



The Genesis Mission: Transforming Science and Energy with AI

DOE Biological and Environmental Research Program Selection Announcement

Summary of BER-supported projects selected in Fiscal Year 2026 under DOE Request for Application DOE-FOA-0003612

The Genesis Mission unites U.S. Department of Energy (DOE) programs, national laboratories, industry, academia, and partner agencies to harness AI for scientific breakthroughs in energy dominance, discovery science, and national security. This coordinated national effort builds upon DOE's years of investment in advanced computing, scientific datasets, and research infrastructure with the goal of doubling research productivity within a decade.

As part of the Genesis Mission, DOE established a set of national cross-government science and technology (S&T) challenges. In Fiscal Year 2026, DOE's Biological and Environmental Research program (BER) selected 58 projects to tackle six program-relevant challenges. These S&T challenges draw on BER's fundamental science strengths in energy, environment, genomics, and biosystems design.

BER solicited proposals from interdisciplinary, multi-institutional teams spanning academia, industry, and the national laboratories. The selected projects will leverage AI to accelerate the pace of scientific discovery. Resulting AI models, workflows, and data will be integrated into the larger Genesis Mission capabilities.

Selection for award negotiations is not a commitment by DOE to issue an award or provide funding. Before funding is issued, DOE and applicants will undergo a negotiation process, and DOE may cancel negotiations and rescind the selection for any reason during that process.

BER-Relevant Science and Technology Challenges

In February 2026, DOE announced an initial 26 national science and technology (S&T) challenges to accelerate U.S. innovation and leadership. These challenges will deliver tangible benefits for the American people through the Genesis Mission's AI platforms and partnerships. BER-relevant S&T challenges include:

-  Scaling the Biotechnology Revolution
-  Predicting U.S. Water for Energy
-  Securing America's Critical Minerals Supply
-  Unleashing Subsurface Strategic Energy Assets
-  Scaling the Grid To Power the American Economy
-  AI in Fluid Flow for Energy Components and Technologies



U.S. DEPARTMENT
of **ENERGY** | Office of
Science

Biological and Environmental Research Program



Websites



DOE Office of Science
energy.gov/science



Genesis Mission
energy.gov/genesis-mission



BER Program Managers

Dawn Adin
dawn.adin@doe.gov

Paul Bayer
paul.bayer@doe.gov

Brian Benscoter
brian.benscoter@doe.gov

Justin Hnilo
justin.hnilo@doe.gov

Resham Kulkarni
resham.kulkarni@doe.gov

Shaima Nasiri
shaima.nasiri@doe.gov

Pablo Rabinowicz
pablo.rabinowicz@doe.gov

Vijay Sharma
vijay.sharma@doe.gov

Daniel Stover
daniel.stover@doe.gov

Boris Wawrik
boris.wawrik@doe.gov

Scaling the Biotechnology Revolution Selections

Designing biology on demand to innovate biotechnology and re-establish U.S. leadership will require accurately defining essential governing principles. The nonlinear complexity of biological systems and long, costly development cycles hinder biotechnological innovations, impeding delivery of crucial biofuels, biochemicals, and bioproducts vital for American economic prosperity and energy independence.

AI will integrate and interpret genomics, multiomics, imaging, dynamics, and phenomics data into embedded models to establish genotype–phenotype relationships, identify causal control points, and guide autonomous experimentation. This will advance AI's ability to reason about long causal chains under uncertainty, integrate across multiple scales, and learn from sparse, noisy data. Tools like digital twins could be used to de-risk process development, optimization, integration, and scaleup, thereby speeding up industrial production of biofuels, biochemicals, and bioproducts by orders of magnitude.

