initiatives are vital steps toward ensuring that AI technologies can be successfully and safely integrated into the nuclear industry.”

— Rick Vilim, Argonne senior nuclear engineer

Together, the projects link regulated AI validation, faster licensing, and earlier fault detection in one deployable pathway.

AI4HPC: AI Workflows for Scientific Software Development

Motivation

Scientific discovery increasingly depends on large, complex software systems running on leadership-class supercomputers. These codes, often millions of lines of legacy software, must evolve to incorporate new models, architectures, and algorithms. Meanwhile, advances in AI-assisted coding are transforming software development but do not directly translate to scientific computing, where correctness, reproducibility, and extreme performance are essential.

Approach

AI4HPC is developing an integrated ecosystem of AI tools that embed assistance across the scientific software lifecycle:

- Generation: creating and integrating new capabilities
- Translation: modernizing large legacy code bases
- Porting: adapting applications to new hardware platforms (e.g. GPUs)
- Optimization: improving performance through AI-assisted algorithmic engineering

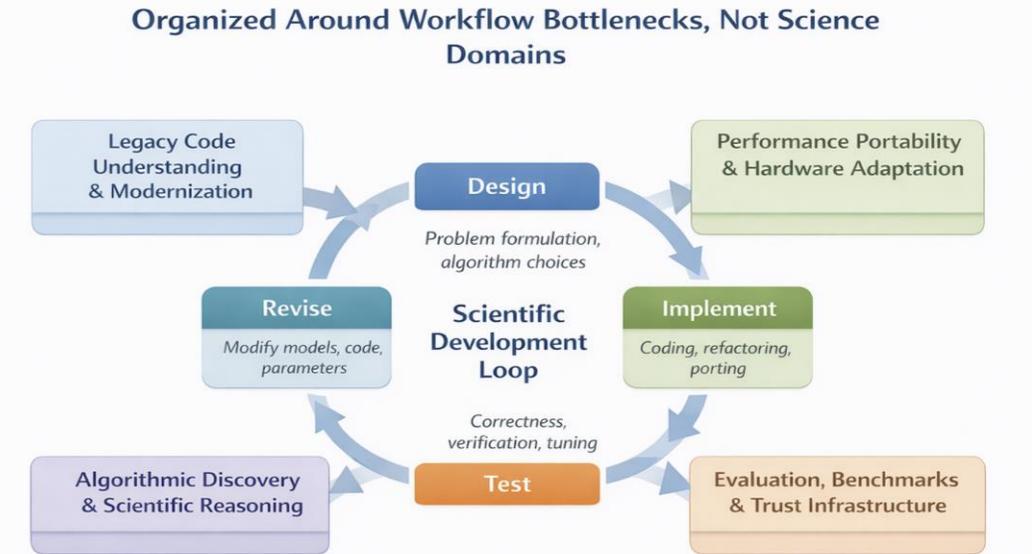
Impact

AI4HPC increases developer productivity and accelerates software evolution.

- Dramatically less effort to generate, modernize, port, and optimize code
- Faster adoption of new architectures
- Shorter time-to-solution for critical problems

S&T Challenge

Aligns with “Securing Leadership in U.S. Data Centers”



AI integrated across the scientific development loop, addressing critical bottlenecks in modernization, porting, optimization, and validation.

In addition to improving existing workflows, AI4HPC is exploring AI-assisted algorithmic discovery and co-design, enabling new approaches to scientific computing beyond code generation and optimization. These capabilities are being developed as an integrated ecosystem within the Genesis platform, leveraging the Genesis Model Consortium and the American Science Cloud for access to tools, hardware, and advanced AI models.

Classifying Heterogeneous Dynamics in Quantum Materials

Motivation

Understanding the spatially non-uniform material evolution under external excitations requires identifying regions with different dynamics

Approach

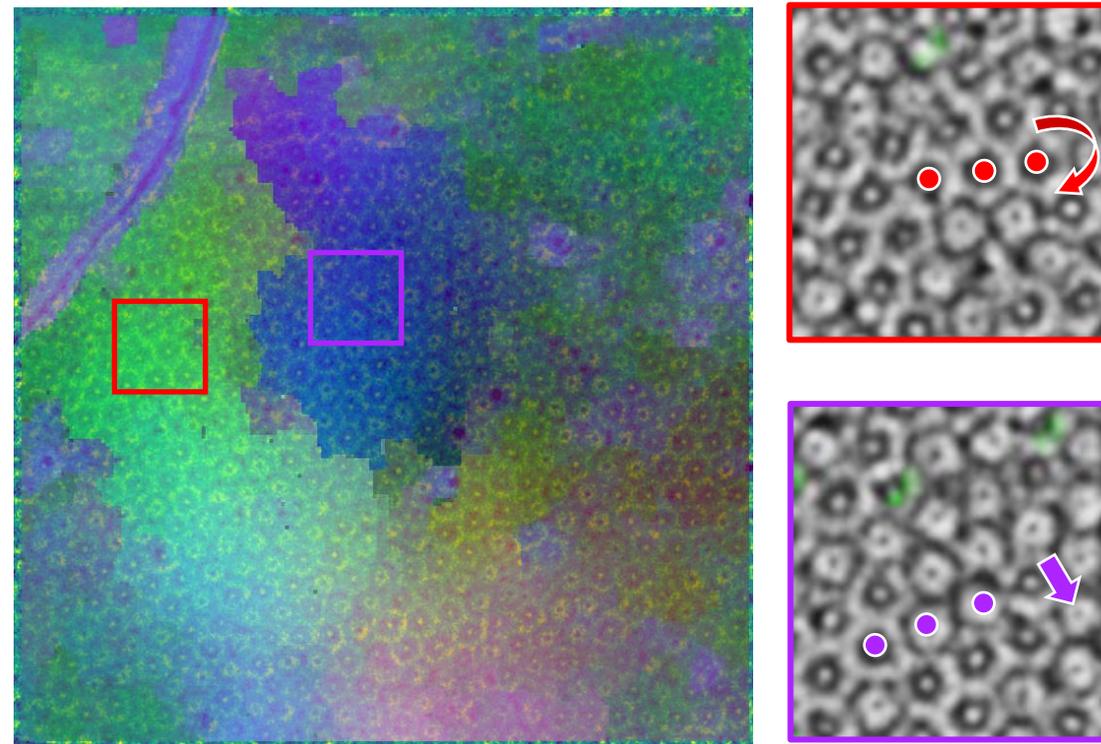
AI model with attention tuned to the time domain and combining featurization and dimensionality reduction allows researchers to understand high-throughput microscopy data near real-time

Impact

- AI acts as a co-scientist and guides decision-making as data comes in
- AI enables discovering regions of the material with novel or desired functionalities for microelectronics
- Approach applicable across microscopy modalities and transferrable to other realms of science with spatially dependent dynamics e.g., critical materials extraction

Aligns with S&T Challenges

- Designing Materials with Functional Predictability
- Recentring Microelectronics in America



AI helped to resolve spatially different magnetic skyrmion dynamics that cannot be identified in conventional approaches.

AI Co-Scientist for Industrial Carbon Conversion

Motivation

Industrial Resources-to-Products (R2P) processes – such as carbon conversion, and reactive media design – remain energy-intensive and empirically driven. Current workflows cannot efficiently connect **molecular chemistry, process dynamics, and real-world operational data**, which slows discovery and optimization.

Approach

- Developed FM4R2P, a **multi-agent AI Co-Scientist** that reasons across scales to accelerate materials and process design while reducing time and cost
- Autonomously analyzes process data (e.g., from National Carbon Capture Center) and builds surrogates of the process
- Building out to launch MD simulations on ALCF infrastructure for autonomous numerical experimentation for hypothesis validation

Impact

- Up to **5x reduction in process-development time** through autonomous hypothesis generation and simulation
- Expected **>15% improvement in energy and material efficiency** via optimized solvent selection / process control
- Improvement in media for industrial separation

Capability Leap

- Faster design and synthesis of carbon conversion media, with AI driven research and simulations, which improves industrial efficiency

S&T Challenge

Aligns with “Reenvisioning Advanced Manufacturing and Industrial Productivity”

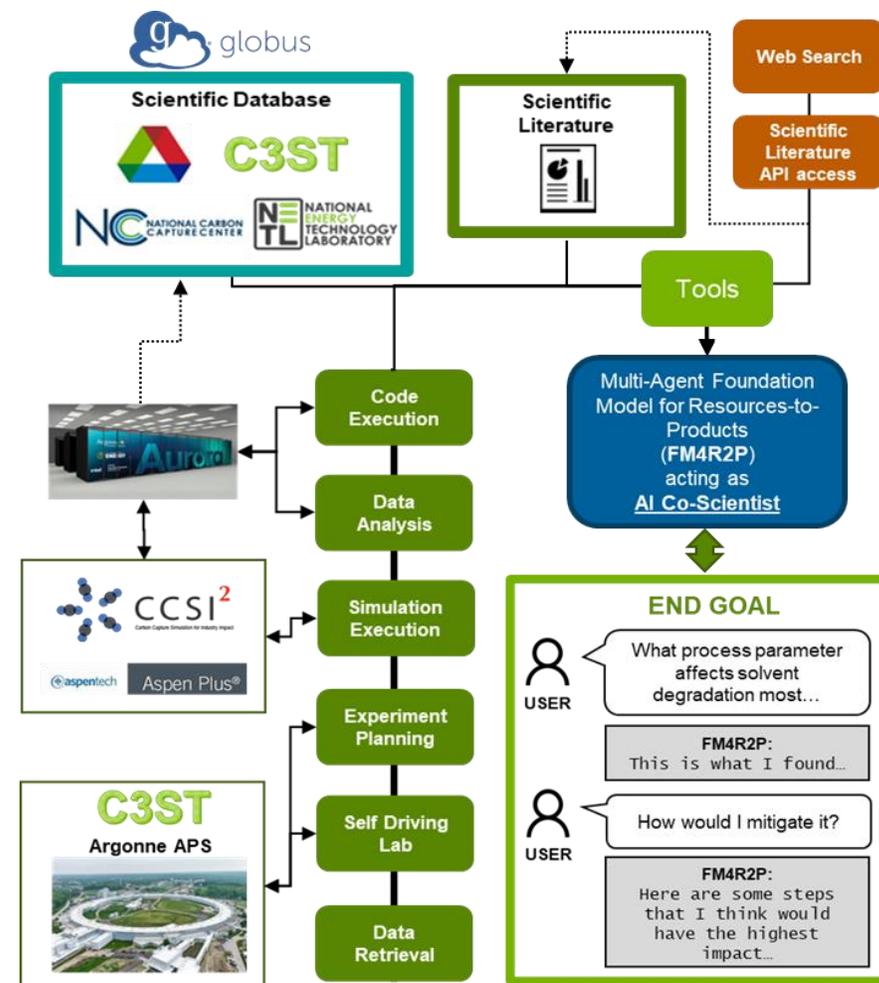


Figure: A schematic of the FM4R2P platform

Workforce Development

Building an AI-enabled STEM workforce

Our Future Focus for AI Integration into Workforce Development

STEM career pathways, from middle school to grad school

Community College Collaboration with 9 local institutions

Activating the regional STEM ecosystem of educators

Argonne's FY25 STEM engagements

- 8,500+ students and family members involved through STEM
- 6,300+ outreach
- 1,200+ pre-college students had skill-building experiences
- 2,900+ college interns worked at Argonne
- college students in Argonne-managed DOE competitions



Argonne's March 2026 AI STEM jam brought together 125 educators to experience firsthand how AI is transforming science and technology and to advance their thinking about the resulting needs in STEM education. Pictured speaker is John Domyancich, Argonne's Manager of STEM Education Learning.

Scientific Imaging Beyond Current Limits

Motivation: To understand today's complex materials and biological systems, we need imaging tools that go beyond current technical limits — extracting insight from incomplete data and integrating knowledge across techniques and disciplines.

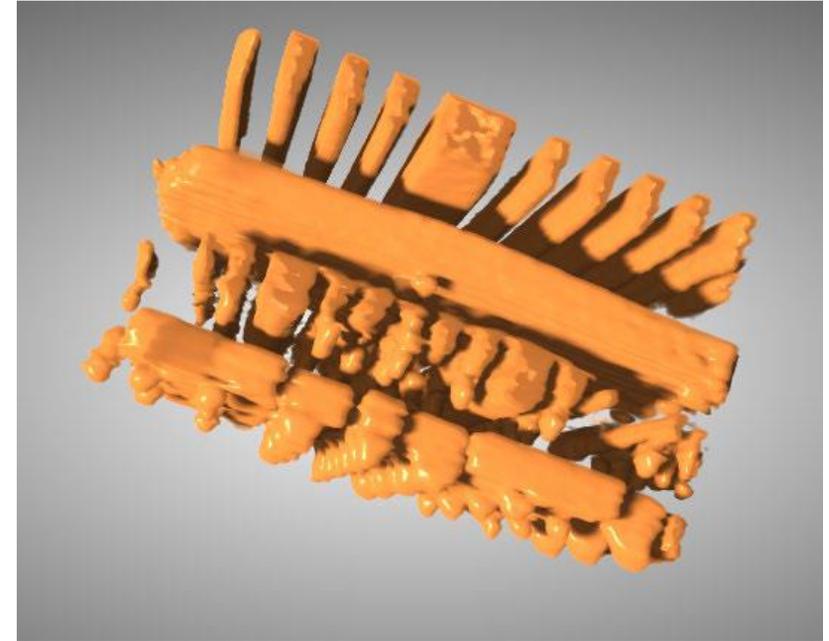
Approach: AI models trained on scientific data allow researchers to study materials and biological systems in real time while experiments are underway.

Impacts:

- Extract meaningful information even under the most challenging experimental conditions
- Reveal patterns and relationships that are difficult for humans to detect
- Accelerate discovery through increased productivity and efficiency

S&T Challenges

Aligns with “Enhancing Particle Accelerators for Discovery” and “Recentering Microelectronics in America”



Advanced microelectronics fabricated by IBM and studied at BNL's light source

AI Advancements at DOE Light Sources

- 10× faster X-ray image reconstruction using AI
- 10× reduction in data required for 3D X-ray imaging
- X-ray image analysis goes from weeks of human effort to minutes of computing time

Powering the Future of Biotechnology

Motivation

Drug discovery relies on manual experiments and slow data collection. AI models lack sufficient high-quality structural data to make reliable predictions.

Approach

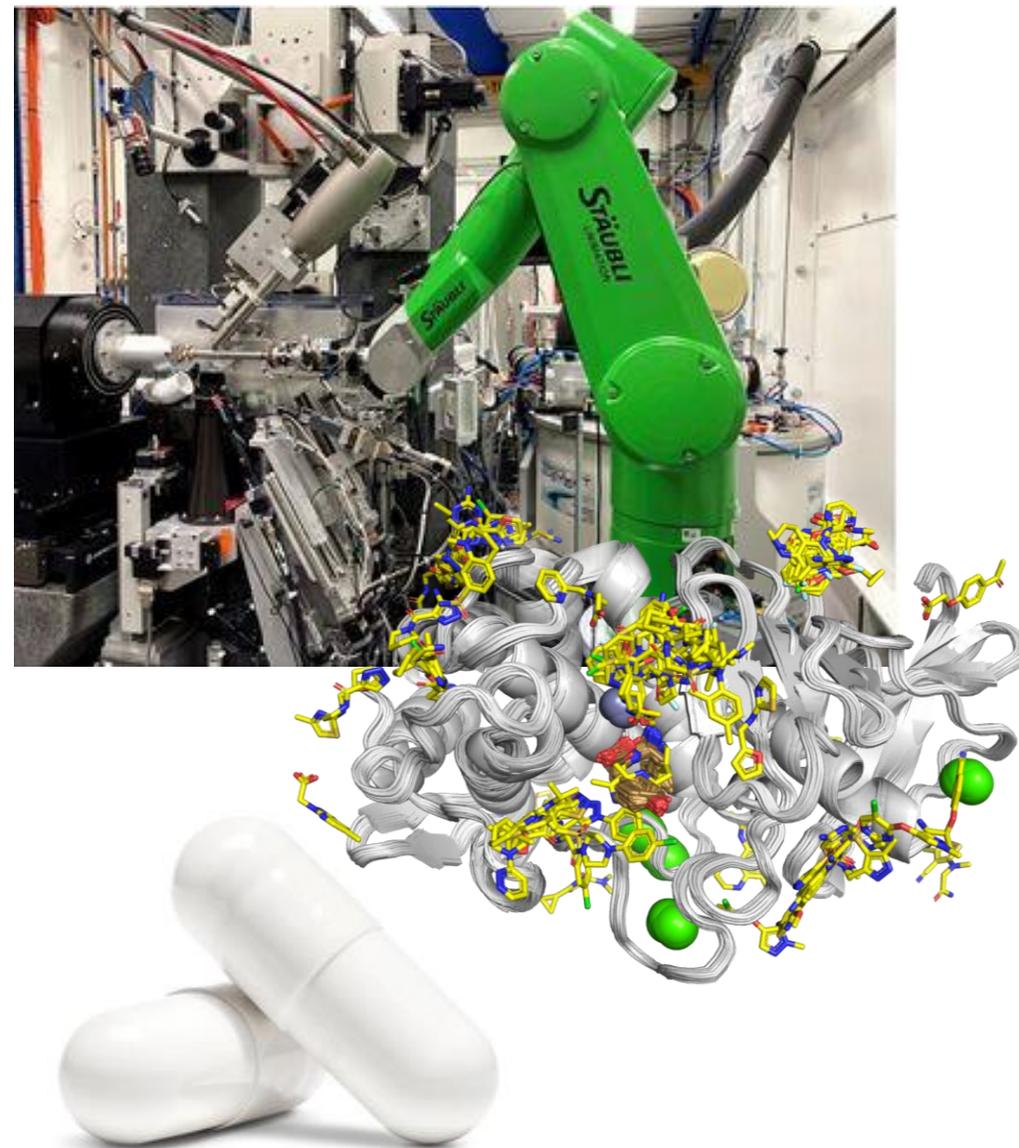
- Integration of AI, robotics, and X-ray science into a continuously operating discovery platform
- 50% throughput increase in FY25 vs. FY24 (70,000+ samples/year), with a clear pathway to an additional 50% gain

Impact

- Generate large, high-quality datasets to train AI models
- Accelerate development of treatments for cancer and other diseases
- Strengthen U.S. leadership in biotechnology and artificial intelligence

S&T Challenges

Aligns with “Enhancing Particle Accelerators for Discovery” and “Scaling the Biotechnology Revolution”



AI-Accelerated Grid Expansion and Interconnection of Large Loads

Motivation: Need to drastically increase our ability to calculate grid expansion and interconnection scenarios (e.g., for large data center loads), which are limited today to 1-2 location scenarios due to computational constraints of existing power flow solvers.

Approach:

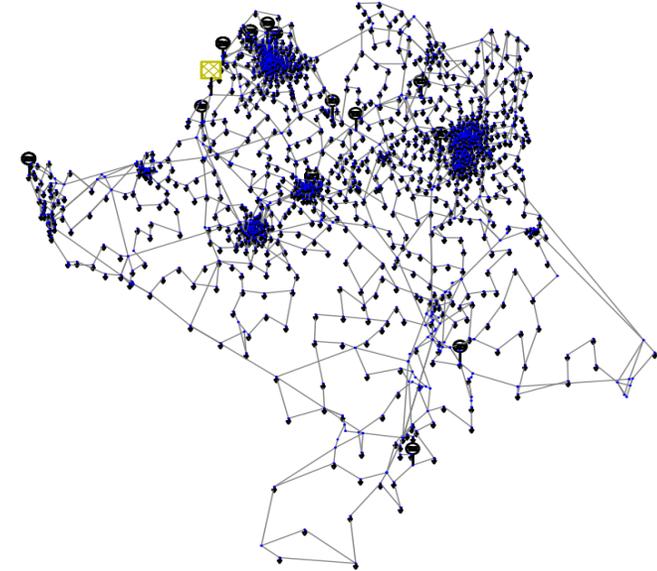
- Use GridFM AI emulator to pre-screen all load connection points, accelerating power flow calculations by >100x
- Run for each load connection point an N-1 contingency analysis
- Identify locations with the minimum violations

Impact:

- Significant acceleration in the interconnection process
- Cost reduction (avoidance of large infrastructure investments)
- Improved grid reliability and resiliency

S&T Challenges

Aligns with “Securing U.S. Leadership in Data Centers” and “Scaling the Grid to Power the American Economy”



Example: Synthetic Texas Transmission Grid with 2000 nodes

- 1124 potential connection points
 - ~ 4000 contingencies
- 10 different 24-hour load scenarios at 5 mins
 - ⇒ ~ 10B power flow simulations

Today: ~ 20 years of calculations

AI emulation: 2 months of calculations

Transforming Materials Discovery with AI

Motivation

Discovering new materials for energy, quantum technologies, and advanced electronics requires testing thousands of material variations within large sample libraries. Today, samples are measured one at a time using multiple techniques. Mapping a single library can take months.

Approach

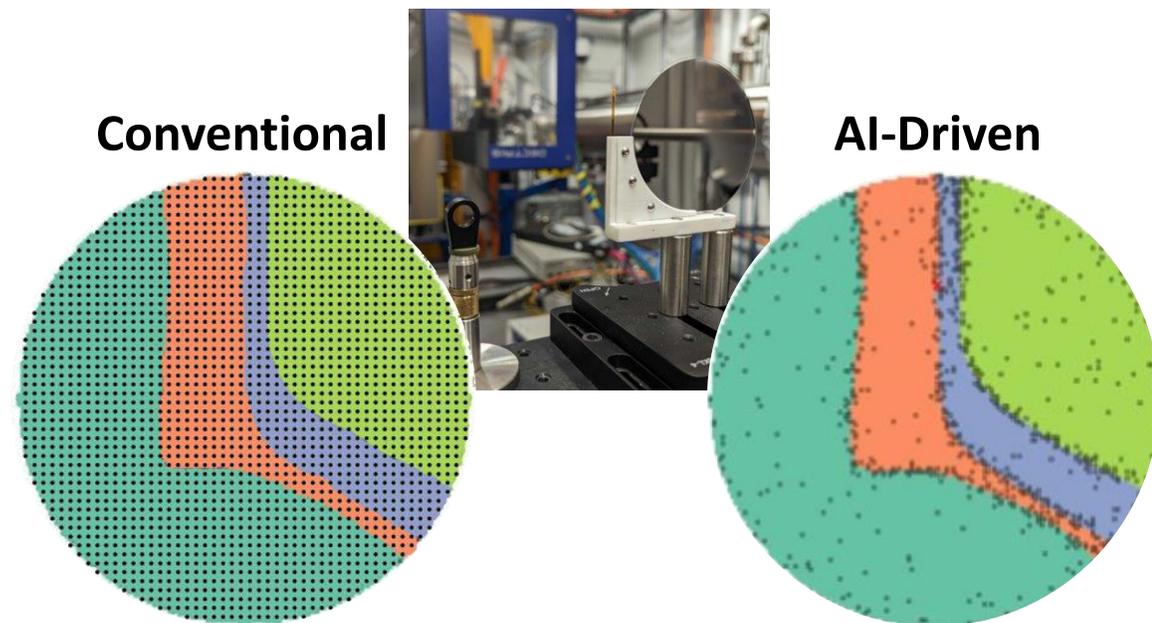
- AI coordinates two X-ray experiments simultaneously, with one guiding the other
- Real-time data analysis focuses only on the most informative measurements

Impact

- Achieves over 99% accuracy using only 5% of the data
- Reduces experiments from months to days
- Deployable across DOE X-ray facilities

S&T Challenges

Aligns with “Enhancing Particle Accelerators for Discovery” and “Designing Materials with Predictable Functionality”



Mapping Compositions in a Sample Library

Conventional: ~ 16,000 measurements | ~2 months

AI-guided: ~ 400 measurements | ~2 days

Colors indicate different material structures.

Achieving AI-driven Biotechnology to Secure America's Critical Minerals Supply

Motivation: Reliance on foreign supply chains for critical minerals and materials (CMM) threatens national security, economic competitiveness, and energy independence. Current domestic production is **costly and complex**, requiring >500 steps and toxic, expensive chemicals to extract, separate, and concentrate materials from ores, tailings, or waste.

Approach:

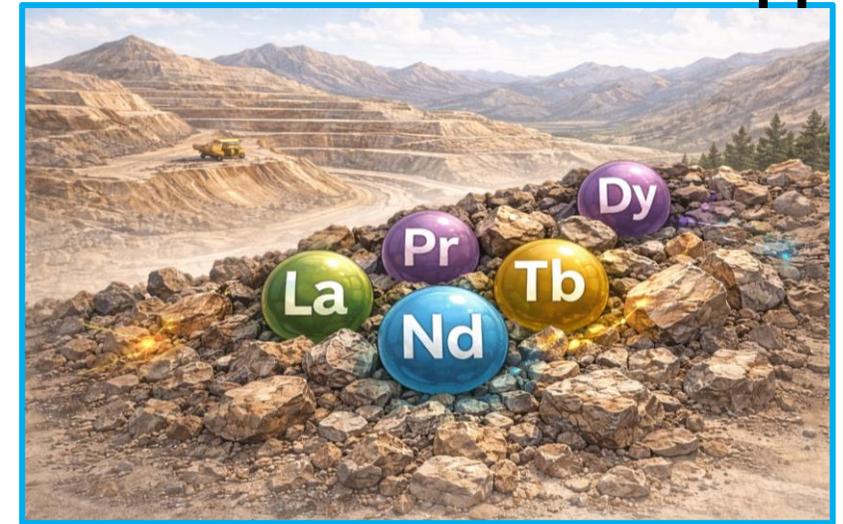
- Develop AI models and agents to design CMM-capture biomolecules and cells.
- Develop AI-driven closed-loop autonomous laboratory to accelerate the discovery of the CMM separation science.

Impact:

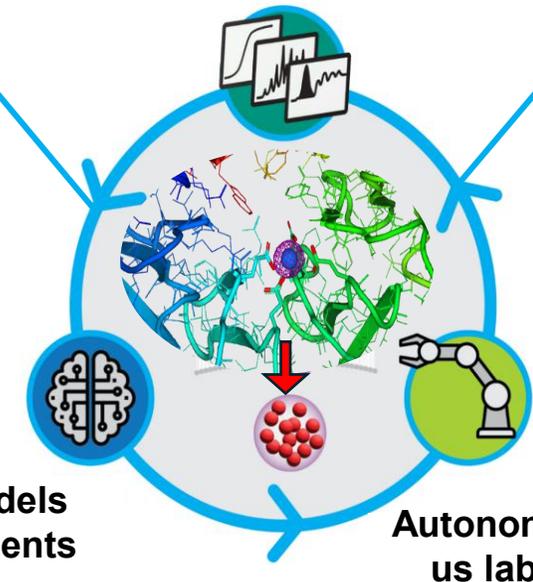
- AI-supported biotechnology is anticipated to **reduce more than 500 steps to less than 5 steps** with the same or higher purity, using mild and inexpensive biological solutions at ambient conditions
- AI speeds up the discovery-to-market cycle by 10x+
- Strengthened domestic CMM supply chain security and sustainability
- Decreased chemical intensity and reduce waste generation

S&T Challenge

Aligns with “Securing America’s Critical Minerals Supply”



CMM separation science



Example: AI helped to discover a new biomolecule that captures a critical rare earth element, terbium.

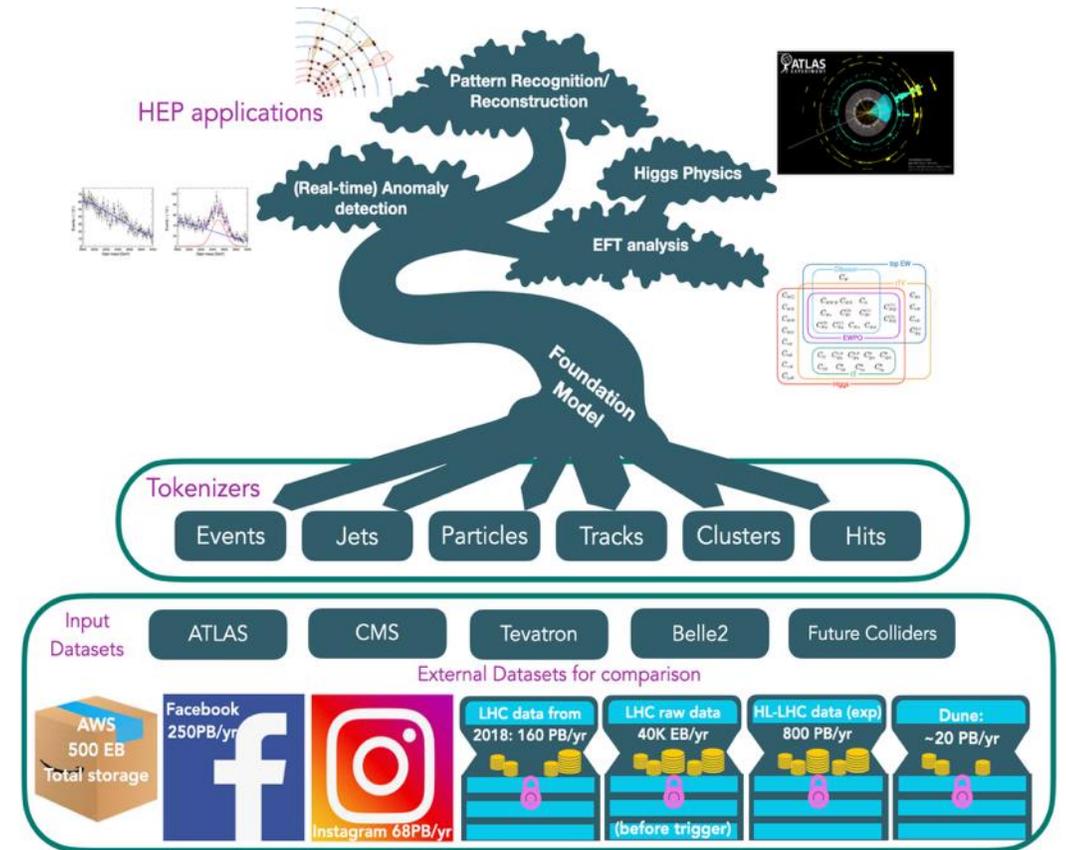
Cross-Experiment Discovery for Particle Physics

Motivation: AI and cross-experiment data integration enable a step-change in precision and discovery potential by combining current and legacy collider datasets.

Approach: We standardize and unify heterogeneous detector data and develop physics-informed foundation models (e.g., hierarchical tokenization and pretraining) to enable AI-driven reasoning across experiments.

Impact: This approach unlocks deeper insights into fundamental physics—enhancing Higgs measurements, enabling novel searches, and increasing sensitivity to new physics beyond the Standard Model.

S&T Challenge: Aligns with “Unifying Physics from Quarks to the Cosmos” through efforts like **TREASURE**, with early results demonstrating AI-ready datasets and scalable platforms on the American Science Cloud and Fermilab data infrastructure.



Example – **TREASURE** project
Tokenized Representations for Energy-frontier AI Searches via Understanding and Reasoning

From isolated experiments to integrated,
AI-driven discovery across the global collider physics program



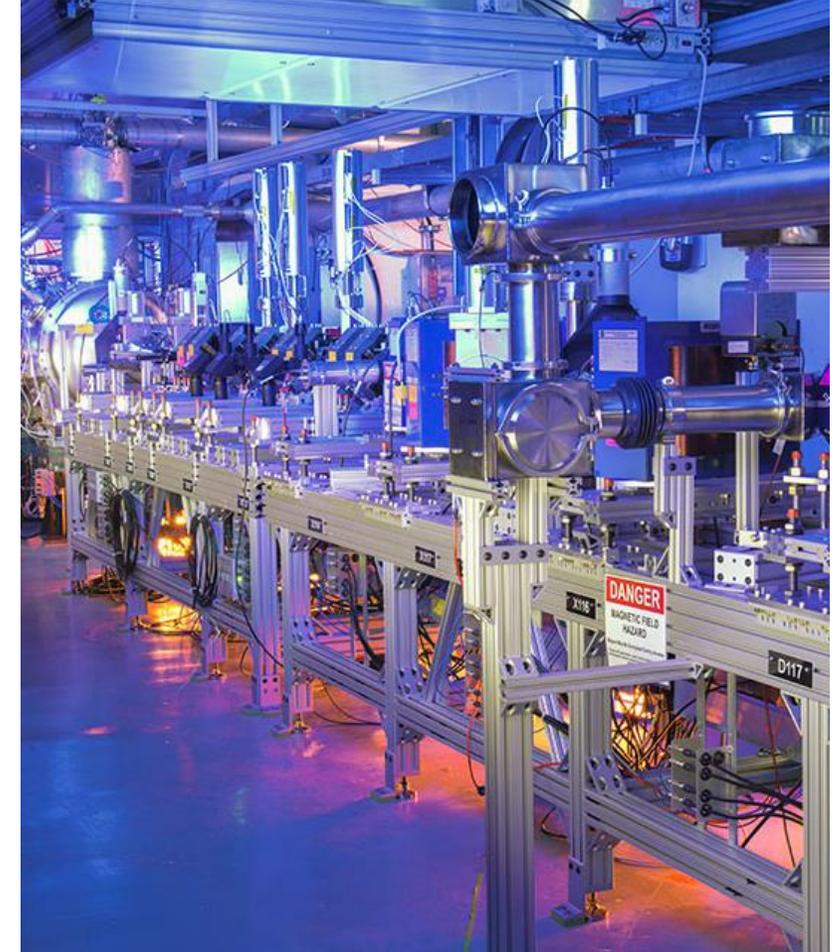
AI-Driven Transformation of Particle Accelerator Design & Operations

Motivation: Particle accelerators are critical DOE infrastructure whose growing complexity limits performance and discovery, creating an opportunity for AI-driven transformation.