Under this S&T challenge, BER is involved in four focus areas:

- Biomolecular Science
- Genotype to Phenotype
- Predictive Engineering of Microbial Communities
- Biodesign

Biomolecular Science

Predictive AI to Map Point Mutation Effects on Protein Function: Measurement and Biosynthesis of Isoprenoids

- **Lead PI:** Caroline Ajo-Franklin, Rice University
- **Collaborating Institutions:** Argonne National Laboratory; Lawrence Berkeley National Laboratory

Isoprenoids represent an important class of chemicals, including drop-in replacements for solvents and jet fuels, but these chemicals are not currently produced in the United States. Producing isoprenoids with microbes is economically uncompetitive due to a fundamental bottleneck: the absence of high-throughput tools for measuring and optimizing production. This project aims to break this bottleneck by deploying an AI-driven engineering platform that simultaneously advances isoprenoid biosynthesis and the frontier of protein prediction. By pairing engineered transcriptional biosensors with terpene synthases in a microbial testbed, the research team will generate unprecedented genotype-to-phenotype datasets linking variation in protein sequence to four critical protein properties—ligand binding, DNA recognition, folding stability, and enzyme activity—that no existing AI model can currently predict. These datasets will enable next-generation structure prediction models capable of resolving mutational effects on protein–terpenoid and protein–DNA interfaces, validated prospectively through high-throughput biosensor activity characterization, automated X-ray crystallography and molecular dynamics simulations at Argonne National Laboratory, and metabolomics-based enzyme characterization at Lawrence Berkeley National Laboratory. The result will be a

scalable, self-reinforcing AI-experimental loop that transforms isoprenoid enzyme engineering from a low-throughput empirical craft into a data-driven, AI-accelerated science—directly advancing DOE's mission to scale the biotechnology revolution for energy-relevant biomanufacturing.

Structural Dynamics Enable Enzyme Design for Useful Biotechnology

- **Lead PI:** James Fraser, University of California–San Francisco
- **Collaborating Institutions:** SLAC National Accelerator Laboratory; Profluent Bio; Radial, Astera Institute; University of California–Santa Barbara

Useful biotechnology depends on catalysts that are active, programmable, controllable, and adaptable to new chemistries. Natural enzymes are remarkable, but they evolved for biological needs rather than for the full range of reactions needed in manufacturing, bioenergy, biomass conversion, lignin valorization, and waste detoxification. Today, many useful chemical transformations still rely on costly and burdensome processes. This project will develop a new AI and experimental platform for designing enzymes that can catalyze useful reactions and ultimately be controlled by small molecules, ions, or changes in their environment.

The central idea of this project is “structure-seeded enzyme design.” Rather than asking AI to design an enzyme entirely from an idealized model of a reaction, the research team will first use AI to generate large numbers of stable protein scaffolds and then use high-throughput crystallography at SLAC National Accelerator Laboratory to determine how small substrate-like and transition-state-like molecules actually bind inside those proteins. These experimentally observed binding geometries will

then seed the next round of AI design, in which catalytic residues, metal sites, cofactors, and regulatory features are built around real molecular poses.

This project will use specific AI methods matched to each stage of the problem. Diffusion and flow-matching generative models will design new protein scaffolds and alternative conformational states. Message-passing neural network sequence design models will design protein sequences that stabilize active sites around bound ligands. Deep learning-guided multistate design will create proteins whose conformational changes can be linked to substrate binding, product release, or allosteric control. These methods are well suited to enzyme design because catalysis depends not only on a stable structure but also on precise atomic geometry and coordinated motion.

All models, chemical libraries, structures, screening results, design trajectories, benchmarks, and workflows will be deposited through the Genesis Mission American Science Cloud and Transformational AI Models Consortium infrastructure so outputs are reusable by other researchers across the Genesis Mission ecosystem. This project will also align datasets and validation workflows with BER's Lakehouse-enabled AI-ready Multimodal Bioimaging Data Architecture (LAMBDA), helping establish structure-seeded AI design as a shared national laboratory capability for AI-ready experimental validation. If successful, this work will provide a foundation for designing useful enzymes more rapidly, more predictably, and with greater control than is possible with conventional enzyme engineering alone.

Agentic AI-Driven Discovery of Enzyme Conformational Dynamics in Microbial Systems*

- **Lead PI:** Travis Wheeler, University of Arizona
- **Collaborating Institution:** Pacific Northwest National Laboratory

Living cells run on proteins. The most important of these are enzymes, the molecular machines that carry out the chemical reactions of life. A protein is often presented as having a static shape, but in fact, proteins are in constant motion. They flex, bend, and switch between different conformations, and that motion is usually what lets an enzyme do its job: perform a reaction, turn activity on or off, or respond to a signal from its surroundings. Knowing the shape of a protein as well as how it moves is necessary to understand what an enzyme does. This knowledge is especially important for redesigning microbes to produce fuels, chemicals, and materials or to support healthy soils and plants.