Approach: We deploy AI—from agentic systems like **Osprey** (now operating across 8 DOE facilities with $>100\times$ faster operator interaction) to the **Genesis** platform—to enable autonomous operations, intelligent design, and a unified accelerator ecosystem.

Impact: AI is accelerating time to discovery, improving beam performance and reliability, and scaling innovation across DOE facilities with broad benefits to science, industry, and medicine.

S&T Challenges: Aligns with DOE’s “Enhancing Particle Accelerators for Discovery” priority while enabling cross-cutting impact across “Quark-to-Cosmos,” fusion & fission energy, quantum, microelectronics, and autonomous labs.



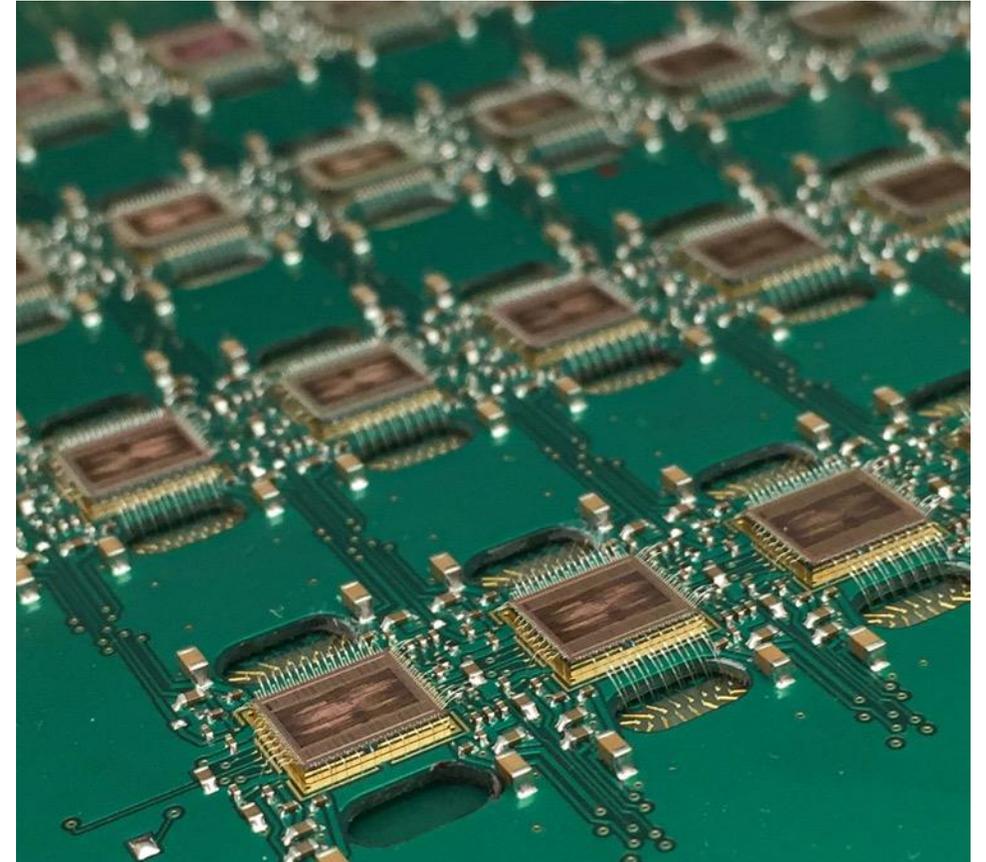
Novel Sensing and Control in Extreme Environments

Motivation: DOE science relies on advanced microelectronics that must operate in extreme environments (radiation, temperature), where traditional design approaches limit performance and slow innovation.

Approach: Develop AI-driven surrogate models and agentic design tools to automate and accelerate microelectronics design across the full stack for extreme-environment applications.

Impact: With the use of AI, enable orders-of-magnitude acceleration in chip design (e.g., **500× faster** exploration for qubit readout circuits), unlocking new capabilities for quantum systems, accelerators, fusion, and beyond

S&T Challenge: Aligns with “Recentering Microelectronics in America” while enabling cross-cutting impact across quantum science, particle accelerators, fusion energy, and materials discovery



**From slow, siloed design to AI-driven microelectronics innovation
for the most demanding environments in science**



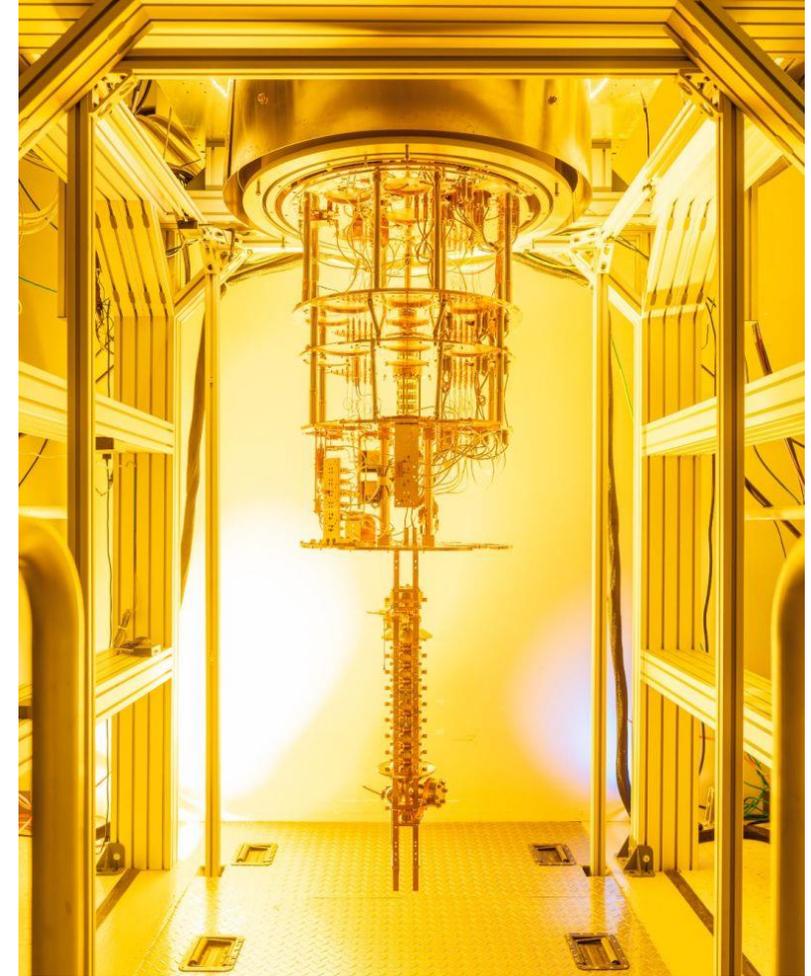
AI-Ready Quantum Data Ecosystem for Discovery at the SQMS Center

Motivation: DOE's National Quantum Information Science Centers generate unique, multi-modal data (materials → devices → systems → algorithms) that remain fragmented with isolated access, which limits their full potential for accelerating discovery.

Approach: A coordinated, cross-center effort to curate, standardize, and integrate quantum data into an AI-ready ecosystem—enabling foundation models and advanced analytics across the full quantum stack

Impact: A unified quantum data resource will accelerate breakthroughs in materials, devices, and algorithms, position DOE as a global leader in AI-driven quantum discovery, and enable broad community access and innovation

S&T Challenge: Advances DOE priorities in “Quantum Information Science” and “AI for Science,” establishing a national-scale capability that connects quantum research across centers, disciplines, and applications.



From fragmented quantum efforts to a unified,
AI-powered national discovery ecosystem

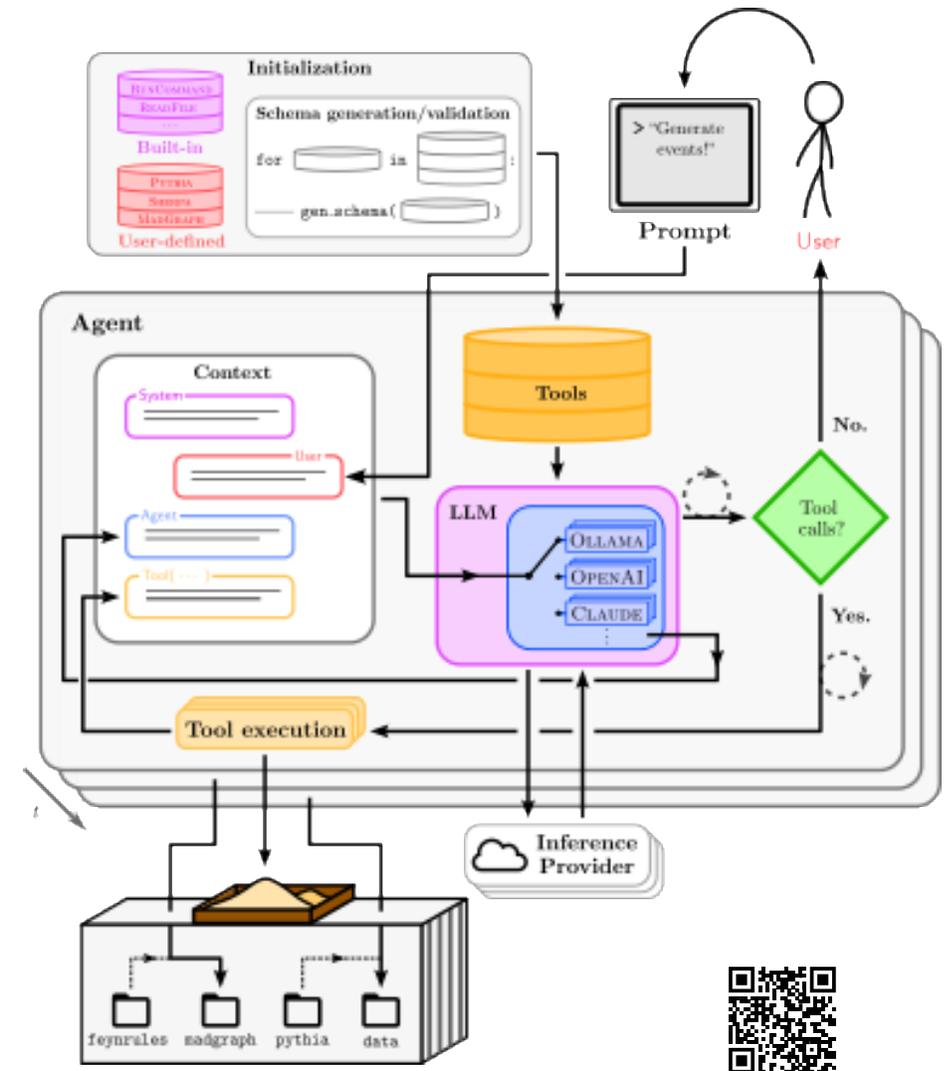
Automating Complex Physics Workflows

Motivation: Flagship physics experiments require complex, multi-step workflows across theory and experiment that are currently manually orchestrated, limiting speed, scalability, and reproducibility.

Approach: Develop agentic AI systems that plan, execute, and optimize end-to-end physics workflows while preserving full provenance for reproducibility and continuous learning.

Impact: Agentic automation accelerates discovery at major experiments (e.g., DUNE, LHC) by reducing human bottlenecks and enabling scalable, iterative exploration of new physics models

S&T Challenges: Advances “Unifying Physics from Quarks to the Cosmos” priorities and extends to other domains (e.g., quantum algorithms and error correction), demonstrating a unified AI framework for automating complex scientific workflows.



AI-ready Scientific Data Services

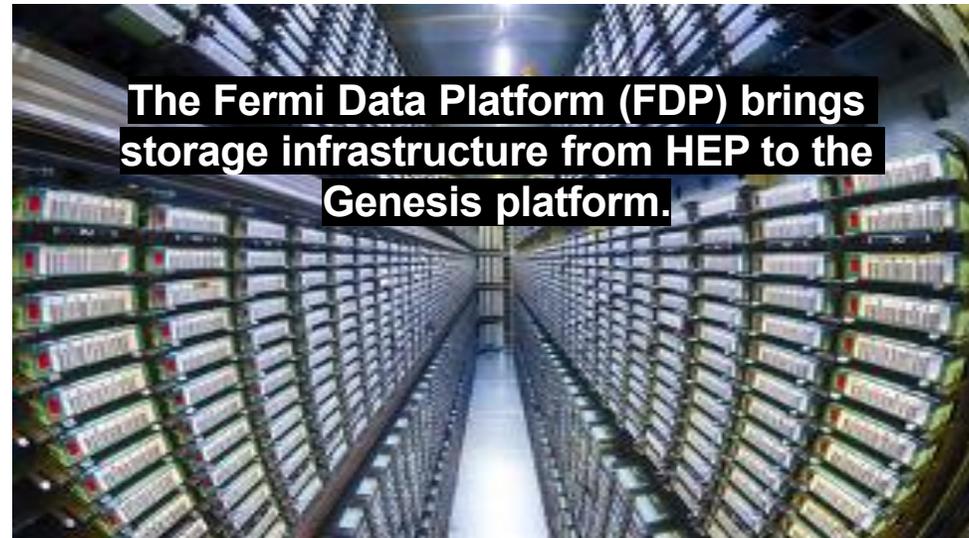
The Fermi Data Platform as an American Science Cloud Infrastructure Partner

Motivation: AI-driven discovery requires scalable, distributed, and accessible data infrastructure across the DOE complex, which traditional systems cannot fully support.

Approach: The **Fermi Data Platform (FDP)** provides WAN-scale, AI-ready data services—leveraging high-speed ESnet connectivity and integration with the American Science Cloud to enable seamless, cross-lab data access for Genesis.

Impact: FDP is already enabling AI-driven science across multiple programs (e.g., TREASURE, AI Universe, AXESS), establishing a shared data foundation that scales discovery across experiments and laboratories

S&T Area: Advances DOE capabilities in data infrastructure, computer science, and high energy physics, positioning Fermilab as a backbone provider for AI-enabled science within the American Science Cloud ecosystem.



The Fermi Data Platform (FDP) brings storage infrastructure from HEP to the Genesis platform.

AI Datasets
and Models

Public
Science Data
Releases

Private
Science Data

Data Tools, Data Formats, and APIs

Fermilab Mass Storage Resources
and Network Connectivity

TMAP8 Validation Case

Motivation: Building and validating physics simulation input cases for nuclear codes like TMAP8 requires deep domain expertise and tedious manual effort - digitizing legacy data, interpreting experimental reports, and iteratively authoring input files.

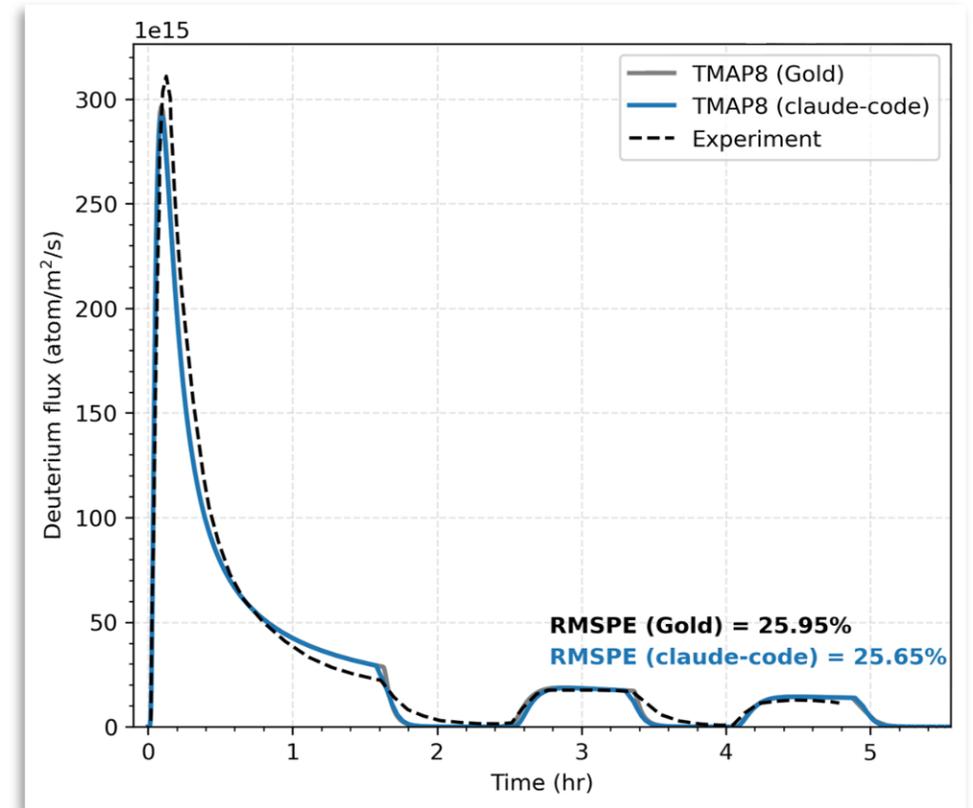
Approach

- Claude Code read legacy PDF documentation and automatically digitized experimental data
- Agent-guided planning elicits domain knowledge interactively, then autonomously scaffolds and iterates TMAP8 input files
- Human expert provides occasional oversight (~6 hours of babysitting) while the agent handles implementation, debugging, and result comparison

Impact

- AI-generated TMAP8 results (RMSPE = 25.65%) match or exceed the gold-standard human implementation (RMSPE = 25.95%)

Aligns with S&T Challenge: "Accelerating Delivery of Fusion Energy"



Conventional approach: **4 days of fusion science expert**
AI-driven approach: **6 hours of general materials staff**

The AI-generated TMAP8 input achieved lower prediction error than the hand-crafted gold standard - AI code agents can match expert-level simulation quality in a fraction of the time.

Developing a Semi-Autonomous Laboratory

Motivation: Current post-irradiation examination (PIE) workflows utilizing robotic systems are hindered by human operation or time-intensive programming resulting in delays and reduced throughput.

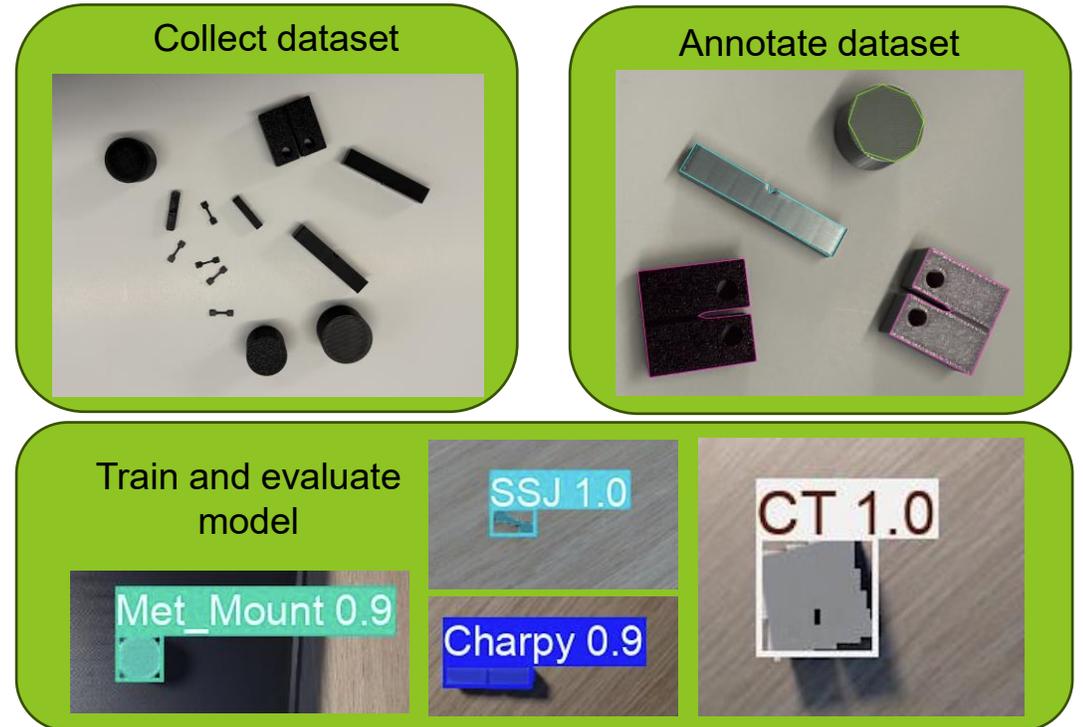
Approach

- Develop computer vision model for various material testing specimens
- Integrate model with robotic arm for automated path planning
- Demonstrate an automated PIE workflow

Impact

- 30%-time reduction compared to traditional methods
- Functional prototype system for high-consequence nuclear environments

Aligns with S&T Challenge: "Achieving AI-Driven Autonomous Laboratories"



Conventional approach: **Manual operation or manual programming of the robotic systems requiring subject matter expertise**

AI-driven approach: **Automated identification and operation of the robotic systems resulting in at least a 30%-time reduction**

Wisdm: wireless signal detection with a distributed energy-efficient model

Motivation

Traditional approaches fail to detect unknown or EM signals at low SNR, leaving critical gaps in real-time spectrum surveillance for national defense and public safety.

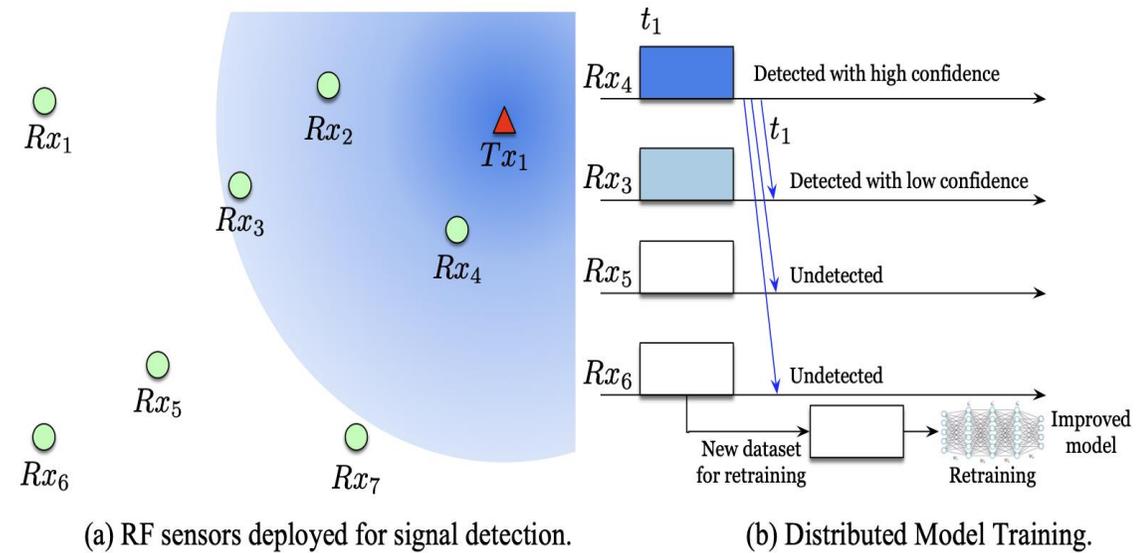
Approach

- Distributed peer labeling for unknown waveform training
- Lightweight transformer model for real-time detection
- Quantization aware pruning for edge deployment

Impact

- Detects unknown signals at -20 dB SNR
- Fewer sensors needed for equivalent coverage
- Strengthens U.S. SIGINT and spectrum defense

Aligns with S&T Challenge “Accelerating Nuclear Threat Assessment, Preparedness, and Response”



Conventional approach: **Static ML model require prior signal knowledge and fail at low SNR.**
AI-driven approach: **Detects unknown waveforms at -20 dB SNR on edge hardware.**

Digital Engineering and Autonomous Robotics

Motivation

Hanford's 177 radioactive waste tanks require costly, hazardous manual pit inspections every three years, with no autonomous alternative for confined, access limited environments.

Approach

- Sim to real AI training for autonomous pit navigation
- Digital twin linking robot, sensors, and environment
- ROS2 and DeepLynx for deployment and feedback

Impact

- Eliminate costly concrete cap removal for inspections
- Reduce worker exposure to radioactive pit environments
- Accelerate Hanford cleanup toward DOE's 2040 goal

Aligns with S&T Challenge: “Transforming Nuclear Cleanup and Restoration”



Conventional approach: **Manual pit inspections require removing concrete caps, exposing workers to radiation every three years.**

AI-driven approach: **Autonomous robots inspect through 4-inch viewports, eliminating cap removal and worker exposure.**

Cognito – Responsible AI Implementation Guide

Motivation

Cognito gives utility leaders a clear, defensible way to decide which AI capabilities are safe and mission-aligned for critical grid operations.

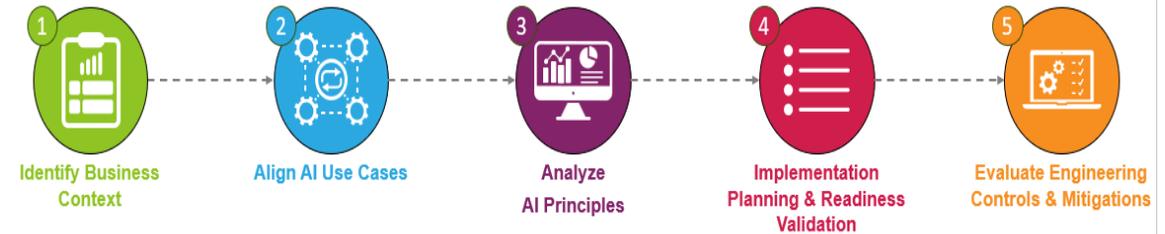
Approach

Cognito uses a structured, consequence-driven, five-step process to guide utilities from exploratory AI ideas to risk-vetted, implementation-ready plans.

Impact

- Accelerate safe, defensible AI adoption
- Improve leadership decision-making
- Enhance operational assurance and reduces system risk

Aligns with S&T Challenge: Developing AI that remains reliable and risk-aware in operational, high-consequence environments where traditional assessment methods fall short



Conventional approach: **Months of manual, fragmented readiness checks.**

Cognito approach: **Days to generate a structured, traceable risk and readiness assessment.**

Cognito speeds AI-readiness assessment, cuts operational uncertainty, and accelerates safe AI adoption in critical environments.

Light Water Reactor Sustainability Program

Motivation

AI can improve the safety, efficiency, and cost-effectiveness of the current U.S. nuclear fleet while expanding its capacity and extending its operational lifespan

Approach

- Prioritize automation of manual, labor-intensive tasks across functions and organizations
- Develop tailored AI methods and tools through advanced research and development (R&D)
- Disseminate outcomes via pilot demonstrations, guidance reports, and licensing of AI models and code

Impact

- Optimize plant operations and enhance efficiency
- Accelerate capacity expansion
- Support long-term plant performance and reliability
- Strengthen facility protection and security

Aligns with S&T Challenge: Quadrupling of nuclear power by 2050



Conventional approach: Plant staff manually collect, prepare, and analyze data to make decisions across all plant organizations.
AI-driven approach: Streamlined, automated decision-making across plant functions and organizations.

Enabling Risk-aware Biochemical Process Scale-up through AI Assisted Design

Motivation

Scaling up cellulosic biorefinery technologies based on limited laboratory and pilot plant data is technically complex, costly, and risky, resulting in limited investments in FOAK projects.

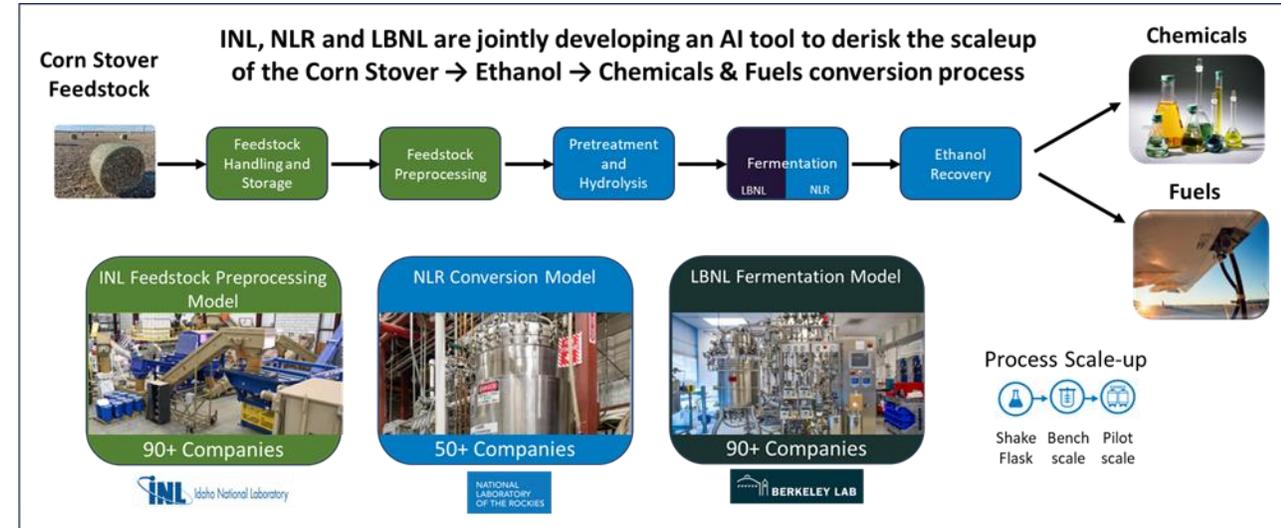
Approach

- Physics informed AI for biomass processing and conversion
- Trained on available data from National User Facilities and Process Demonstration Units
- Uncertainty-quantified predictions

Impact

- Inform process design for variable feedstocks
- Reduce the risks associated with biomanufacturing
- Enable reduction of foreign import of fertilizers and chemicals

S&T Challenge: Scaling the Biotechnology Revolution



Expected reduction in development time of **20 – 40%** with AI based tools when compared to conventional methods. Translating to cost savings of **30 – 50%** due to streamlined testing and experimentation.

MOOSEng: Lowering the Barrier to Multiphysics Modeling and Accelerating Technology Transfer

Motivation

High-fidelity multiphysics modeling is essential for reactor design, safety, and licensing, but these tools can have a steep learning curve.

Approach

- Domain-specific AI agent with retrieval-augmented generation assists with tool use.
- Complex inputs can be automatically repaired and validated thanks to the use of tuned AI agents.
- Runtime-in-the-loop execution and evaluation use real MOOSE feedback to iteratively improve outputs and measure reliability.

Impact

- Agent evaluation across 175 prompts / 7 physics families: 158/175 runnable cases (**90%**) vs 10/175 (**6%**) for ChatGPT 5.2 API.
- Speeds user adoption, broadens access to advanced modeling, and creates a reusable pathway for applying AI across the MOOSE ecosystem.

S&T Challenge: Supporting AI-powered nuclear design in Prometheus

MOOSENGER

MOOSEng
Chat * Reason * Simulate
(powered by DOE innovation)

Notice: These data were produced by Battelle Energy All No. DE-AC07-05ID14517 with the Department of Energy. Du commercialization or such other time period specified b Energy, the Government is granted for itself and others nonexclusive, paid-up, irrevocable worldwide license in prepare derivative works, and perform publicly and disp behalf of the Government. Subsequent to that period the for itself and others acting on its behalf a nonexclusi worldwide license in this data to reproduce, prepare de distribute copies to the public, perform publicly and d permit others to do so. The specific term of the licens inquiry made to the Contractor or DOE. NEITHER THE UNIT STATES DEPARTMENT OF ENERGY, NOR ANY OF THEIR EMPLOYEES EXPRESS OR IMPLIED, OR ASSUMES ANY LEGAL LIABILITY OR R ACCURACY, COMPLETENESS, OR USEFULNESS OF ANY DATA, APPA PROCESS DISCLOSED, OR REPRESENTS THAT ITS USE WOULD NOT RIGHTS.