While biologists can now predict a protein's static shape quickly and accurately, predicting a protein's motion is one of the hardest problems in biology. The most reliable method, molecular dynamics simulation, is essentially a physics-based movie that tracks every atom in a protein as it jiggles over time. These simulations are accurate but computationally intensive, often taking

days or weeks of supercomputer time for a single protein. A microbe contains thousands of proteins, so simulating all of them remains out of reach and limits biologists' ability to predict its motion at scale.

This project will use AI to address that limitation. The research team will train AI models to learn the patterns of protein motion from a large collection of existing simulations. The models will learn to predict how a protein will move and how its motion changes when a drug-like molecule binds to it, when it is chemically tagged, or when its sequence is altered. The goal is to turn a calculation that takes weeks into one that takes seconds—roughly a thousandfold speedup—while keeping the predictions accurate enough to be scientifically useful.

A central part of the work is building the data itself. The research team will gather and standardize the largest ready-to-use collection of protein-motion data assembled to date, integrating large public datasets into MDRepo, an open, community-contributed repository. This resource will be valuable in its own right. Much as the public archive of protein shapes made today's shape-prediction tools possible, a clean and unified archive of protein motion is a prerequisite for the next generation of predictive models. This project will make MDRepo freely available to the research community.

This work brings together two complementary teams. Researchers at the University of Arizona will develop the data infrastructure, AI models, and methods for summarizing protein motion. Researchers at Pacific Northwest National Laboratory will run high-quality simulations of real bacterial proteins important to DOE's goals in bioenergy and environmental science and will help choose the biological targets that matter most. Early targets include enzymes from soil bacteria whose shape changes control the flow of chemistry through their metabolism.

If this project succeeds, it will give scientists a fast and inexpensive route from a protein's static structure to a picture of how it behaves. That capability would speed the design of microbes and enzymes for producing fuels and chemicals, sharpen understanding of the microbial communities that shape soils and support crops, and advance the national goal of using AI to accelerate scientific discovery.

* BER is lead; co-funded with the DOE Advanced Scientific Computing Research (ASCR) program

Genotype to Phenotype

IMAGINE-AI: Integrated Microbiome Analytic Generator of Interaction Network Evolution

- **Lead PI:** Ember Morrissey, West Virginia University
- **Collaborating Institutions:** Johns Hopkins University; Pacific Northwest National Laboratory; Lawrence Livermore National Laboratory

The untapped metabolic potential of microbial communities could revolutionize domestic bioenergy production and ensure U.S. leadership in biochemical and bioproduct innovation. Yet realizing this potential is blocked by a fundamental barrier: an organism's genes alone (genotype) do not reliably predict its function (phenotype). In multispecies communities, phenotypes emerge from seemingly random and unpredictable interaction networks that defy traditional modeling. For example, biological nitrogen fixation, the natural process by which certain bacteria convert nitrogen gas from the air into plant-available nitrogen, is influenced by neighboring species that share nutrients and chemical byproducts or consume oxygen that would otherwise disrupt nitrogen fixation. Multispecies interactions such as these that influence gene expression make genotype-to-phenotype predictions challenging. Predicting how interactions evolve and reshape emergent phenotypes in dynamic environments is even more difficult. However, deep learning has repeatedly solved biological prediction problems once considered unsolvable. A rapidly growing body of work shows the same revolution is now within reach for genotype-to-phenotype inference, even in the context of complex microbial networks.

This project's objective is to develop an advanced deep learning AI model called the Integrated Microbiome Analytic Generator of Interaction Network Evolution-AI (IMAGINE-AI) that can predict phenotypes for individual strains and multispecies communities. While the AI that the team aims to build will be versatile and applicable to any microbial system, this project will focus model testing and refinement on nitrogen-fixing microbial communities because engineering and optimizing nitrogen fixation would revolutionize bioenergy and agriculture for domestic food and energy security. To this end, this unique team of experts will integrate large-scale measurements of genes, molecules, and traits with deep learning to develop an AI advantage in genotype-to-phenotype prediction.