Test problem	MOOSEng (Pass)	Baseline (Pass) ChatGPT 5.2 API
Diffusion	25/25 (100%)	9/25 (36%)
Transient heat conduction	23/25 (92%)	0/25 (0%)
Solid mechanics	24/25 (96%)	0/25 (0%)
Porous flow	23/25 (92%)	1/25 (4%)
Navier-Stokes	21/25 (84%)	0/25 (0%)
Phase-Field	21/25 (84%)	0/25 (0%)
Plasticity	21/25 (84%)	0/25 (0%)
Overall	158/175 (90%)	10/175 (6%)

Conventional approach: expert-mediated tool selection, manual input development, iterative debugging | **weeks** to first model
MOOSEng approach: AI-guided workflow selection, guarded input generation, solver-backed repair | **hours to days** to runnable simulation

Turns expert-dependent multiphysics modeling into a guided, auditable workflow—reducing onboarding friction and broadening access to DOE simulation capabilities.

Trustworthiness Assessment of Grid Foundation Models

Motivation

- Grid topologies are changing rapidly; current tools are **too slow** or **too inaccurate**.
- AI models exist but **don't generalize** to new grid structures and **lack uncertainty quantification**, limiting trust and deployment.

Approach

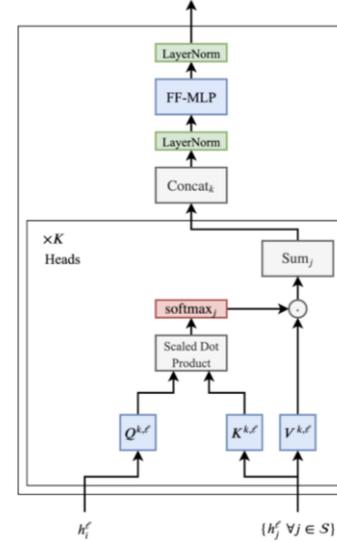
- **Micro-foundational Graph AI models** trained across many grid topologies.
- Embedded uncertainty quantification for operator-trustworthy predictions.
- Distributed GPU training, enabling large-scale model development and rapid fine-tuning for specific utilities.

Impact

- **Up to 20× accuracy improvement** over DC approximations and **~10× faster** than traditional workflows.
- Supports DOE missions in **grid reliability**, clean-energy integration, and accelerated interconnection studies.

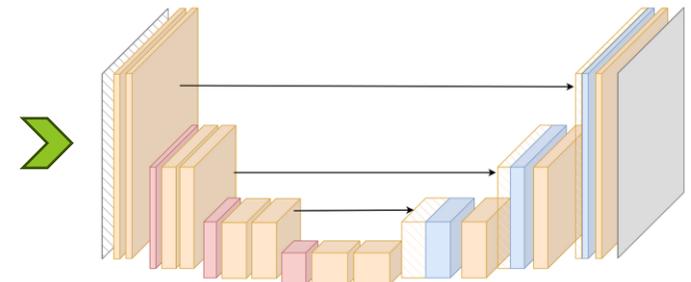
S&T Challenge: Scaling the Grid to Power the American Economy

Graph Transformers



Incorporate attention

Graph U-Net with Bayesian LoRA



Generalize across topologies and provide UQ

This effort can reduce the latency of real-time grid contingency analysis and remedial action generation from 15-45 minutes to under 5 minutes, while maintaining or exceeding the physical safety standards required by NERC reliability criteria.

Digital Twin Development for Remote Operations

Motivation

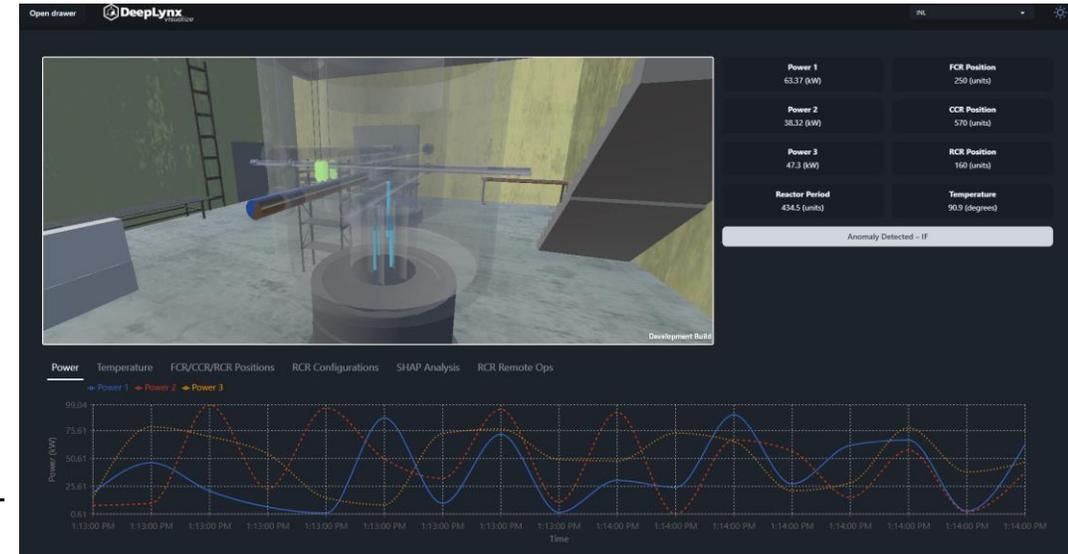
Current reactor operations require on-site personnel, limiting deployment in remote locations and preventing multi-reactor consolidation. Remote and autonomous operations could reduce overhead costs, enhance safety, and enable microreactor and SMR deployment in previously inaccessible markets – yet remain largely undemonstrated.

- Remote Control Rod (RCR) integrated with a Data Acquisition System (DAS)
- Physics (PKE/IPKE) and ML methods (LSTM, Isolation Forest) to predict power, detect anomalies, and recommend RCR control actions
- DeepLynx Nexus digital twin to synthesize real-time reactor data, physics, and ML

Impact

- First remote operations demonstration of the AGN-201 using a digital twin
- Reusable DT framework for flexible remote operations available to industry and academia accelerating the path to autonomous nuclear operations across the U.S.

S&T Challenge: Aligns with "Delivering Nuclear Energy that is Faster, Safer, Cheaper"



Mockup UI for Remote Operations via Nexus

Conventional approach: **Requires on-site operator presence | Manual monitoring and control**
AI/Digital twin-driven approach: **Remote operations from any location | Automated anomaly detection (IF, ~5% rate), power prediction (LSTM, $R^2 \sim 0.9$), and RCR control recommendations (IPKE, surrogate model)**

Demonstrating Autonomous Operations at the Neutron Radiography Reactor using AI

Motivation

Autonomous operations is desired by industry and research groups due to the ability to reduce operational costs, operational errors, and increase efficiencies.

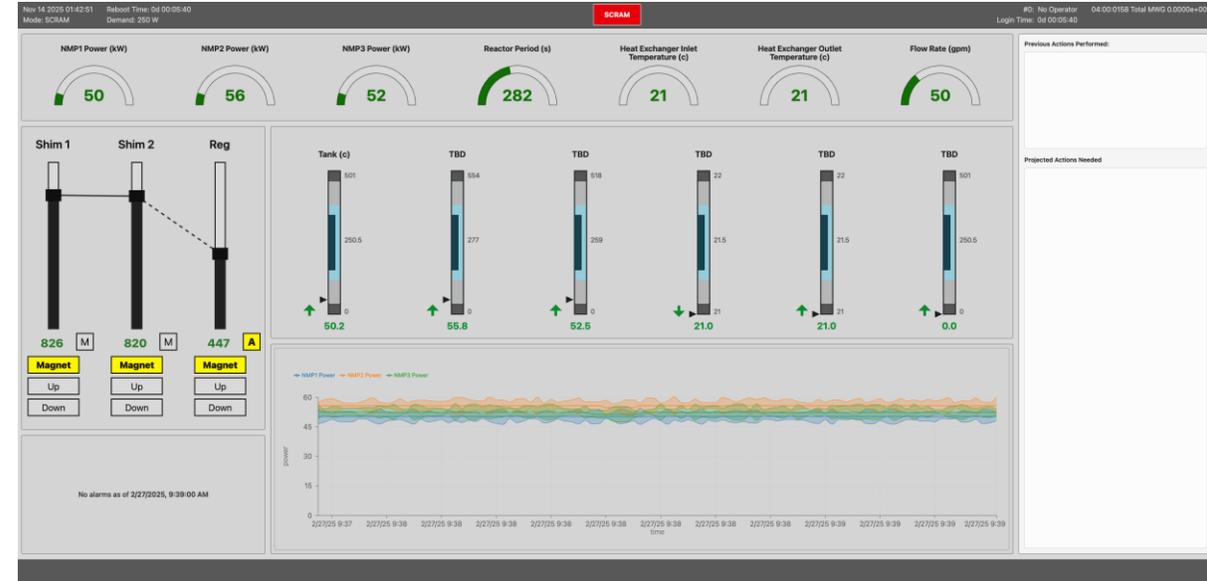
Approach

- Develop a digital twin framework for ingesting, analyzing, and determining actions for a reactor to take.
- Utilize model predictive control, reinforcement learning, and PID controller for determining control actions
- Combine multiple lab-developed software to accelerate and stress-test them in a real-world environment

Impact

- Assess methodologies, de-risk future demonstrations, and demonstrate autonomous operations
- Determine where deterministic and AI method best interface

S&T Challenge: Delivering Nuclear Energy that is Faster, Safer, Cheaper



Conventional approach: **Manually select and move control rods and continually interface to meet operations goals**

AI-driven approach: **Provide overarching goal(s) to system and monitor to verify goals are met**

Demonstration of autonomous operations using AI for full demonstration of startup, operations, and shutdown.

Safeguarding Advanced Microreactors with Sensor-Informed Digital Twins

Motivation

The IAEA anticipates global nuclear capacity could triple over the next 50+ years. With limited resources, new technologies will be imperative to ensuring safeguards verification for the expansion of nuclear worldwide.

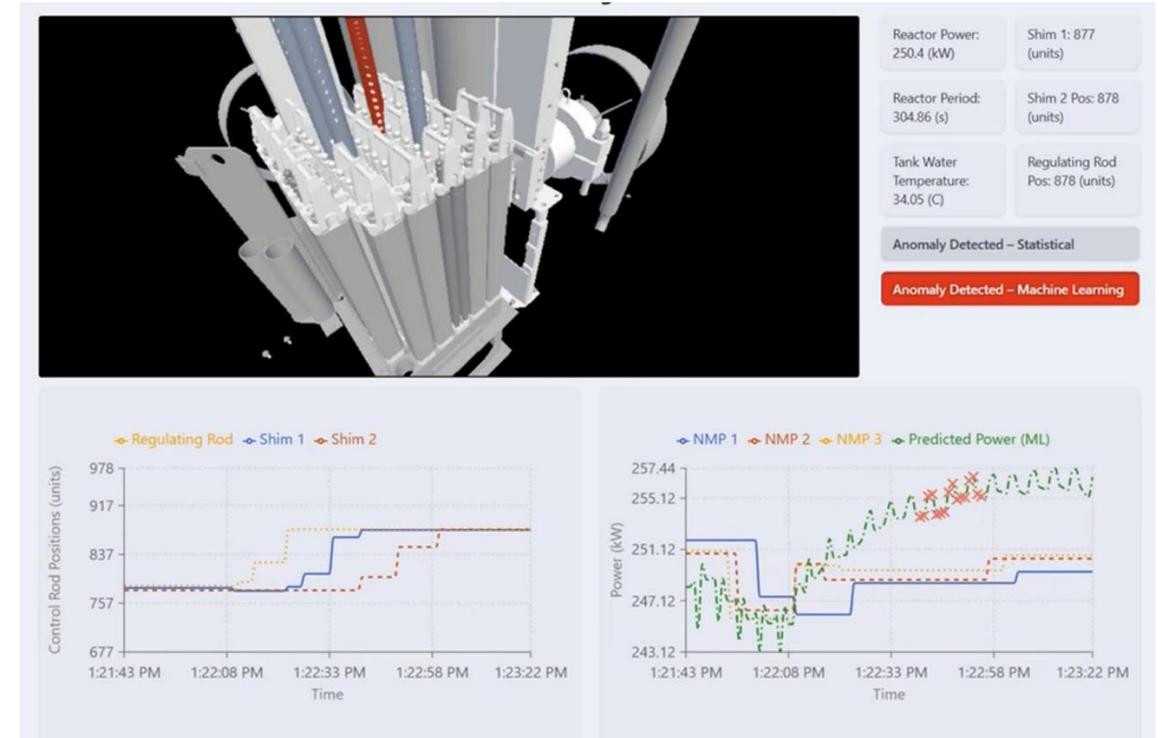
Approach

- Digital twin technologies help synthesize vast amounts of sensor data ingested by an operating reactor
- Combine statistical analysis and machine learning to automatically detect and flag anomalies

Impact

- Demonstrate data ingestion and anomaly detection with an operating research reactor
- Ensure subject matter experts can review and make safeguards conclusions

S&T Challenge: Delivering Nuclear Energy that is Faster, Safer, Cheaper



Current Approach: Humans currently examine gigabits of data to draw safeguards conclusions.

AI-Driven Approach: Digital twin collects, analyzes, and provides areas of interest for an inspector to examine allowing the human spend **more time** on drawing **conclusions** and less time on sorting data.

ACORN: Autonomous Control for Reactor techNologies

Motivation

Demonstrating artificial intelligence (AI)-guided autonomous operation is an economic necessity for advanced reactor faster and safer deployment.

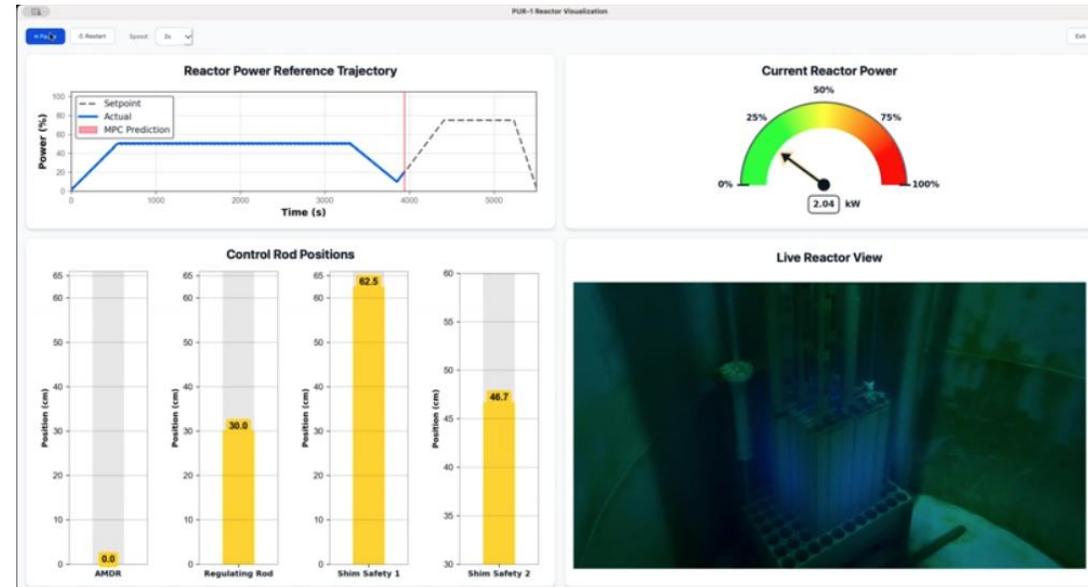
Approach

- Autonomous Control for Reactor Technologies (ACORN) leverages physics- and data-driven model predictive control.
- Agnostic to reactor technologies.
- Demonstration planned at Purdue University Reactor One (PUR-1) and at INL's Neutron Radiography Reactor.

Impact

- Provides regulators with critical evidence, through practical demonstration that autonomous operation enhances and does not undermine nuclear safety.
- Benchmarked in a virtual environment for different microreactor technologies.

S&T Challenge: Delivering Nuclear Energy that is Faster, Safer, Cheaper



ACORN interface developed for PUR-1

- Conventional controllers are reactive.
- ACORN is a **predictive** controller – anticipates reactors operational state ahead of time and recommends an optimal control action.

VIPER: Visualization of Predictive MaintEnance Recommendations

Motivation

To economically uprate U.S. nuclear power plants while strengthening system reliability, the industry needs a single, nuclear-specific artificial intelligence (AI) platform.

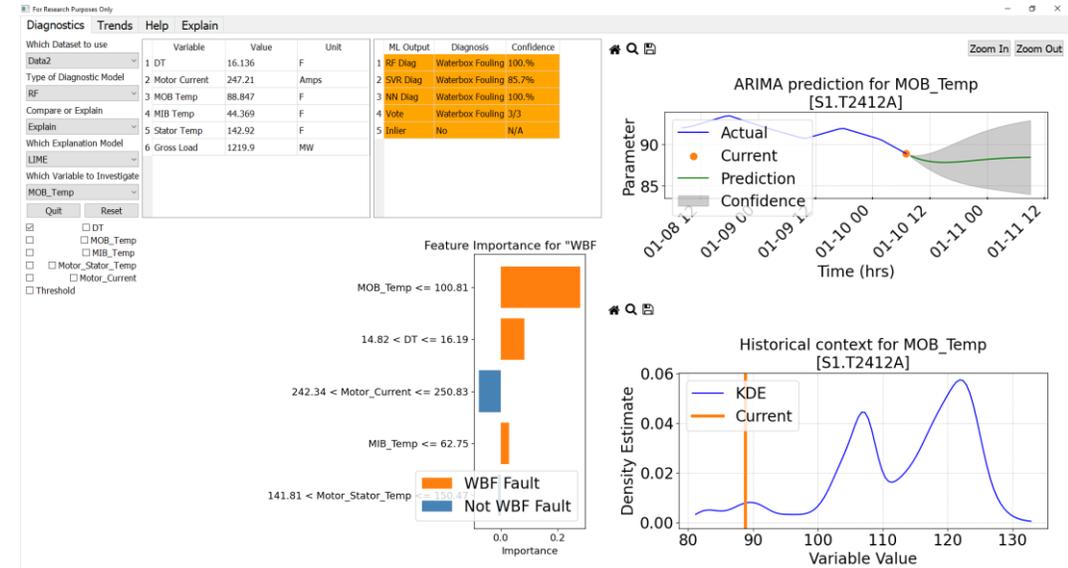
Approach

- VIPER — an AI platform combining ML fault diagnostics, explainable AI, and an RAG-powered nuclear knowledge chatbot.
- VIPER predicts and recommends actions based on system reliability to maximize plant operational efficiency.

Impact

- The integrated VIPER platform directly addresses the nuclear industry's dual challenge of workforce reduction and aging asset management.

S&T Challenge: Delivering Nuclear Energy that is Faster, Safer, Cheaper



VIPER interface developed for in collaboration with the Nuclear Power Plant

- Conventional monitoring and diagnosis at nuclear power plants is **labor-centric business model**.
- The AI-driven predictive analysis enables transitions to a **technology-centric business model**.

Physics-Informed AI for Critical Mineral Orebody Mapping

Motivation

Domestic uranium production requires faster, more reliable identification of roll-front orebodies; current manual workflows are slow, subjective, and limit timely decisions critical to U.S. energy and mineral-security priorities.

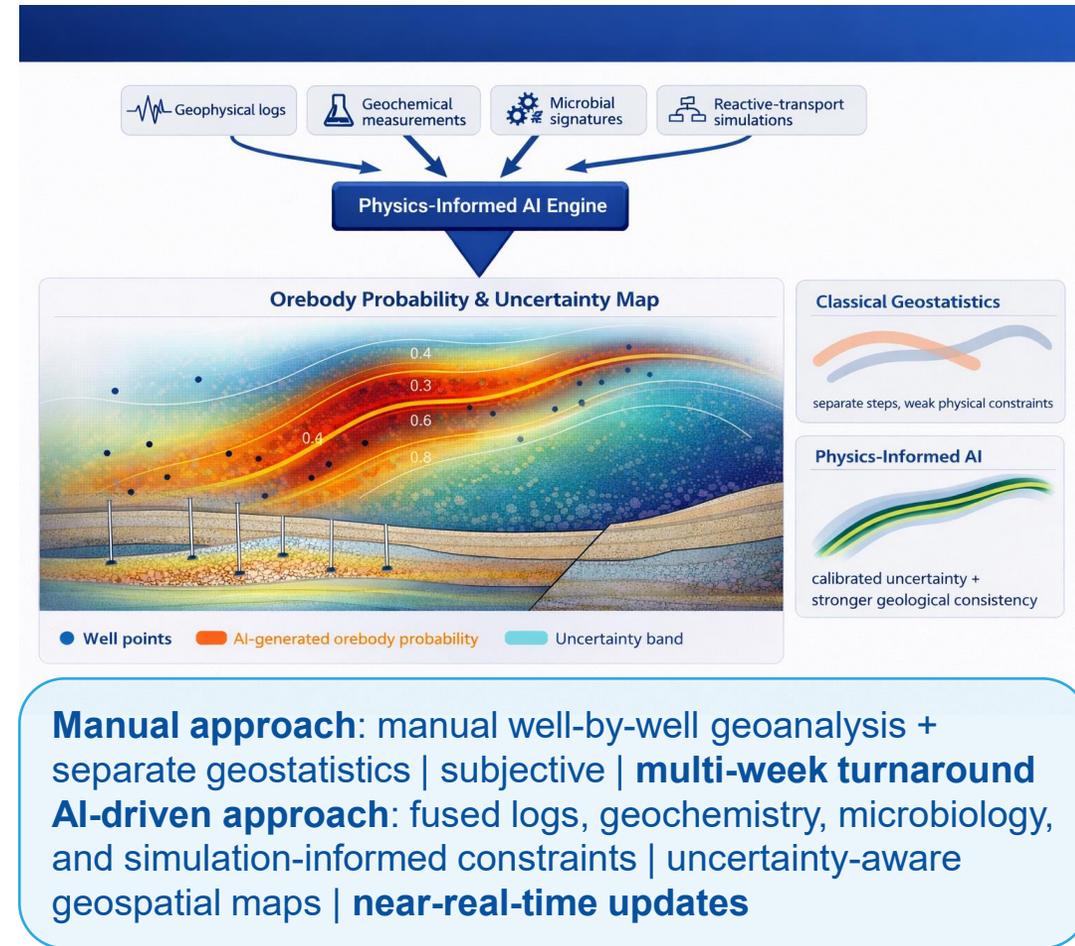
Approach

- Fuse geophysical, geochemical, microbiological, and simulation data in a physics-informed AI engine.
- Automate subsurface correlation and mineralization prediction with uncertainty propagation.
- Produce geospatial orebody probability maps benchmarked against traditional tools.

Impact

- Shrinks multi-week expert interpretation to rapid, reproducible field-scale mapping.
- Improves drilling and mine-planning confidence with uncertainty-aware targeting.

S&T Challenge: Aligns with "Securing America's Critical Minerals Supply".



FORESIGHT: AI Enabled Supply Chain Analysis

Motivation

Increased trade globalization has led to increased supply chain complexity reducing effectiveness of 'traditional' supply chain analytics techniques. Advancements in AI have enabled new approaches to be tested and deployed expanding the boundary of what is possible in supply chain analysis.

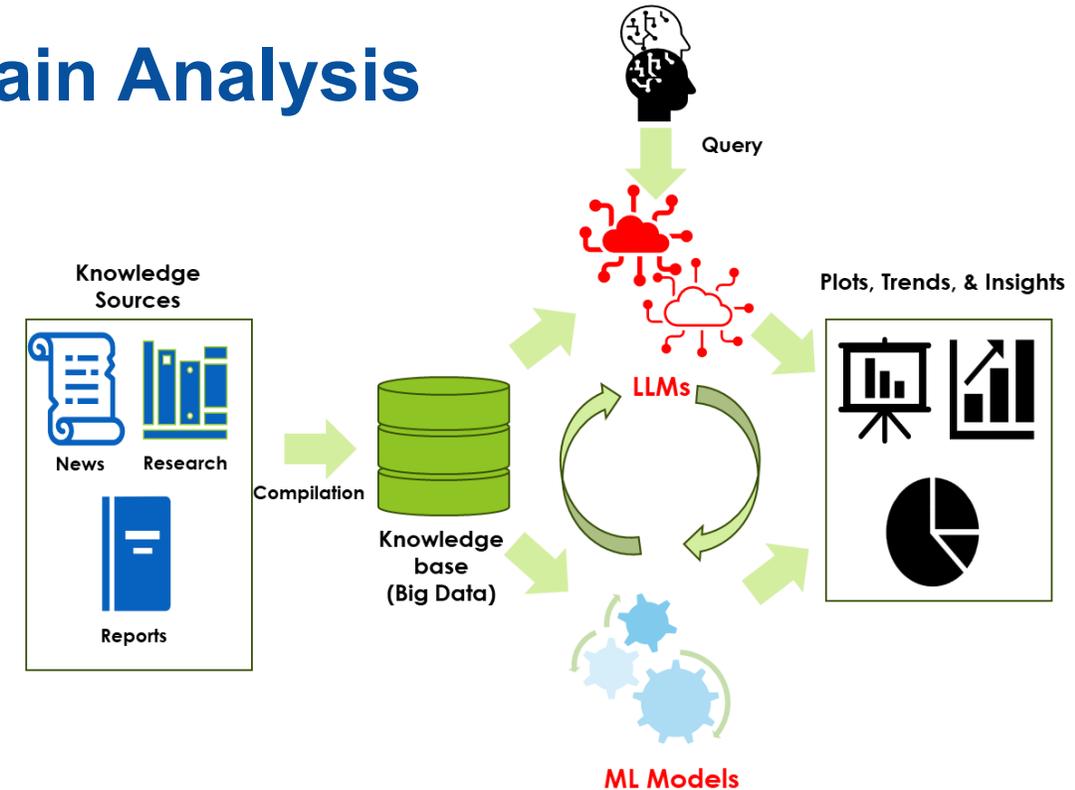
Approach

- Predict challenges and disruptions within the critical materials supply chain through big data and AI/Machine Learning to develop correlation/causation for supply chain events/entities

Impact

- Enable customers such as government policy makers, warfighters, and industry leaders to stay ahead of supply chain challenges and favorably position the U.S. ahead of incoming disruptions
- Dual capability —semantic synthesis and scenario prediction— empowers resilient supply chains and innovation-ready material systems

S&T Challenge: Aligns with "Securing America's Critical Minerals Supply"



Conventional approach: **Simulation models | 12-18 months**

AI-driven approach: **AI predictive models | 1 day**

Completion of AI model results in rapid supply chain analyses capable of continuous assessment based on real-time events.

AI to Predict Electrophoretic Transport Properties of Lanthanide-Ligand Complexes for Separations

Motivation

Efficient separation of lanthanides is a critical bottleneck for rare earth supply chains. Minimal chemical differences across the series make separation difficult, and current methods rely on complex, multi-step processes with limited predictive capability.

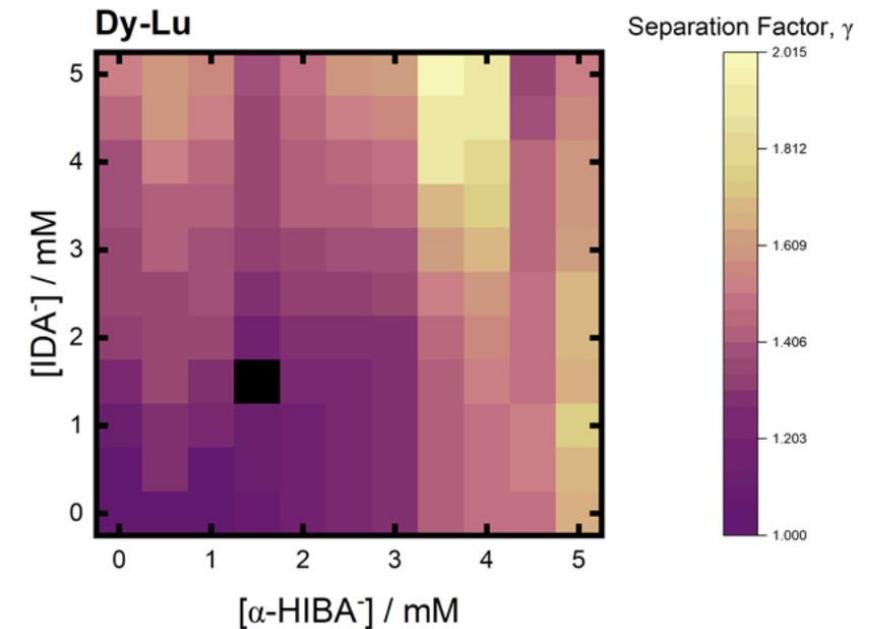
Approach

- Develop AI models to quantify how ligand chemistry and concentration control separation behavior across the lanthanide series.
- Integrate mobility and ligand descriptors to build a predictive framework for separations.

Impact

- Reveals a transition from mobility-driven to thermodynamic and cooperative separation regimes.
- Enables predictive ligand selection and reduces experimental trial-and-error.

S&T Challenge: Aligns with "Securing America's Critical Mineral's Supply"



Conventional approach: Experimental mapping of resolution surfaces for each ligand system | extensive testing

AI-driven approach: Predictive modeling of ligand performance and optimal conditions | few minutes

Agentic AI for Critical Minerals & Materials Resource Estimation

Motivation

U.S. competitors have built a formidable advantage in critical minerals and materials through four decades of sustained R&D investment — a gap that traditional methods cannot close without AI.

Approach

- Model Context Protocol (**MCP**) servers acquiring data from several CMM data servers (USGS, BGS, OSTI.gov, etc.)
- LangGraph agents calling tools and models to perform prediction/inference
- Models: A mixture of classical surrogate models and pretrained foundation models depending on the downstream process.

Impact

- Exploration cost reduction, 1.5-2x increased recovery of rare earths.

S&T Challenge: Aligns with "Securing America's Critical Minerals Supply"



Conventional approach: 10+ years
AI-driven approach: <5 years

BIOWASTE-AI: LLM-Driven Knowledge Extraction to Guide Biomass & Waste Processing System Design

Motivation

Designing optimal preprocessing systems for biomass and waste feedstocks demands understanding how feedstock properties govern equipment selection, process performance and optimization, and cost.

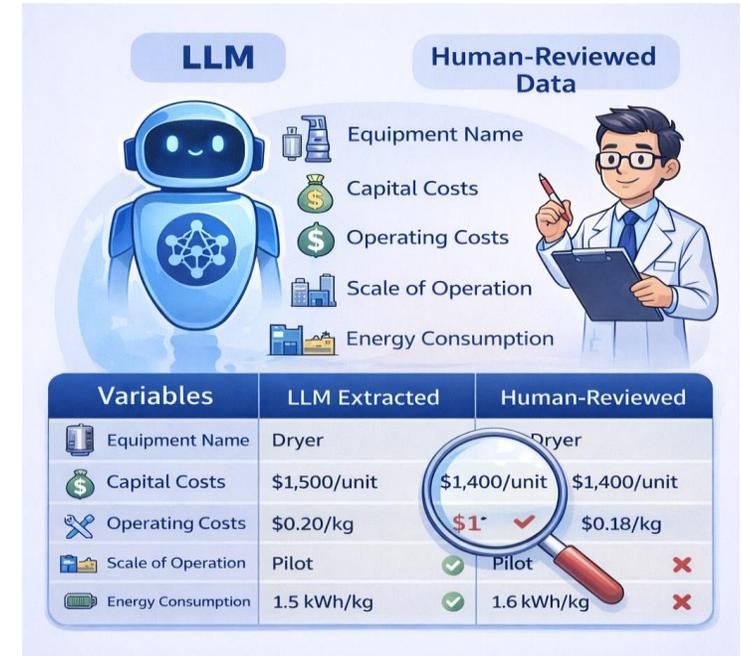
Approach

- LLMs were used to extract and structure property-to-performance-cost knowledge from 110+ biomass and waste preprocessing papers
- Structured outputs mapping feedstock characteristics to process steps, equipment, and techno-economic parameters

Impacts

- ~99.5% accuracy achieved in LLM-extracted equipment-specific parameters
- Transferable methodology established for LLM-driven process optimization, feedstock selection, and product manufacturing
- Rapid, scalable knowledge synthesis reduces months of manual expert effort to hours, supporting U.S. energy development

S&T Challenge: Aligns with “Reenvisioning Advanced Manufacturing and Industrial Productivity”



Conventional approach: Manual literature synthesis to extract feedstock-processing relationships | months of expert effort

AI-driven approach: LLM-driven knowledge extraction linking feedstock properties to system design and cost | hours to days



PROMETHEUS

Prometheus: AI-Assisted Power Plant Design & Autonomous Operations

Motivation

The U.S. seeks to quadruple nuclear capacity by 2050, but traditional methods in nuclear power plant deployment can't get the country to meet today's demand fast enough.

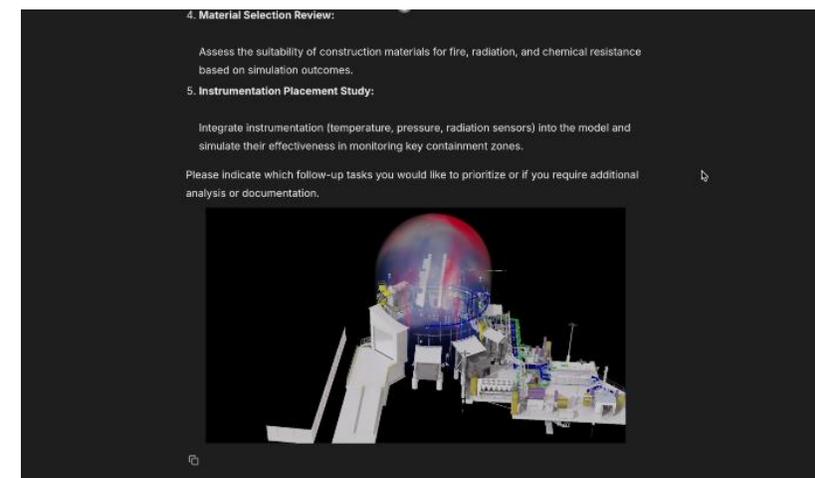
Approach

- Agentic AI capable of orchestrating engineering design, licensing, manufacturing, and construction workflows, connected to DOE codes, DOE models, and commercial tools
- AI surrogate models for key nuclear physics (neutronics, T/H, etc.) enabling rapid iteration and design optimization
- Comprehensive data management that includes historical, curated data sets for additional context/training and in-process artifact storage for project-specific context

Impact

- 2-3x near-term and up to 5x long-term plant deployment acceleration
- 50% operating cost reduction for microreactor/SMR technology
- Enable the U.S. to both 1) quadruple (400 GW) nuclear capacity by 2050 and 2) meet today's energy demand from data centers

S&T Challenge: Delivering Nuclear Energy that is Faster, Safer, Cheaper



Optimizing Sensor Placement for Nuclear Digital Twins

Motivation

Nuclear assets lack the luxury of immersive sensing networks because of harsh conditions, physical inaccessibility, and budget limitations. This motivates the urge to optimally place and learn from such sensors.

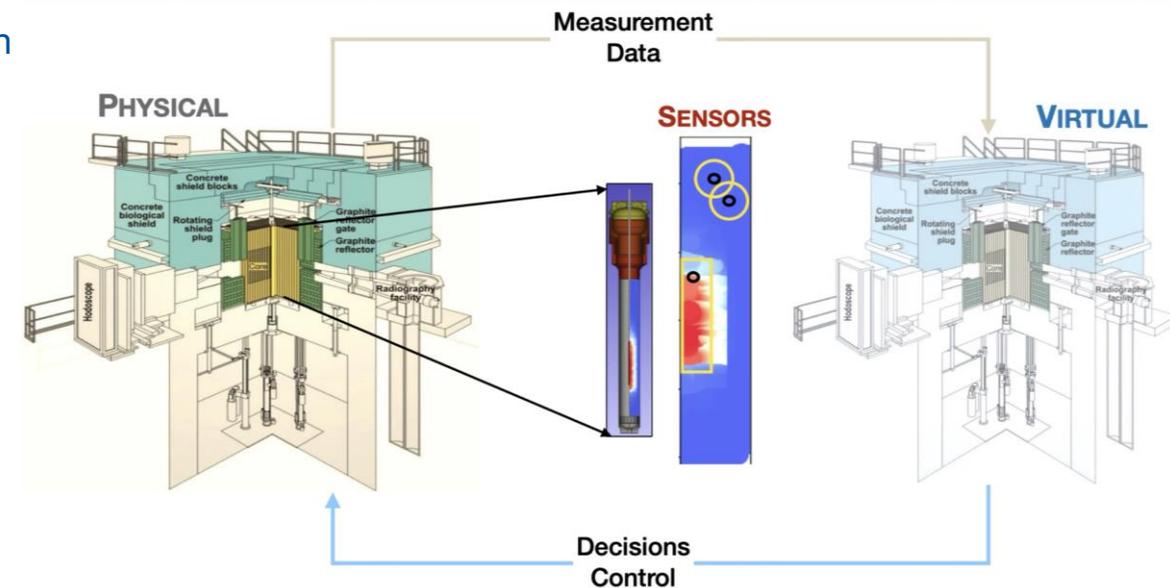
Approach

- AI-based decompositions, bases and surrogate models
- Greedy optimizers with physical and cost constraints for both reconstruction and off-normal conditions detection
- In-house code RAVEN and open-source code PySensors

Impact

- X optimal sensors outperforms 4X random sensors by several orders of magnitudes in reconstructing field of interest.
- Mitigates sensors drifting due to high temperature and act as virtual sensors to reflect real-time reactor state
- U.S. is heading towards next gen reactors including small modular reactors in remote setups which dictates minimal human intervention which necessitates this technology.

S&T Challenge: Aligns with "Delivering Nuclear Energy That is Faster, Safer, Cheaper"



Conventional approach: **20 measurements | 2 months**
AI-driven approach: **5 measurements | 1 day**

Multi-modal Foundation Models for Power Plant Anomaly Detection

Motivation

Nuclear safety depends on identifying issues before failures occur. Modern plants produce vast amounts of sensor data, logs, imagery, and video – far exceeding the capabilities of traditional monitoring and rule-based analytics.

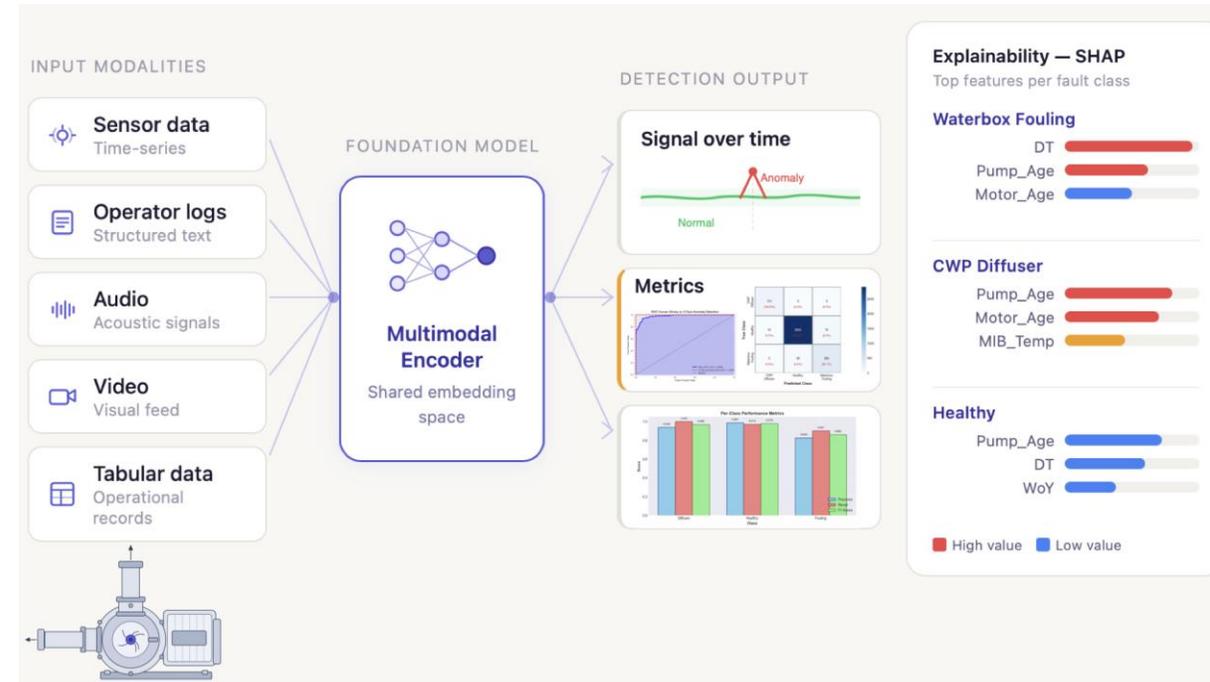
Approach

- Digests various data modalities (sensor data, textual logbooks, images, etc.)
- Foundation model are NOT retrained or finetuned but use in-context learning for nuclear applications.
- Model has been assessed against industrial benchmarks and outperformed all classical methods.

Impact

A step towards autonomous operations, diagnostics and prognostics and remote monitoring and control based on AI decision making.

S&T Challenge: Aligns with "Delivering Nuclear Energy That is Faster, Safer, Cheaper"



Conventional approach: Data analyzed separately, triggers via **simple rule violations. [takes weeks, detects minutes before the anomaly if ever detected]**

AI-driven approach: **All modalities ingested simultaneously [analysis takes minutes, detection can happen weeks ahead].**

SYNAPS-I (AI for Light and Neutron Sources)

Motivation: Advanced light and neutron sources generate petabytes of imaging data, yet most of it remains underutilized. Extracting actionable knowledge at scale requires AI capabilities and close integration with beamlines and HPC.

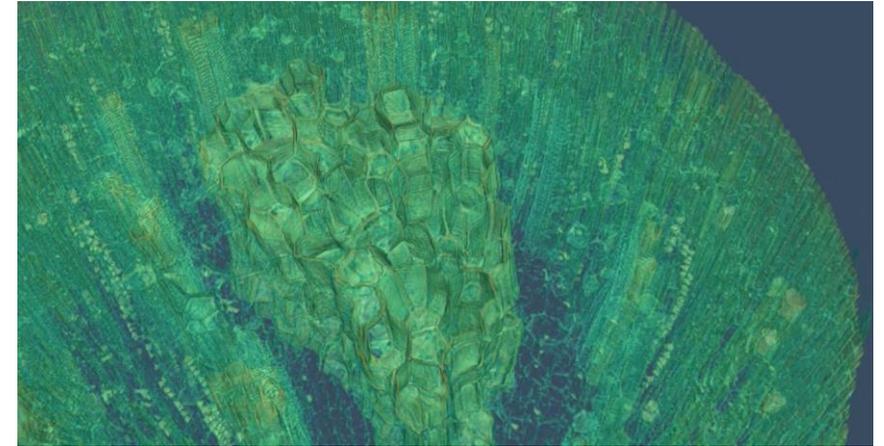
Approach

- Integrate light source beamlines with state-of-the-art industry foundation models—fine-tuned by SUF users and staff—and deploy across thousands of CPUs and hundreds of GPUs on HPC
- Close integration with the American Science Cloud

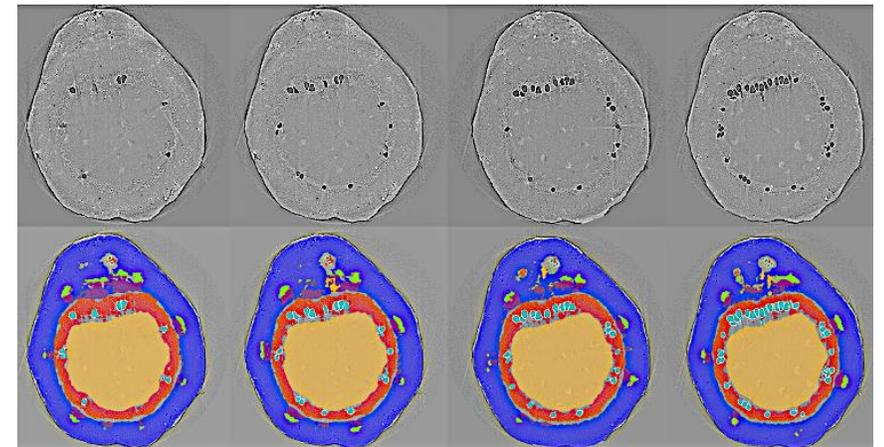
Impact

- Providing semantic segmentation of large datasets within **15 minutes rather than months**.
- Enables unified, cross-facility AI analysis spanning five national laboratories
- Strengthens U.S. leadership in AI-accelerated scientific discovery and critical technologies

S&T Challenge: Aligns with “Designing Materials with Predictable Functionality” and “Achieving AI-Driven Autonomous Laboratories”



3D volume of a grapevine stem



Segmented time series. (factor **3000 speed** increase of manual)

American Science Cloud – AI Universe



The LBNL-led **AI Universe** project unites five US National Labs to build an integrated, AI-ready Cosmic Frontier data repository and use foundation models to advance discovery of dark energy, dark matter, and cosmic structure through multimodal astrophysical data analysis.

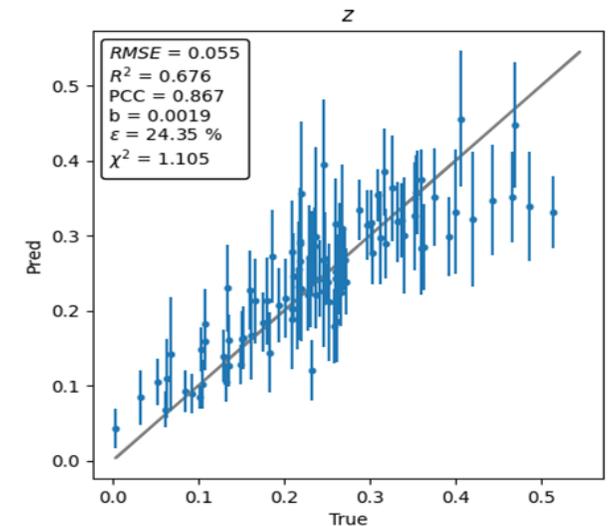
Launched October 2025. PIs: Simone Ferraro & Uros Seljak, co-leading with SLAC

Data curation:

- A data repository for the AI Universe, with particular focus on flagship HEP-funded surveys (DES, Rubin-LSST, DESI) and precursor experiments.
- Standards and workflows for AI-ready CF data.
- Enabling new Cosmic Frontier science with map-centric products, traditional “catalogs”

Scientific Discovery with Foundation Models applied to:

- Photometric redshift estimation
- Galaxy parameter estimation
- Strong lensing detection
- Weak gravitational lensing



FORUM-AI: Intelligent, Scalable Agents for Materials Science & Discovery



Motivation: To solve complex materials science and materials discovery problems by developing agents that autonomously integrate data, simulation tools, and experimental equipment using large-scale HPC resources.

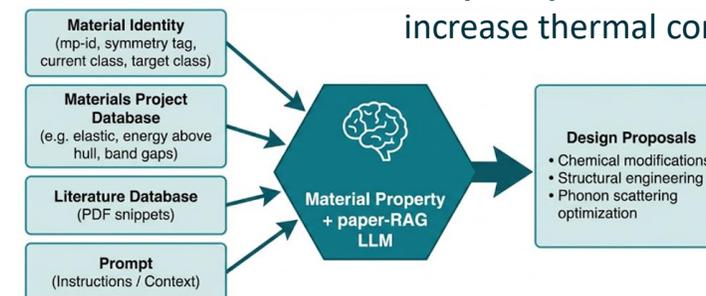
Approach: Develop agents that can reason over existing materials data, make hypotheses, and test and refine those hypotheses by automatically running simulations and experiments to solve complex problems.

Impact

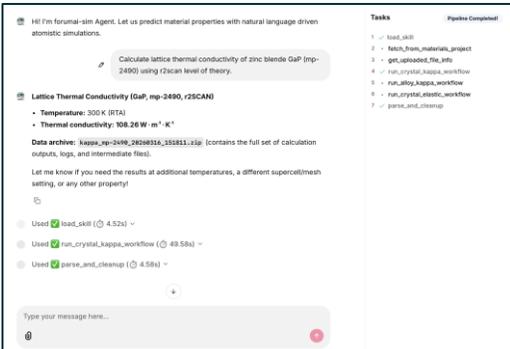
- Agents for data (Materials Project), simulation (atomate2), and robotic experiment (A-lab). Can already custom-tailor some materials properties with simulation (see right)
- Develop fine-tuning strategies on HPC for large foundational potentials to better target specific chemistries
- Extracted new large (>200K curves) spectroscopic data set from the scientific literature for improved reasoning
- Target: hundreds of materials science hypotheses run in parallel over HPC, >2X speedup in solving materials problems

S&T Challenge: Aligns with “Designing Materials with Predictable Functionality” and “Achieving AI-Driven Autonomous Laboratories”

1. Example input: modify GaAs to increase thermal conductivity



2. Orchestration agent designs proposed material (P-doped) by calling upon historical data sets and knowledge



Hi! I'm forum-sim Agent. Let us predict material properties with natural language driven atomistic simulations.

Calculate lattice thermal conductivity of zinc blende GaP (mp-2490) using 23can level of theory.

Lattice Thermal Conductivity (GaP, mp-2490, 23CAN)

- Temperature: 300 K (RT)
- Thermal conductivity: 108.26 W·m⁻¹·K⁻¹

Data archive: kappas_mp-2490_23can_300k_20240115_20240115.zip (contains the full set of calculation outputs, logs, and intermediate files).

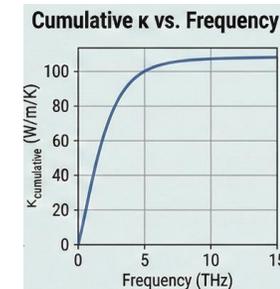
Let me know if you need the results at additional temperatures, a different supercell/mesh setting, or any other property!

- Used load_skel (4.52s)
- Used run_crystal_kappa_workflow (49.58s)
- Used parse_and_cleanmp (4.58s)

Tasks

- load_skel
- load_from_materials_project
- get_unloaded_file_info
- run_crystal_kappa_workflow
- run_alloy_kappa_workflow
- run_crystal_thermal_workflow
- parse_and_cleanmp

3. Tool agents carry out research plan on agent-proposed material



4. Simulation results confirm >2.5X increase in thermal conductivity

Goal achieved in a few minutes w/o user intervention

AI-booster Particle Accelerator Operations & Design

Motivation

Particle accelerators are strategic infrastructure for science, industry and medicine. The largest are among the most complex machines ever built. AI is poised to transform accelerator operations, design, and how research is done at DOE, academia and industry.

Approach

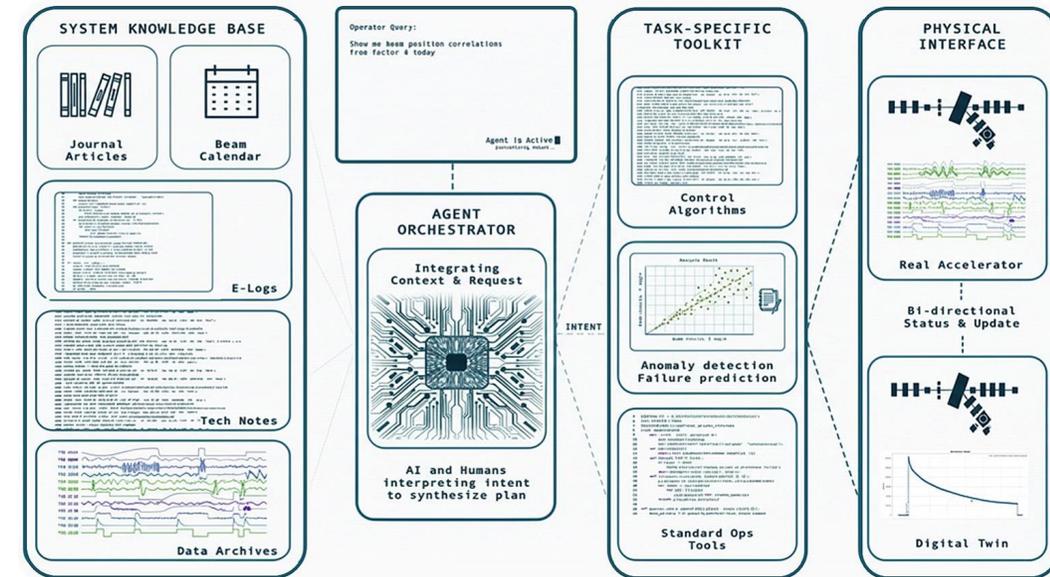
- Develop and implement language interfaces, AI assistants, and information synthesis at unprecedented scale
- Deploy AI-assisted-to-fully autonomous design and operations with preventive fault detection and auto-fault-recovery
- Treats nation's accelerator ecosystem as a single, intelligent system

Impact

- Accelerate discovery, industrial & medical applications with enhanced, intelligent, automated operations and novel accelerator designs that deliver better beams (e.g., higher energy and precision)

S&T Challenge

- Addresses “Enhancing Particle Accelerators for Discovery”
- Impacts/synergies with Quark-to-Cosmos, Scientific Imaging, Fusion & Fission Energy, Microelectronics, Quantum, Autonomous Labs, ...



Agentic AI Osprey for accelerator operations

- Installed at 8 facilities in first phase
- Streamlines access to machine parameters for control (Millions parameters when fully deployed)
- >100x faster operator queries using natural language, e.g., “Show me beam position correlations from sector 4 today.”

One codebase, eight sites, the Genesis Mission's infrastructure, a shared platform for transforming particle accelerators in the age of AI.

Accelerating Biotechnology

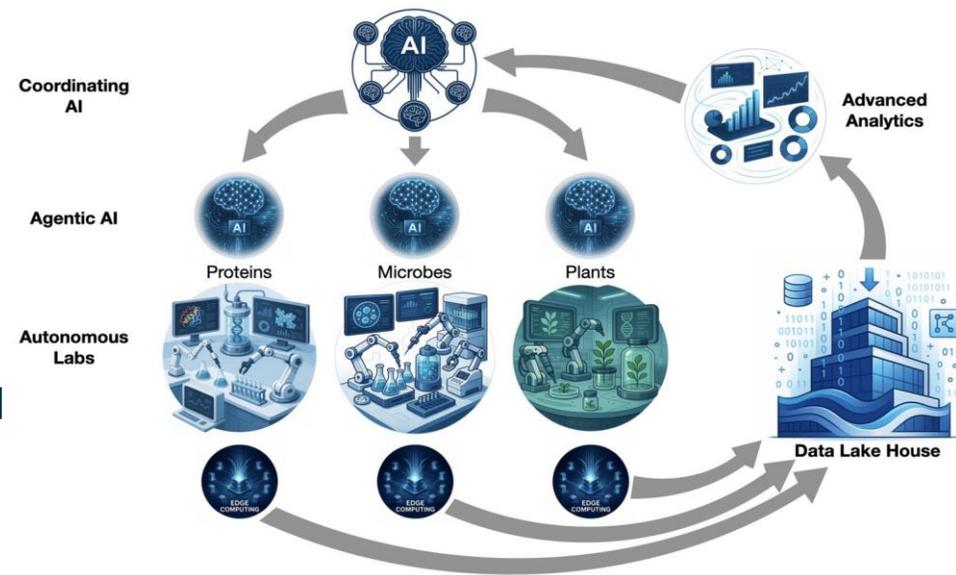
Motivation: Growth of the U.S. economy with biology requires new AI approaches that will accelerate industry adoption.

Approach: The multi-lab OPAL project and the ModCon OPAL-FAMOUS seed project are developing AI models, agentic AI and a network of autonomous laboratories that can learn, adapt, and accelerate breakthroughs across biology, biotechnology, and energy science.

Impact

- AI models redesigned 2 enzymes for critical mineral recovery and bioproduct production.
- AI-guided experiments rapidly doubled the production of the feedstock chemical isoprenol in weeks instead of months.
- Agentic AI provided near-real-time analysis of plant traits, cutting the time from data collection to actionable insight in half.

S&T Challenge: Aligns with “Scaling the Biotechnology Revolution” and “Achieving AI-Driven Autonomous Laboratories”



AI-Powered Rapid Materials Discovery

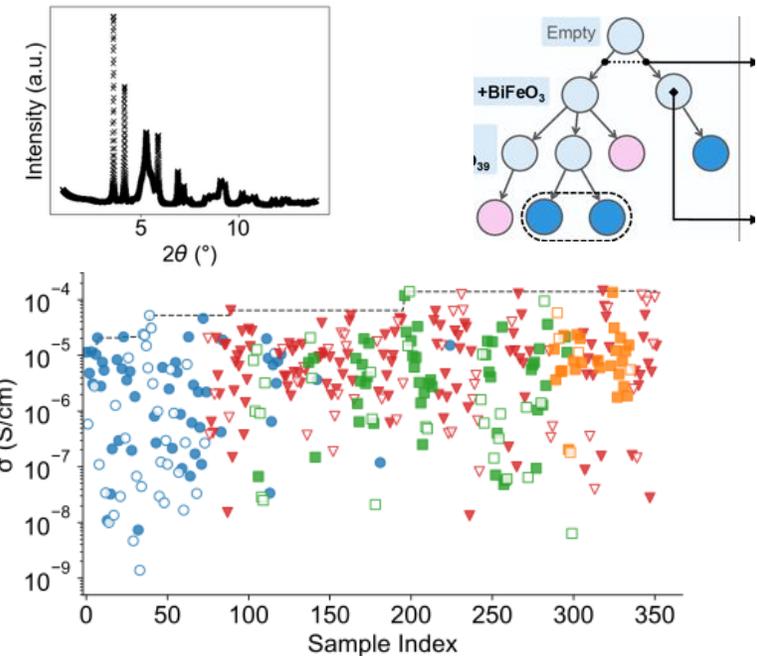
Motivation: New or improved materials empower the US industrial base and its competitiveness, and ensure national security. Materials development is slow and relies on dispersed knowledge and data.

Approach: Autonomous synthesis and testing of materials, integrating LLMs to gather background knowledge, high-throughput computing (Materials Project) for property prediction, and AI for on-the-fly data interpretation and decision making. A-lab is addressable by agentic-AI to be a resource in a variety of discovery campaigns.

Impacts:

- Accelerate materials idea to synthesis and testing by factor of more than 100
- Achieved 24/7 operation on relevant materials chemistries and quantities
- Optimization of materials for energy storage in record time

S&T Challenge: Aligns with “Achieving AI-Driven Autonomous Laboratories” and “Designing Materials with Predictable Functionality”



Transforming Materials Discovery with AI

Motivation: Reliable, systematic data is imperative to build powerful data-driven models. Founded at LBL in 2011, the *Materials Project* pioneered robust web-accessible materials workflows, data, and associated AI/ML analyses algorithms.

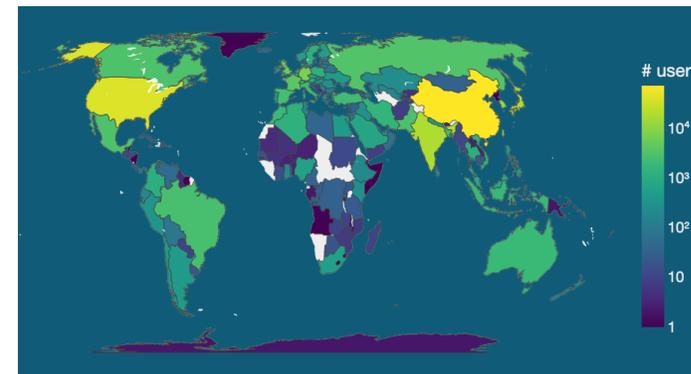
Approach

- The Materials Project combines high-throughput quantum simulations, cloud-based data infrastructure, and AI-enabled tools to compute and organize properties of hundreds of thousands of materials.
- Open APIs enable rapid screening and integration with experiments.

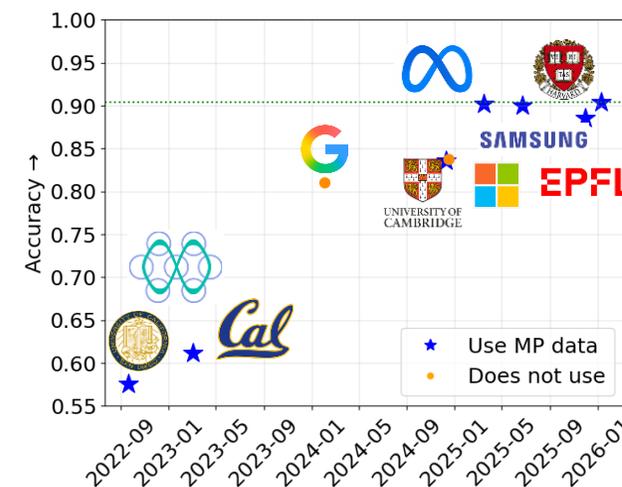
Impact

- Cited over **6 times per day**, 4 times for ML applications.
- Millions of computed materials and experimental data entries.
- Pioneered ML foundation models and standardized benchmarking through unprecedented materials data.
- Acts as innovation multiplier with hundreds of materials discoveries linked to Materials Project data and tools.

S&T Challenge: Aligns with “Designing Materials with Predictable Functionality”



- >700k users across industry + research sectors
- > 1.5k new users per week
- > 4M API requests per week
- >1,000,000 materials data points in AI-ready format



- Trained ML foundation models with >10x speedup over standard methods

PD4M: Advancing Nuclear Physics with Autonomous Data Labeling

Motivation: Advance our understanding of the forces acting in nuclei by having large scale nuclear physics experiments directly inform nuclear models. Provide the data infrastructure to make this possible.

Approach

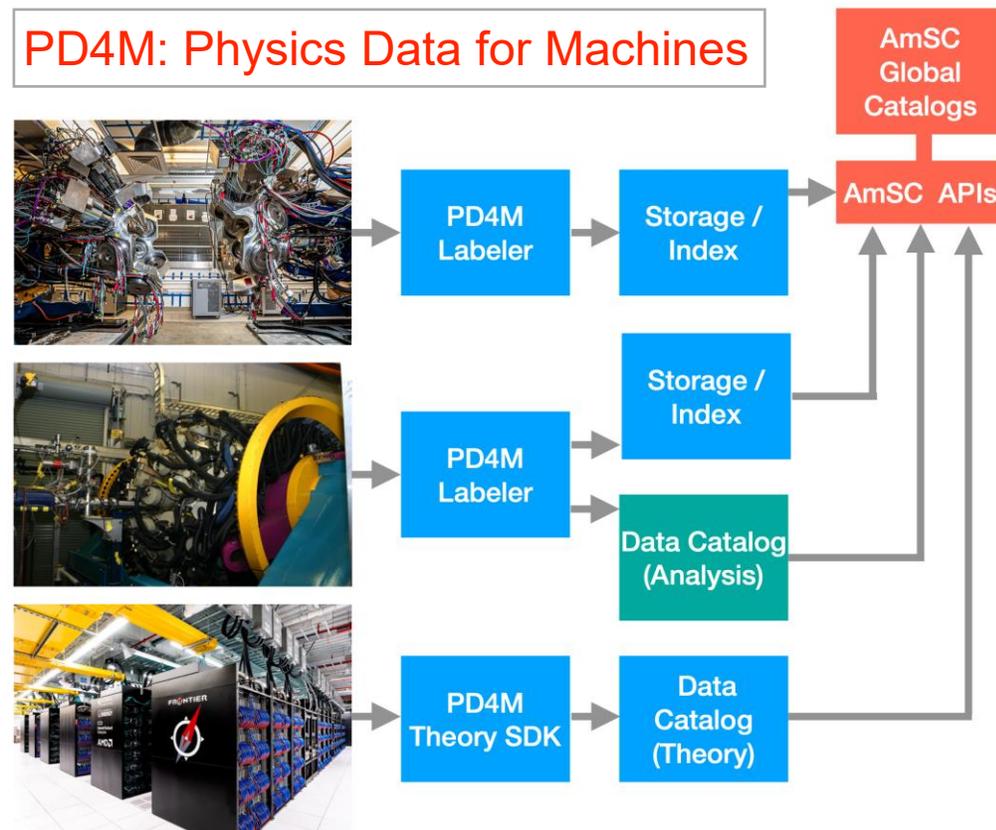
- Live data-taps on operational nuclear physics experiments to stream complete, fully-contextualized data autonomously at scale (~10 Gb/s current → 1 Tb/s projected)
- Interface existing HPC-based nuclear theory codes to AmSC data catalogs

Impact

- Enables autonomous, self-guided experiments in nuclear physics connecting theory and measurement at rates covering all NP devices
- Removes analysis bottlenecks to enable full utilization of data and accelerate analysis by more than a factor of 10

S&T Challenge: Aligns with “Enhancing Particle Accelerators for Discovery” and “Unifying Physics from Quarks to the Cosmos”

PD4M: Physics Data for Machines



Status

- First HPC-based (Frontier) nuclear theory application (IMSRG) already interfaced to prototype data catalog
- Experimental data labeler developed and will be deployed in summer 2026 (GRETA, Gammasphere)

Agentic Advanced Manufacturing for Production Workflows

Motivation: Production environments have long required closed-loop fabrication that measures part quality and adapts the process in real time. This will increase rate production and quality of NNSA-relevant manufacturing, while accelerating the insertion of new designs and materials.