IMAGINE-AI will combine two complementary AI approaches. The first, transformer-based protein and genome language models, works like the AI behind tools such as ChatGPT but applied to DNA and protein sequences instead of text. The second, graph neural networks, models the web of gene-gene and metabolic interactions. This two-pronged approach is well suited to the problem because transformer-based models learn representations directly from genomic sequence, capturing molecular function beyond curated annotations, while graph neural networks explicitly model nonlinear gene-gene and interspecies interactions that are inaccessible to conventional statistical methods.

Together, these models will link intracellular molecular function to community-level processes such as microbial cooperation and competition, enabling the prediction of emergent phenotypes in multispecies systems. Comprehensive public databases will be augmented with newly generated data on bacterial species and microbial communities to train and test IMAGINE-AI.

This project will demonstrate an AI advantage by comparing its novel AI genotype-to-phenotype predictions with traditional statistical approaches. In tandem with AI development, the team will establish high-throughput, automated laboratory protocols that rapidly measure microbial traits. These protocols will capture how bacterial species and communities respond to environmental drivers, further refining IMAGINE-AI. All models, datasets, and workflows will be deployed on Genesis Mission infrastructure to ensure interoperability, reuse, and integration with other projects across the Genesis Mission ecosystem. The proposed IMAGINE-AI model will rapidly advance biotechnology by providing predictive modeling of microbial communities that can be used to fix nitrogen for crops, as well as myriad other applications that ensure U.S. leadership in agricultural, biotechnology, and energy sectors.

AI-Driven Prediction of Emerging Microbial Phenotypes in Electrogenic Consortia Across Scales

- **Lead PI:** Victoria Orphan, California Institute of Technology
- **Collaborating Institutions:** Argonne National Laboratory; University of Georgia; Dayhoff Labs

The expression of a microorganism's phenotype (i.e., its behavior and metabolic activity) depends both on the composition and regulation of its genome and proteome and on external influences, including the physical and chemical environment and behavior of other interacting microbes. Most existing models treat these dimensions separately, which limits how well they predict microbial behavior in complex, real-world settings. This project asks how AI can improve microbial phenotype predictions, moving beyond models that rely on genes alone toward dynamic models that also account for a microbe's environment and its interactions with neighboring species.

The research team will build and test this AI framework using a DOE-relevant yet challenging microbial system: interdependent microbes that live together in structured aggregations in matrices without oxygen, operate on very little energy, and are fueled by methane. These associations consist of ANME (anaerobic methane-oxidizing archaea), which consume methane and transfer electrons extracellularly to a sulfate-reducing bacterial partner, which in turn converts sulfate to sulfide. Their phenotype is shaped by close metabolic cooperation (i.e., syntrophy) through the sharing of electrons and nutrients from one cell to another and by the surrounding chemical and physical environment. These electron-conducting communities are an ideal test case precisely because genes alone cannot predict their phenotype; it also depends on tightly linked cooperation between partners, on chemical pathways that can run in either direction, and on

fine-scale control by the local chemical environment. These same features make them relevant to DOE's interest in how life functions under low-energy conditions and to broader efforts to predict how engineered microbial communities perform and how well they withstand changes.

This project is organized into two main aims. Aim 1 investigates whether combining AI with mechanism-based models can improve phenotype predictions. AI methods that interpret what genes do in context will be paired with mechanism-based simulations of how community-level behavior emerges. The research team will build a new context-aware model that combines two sources of information: predictions of what individual enzymes do and predictions based on which genes tend to occur together across genomes. Gene predictions from this AI-based model will be compared against mechanism-based models informed about energetic limits, reaction rates, and experimental laboratory measurements. These predictions will simulate the phenotype(s) observed in ANME and sulfate-reducing bacterial syntrophs and genetic data from closely related, well-characterized cultured microbes. Repeating this iterative comparative process will continuously improve both the AI and the mechanism-based models, ultimately building toward a unified framework for explaining how these cooperative syntrophic associations work—one that is consistent with both their genomic content and the relevant chemistry. If successful, Aim 1 will transform automated gene annotation from a task based on sequence alone into one that accounts for context and can predict and explain complex, community-level behaviors, such as electron sharing and microbial metabolic cooperation.