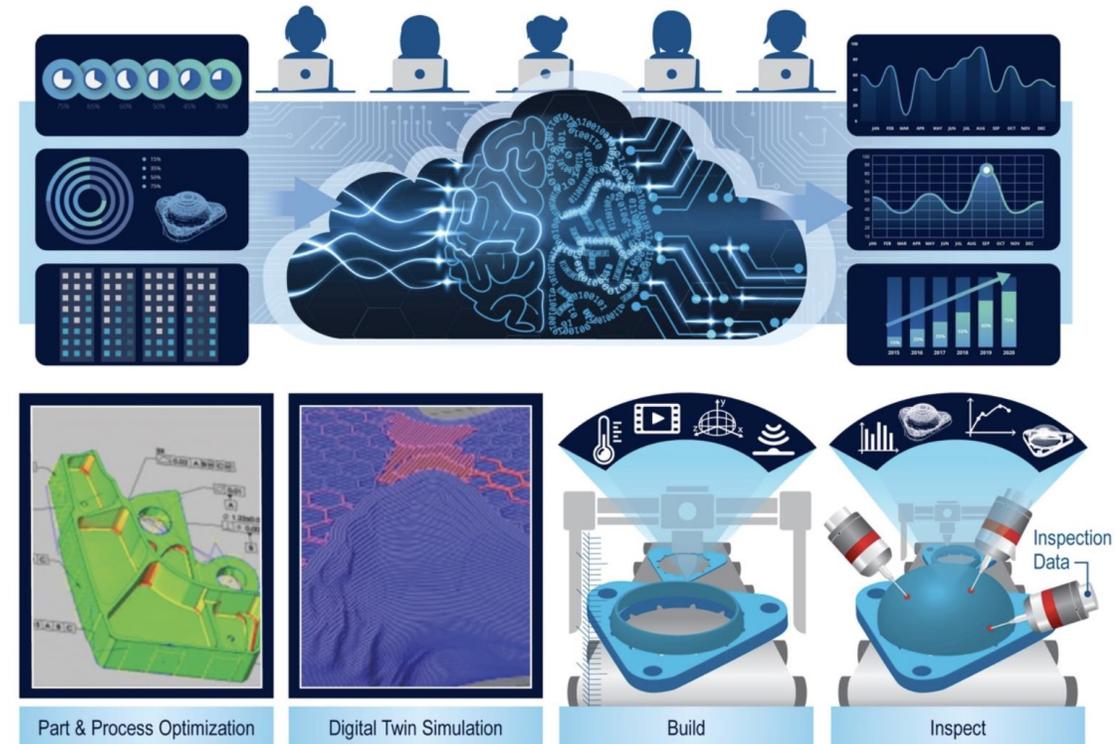
Approach: Build an AI agent that closes the loop for Direct Ink Write (DIW) printing, from job start through measurement, decision, and adjustment.

- *Launch:* Agent starts prints from operator requirements
- *Collect:* Stream images and machine telemetry to a shared data store
- *Analyze:* Apply AI quality measurement, add sensing modalities as needed
- *Decide:* Classify quality and update parameters for next layers or prints, generate as-built digital twins, and optionally launch simulation-based evaluation

Impact:

- *Demonstrate closed-loop DIW printing:* automated print, measure, decide, adjust workflow on a production-relevant DIW platform.
- *Increase reliability and throughput:* detect quality issues early and adapt process parameters in real time, reducing manual intervention and rework.
- *Enable as-built digital twins:* generate process and part representations from on-machine data to support qualification and simulation-ready evaluation.

S&T Challenges: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”





Fusion Experiment Design at the National Ignition Facility

Motivation: After the 2022 fusion net energy gain result, the focus is *robust, repeatable, higher-yield* ICF designs. This requires navigating a *high-dimensional space* of coupled parameters where brute-force search is impractical. AI focuses simulation and experimental effort on the most promising regions.

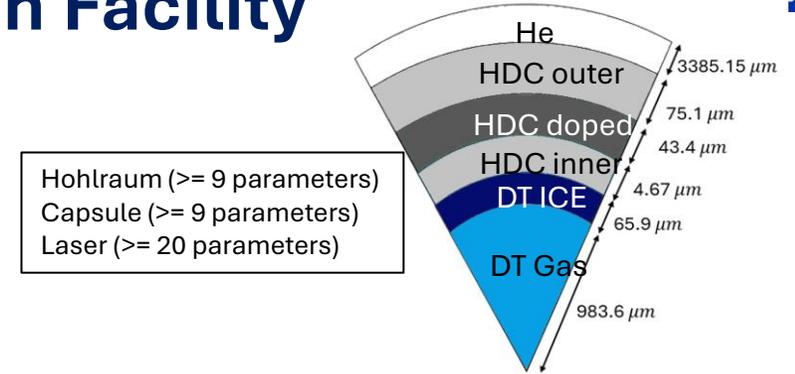
Approach:

- *Agent-based workflows* run simulation campaigns, analyze outcomes, and choose the next designs to evaluate.
- *Surrogate models (emulators)* accelerate costly physics in rad-hydro and kinetics, including *NLTE*.
- *LLM interfaces* coordinate workflows and support designer training via natural-language interaction with models and results.

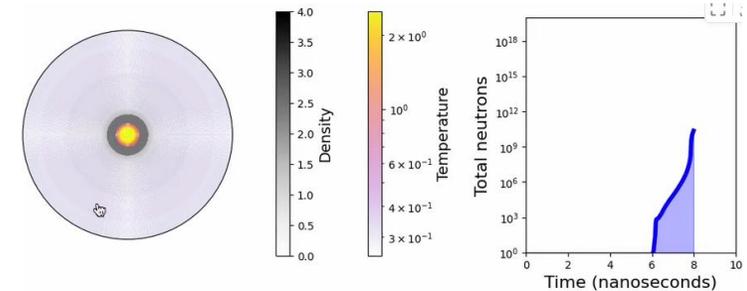
Impact:

- *AI-informed NIF designs* are moving to manufacture for an upcoming shot campaign.
- The workflow scales toward *next-generation, higher-yield facilities* with larger design space.
- Lessons learned support the *NNSA mission*, including *stockpile modernization* through improved predictive design.

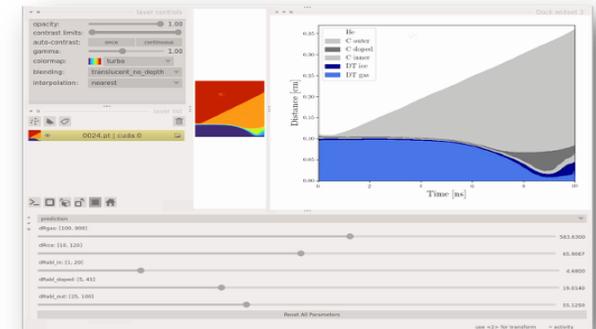
S&T Challenges: Aligns with “*Increasing Experimental Capacity and Nuclear Research Facilities*” and “*Accelerating Delivery of Fusion Energy*”



ICF capsule design spans a high-dimensional, coupled space that cannot be brute-forced, and intuition can take years to build.



Near-real-time simulation feedback, including predicted neutron yield, for AI-prioritized design options.



Interactive exploration links early design choices to late-time performance, accelerating intuition and iteration.

SPARK: Accelerating Alloy Discovery with AI and Self-Driving Labs

Motivation: Critical defense, energy, and aerospace systems operate in extreme environments—shock, high temperature, and radiation—where today’s materials fail. Discovering and qualifying a new alloy takes 5–10 years. This timeline is too slow for national security needs.

Approach:

- Trained models on decades of experimental data to design new alloys from scratch—including how to actually make them.
- Robotic systems autonomously synthesized bulk materials and tested their mechanical performance.
- Experimental results fed directly back into the AI, continuously improving predictions.

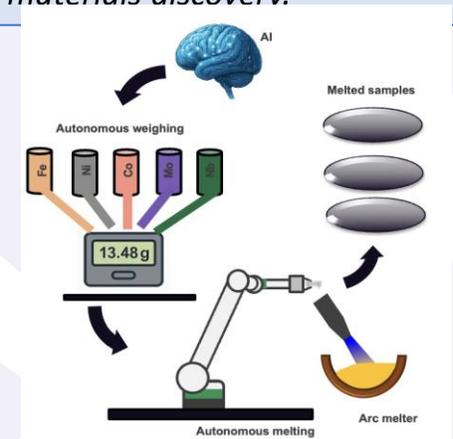
Impact:

- Compress discovery timelines from decades to months.
- Rapidly test materials in relevant environments, closing the gap between discovery and deployment.
- Build a first-of-its-kind high-quality materials database—critical for training next-generation AI models.

S&T Challenge: Aligns with “Designing Materials with Predictable Functionality,” “Achieving AI-Driven Autonomous Laboratories,” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”



Autonomous arc melting, which Los Alamos has demonstrated, is a step toward physical AI in materials discovery.



AI for Structured Composite Development

Motivation: Structured composite materials combine two or more materials for superior mechanical performance. These materials are challenging to design: They are inherently non-homogenous and anisotropic, often exhibiting complex failure modes.

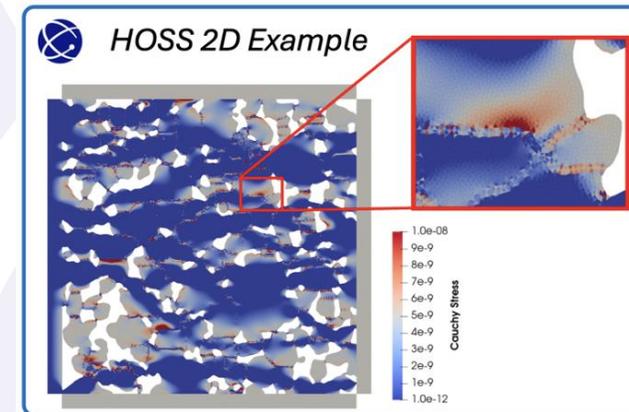
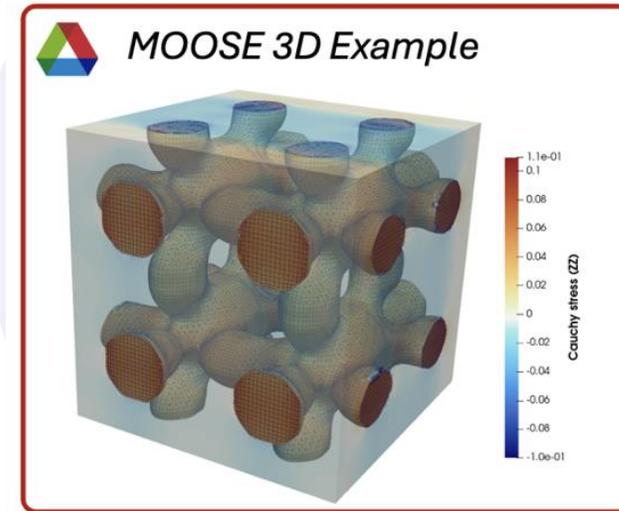
Approach:

- Combined data generated by Los Alamos codes (HOSS) and Argonne codes (MOOSE) via foundation models to determine fracture patterns.
- 3D printed and tested new composite patterns at Argonne.

Impact:

- Develop new stockpile materials with desirable stress/strain/fracture properties.
- Close the loop between AI predictions, simulation, and experiments to develop new agile manufacturing processes.
- Link these workflows with large-scale experiments to accelerate the evaluation of AI-generated hypotheses.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Designing Materials with Predictable Functionality”



Simulation tools used to generate the AI datasets: MOOSE 3D example showing the local stress field under periodic mechanical load; HOSS 2D simulation showing the stress load-bearing paths through the material.

AnnoMate & MicroSentryAI: AI-Assisted Inspection for “Bowtie” Production

Motivation: Manual inspection of precision components is time-intensive and relies heavily on expert judgment. Inspector fatigue and variability can affect consistency, while growing production demands require faster, more reliable inspection workflows.

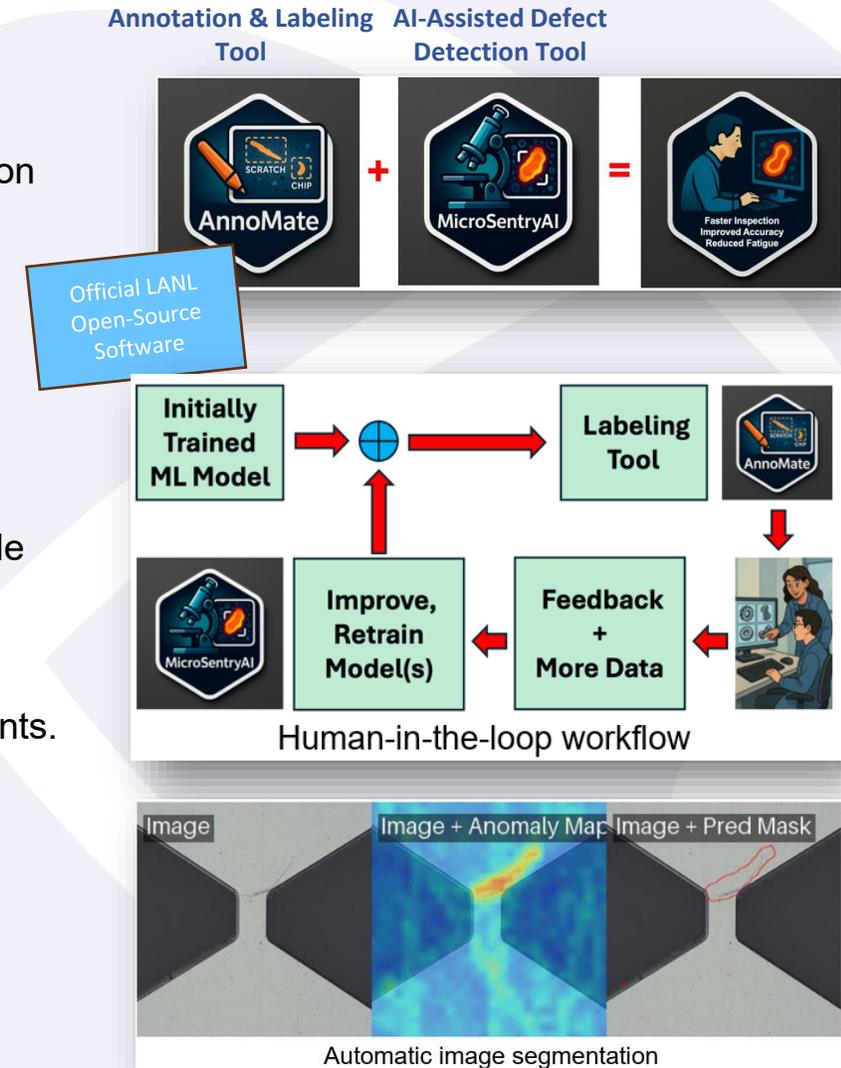
Approach:

- AnnoMate enabled inspectors to label defects by precise location and type
- MicroSentryAI performed AI-assisted defect classification and segmentation
- Human-in-the-loop workflow continuously improves models with new data

Impact:

- Prototype unified inspection tool delivered to users improves inspection consistency while keeping expert inspectors in the decision loop.
- Captures expert knowledge through labeled datasets to support inspector training.
- Supports training of new inspectors using curated inspection data.
- Strengthens manufacturing quality assurance and reliability for mission-critical components.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”



CHELATRON: AI-Driven Molecular Discovery for Cancer Treatment

Motivation: Targeted alpha therapy uses radionuclides such as ^{225}Ac Actinium to treat cancer. Molecules called chelators bind and deliver these radionuclides to target sites. New chelators with stronger binding affinities can bolster treatment effectiveness and limit collateral damage to healthy cells.

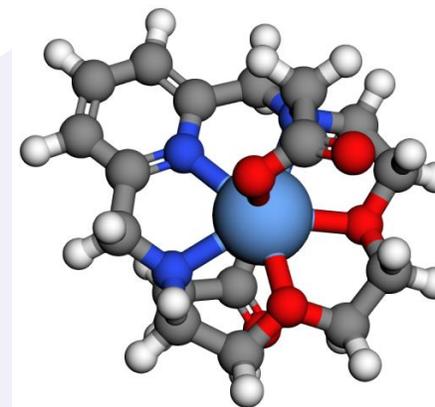
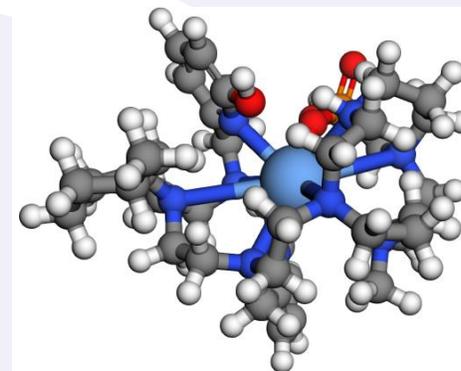
Approach:

- Developed interpretable chemical hypotheses using NVIDIA's Nemotron LLM.
- Turned hypotheses into molecules by sampling A/B group chemical spaces with NVIDIA's GenMoL.
- Estimated binding affinity of complexes using LANL's Architector complex constructor tool and high-throughput quantum calculations. Feed assessments back into LLM for iterative hypothesis refinement.

Impact:

- Hundreds of new chelators can be screened overnight—more in 13 hours than in 50 years of research.
- Experimental evaluation of candidate chelators is underway.

S&T Challenge: Aligns with “Securing America’s Critical Minerals Supply” and “Achieving AI-Driven Autonomous Laboratories”



These chelators are predicted to have a stronger binding affinity than Macropa (the state-of-the-art chelator in nuclear medicine).

URSA: Universal Research and Scientific Agent

Motivation: URSA creates an ecosystem where specialized autonomous agents collaborate to “do science.”

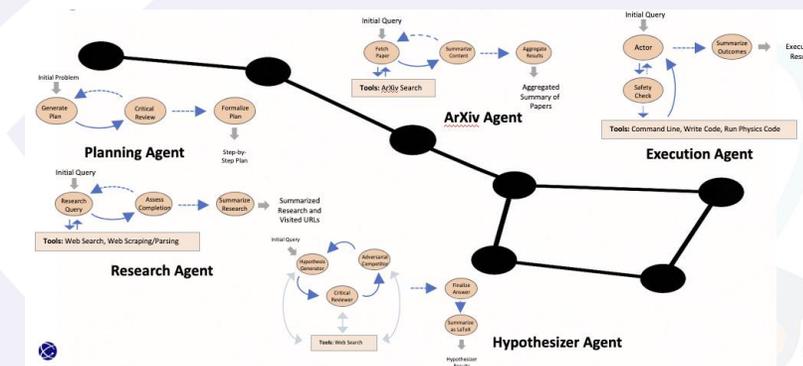
Approach: Composable agents each handle a specific subtask of the research workflow:

- Planning Agent decomposes high-level scientific objectives into actionable steps.
- Execution Agent generates and runs code and orchestrates computational experiments.
- arXiv Agent retrieves, downloads, and summarizes scientific papers (PDFs) to inform research context.
- Web Search Agents performs online searches to gather domain-relevant knowledge.
- Hypothesizer Agent proposes and iteratively refines hypotheses.

Impact:

- Enables automated materials discovery using an agentic workflow that begins with a high-level objective.

S&T Challenge: Aligns with “Designing Materials with Predictable Functionality,” “Achieving AI-Driven Autonomous Laboratories,” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”



The URSA workflow. URSA has been used for materials discovery, ICF capsule design, and more.

AI-Based Monitoring of Manufacturing Processes Using Electrical Signals

Motivation: Traditional defect detection relies on post-process inspection and additional sensors. A scalable, real-time monitoring method is needed to detect anomalies during welding and machining without increasing instrumentation complexity.

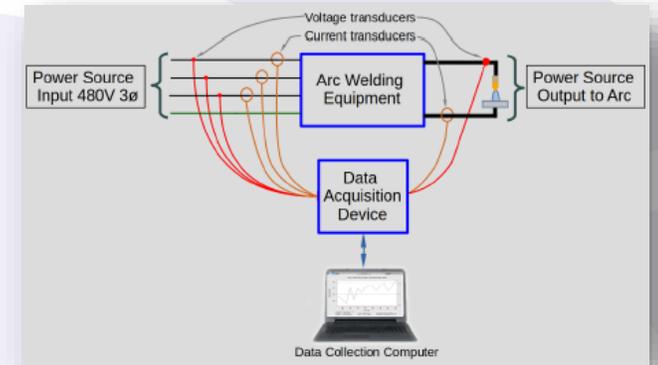
Approach:

- Analyzed electrical load signals to detect anomalies without additional sensors.
- Evaluated multiple AI architectures to distinguish nominal vs. off-nominal welds
- Deployed lightweight, low-parameter models suitable for edge environments

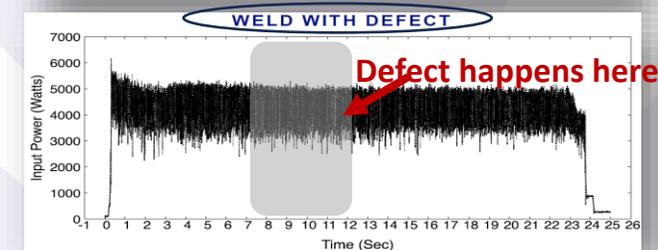
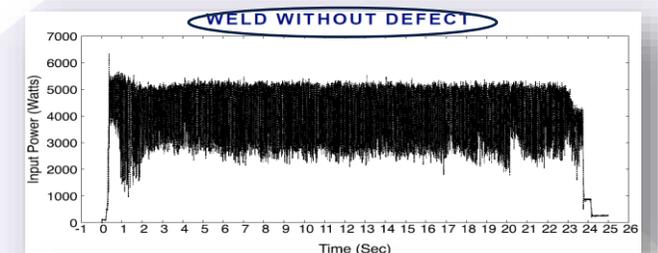
Impact:

- Demonstrates feasibility of non-intrusive-load-monitoring-based real-time anomaly detection.
- Enables early detection of welding anomalies before post-process inspection.
- Supports more reliable manufacturing by identifying off-nominal process behavior in real time.
- Reduces rework, material waste, and downtime while accelerating quality assurance.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”



NILM monitors power consumption without modifications, collecting input data.



AI for Materials Fracture Research

Motivation: Most catastrophic material failures occur through a sudden brittle fracture. It is imperative to develop materials that do not fail under specific conditions such as extreme temperatures, loading, and impact.

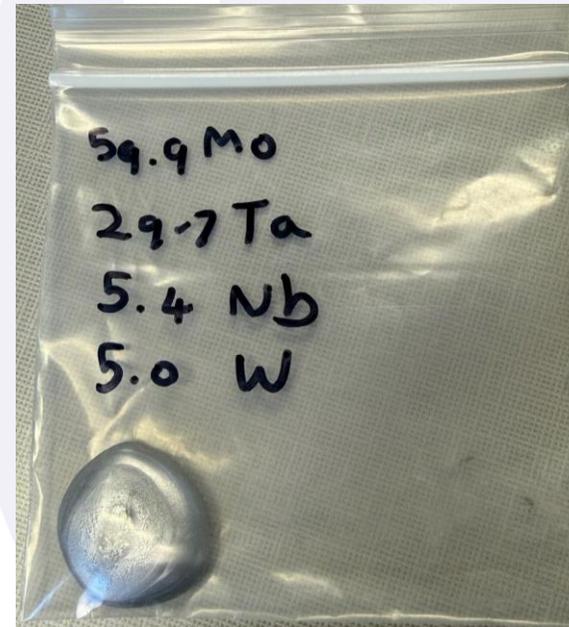
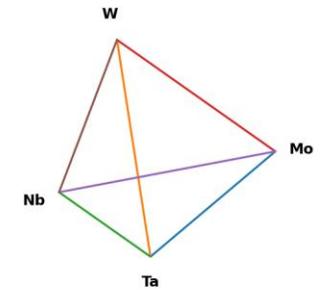
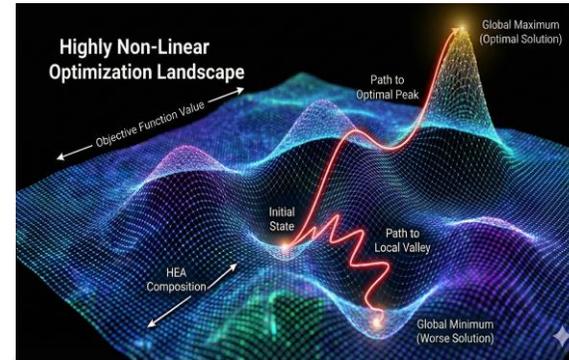
Approach:

- Developed an AI-guided workflow that fuses literature mining, simulations, and optimization.
- Robotically manufactured high-entropy alloys for testing and evaluation.

Impact:

- Accelerate discovery by increasing productivity and efficiency by 100x.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence”



Element	Measured %
Mo	59.9
Ta	29.7
Nb	5.4
W	5.0

Above, a sample of robotically manufactured AI-optimized material.

Accelerating the Design/Test Loop of Additively Manufactured Foams, Polymers, and Composites

Motivation: Designing silicone foams for mission applications traditionally requires slow trial-and-error experimentation. Additive manufacturing enables precise control of foam structure, but identifying optimal process parameters remains time-intensive.

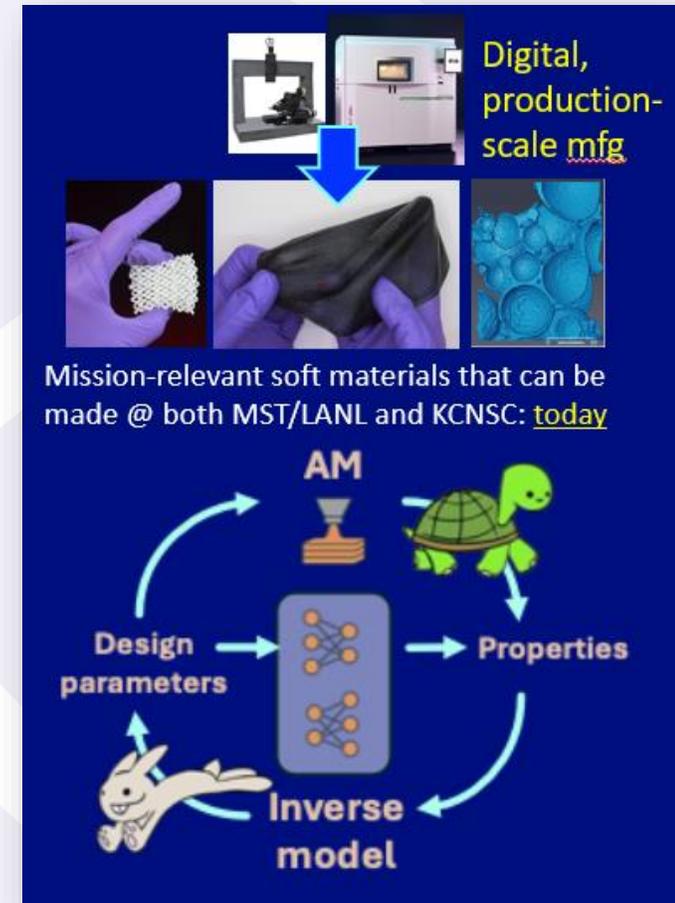
Approach:

- Developed Digital Foam, an inverse ML model linking direct-ink write print parameters to final foam performance
- Combined additive manufacturing, high-throughput testing, and ML to learn process-structure-property relationships
- Collected thermal imaging data to support AI-based defect detection for selective laser sintering.

Impact:

- Reduced print iteration cycles (by at least a factor of two) and accelerated silicone foam design workflows.
- Generated mission-relevant datasets enabling ML-driven materials and manufacturing optimization.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”



AI-Enhanced Monitoring for Airborne Plutonium Detection

Motivation: Continuous air monitors (CAMs) protect workers by detecting airborne radioactive material. However, environmental noise and statistical fluctuations can produce false alarms, interrupting operations and requiring costly investigation.

Approach:

- Trained physics-guided AI models using synthetic radiation spectra from known isotope energies.
- Classified borderline detection events to distinguish real signals from background noise.
- Continuously improved model performance using operational monitoring data.

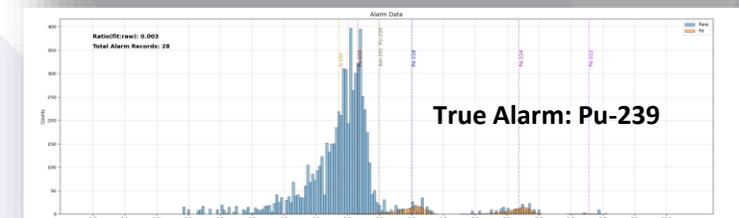
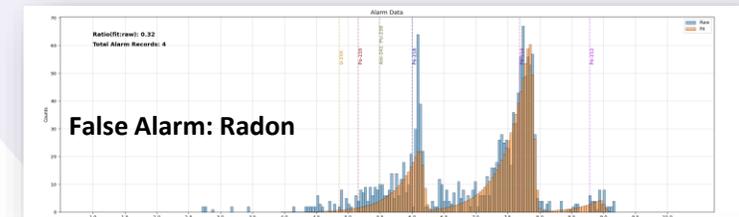
Impact:

- Reduces false alarms while maintaining reliable airborne contamination detection.
- Improves operational efficiency in high-security manufacturing environments.
- Enhances worker safety through faster interpretation of radiological monitoring data.
- Provides a scalable, cost-effective enhancement to existing CAM systems.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Reenvisioning Advanced Manufacturing and Industrial Productivity”



A CAM alarm.



AI-Guided Hydrogen Sink Discovery

Motivation: Hydrogen sink materials are key system components. Traditional sink discovery is a trial-and-error process.

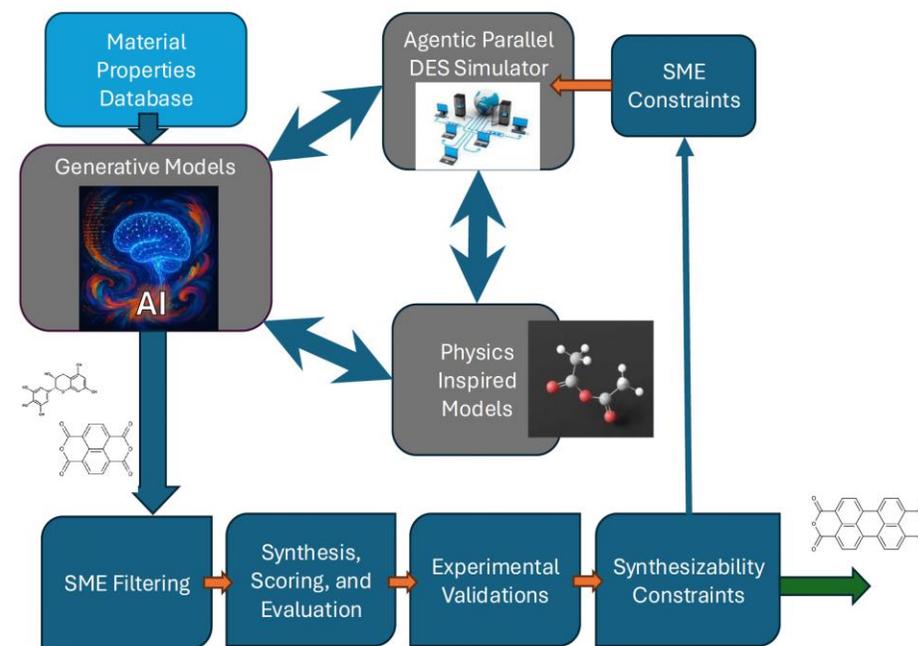
Approach:

- Developed a parallel discrete event simulation based on an agentic AI framework that coordinates physics modes, surrogate AI/ML models, SME inputs, and generative AI.
- From 50k AI-proposed H₂ sink candidate molecules, SMEs down-selected 300 with uptake capacity ranging from 125% to 200% of the capacity of legacy materials.
- SMEs further down-selected 29 molecules for synthesis. Experimental results are forthcoming.

Impact:

- Extends system lifetime by slowing the aging process of materials.
- Serves as a bridge for transitioning AI-guided predictions into synthesis and performance validation.

S&T Challenge: Aligns with “Accelerating Materials Discovery, Production, and Qualification for Strategic Deterrence” and “Designing Materials with Predictable Functionality.”



Hydrogen sink materials are used for absorbing and potentially recycling various hydrogen isotopes. Generative AI can dramatically shorten the discovery process.

CM²US

Motivation: To unlock the full potential of American energy, the nation needs a secure, reliable and resilient supply of critical minerals and materials (CMM).

Approach: NETL researchers are co-leading this national initiative to enhance CMM discovery, development and supply chain resilience.

Impacts: NETL contributes specialized expertise and tools, including:

- **AI-Accelerated Data Management:** Utilizing platforms like NETL's EDX ClaiMM to provide essential data.
- **Advanced Modeling:** Using the Unconventional REE and Critical Minerals modeling suite to forecast resource potential.
- **Partnership Development:** Co-leading efforts to engage commercial partners to test and accelerate capabilities.



Genesis Alignment:

The Genesis Mission directs targeted investment and advancement in critical areas like AI and CMM. Mobilizing technologies in these areas can spur industry development and secure national economic and scientific advantages.

American Science Cloud



U.S. DEPARTMENT
of ENERGY



AmSC

Motivation: To provide DOE and the National Laboratories with a platform to collaborate on crucial, cutting-edge scientific research.

Approach: NETL's Energy Data eXchange® (EDX) is a key infrastructure partner in the AmSC, providing crucial data, expertise and infrastructure to this national scientific computing resource.

Impacts:

- The EDX platform allows personnel across the DOE and National Laboratories to securely and efficiently connect to AmSC and its growing repository of data, models, AI agents and capabilities.
- NETL's EDX team is actively involved in the AmSC co-design process, transferring established data governance, technical capabilities and institutional knowledge to align with AmSC's broader platform requirements and scientific mission.



Genesis Alignment:

The Genesis Mission aims to build a platform for the world's most complex and powerful scientific discoveries.

EDX CLaiMM

NETL launched EDX CLaiMM, the first digital platform dedicated to supporting U.S. research and innovation across the critical minerals and materials (CMM) value chain, unlocking digital tools and data from resource to supply. Hosted on NETL's EDX platform, ClaiMM gives CMM researchers, scientists and stakeholders access to curated CMM datasets, databases, models, tools and services within a secure, authoritative and adaptive digital research ecosystem.



MapIT

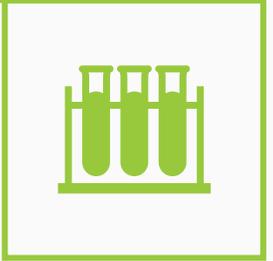
The Subsurface Energy Systems Mapping Inquiry Tool (MapIT) is a published, online web mapping tool that expertly aggregates millions of features and attributes from published, authoritative sources, into an AI-ready digital collection. MapIT is designed to help users discover available public-sourced subsurface data to inform discovery and innovation for subsurface energy exploration and characterization efforts for resource identification (e.g., critical minerals, hydrocarbons or geothermal).

NEWTS Database Dashboard

The National Energy Water Treatment and Speciation (NEWTS) database and dashboard displays sites across the nation where energy-related wastewater stream samples and composition data have been collected. To date, these digital products have been downloaded for use over 32,000 times by a range of stakeholders. Using the database and dashboard, users are able to quickly explore and obtain sample data from locations where various government agencies collect and analyze water samples from energy-related wastewater streams.

FlowDash Geothermal Energy Enhancer

An NETL-led team developed FlowDash Geothermal Energy Enhancer, a new machine learning (ML) tool that can reduce risks and costs for geothermal energy operations and enhance safety for neighboring communities. The award-winning FlowDash tool demonstrated superior predictive accuracy, creativity and exploration indications.



SAMI

NETL established the Science-based Artificial Intelligence and Machine Learning Institute (SAMI), which leverages science-based models, AI/ML methods, data analytics and high-performance computing to accelerate applied technology development. These advances empower NETL to push the frontiers of AI technology and create next generation architectures, tools, and approaches; expand partnerships to advance the development and adoption of AI; and foster AI workforce development.

Wafer-Scale Engine

NETL researchers, in partnership with a key pioneering computing innovator, took a significant step forward in harnessing the power of the world's largest computer chip — the Wafer-Scale Engine — by using an application programming interface designed in-house to connect computational fluid dynamics software through data-file sharing with the next-generation computing technology. Efforts have recently included development of subsurface modeling capabilities that are more energy efficient, computationally efficient, and AI-optimized using the wafer scale engine computational architecture.

GISA

NETL innovated the Geospatial and Information Substitution and Anonymization (GISA) tool, a published, digital model that leverages AI to analyze and remove sensitive information from products prior to publishing. Datasets shared between partners and DOE can contain sensitive information that may prevent or delay data from becoming public or shared with other entities. GISA helps the producers of those data prepare anonymized versions for public use and reuse.



URC Model & Tool

The Unconventional Rare earth element and Critical minerals prospectivity analysis tool, or the URC model and Tool, was developed to reduce U.S. reliance on foreign rare earth elements and critical minerals producers by systematically identifying new, unconventional domestic sources. The URC model and tool was developed by NETL using an AI-driven, big data modeling approach that forecasts and identifies new sources of critical minerals and rare earth elements (REEs) in unconventional deposits like coal ash, mine waste, and *in situ*, but non-traditional geologic sources.

AI-accelerated Permanent Magnet Discovery

Motivation: New low-cost strong permanent magnets are essential to the continued growth of the U.S. economy and reduce dependence on foreign sources of rare earths and rare earth magnets.

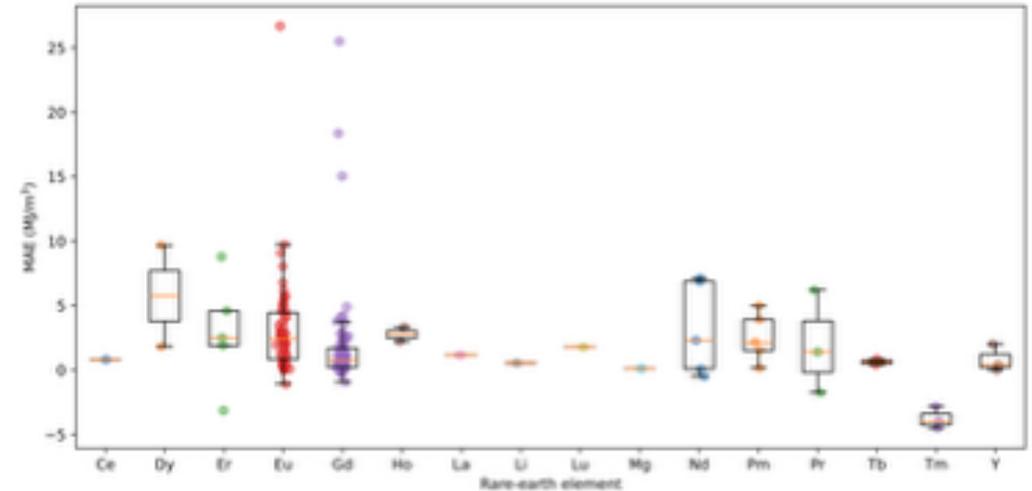
Approach: ORNL trained AI models on structural and compositional data of calcium pentacopper–structure rare earth materials for discovery of new high-magnetic anisotropy materials.

Impacts:

- 5-10× accelerated discovery of new permanent magnet materials (correlating to months instead of years).
- Decreased reliance on critical rare earths neodymium and dysprosium for generation of coercivity (ability to resist demagnetization) in permanent magnets.

S&T Challenge:

Aligns with “Securing America’s Critical Minerals Supply” and “Designing Materials with Predictable Functionality”



Distribution of predicted magnetic anisotropy values by rare earth constituent, with positive values indicating highly desirable uniaxial anisotropy. Anisotropic materials exhibit directionally dependent physical properties, allowing retention of strong, directional magnetism.

AI Advancements in Magnet Discovery

Over 70,000 prototypes theoretically studied

Over 400 thermodynamically stable compositions identified

171 materials identified with single-axis anisotropy, which provides directional magnetism

AI-Accelerated Nuclear Reactor Safety Analysis to Enable New Nuclear Energy Deployment

Motivation: Current NRC safety reviews are expert-intensive and slow, a process inadequate to meet the increased demand of advanced reactor applications.

Approach: LLMs used in this workflow are accessed from the Genesis Mission Platform. This work leverages years of DOE and NRC investments in software tools for confirmatory analysis of reactor regulation (SCALE, MELCOR, BlueCRAB).

S&T Challenge: Aligns with “Delivering Nuclear Energy That is Faster, Safer, Cheaper” and “Harnessing America’s Historic Nuclear Data and Research”

Impact: Goal of reducing reviews from 2 years to 2 weeks with a SCALE-based confirmatory analysis; verified by assessing Kairos Power’s Hermes 1 construction permit through an analysis of power profiles and operating parameters.

NRC Impact: Consistent, AI-assisted workflows, reducing subjectivity and staff burden.

Industry Impact: Pre-submittal access to NRC agentic workflows through Genesis will improve license applications.

DOE Impact: Workflow adoption and use will influence future resource investments to fill gaps through experiments and/or computational modeling.



3D visualization of Kairos Hermes prototype reactor

AI-Driven Autonomous Manufacturing Accelerates Materials Production and Qualification

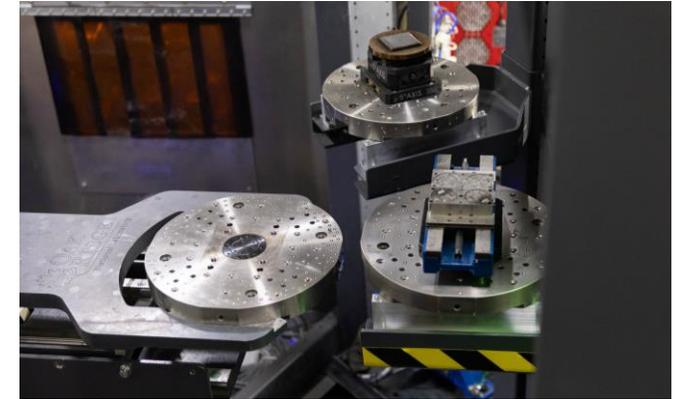
Motivation: Advanced manufacturing processes require extensive trial-and-error experimentation, manual process tuning, and slow inspection cycles, delaying deployment of critical materials and components for energy, defense, and infrastructure systems.

Approach: ORNL integrated machine learning, reinforcement learning control, and AI-enabled imaging across manufacturing systems to autonomously optimize process parameters, detect defects, generate toolpaths, and correct deviations in real time. Demonstrations include:

- AI-optimized carbon fiber and silicon carbide fiber production, wire-arc additive manufacturing, and automated toolpath generation.
- AI-accelerated X-ray CT inspection.

Impacts: Autonomous experimentation reduced manufacturing cycle times by up to 80%, improved fiber quality by 100%, accelerated inspections, and reduced scrap and qualification timelines.

These capabilities lower production costs, strengthen domestic supply chains, and accelerate deployment of advanced materials.



Components move seamlessly and autonomously between manufacturing stages within ORNL's Future Foundries platform, enabling a fully integrated, highly adaptable production workflow.

S&T Challenge: Aligns with “Reenvisioning Advanced Manufacturing and Industrial Productivity,” “Securing America’s Critical Minerals Supply,” and “Delivering Nuclear Energy that is Faster, Safer, Cheaper”

Scalable AI for Real-Time Accelerator Operation and Experiment Steering at Neutron Facilities

Motivation: AI is poised to optimize accelerator operations while reducing neutron experiment time and accelerating scientific discovery for impacts in industry, medicine, and energy.

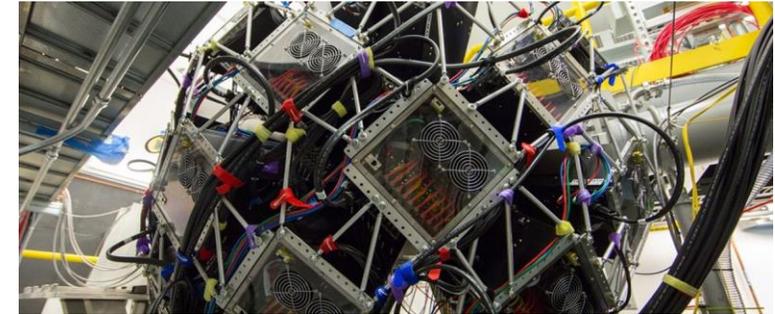
Approach:

- Developed an AI model to better analyze neutron scattering experiments, reducing experiment time.
- Installed and ran Osprey Agentic Workflow package on Enterprise and Accelerator networks with gtp-oss-120b LLM; linked Osprey to the Spallation Neutron Source (SNS) Control System.
- Using ML models to minimize beam losses and detect anomalies before downtime occurs.
- Potential deployment of AI-assisted-to-fully autonomous design and operations with preventive fault detection and auto-fault-recovery.

Impacts:

- By bridging the SNS and the exascale supercomputer Frontier with AI for real-time analysis, scientists can run neutron experiments more efficiently, reducing experiment times by 30-50% via on-the-fly data analysis and adjustments.
- Demonstrated outstanding scalability of the machine learning model on Frontier (scaling up to 4,608 GPUs).
- Continual learning framework will allow ML models to adapt to new accelerator configurations and drifts.

S&T Challenge: Aligns with “Achieving AI-Driven Autonomous Laboratories” and “Enhancing Particle Accelerators for Discovery”



A single crystal diffractometer at ORNL's Spallation Neutron Source.

Agentic AI Osprey for accelerator operations

- Installed at 8 facilities in first phase
- Streamlines access to machine parameters for control (millions of parameters when fully deployed)
- >100x faster operator queries & response for key activities
- One codebase, eight sites, Genesis Mission's infrastructure

Transforming Biodesign for the Nation's Energy and Security

Motivation: Connecting genes to desired plant and microbial traits is complex but essential to designing biotechnologies that meet American needs for energy, critical materials, and manufacturing.

Approach: Aligned with the Genesis OPAL Project, agentic AI leverages the Frontier exascale supercomputer for real-time analyses of data from ORNL's autonomous Advanced Plant Phenotyping Laboratory (APPL), exploring whether roots can “mine” dilute critical minerals from soils.

Impacts:

- Agentic AI monitors, interprets, and recommends on-the-fly adjustments to experiments to accelerate adaptive learning for bioengineering.
- Reduces analysis time of APPL imaging from 168 hours to <1 minute.
- Enables adaptive workflows and smart experiments for bioengineering (enhancing plant uptake of critical minerals).

S&T Challenge: Aligns with “Scaling the Biotechnology Revolution,” “Achieving AI-Driven Autonomous Laboratories,” and “Securing America's Critical Minerals Supply”



ORNL's Advanced Plant Phenotyping Laboratory and Frontier exascale supercomputer provide the ideal environment to apply AI for enhancing plant uptake of CMMs using engineered microbes.

Revolutionizing Quantum Materials Discovery with AI

Motivation: Quantum materials can enable next-generation technologies in computing, communication, and sensing. The Multimodal AI for 2D Quantum Magnets (MAIQMag Pathfinder) ModCon project will accelerate the detailed understanding of underlying interactions in materials required to enable breakthroughs and strengthen U.S. leadership in quantum and advanced AI.

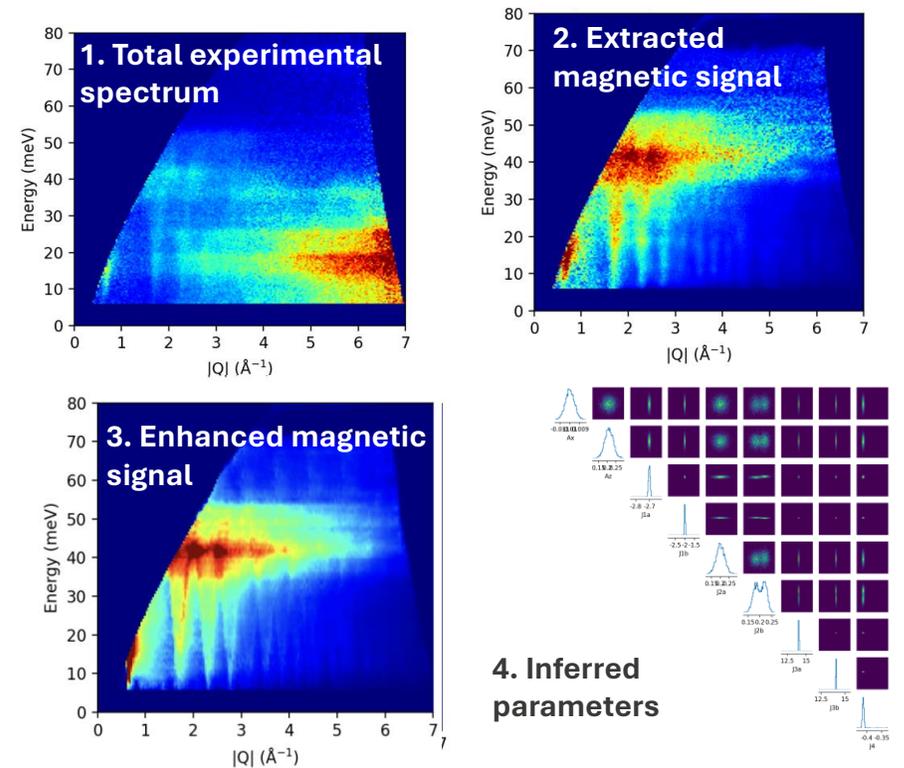
Approach:

- Develop AI methods to rapidly uncover the essential atomic and electronic interactions in complex quantum materials by simultaneously using multimodal data from DOE neutron and light sources. This will transform a slow, manual process into an automatic, precise, and predictive capability, rapidly accelerating the ability to understand new materials and predict their functionality.
- Use AI to propose solutions that align with multiple experimental data from light and neutron sources to determine how energy levels, dynamics, and properties emerge, reducing a process that once took months or years to just days or hours.

Impacts:

- More accurate predictions with reduced uncertainty compared to single-data approaches.
- Deployable across all DOE Basic Energy Sciences facilities.
- Faster scientific discoveries and improved use of DOE facilities.

S&T Challenge: Aligns with “Designing Materials with Predictable Functionality” and “Achieving AI-Driven Autonomous Laboratories”



An AI-enabled workflow for solving quantum interactions.

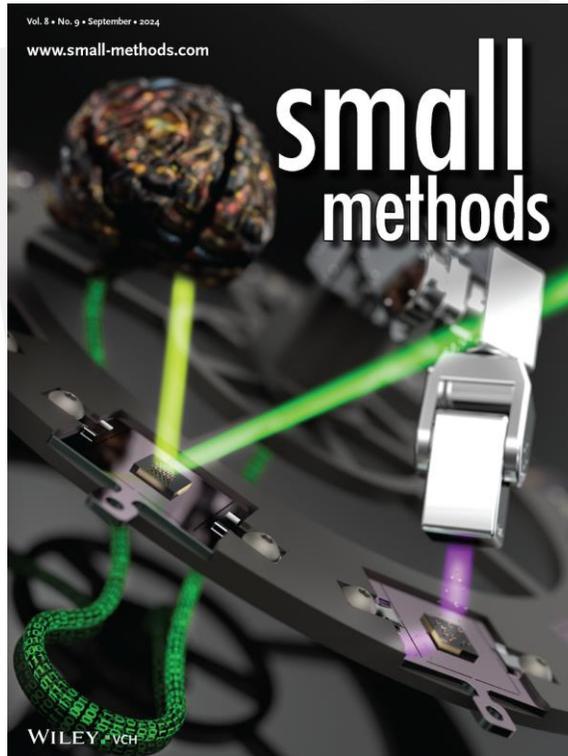
AI Advancements in Quantum Materials

- >10× faster experiment-to-solution time
- Comprehensive solutions by uncertainty-aware multimodal inference
- Scalable to high-dimensional spaces

First-Ever Autonomous Material Synthesis by Pulsed Laser Deposition (PLD)

Motivation: Synthesis of thin films (for basic scientific study of 2D, quantum, and electronic materials) has traditionally relied upon slow, sequential processes with substantial human intervention.

Approach: For the first time, a self-driving experiment grew ultrathin films of WSe_2 . LLMs augmented with literature knowledge are used to generate hypotheses, self-driving synthesis systems autonomously run the experiments, and other AI models correlate complex data to thin film characteristics.



An illustration depicting the novel automated PLD platform developed at ORNL that enables AI-driven optimization of growth conditions for thin films, establishing a novel workflow applicable to any materials that can be synthesized using PLD.

Impacts:

- Results demonstrated 10-100× faster exploration and synthesis.
- The novel approach enables universal accelerated discovery and autonomous growth materials for electronics, energy, etc.

S&T Challenge: Aligns with “Recentering Microelectronics in America,” “Designing Materials with Predictable Functionality,” and “Achieving AI-Driven Autonomous Laboratories”

Enabling Tritium Fuel Production for Domestic Fusion Energy

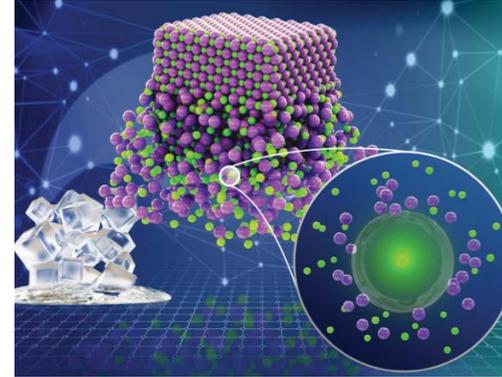
Motivation: To optimize the production and recovery of tritium fuel generated during breeding reactions in a tokamak molten salt blanket material.

Approach: Harness the convergence of agentic AI, quantum computing, and leadership-class HPC to explore a vast chemical space for optimal conditions that maximize tritium production. Large-scale high-fidelity simulations (enabled by AI surrogate models and enhanced by quantum computing results) will drive the discovery process.

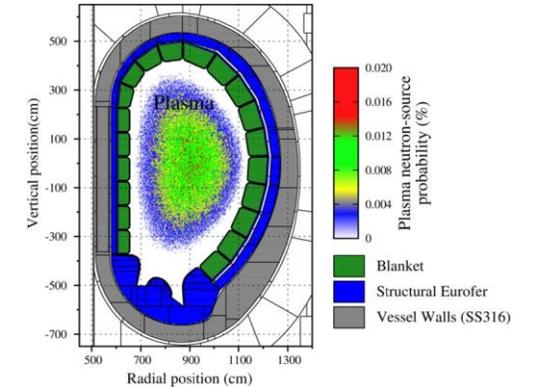
Impacts:

- Acceleration of the design process through rapid sampling of the phase diagram composition space, coupled neutronics models, and high-fidelity simulations to predict blanket properties.
- Extreme-accuracy tritium affinity predictions in the molten salt.
- Guided design of additives that control the tritium chemistry for effective recovery essential for net energy production.

S&T Challenge: Aligns with “Accelerating Delivery of Fusion Energy” and “Delivering Nuclear Energy That is Faster, Safer, Cheaper”



AI-driven simulations of molten salt free energies. Y. Shi, S. Lam, T. L. Beck, Chem. Sci. 2022



Schematic of a blanket structure surrounding the high temperature plasma in the Commonwealth Fusion Tokamak

AI and Quantum Advancements for Fusion Energy

- Agentic AI orchestration of the coupled design cycle
- Advanced surrogate models based on the SuperSalt method for molten salt properties
- AI acceleration of quantum computing results by optimizing quantum circuit design and error mitigation/correction

Fission: The results also apply to tritium removal from molten salt small modular reactors.

Real-Time, AI-Based Water Intelligence for U.S. Energy Systems

Motivation: Water availability limits U.S. energy production and infrastructure resilience, yet current prediction systems are too slow and too coarse for real-time decision-making, constraining hydropower operations and flood response.

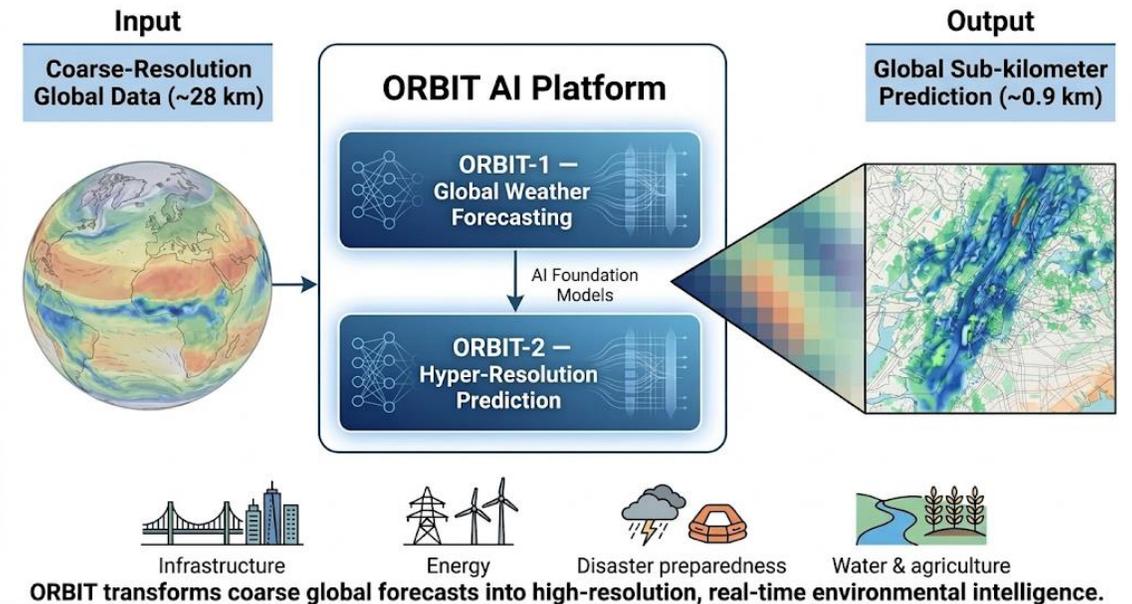
Approach:

- Develop ORBIT, a unified AI foundation platform that transforms global weather forecasts into sub-kilometer environmental intelligence.
- Train on DOE exascale supercomputers and deploy anywhere to deliver real-time predictions.

Impacts:

- Enables real-time water forecasting for energy systems
- Strengthens grid reliability and infrastructure resilience
- Democratizes access to high-resolution prediction
- Turns days of supercomputer simulation into seconds on a standard computer

S&T Challenge: Aligns with “Predicting U.S. Water for Energy”



AI Advancements for Water for Energy

- Global prediction at 0.9 km resolution
- Orders-of-magnitude speedup: days → seconds
- Up to 52% higher accuracy at extended forecast horizons
- Achieve up to 99% accuracy for global prediction
- Being integrated into NVIDIA Earth-2

AI Foundation Models to Drive Fusion Energy Deployment

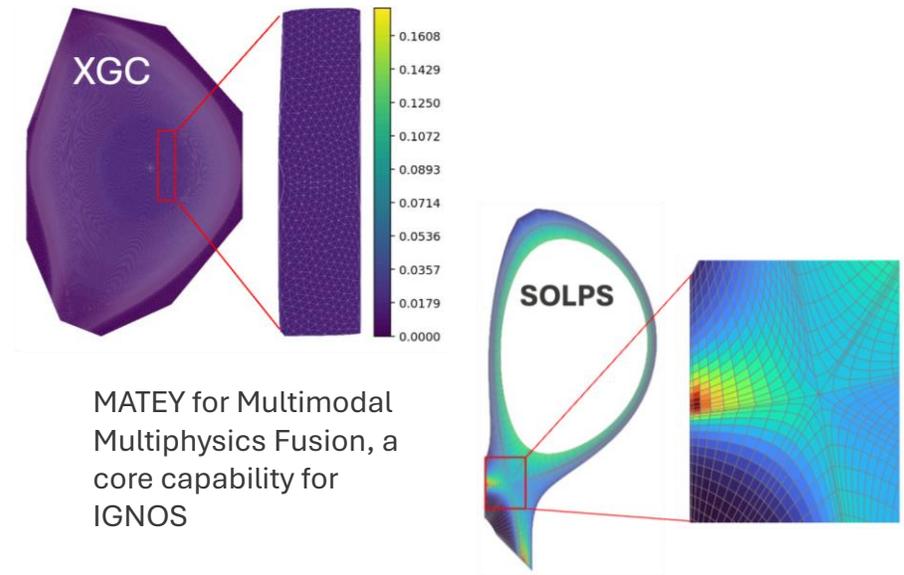
Motivation: Move toward a unified fusion foundation model that can learn across diverse fusion data sources, codes, and plasma regimes. Ingest data from state-of-the-art fusion codes with differing physics systems, fidelity, and format to enable broader review of data.

Approach:

- Built on multiphysics modules to support multimodal fusion data.
- Used CNN and graph convolutional operators for structured and unstructured meshes, respectively, to capture local physics.
- Used attention to capture long-range dependencies across data types.

Impacts:

- Allows one Fusion FM for both structured and unstructured data.
- Expands the amount and diversity of fusion data that can be used for training, improving the path toward more general and robust models.
- Supports the broader DOE mission by enabling foundation models to work across heterogeneous scientific datasets from experiment and simulation.
- Leverages collaboration with PPPL and ANL in FusionFM Seed project.



S&T Challenge: Aligns with “Accelerating Delivery of Fusion Energy” and “Delivering Nuclear Energy That is Faster, Safer, Cheaper”

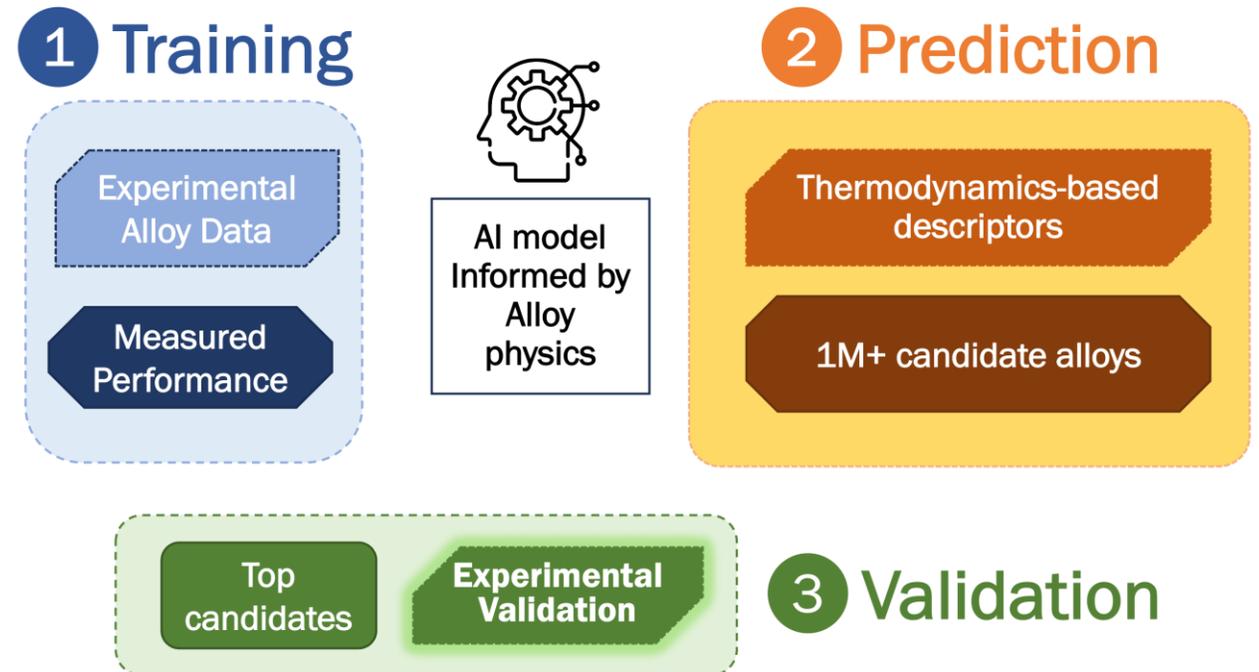
Rapid AI-Guided Design of High-Temperature Alloys

Motivation: Developing alloys for extreme environments depends on complex interactions among composition, phase stability, and processing.

Approach: We combined ORNL’s experimental alloy dataset with AI models informed by thermodynamics-based features to screen more than one million candidate compositions before experimental validation.

Impacts: This approach could enable a 200,000× acceleration in the speed of design for corrosion-resistant, high-temperature steels by reducing the experimental search space. It also improves confidence that the AI predictions reflected underlying alloy behavior, not just statistical correlations.

S&T Challenge: Aligns with “Reenvisioning Advanced Manufacturing and Industrial Productivity” and “Designing Materials with Predictable Functionality”



Workflow combining experimental data, thermodynamics-informed AI, large-scale virtual screening, and targeted experimental validation

AI-Enabled Physics Analyses for Energy, Defense, and Discovery Applications

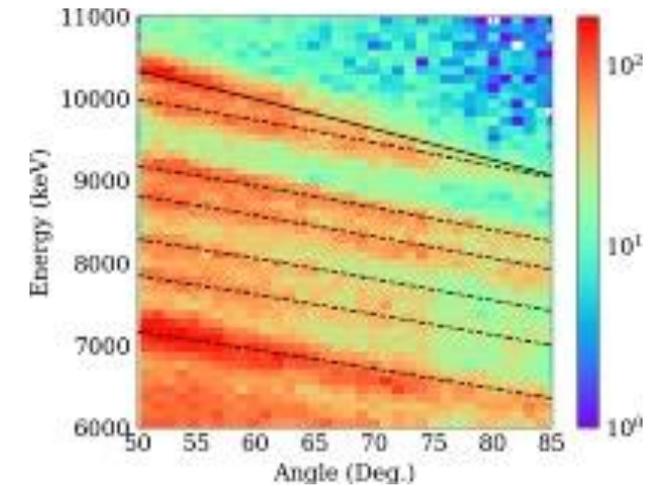
Motivation: Understanding the atomic nucleus, and the reactions that create and destroy nuclei in extreme environments (e.g., nuclear and fusion reactors, thermonuclear weapons, and stellar processes) requires increasingly complex experimental tools and analysis pipelines.

Approach: Algorithms were developed to predict and subtract backgrounds at nuclear physics facilities including ATLAS, FRIB, and SECAR. Bayesian normalization of CMOS images in the recoil separator at FRIB and neural networks for rare event searches in FRIB Decay Station data were used.

Impacts: 20% reduction in experiment time for determination of neutron-induced cross sections and measurements of charge state distributions—results relevant to discovery science, stockpile stewardship, and nuclear energy—corresponding to ~\$200k savings in facility operational costs per experiment.