Aim 2 integrates these advances into the project's fine-scale spatial model to predict microbial phenotype with greater accuracy and computational affordability in the context of a 3D environmental simulation focused on the micron-scale pore spaces where these microbial syntrophs live. Linking a cell's internal metabolism to the movement and reaction of chemicals around it in 3D is typically too computationally expensive to apply in spatially heterogeneous environments. To overcome this, the research team will use AI approximations of the full metabolic calculations. Developing physics-based surrogate models of how chemicals move and react at this fine scale will allow the team to quickly predict how microbial phenotypes shift in the context of local differences in chemical conditions and required nutrient sources.

Combined, these two aims will establish an integrated framework for predicting microbial behavior in natural and engineered systems relevant to bioenergy.

AI-Enabled Multiscale Discovery of Genotype-to-Phenotype Control Points for Lodging-Resistant Bioenergy Sorghum

- **Lead PI:** Maria Salas-Fernandez, Iowa State University
- **Collaborating Institution:** National Laboratory of the Rockies

Bioenergy sorghum is a promising feedstock because it can produce large amounts of plant biomass. However, its tall plants are vulnerable to lodging, which occurs when stems bend or break under wind and other environmental forces. Lodging reduces harvestable yield and increases production risk. At the same time, traits that make stems stronger can make the harvested biomass more difficult to break down into sugars and other products at a biorefinery. Improving bioenergy sorghum therefore requires a way to strengthen plants without reducing their processing efficiency.

This project will develop TITAN, the Tokenized Integrated Twin for Agricultural Resilience, an AI framework that connects genetic variation with plant structure, stalk strength, environmental response, and biomass-processing performance. The goal is to identify the biological control points that determine whether a plant remains standing in the field while still being efficiently converted into fuels and other bioproducts.

TITAN will combine several types of information, including genomics, plant and tissue imaging, cell wall chemistry, mechanical measurements, and environmental data. First, separate AI models will be trained on large existing datasets, even when all measurements were not collected from the same plants. This self-supervised learning approach allows the models to identify useful patterns without requiring a large, fully matched dataset. The project will then generate a smaller, carefully coordinated dataset from fewer than 100 genetically diverse sorghum lines. A transformer-based model will combine the different data types and test whether information from multiple sources improves predictions beyond what can be achieved with any single source.

This project will also create a physics-informed digital twin, or computer representation, of a sorghum stalk. Measurements of cell wall structure and stiffness obtained through atomic force microscopy will be linked with tissue architecture and whole-stalk mechanical behavior. Computer simulations will then predict how stalks bend or fail under wind loading. By combining experimental measurements with these simulations, TITAN will be able to make predictions beyond the conditions directly observed in the laboratory or field while remaining consistent with known biomechanical principles.

In Phase I, the project has four objectives: (1) train data-specific AI models using existing datasets; (2) create an AI-designed, synchronized multiscale dataset and demonstrate the value of combining multiple data types; (3) construct a digital twin pipeline linking nanoscale cell wall properties (e.g., structural and mechanical) to stalk-level mechanics; and (4) demonstrate an integrated

genotype-to-phenotype model capable of identifying genetic and structural changes associated with desired outcomes.

This project is designed to demonstrate a clear AI advantage. The team will target a statistical measure of cross-scale prediction performance, known as R^2 , of at least 0.5 using synchronized measurements from fewer than 100 genotypes. This would reduce the number of costly paired samples needed by approximately 10- to 50-fold compared to conventional supervised approaches. The multimodal model is also expected to improve prediction accuracy by at least 20% compared with models based on a single data type.

Expected outcomes include pretrained AI models for genomics, imaging, and mechanics; a library of models describing the mechanical behavior of biological tissues; a digital twin workflow for sorghum stalks; and an accessible interface for genotype-to-phenotype prediction. These products will be released as reusable modules through the American Science Cloud and Transformational AI Models Consortium platforms.

The broader impact will be a faster and more cost-effective way to design crops that are both resilient in the field and efficient in the biorefinery. This framework will also transfer to other traits, environments, crops, and biological systems. By reducing experimental burden and accelerating the discovery of useful genotype-to-phenotype relationships, this project will support more resilient agriculture, strengthen the U.S. bioeconomy, and advance energy security.