S&T Challenge:

Aligns with "Unifying Physics from Quarks to Cosmos," "Strengthening Deterrence Through Attribution of Nuclear and Radiological Signatures," and "Achieving AI-Driven Autonomous Laboratories"



A random forest algorithm predicts and subtracts backgrounds from data from the ORNL-led experimental campaign at ATLAS without the need for new measurements. This allows for faster determination of neutron-induced cross sections relevant to astrophysics and myriad applications.

Optimizing Microorganisms to Accelerate Commercial Readiness of Biotechnologies

Motivation: Biomanufacturing contributes to economic prosperity and energy independence. Optimizing the microorganisms that bolster production of high-value fuels, chemicals, and materials accelerates these innovations.

Approach:

- PNNL advanced the open-source BacterAI platform to predict when microbial growth will stop and identify where more detailed analysis is needed.
- Reinforcement learning connects AI models for biological discovery to lab automation such as liquid handling to significantly increase experimental throughput.

Impacts:

- Algorithm explores large parameter space efficiently and quickly to build comprehensive data that is AI-ready.
- Continuous closed-loop experimentation that analyzes results and issues new experiments in minutes, compared to hours or days in human-oriented workflows.



S&T Challenge: Aligns with *Scaling the Biotechnology Revolution and Achieving AI-Driven Autonomous Laboratories*

AI Analysis of Nanoribbon Assembly Reveals Protein Design Principles

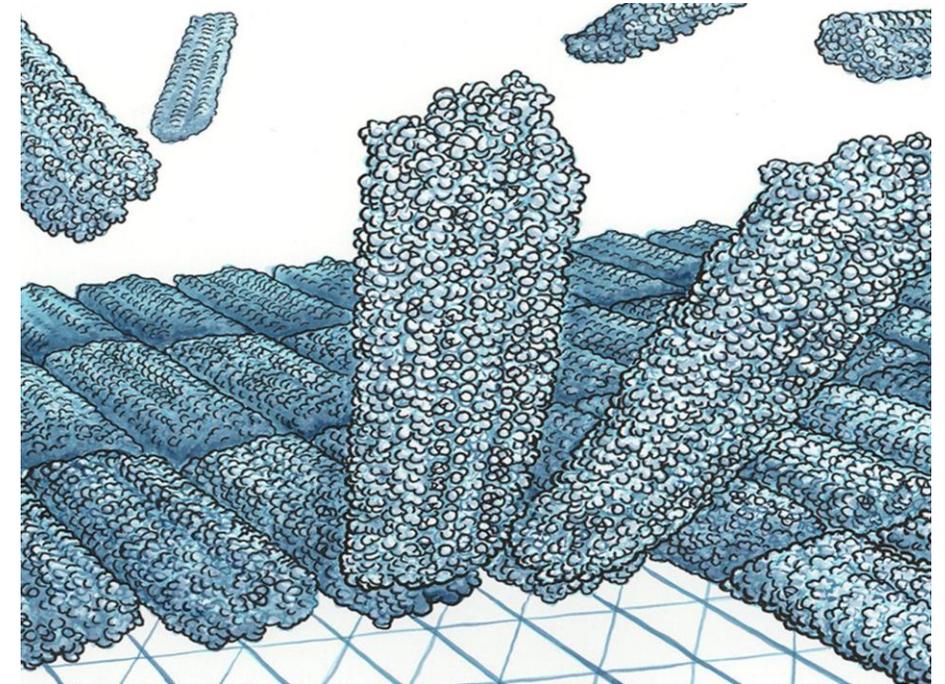
Motivation: Inorganic materials that mimic hierarchical structures found in nature could be designed and produced with desirable properties for catalysts, biosensors, critical element recovery, and biomedical devices.

Approach:

- Analyzed atomic force microscopy to quantify nanorod orientation and organization enabled by AtomAI computer vision model.
- Explored how protein nanoribbons assemble on mineral surfaces after observing unexpected organization.

Impact:

- Machine learning algorithms are enabling studies that otherwise would be impossible due to imagery complexity.
- Discovered that physical forces are missing from protein design algorithms; physics-informed machine learning is needed to account for solvent effects in protein design.



S&T Challenge: Aligns with *Designing Materials with Predictable Functionality*

Enabling the Recovery of Critical Materials from Complex Feedstocks via Selective Precipitation

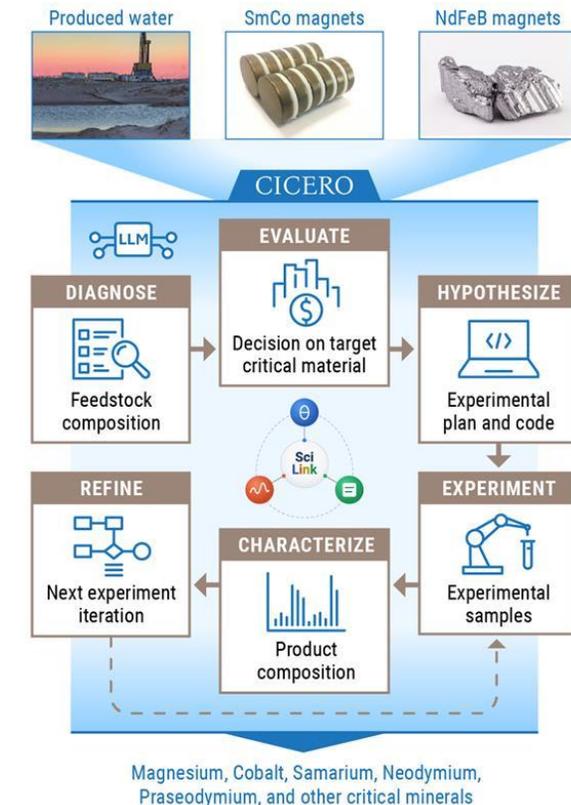
Motivation: The ability to efficiently and cost-effectively recover critical materials from complex feedstocks and waste streams addresses an urgent and growing need for domestic supplies of elements for energy technologies.

Approach:

- Designed an agentic AI-guided closed-loop workflow that autonomously plans, evaluates, and optimizes selective precipitation campaigns for critical materials recovery.
- Builds on PNNL's SciLink LLM multi-agent automation platform.

Impact:

- Recovered magnesium from produced water, samarium from SmCo magnets, and rare earth elements from NdFeB magnets in high purity and yield.
- Shortens timeline for development of efficient, adaptable, and scalable separations from months or years to days.



S&T Challenge: Aligns with *Securing America's Critical Minerals Supply*

Speeding Up Essential Groundwork for Quantum Computing

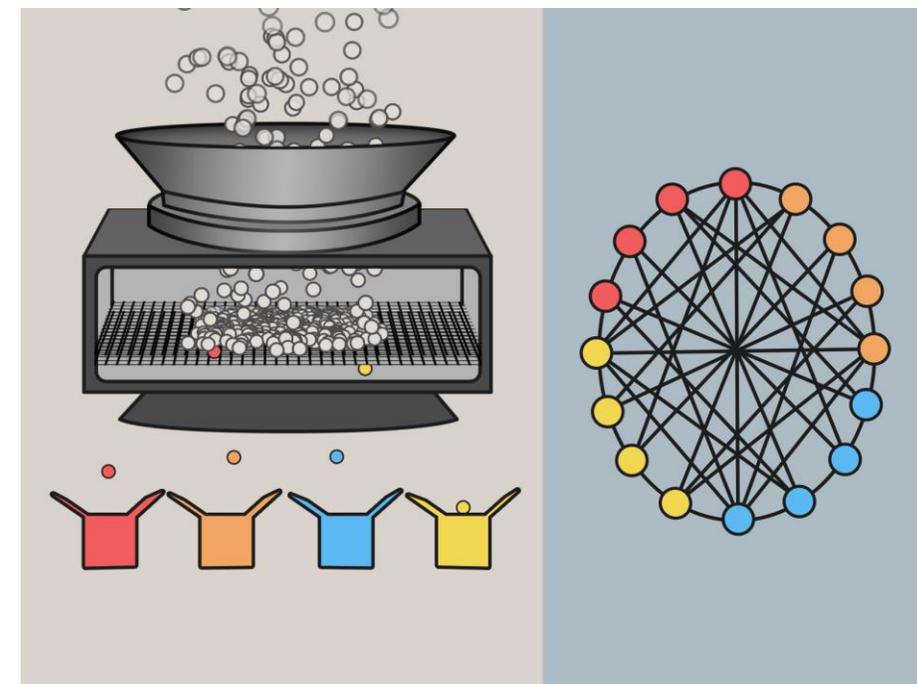
Motivation: Quantum computers promise to speed up complex calculations dramatically, but data must be packed differently to avoid bottlenecks and gain full benefit.

Approach:

- PNNL developed “Picasso,” a parallel algorithm for efficient hybrid computing that uses classical computation to prepare quantum data for quantum computing.
- Leverages AI-augmented graph coloring to compute space-efficient clique partitioning that quickly sorts terms into as few groupings as possible for computation on a quantum computer.

Impact:

- Solved a problem with 2 million Pauli strings and a trillion-plus relationships in 15 minutes using only 40 GB GPU memory while current tools are typically limited to systems with tens of thousands of Pauli strings.



S&T Challenge: Aligns with *Discovering Quantum Algorithms with AI and Realizing Quantum Systems for Discovery*

Autonomous Synthesis for Redox Flow Battery Experiments

Motivation: Redox flow batteries enable grid-scale energy storage that supports grid flexibility and energy reliability.

Approach:

- Integrated robotics and AI into an autonomous workflow to optimize synthesis of fluorenone derivatives, a key component of redox flow batteries.
- Leveraged flexible-batch Bayesian optimization, an active learning method that can suggest multiple paths per experimental round.

Impact:

- By iteratively varying reactant concentrations, reaction time, and temperature, PNNL researchers discovered the conditions that resulted in over 90 percent yield of the fluorenone product at relatively mild temperatures.
- Exploring same options with traditional methods would take about 15 years of continuous data collection.



S&T Challenge: Aligns with *Achieving AI-Driven Autonomous Laboratories, Scaling the Grid to Power the American Energy Economy, and Designing Materials with Predictable Functionality*

Artificial Intelligence Accelerates Nuclear Forensics Analysis

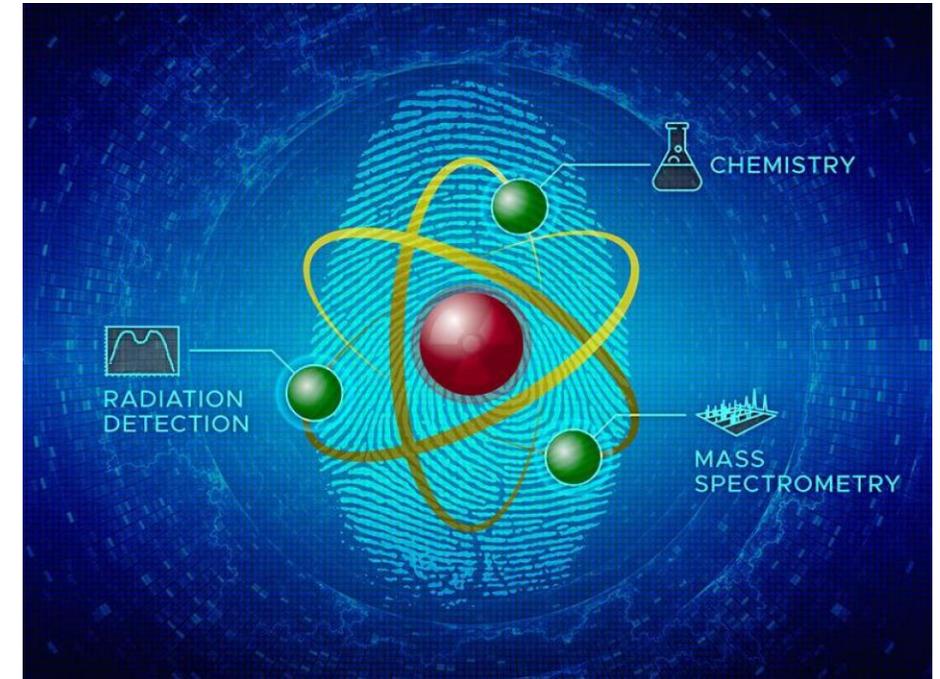
Motivation: Determining the fingerprints of nuclear debris helps understand and protect against nuclear threats.

Approach:

- Developed AI-driven method for understanding complex mixes of radioactive debris, advancing AI+HPC for chemical simulation through PNNL-Microsoft partnership on Discovery Platform.
- Simulated complex computational chemistry to prioritize and target chemical separations needed to investigate an event and identify materials' origins.

Impacts:

- Chemical separation modeling calculates properties of possible molecular combinations, significantly reducing experimental validation timelines.
- A new workflow demonstrates integration of traditional chemical simulations with generative AI, reducing computational workload to respond to nuclear events.



S&T Challenge: Aligns with *Accelerating Nuclear Threat Assessment, Preparedness, and Response; Safeguarding Nuclear Materials from Proliferation Threats; and Strengthening Deterrence Through Attribution of Nuclear and Radiological Signatures*

Shortening Environmental Cleanup Mission Timeline by Optimizing Glass Formulations

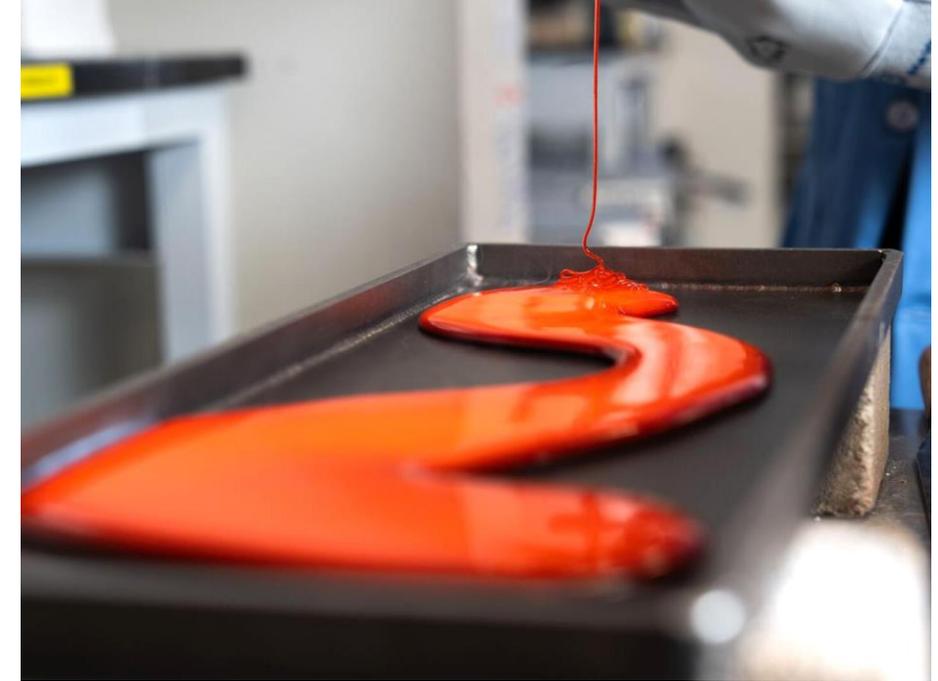
Motivation: Millions of gallons of low-activity, radioactive waste from the Manhattan Project and Cold War is being converted into glass for safe long-term storage. Optimizing glass formulas increases waste loading and process efficiency, reducing costs and mission duration.

Approach:

- Uses computer modeling to predict “recipes” for glass design that accounts for each tank’s chemically unique and nonuniform waste.
- Computer models are trained on hundreds of historical testing results. Waste analysis data is fed into an algorithm that generates the corresponding glass design.

Impact:

- Glass design process at Hanford now takes less than 120 minutes and is getting faster, compared to the traditional approach used at the West Valley Demonstration Project that took about a decade.



S&T Challenge: Aligns with *Transforming Nuclear Cleanup and Restoration*

Accelerating Delivery of Fusion Energy: Taming instabilities through real-time control

Motivation: Fast surrogate physics models are required to predict, mitigate and control problematic instabilities in fusion systems.

Approach – raw AI from data

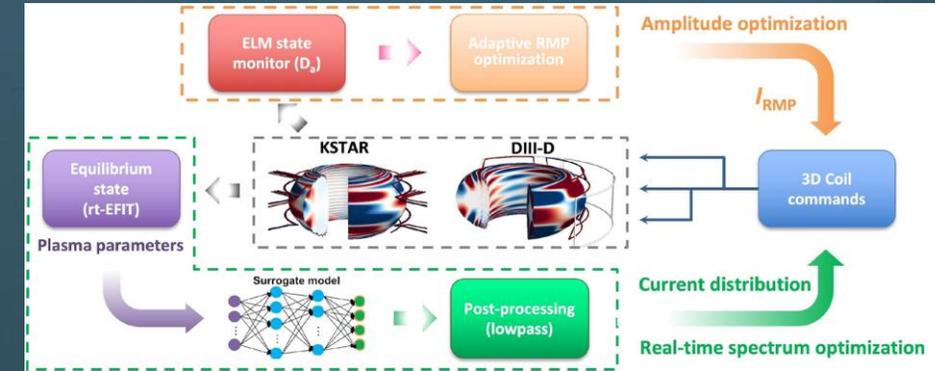
- Plasma state detection from data.
- AI-surrogate physics model to:
 - Predict instability onset
 - Control evolution to avoid instability or trigger mitigation

Impact

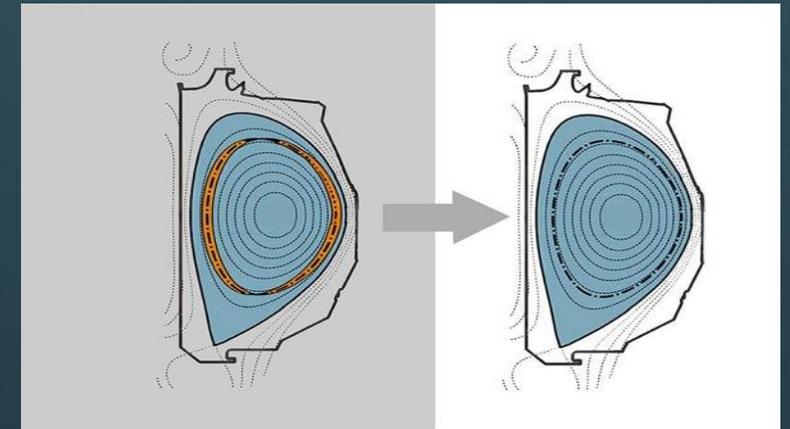
- Autonomous AI control successfully deployed in major tokamak facilities, DIII-D at General Atomics and KSTAR in Korea – record performance achieved!

S&T Challenge

- Robust systems for future devices.



[S.K.Kim, R. Shousha, S.M. Yang et al., *Nat. Comm.* 2024]



Seo et. al. *Nature*, 626, 746(2024)

AI Foundation Models for Fusion Energy

Motivation

- Optimize fusion energy systems utilizing validated foundation model

Approach

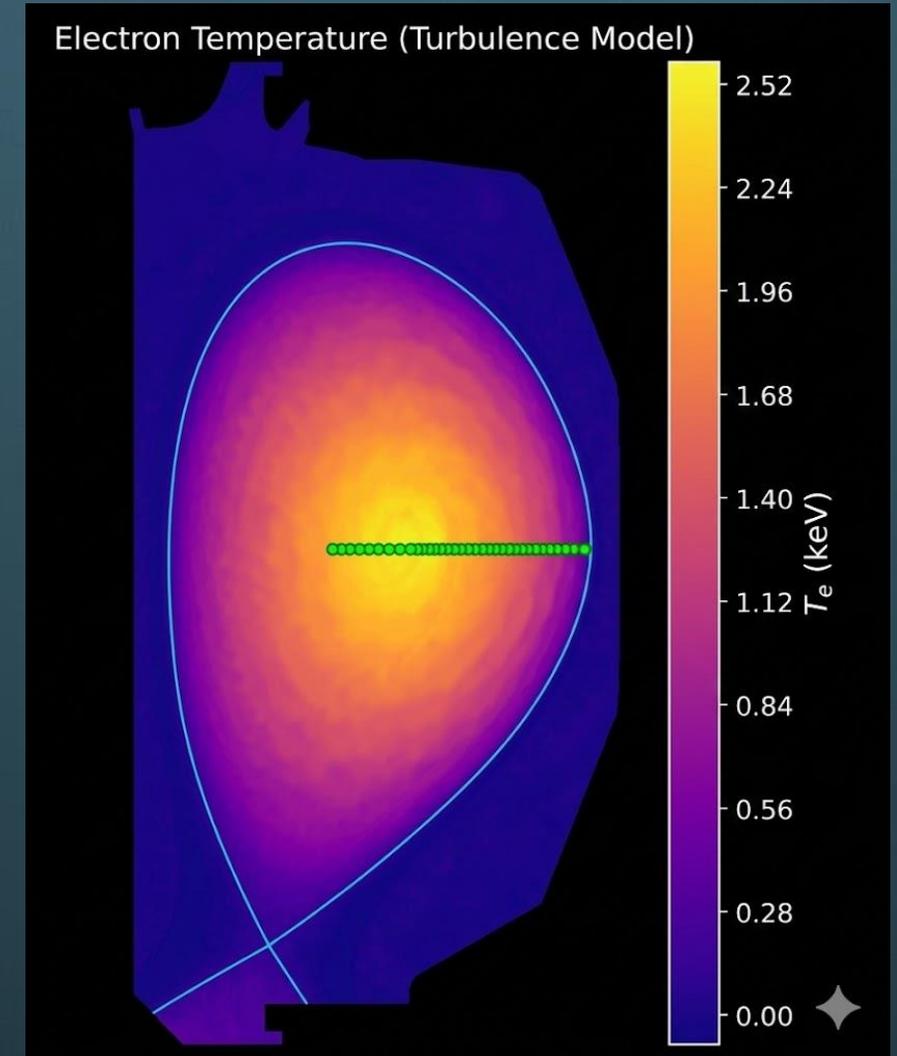
- Ingest structured and unstructured data from experiments and codes.
- Use spatiotemporal attention to capture long-range dependencies across data types.

Impact

- Collaboration with PPPL, ORNL, NVIDIA and ANL in Fusion FM Seed project – early results by July.
- Early success with Simons Foundation funded *Hidden Symmetries* project on Stellarator optimization.

S&T Challenge:

- Accelerate delivery of fusion power plants.



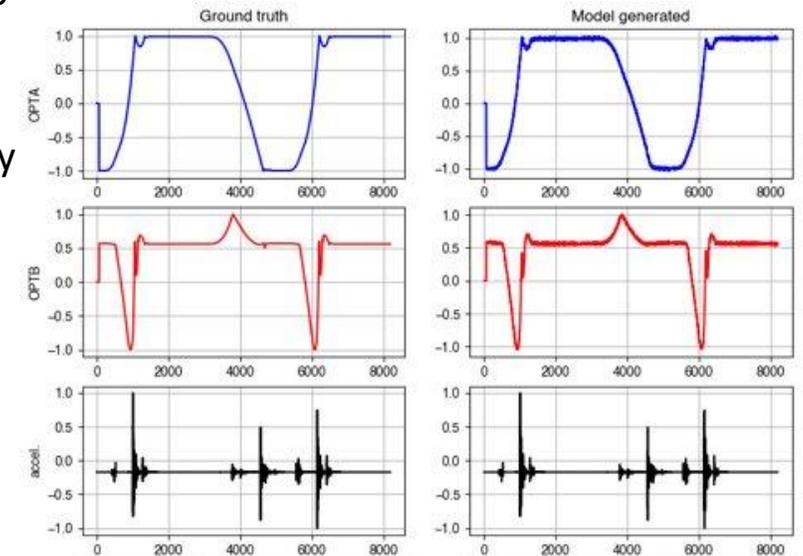
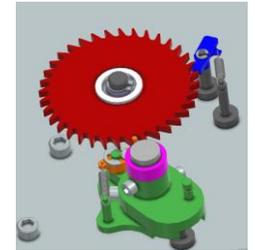
BANYAN: Scientific Foundation Models for National Security Applications

Motivation: Destructive testing is definitive but costly and slow; we instead leverage non-destructive video and sensor data to predict destructive outcomes. A multi-modal vision transformer trained on $\sim 2,000$ destructive tests predicts results with conformal-prediction uncertainty, enabling non-destructive diagnosis of stronglink defects on real weapons data in collaboration with Kansas City.

Approach: Use a multi-modal (video plus sensor data) variant of a vision transformer to generate results of destructive tests (with non-destructive data). Use data from $\sim 2k$ destructive tests for training.

Impact: Demonstrated vision transformer approach can predict and provide uncertainty estimates for destructive tests results. Generative AI approaches for time-series data enable new, non-destructive ways to diagnose stronglink defects. In collaboration with Kansas City, this work can operate directly on real weapons data, and it introduces a first-of-its-kind uncertainty estimation method for generative AI using conformal prediction.

S&T Challenge: Aligns with “Integrating Design and Production Operations for Nuclear Deterrence” and “Streamlining Production, Removing Red Tape, and Ensuring Safety in Nuclear Enterprise.”



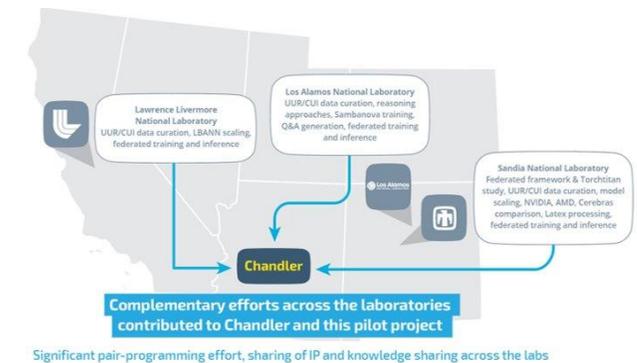
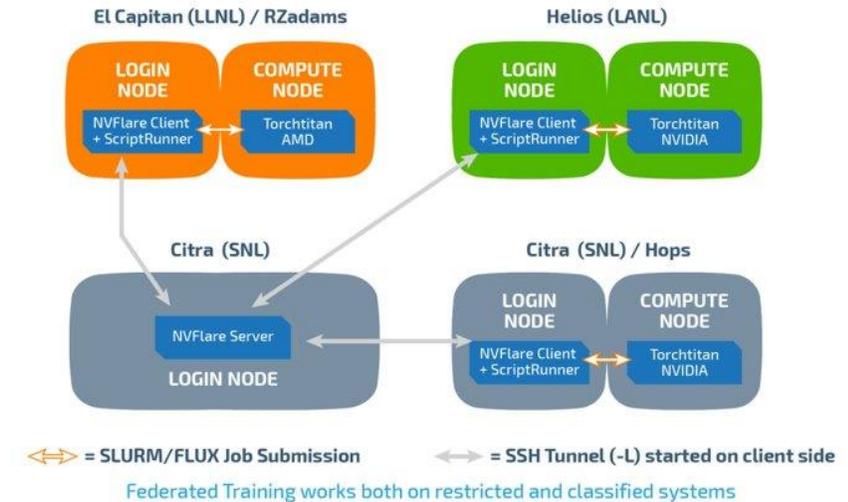
CHANDLER: A NNSA TriLab Federated Language Model

Motivation: Tri-lab AI should benefit from distributed data without moving sensitive datasets, but federated-learning tools historically didn't scale to LLMs. We worked with vendors using NVFlare and TorchTitan and fine-tuned Meta open-weight models to enable sharing model weights across the Tri-labs; through CHANDLER, this improved LLM-scale support in NVFlare, Flower, and APPFL and sped up NNSA model quality gains from months to days.

Approach: Collaborate with vendors to address gaps in federated-learning frameworks, training models using NVIDIA NVFlare and Meta's TorchTitan. Fine-tune existing open-weight models from Meta.

Impact: Created shared models without sharing large amounts of data cross the Tri-labs and sharing just model weights. Vendor federated-training libraries initially couldn't scale to LLM-sized models, but this has now been addressed in NVFlare, Flower, and APPFL through the CHANDLER collaboration. We also developed new data-processing tools that support other efforts (e.g., Genesis Mission) and demonstrated that NNSA models can achieve clear output-quality improvements in days rather than months.

S&T Challenge: Aligns with "Accelerating Nuclear Threat Assessment, Preparedness, and Response," "Harnessing America's Historic Nuclear Data & Research," and "Integrating Design and Production Operations for Nuclear Deterrence."



Accelerating Multiscale Materials Modeling with Machine Learning

Motivation: Density Functional Theory (DFT) is the gold standard for first-principles electronic-structure calculations, but its computational scaling limits practical simulations to only a few hundred atoms. Sandia's mission problems increasingly require accurate electronic-structure modeling at much larger length/time scales, motivating a machine-learning surrogate approach such as MALA.

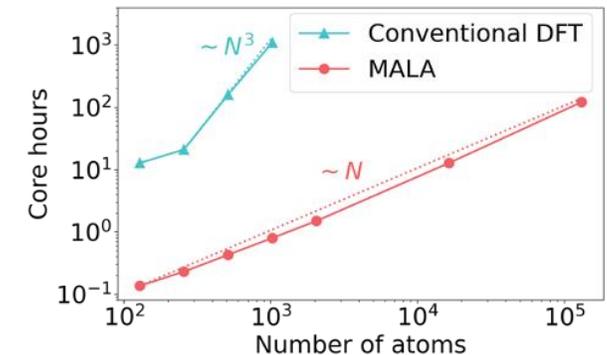
Approach: Machine learned surrogates within MALA is accurate within chemical accuracy and scales computationally as $O(N)$. MALA is a software framework that uses machine learning to predict the electronic structure of materials and enables these calculations at length and time scales that were previously unfeasible.

Impact: MALA can scale up to mesoscale ($10^4 - 10^5$ atoms), which is beyond the reach of traditional DFT. MALA results in $> 5000x$ speedup on 2000 atom systems compared to DFT.

S&T Challenge: Aligns with "Designing Materials with Predictable Functionality"



MALA has been selected for the 2023 R&D Award



Transforming Nuclear Cleanup for Restoration and Revitalization

Motivation: The Department of Energy faces an immense environmental cleanup challenge: an estimated *\$540 billion* liability spanning *eight decades*, involving *90 million gallons* of highly radioactive waste, *6.5 trillion liters* of contaminated groundwater, and *40 million cubic meters* of contaminated soil. Current remediation efforts are often manual, impeded by complex scientific unknowns, and burdened by extensive regulatory, permitting, and operating procedures.

Approach: Leveraging DOE-EM's 30+ years of hydrogeologic and operational data from large-scale facilities, AI technologies will be rapidly deployed at EM sites to accelerate repetitive workflows and foundation models will be trained to accelerate design and implementation of cleanup remedies at complex sites.

Impact:

- Achieve DOE-EM's 2040 vision, substantially accelerating the cleanup mission
- Reduce DOE's environmental cleanup liability by 50%
- Enable recovery of critical material byproducts (e.g., medical isotopes) from legacy nuclear materials
- Unlock valuable land and energy infrastructure for American energy dominance

S&T Challenge: Aligns with "Transforming Nuclear Cleanup and Restoration" and "Reenvisioning Advanced Manufacturing and Industrial Productivity"



Example: AI will augment traditionally human-only workflows to optimize remediation design across EM mission areas, resulting in significant cost savings and schedule acceleration.

AI Enabled Accelerators

Motivation: Particle accelerators are large complex systems. SLAC hosts several major accelerator facilities which will benefit from incorporation of AI techniques and tools into operation and experiment design. AI tools will enable improved reliability and performance from today's accelerators and design tomorrow's accelerators.

Approach: Integrated agentic AI workflows for operations tasks into SLAC's accelerator control system, demonstrated continual learning and digital twin workflows, demonstrated real-time integration between HPC and the LCLS accelerator for several key AI/ML applications, improved advanced ML system modeling and data standardization, and demonstrated progress in autonomous accelerator tuning and complex diagnostic analysis workflows (e.g. advanced 6D phase space reconstruction)

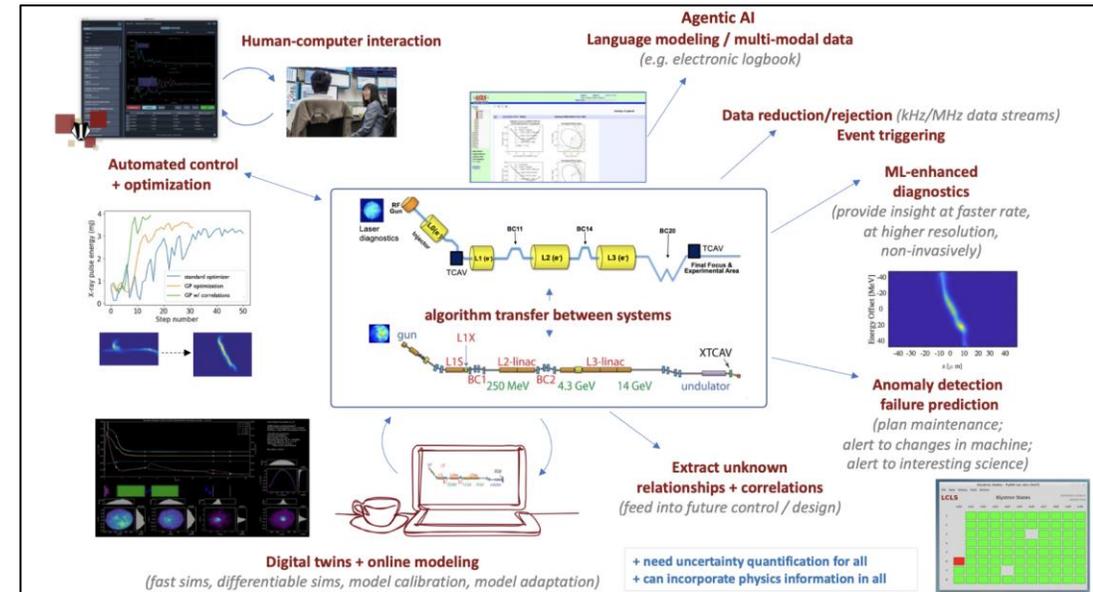
Impacts:

- Higher beam quality and on-demand customization of complex beams for scientific user experiments and commercial accelerator applications in spanning mining to medicine
- Increased scientific throughput and greater facility automation to accelerate scientific discovery
- Fundamental improvements to AI/ML for complex system modeling and control: physics-informed modeling, continual learning, uncertainty quantification

S&T Challenge:

Addresses "Enhancing Particle Accelerators for Discovery"

Impacts/synergies with Quark-to-Cosmos, Scientific Imaging, Fusion & Fission Energy, Microelectronics, Quantum, Autonomous Labs, etc.



SLAC is linking together numerous aspects of AI/ML for accelerator physics and facility operation

AI Impacts

- 10^6 X improvement in prediction speed for digital twins
- Orders of magnitude improvement in speed of beam customization
- Better performance period (pulse intensity, emittance, & acceleration efficiency)

Experiment Steering with Digital Twins unlocks informed experiments @ KHz rep-rate

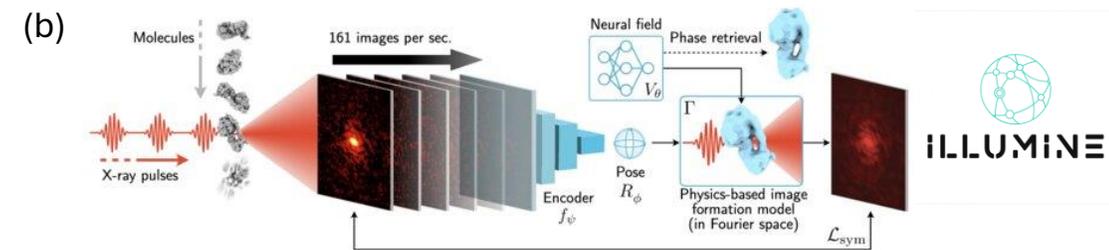
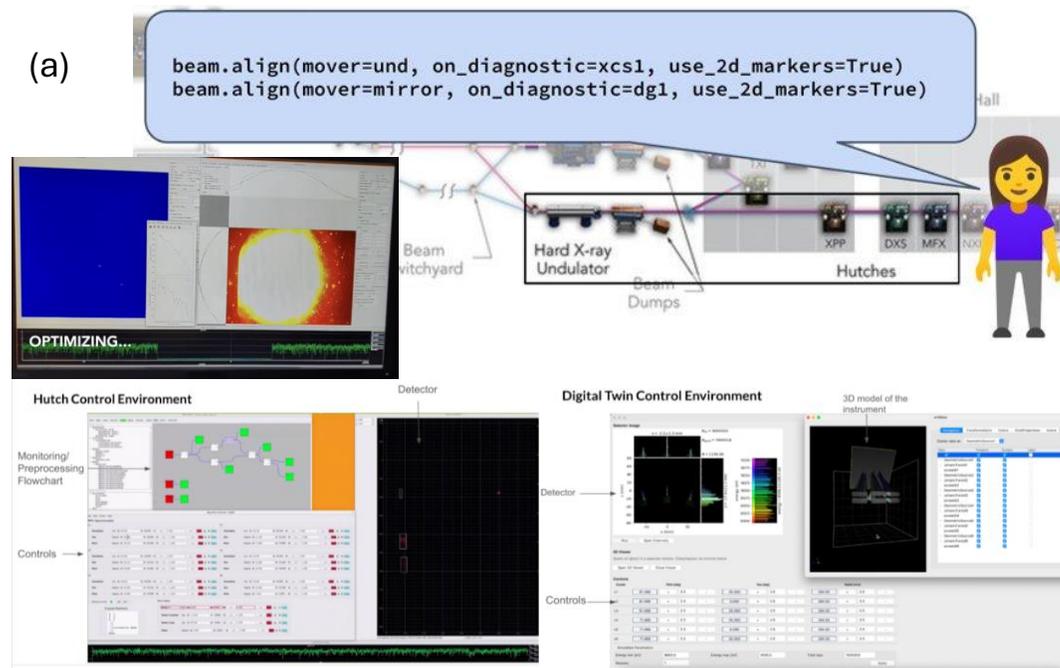
Motivation: Accelerate discovery and reduce time to science through automation, improved real-time feature extraction, AI-assisted decision-making, and increased instrument uptime.

Approach: Build an interoperable experiment steering framework for light sources that enables plug-and-play of algorithms, agents, and Digital Twins. Develop 1) a platform that connects Instrument (Bluesky) and Accelerator (Xopt) control systems to close the information/steering loop; 2) Digital Twin of MFX instrument to enable testing of steering; 3) Digital Twin of the sample with generative AI prototypes; and 4) Implement with closed loop feedback.

Impacts:

- Less tuning time, more science: +20% science measurement time per shift.
- Live, data-driven experiment feedback on data quality.

S&T Challenge: Achieving AI-Driven Autonomous Laboratories



- (a) Bluesky+Xopt framework equips MFX with automated beamline alignment tools and digital twin.
- (b) Prototype for live molecular reconstruction from diffraction image stream at LCLS. [doi:10.1038/s41467-025-62226-7](https://doi.org/10.1038/s41467-025-62226-7)

Physical AI for Dynamic Reactor Control Systems

Motivation

Establish a template for autonomous control of complex reactor systems using tight-coupling of live diagnostics, FPGA edge processing, HPC model training, and real-time control systems.

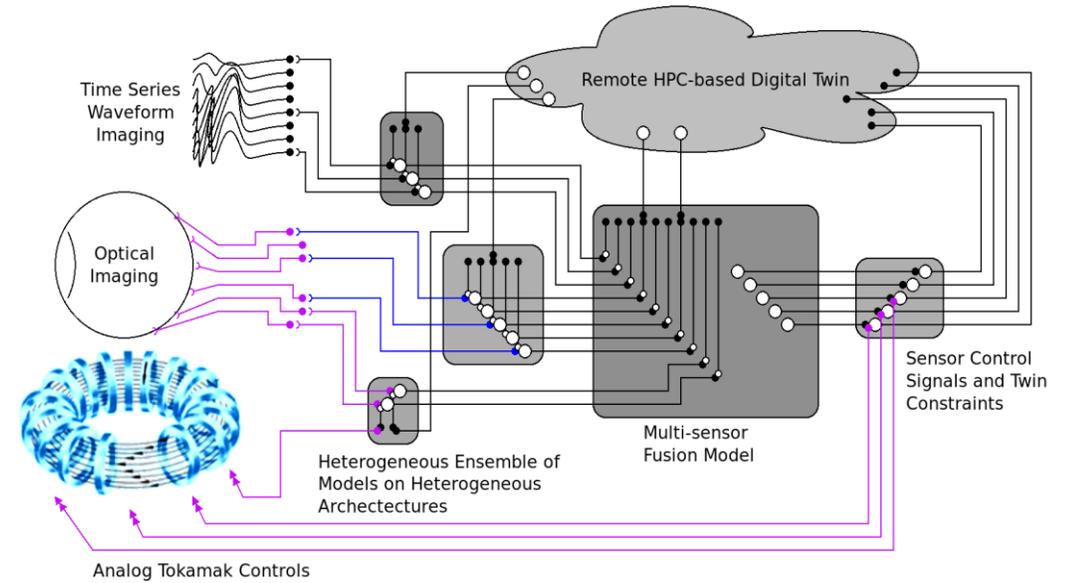
Approach

- Connect to remote-hosted AI surrogate models that will continuously learn from each operational cycle (shot) of a tokamak
- Minimize latency using streaming measure, diagnose, act cycle.
- Adapt operation controls to explore plasma state-space via ensembles of surrogate models

Impact

- Operationalize ML-based strategies to predict and mitigate plasma instabilities
- Self-improving control ecosystem that brings long-duration, stable magnetically confined plasma burns to reality
- Incorporate fusion and fission into a diversified, resilient, adaptive power grid infrastructure.

S&T Challenge: Aligns with “Accelerating Delivery of Fusion Energy” and “Scaling the Grid to Power the American Economy”



Microseconds “instinct” loop control system integrated with seconds “reasoning” loop in AmSC/Genesis ecosystem

AI impacts:

- From minutes to microseconds for live reconstruction and continuous autonomous control
- 500x data reduction... signals to controller tokens

Agentic AI Decodes Catalytic Selectivity and Active-State Evolution

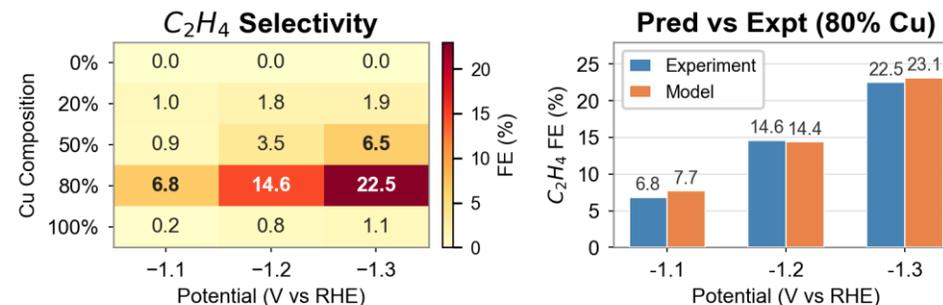
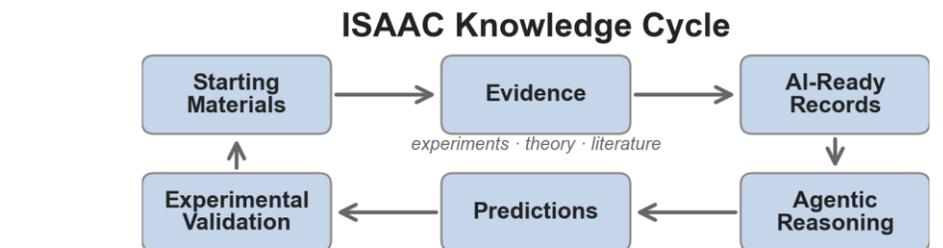
Motivation: Catalysis drives 90% of chemical manufacturing and 35% of global GDP. Predictive catalyst design is essential to American competitiveness and prosperity, yet it demands multi-scale understanding. DOE offers unmatched characterization tools, but they remain fragmented and in isolation.

Approach: Agentic AI converts siloed measurements, theory, and literature into integrated knowledge across National Labs. Ontology-controlled records enable agents to fuse and reason heterogeneous evidence into mechanistic hypotheses.

Impacts: Predictive selectivity across composition–structure–condition space. A closed cycle from starting materials through active-state evolution to targeted outcomes, replacing trial-and-error with physics-informed design.

S&T Challenge:

Achieving AI-Driven Autonomous Laboratories; Designing Materials with Predictable Functionality



CO₂ reduction on Cu|Au striped electrodes: AI-driven knowledge cycle integrates multi-modal evidence to predict composition- and potential-dependent selectivity.

AI impacts

- 6 national labs unified under one ontology
- 100% deterministic inference via schematic layers
- Months → hours for mechanistic interpretation (100x)

Accelerating Structural Biology with PeakNet

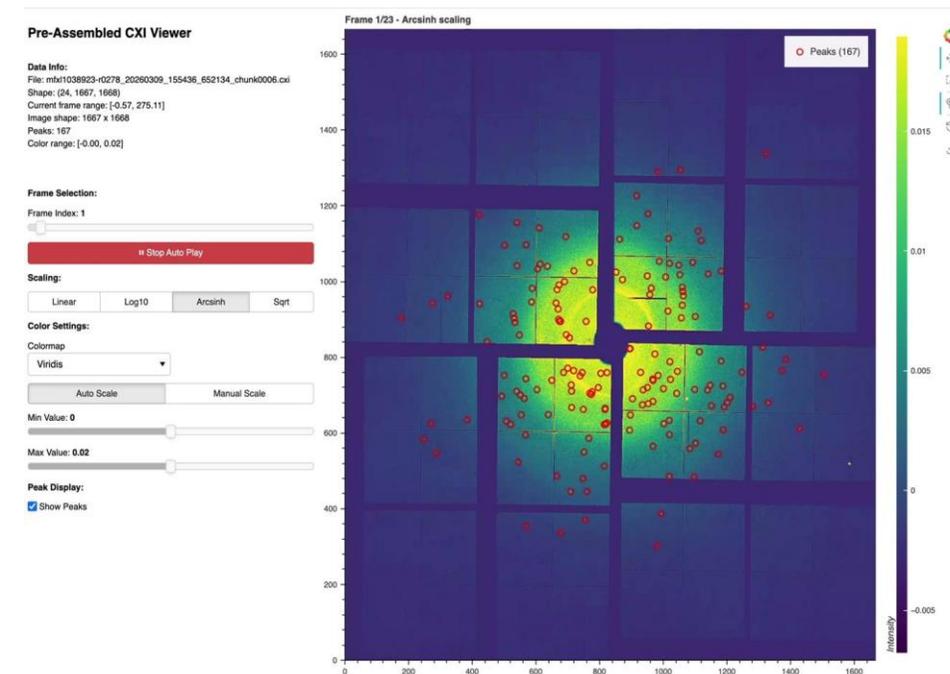
Motivation: To process diffraction images as fast as they are generated without human intervention, enabling faster scientific discovery and experiment steering.

Approach: Deploy PeakNet, a supervised deep neural network to perform real-time segmentation of Bragg Peaks: 1) **Autonomous Learning:** Adapts to real-time background fluctuations without manual tuning; 2) **Scalable Architecture:** Built on a producer-consumer model that connects SLAC's data center to remote HPC to scale inference across remote HPC GPU/CPU; 3) **Massive Training:** Initially trained on 20,000 images with a roadmap to ingest 7 PB of LCLS data.

Impacts: Outperforming the expert

- Higher yield of indexable hits.
- Operational Efficiency: Eliminates the stop-and-tune cycle allowing downstream processing and real-time feedback to proceed without interruption.

S&T Challenge: Biotechnology and AI-driven laboratories



PeakNet learns autonomously to identify Bragg Peaks transforming raw data into scientific insight at the speed of experiment.

- **Speed:** Live 1 MHz autonomous processing.
- **Quality:** Higher index rate (72.7%+) compared to human baseline; catches weak signal missed by traditional algorithms
- **Scalability:** Seamless integration between LCLS data streams and remote HPC facilities.

Hardware-Accelerated AI-Based Distributed Inverse Solver

Motivation

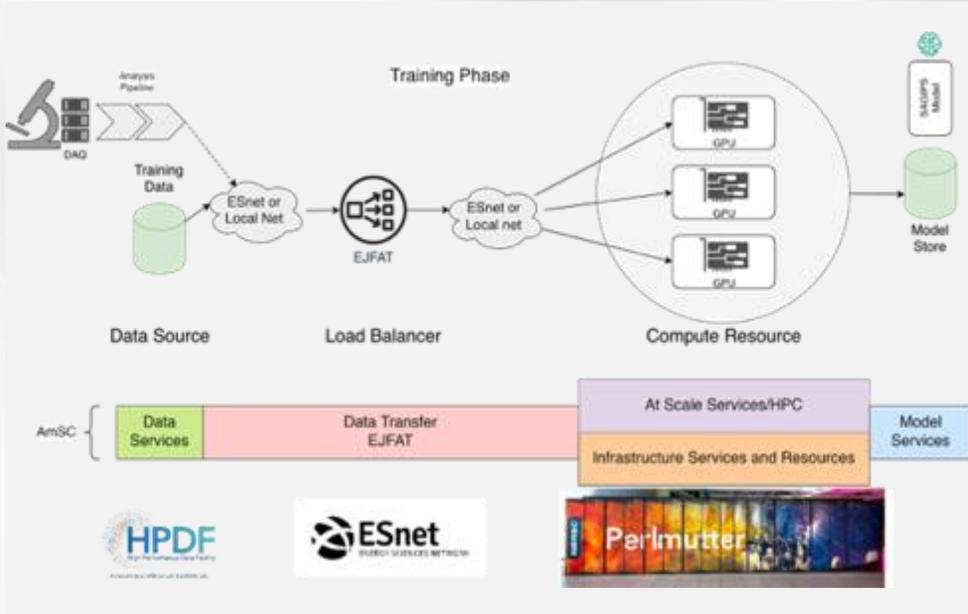
Extracting fundamental physics observables from large-scale nuclear physics data sets is time consuming, typically tailored to one specific experiment, and invasive of the data set.

Approach

- Combine streaming data service (EJFAT) among different facilities with scalable inverse problem solver (SAGIPS).
- Efficiently feed continuous data stream into solver.
- Explore solution space via generative AI.
- Use full computational representation of experimental pipeline.

Impact

- Extract physics observables non-invasively.
- Analyze multiple datasets across facilities simultaneously.
- Reduce time to scientific discovery across domains while enhancing the quality of extracted results.



Schematic representation of HADIS: Data coming from an experiment located at Jefferson Lab is sent via EJFAT to Perlmutter where SAGIPS extracts unknown observables

S&T Challenge

Aligns with “Unifying Physics from Quarks to Cosmos”

Automating Particle Detector Calibration with RoboCDC

Motivation

Drift Chamber detectors are a core technology used in accelerator-based nuclear physics experiments, but accelerator beam and atmospheric conditions affect their operation.

Approach

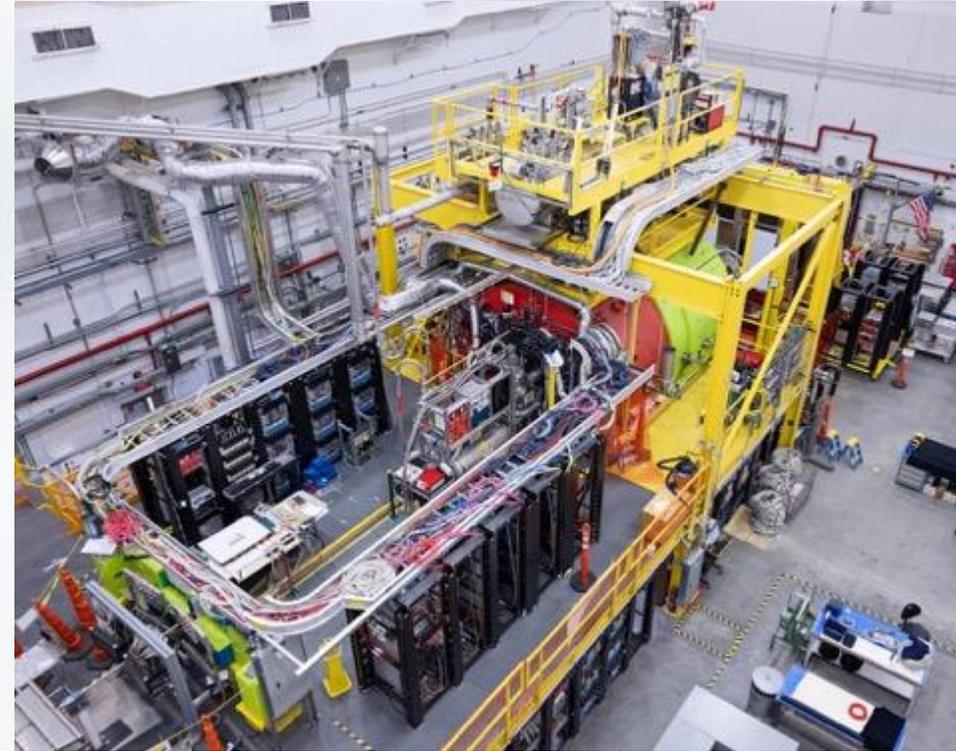
- Automate fine adjustments of the detector's high voltage input using AI to stabilize operation, vastly reducing the need for calibration of the data after it is acquired.

Impact

- Calibration time for the GlueX Central Drift Chambers reduced from months to moments.
- Better data stability along with a reduction in latency between data acquisition and physics analysis.

S&T Challenge

Aligns with "Achieving AI-Driven Autonomous Laboratories" and "Unifying Physics from Quarks to the Cosmos"



The GlueX Central Drift Chamber (installed inside the cylindrical red and lime-green magnets) enable studies of the particles and forces that build our visible universe.

AI Data Quality Monitoring with Hydra

Motivation

Collection of high-speed experimental nuclear physics data must be monitored to ensure proper particle detector operation during valuable research accelerator beam time.

Approach

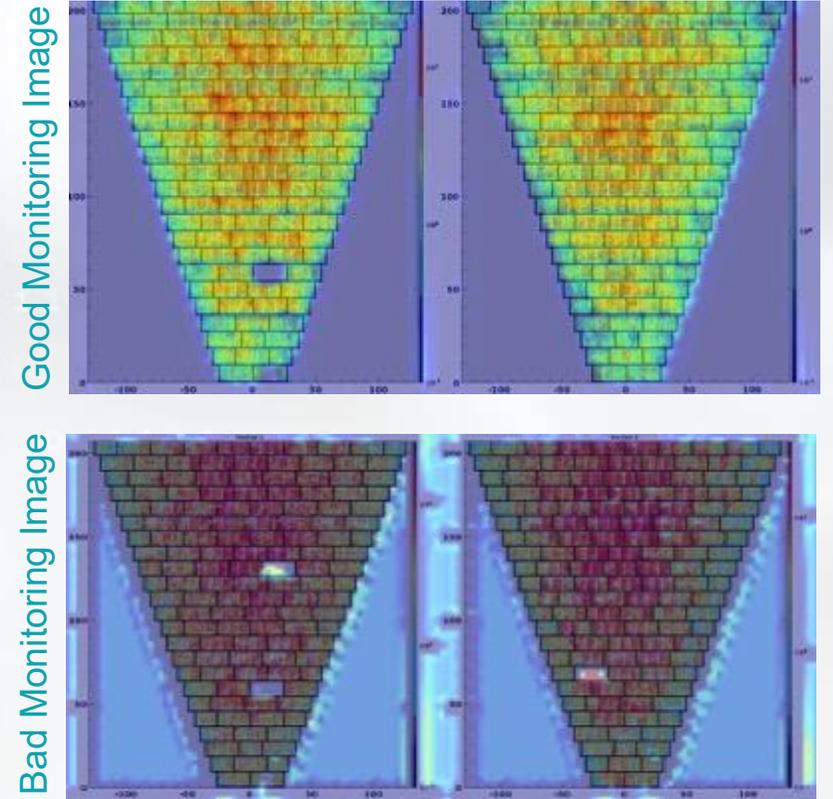
- Train AI models to view monitoring plots to detect anomalies with high precision and at high rates.
- Visually identify location of anomalies by inspecting inner layers of AI models.

Impact

- Problems identified much quicker (minutes vs. hours) and much more consistently compared to current monitoring.
- Allows scientists to focus on new data rather than operations.

S&T Challenge

Aligns with “Achieving AI-Driven Autonomous Laboratories” and “Unifying Physics from Quarks to the Cosmos”



The Experimental Physics Software and Computing Infrastructure (EPSCI) group develops and maintains Hydra, an AI-driven monitoring framework for image-based data quality management in particle accelerator and experimental facilities.

Unleashing AI for More Powerful and Efficient Particle Accelerators

Motivation

Powerful research particle accelerators provide insight into our visible universe, but these highly sophisticated machines are difficult to operate and maintain at peak efficiency.

Approach

- Identifying and predicting faults before they interrupt operations.
- Lowering levels of a damaging type of radiation.
- Load-balancing components to prevent anomalies.

Impact

- Increasing uptime and reliability of research accelerators.
- Revolutionizing design of next-gen machines for applications in medical care, materials science, and energy production.

S&T Challenge

Aligns with “Enhancing Particle Accelerators for Discovery” and “Unifying Physics from Quarks to the Cosmos”



Three of the nation's most powerful particle accelerators rely on superconducting radiofrequency (SRF) accelerator technology, including the Continuous Electron Beam Accelerator Facility, shown here. This research is ensuring such machines operate at peak efficiency now and into the future.

Clean Particle Tracking in Nuclear Physics Experiments

Motivation

In nuclear physics experiments, sensitive particle detectors collect hits from real particles as well as background (X-rays or electronic noise) that is not of interest. It is imperative to remove the background noise for maximum efficiency and resolution.

Approach

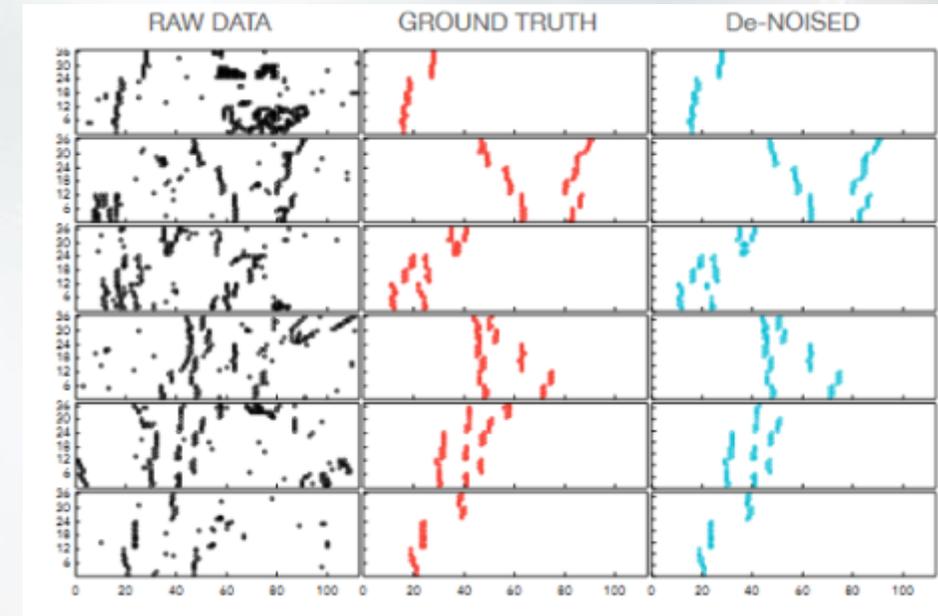
- Train an AI auto-encoder model to filter identified background hits, while allowing all real signal hits from tracks to pass.
- Identify real particle tracks through the detector using a second AI model.

Impact

- Delivers cleaner data sets for analysis, saving research time.
- Strengthens the link between AI model development and scientific outcomes.

S&T Challenge

Aligns with "Unifying Physics from Quarks to the Cosmos"



These data plots show what the raw data look like compared to real signal hits as cleaned manually by a nuclear physicist (orange) and the same raw data with background removed by a trained AI model (teal).

Norfolk Street Flooding Predictions With Machine Learning Models

Motivation

Traditional deep learning models for predicting flood risk in flood-prone areas meet the computational demands of physics-based models, but their predictions lack corresponding uncertainty.

Approach

- Develop deep learning models with distance-aware uncertainty quantification.
- Enable approaches that support accurate confidence intervals and distance-aware behavior, which are critical for risk-aware decision support.

Impact

- Enables 3-day-ahead forecasting and real-time, risk-aware flood forecasting where physics simulations alone are too slow.
- Supports city-level decisions on street closures, stormwater pump operations, and flood-stage monitoring.
- Strengthens resilience for critical infrastructure near water.



This illustration visualizes the predictions of Norfolk, Va., street-scale flooding by machine learning surrogate models.

S&T Challenge

Aligns with “Predicting U.S. Water for Energy”

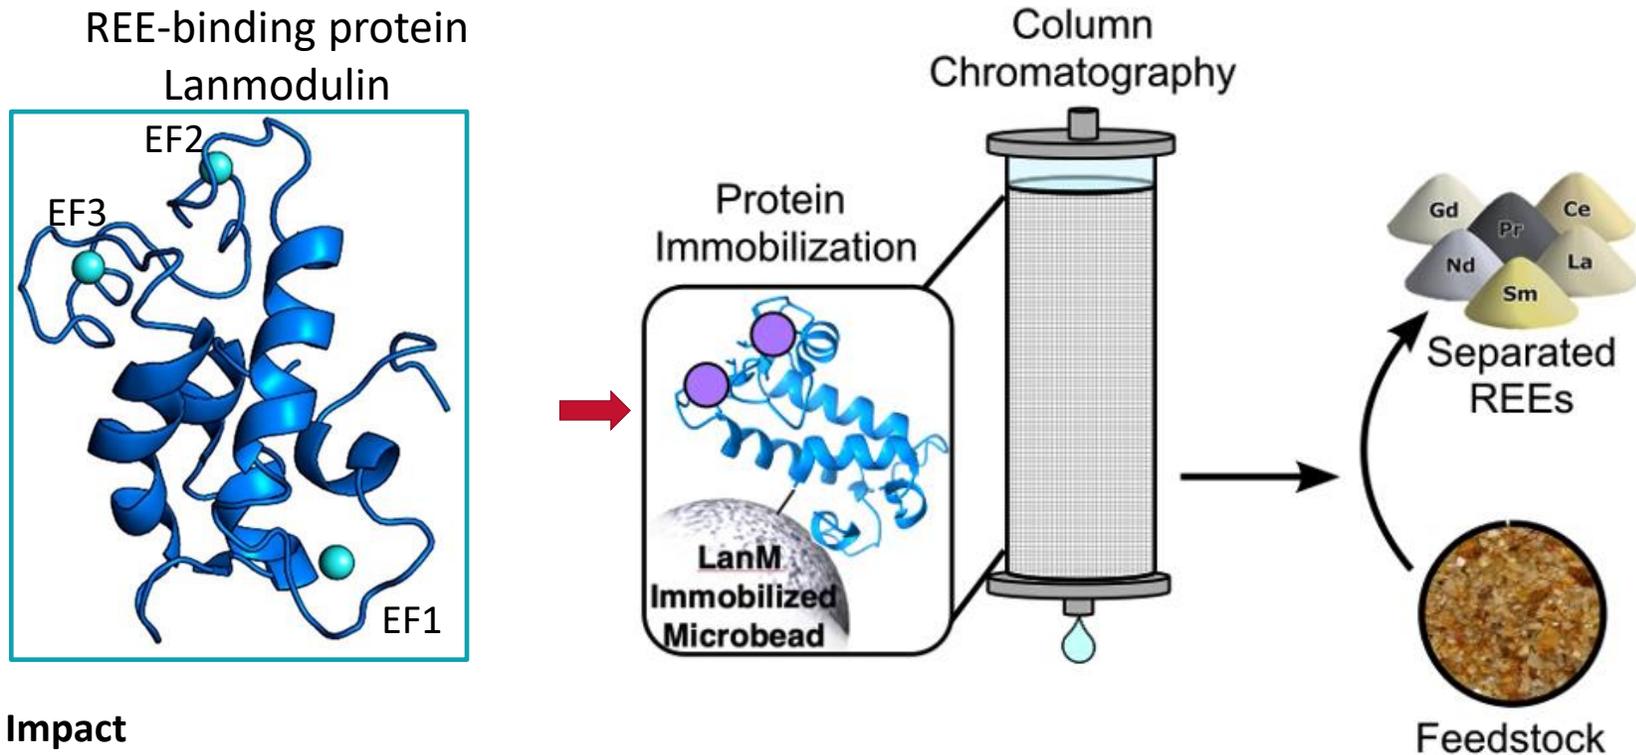
AI Enables Predictive Design for Rare-Earth Elements (REE) Separation

Motivation

- Current REE separation relies on organic solvents and non-selective ligands
- There is a critical need for predictive ligand design to enable efficient REE separation with complex feedstocks

AI-enabled Approach

- Molecular level: infer protein sequence–structure–function relationships to design and prioritize desirable variants
- Process level: model relationships between resin architecture, buffer chemistry, feed composition, and mass transfer to predict separation performance and economics



Impact

- Shift in R&D paradigm in AI-generated or prioritized candidate selection, multivariate co-optimization, and learning from realistic feed conditions
- Build on validated advances in protein engineering and REE bioseparation (ACS Central Science 2021; Nature 2023; Nat Chem Bio 2026)
- Technology licensed to Alta Resource Technologies and advancing toward pilot-scale deployment at LLNL and Alta

S&T challenges: Align with “Scaling the Biotechnology Revolution” and “Securing America’s Critical Minerals Supply”