Foundation Models for Metabolic Engineering

- **Lead PI:** Huimin Zhao, University of Illinois Urbana-Champaign
- **Collaborating Institutions:** Argonne National Laboratory; Lawrence Berkeley National Laboratory

Developing metabolic engineering foundation models (MEFMs) is critical for addressing the complexity and scale of biological design problems. Engineering microbial metabolism requires navigating an enormous combinatorial space of genes, enzymes, regulatory interactions, and environmental conditions. Traditional modeling and machine learning (ML) approaches are typically trained for narrow tasks and rely on limited labeled datasets, making them difficult to generalize across organisms and pathways. In contrast, foundation models—large AI models pretrained on diverse biological datasets such as genomic sequences, metabolic networks, protein structures, and multiomics measurements—can learn transferable representations of biological systems and be adapted to many downstream tasks with relatively little additional data. So far, no foundation model has been developed for metabolic engineering.

The Zhao, Martin, and Ramanathan groups are pioneers in developing AI models for metabolic engineering. Building upon previous accomplishments, this project proposes to develop MEFMs through the following specific aims:

- Extend the TorchCell framework developed by the Zhao group to facilitate the development of MEFMs by expanding the

TorchCell database and improving metabolic engineering model architecture.

- Develop an improved genome large language model based on GenSLM (Genome-Scale Language Model), developed by the Ramanathan group for metabolic engineering, and integrate it with TorchCell.
- Validate AI model predictions and generate large amounts of data in a high-throughput manner for fine-tuning and training these models by using biofoundries, including (1) the iBioFAB, National Science Foundation's iBioFoundry, and iCloudBiofoundry at the University of Illinois Urbana-Champaign; (2) DOE's Orchestrated Platform for Autonomous Laboratories (OPAL); and (3) the IDeA (Intelligent Design Assistant) co-scientist infrastructure at Argonne National Laboratory.

If successful, the resulting MEFMs would enable predictive design of metabolic pathways, enzyme functions, and strain engineering strategies across multiple organisms, significantly accelerating the design-build-test-learn (DBTL) cycle that underpins metabolic engineering. More importantly, these MEFMs can transform metabolic engineering into a more predictive and scalable discipline, enabling rapid development of biomanufacturing platforms for chemicals, fuels, and materials.

Predictive Engineering of Microbial Communities

MOBILE-PATH: Machine learning Optimization of Biological Interfaces for Lateral Gene Exchange Pathways

- **Lead PI:** Joshua Atkinson, Princeton University
- **Collaborating Institution:** Oak Ridge National Laboratory

Microbiome engineering faces a central challenge: scientists have yet to understand how genetic information flows through microbial communities via horizontal gene transfer (HGT). The field lacks a comprehensive framework for anticipating how mobile genetic elements (MGEs) traverse and persist within diverse taxa and how environmental factors shape these processes. This knowledge gap makes it challenging to engineer microbiomes by deploying nonnative or engineered microbes in the context of native microbiomes or to tailor the activity of native microbiomes by delivering functional cargos through HGT. To address this challenge, this project will develop (1) an ML-guided model for anticipating MGE transmissibility pathways, (2) an automated framework to parameterize models using environmental metagenomes, and (3) an agentic decision-making interface to identify optimal MGEs to deliver functional cargo to specified target taxa and minimize the spread of genetic information to nontarget hosts. This project's goal is to develop a holistic framework to predict MGE trajectories within microbial communities. To achieve this goal, the research team will leverage publicly available longitudinal datasets from the human microbiome and a plant rhizosphere to parametrize a graph-representation of HGT in microbial

communities. This work will treat MGE movement as a pathway-tracing problem on a dynamic network and leverage higher-order network models and causality-aware graph neural networks to predict the trajectory of MGEs through microbial communities. To extend beyond these two environmental contexts, the team will develop an automated parameterization tool that can utilize arbitrary sets of metagenome-assembled genomes to generate HGT networks with MGE transfer likelihoods based on predictions of HGT machinery and defense mechanisms.

The proposed work will reveal fundamental principles governing MGE movement across microbial taxa to achieve a predictive understanding of HGT that is generalizable across environments and microbiome functions. The MOBILE-PATH (Machine-learning Optimization of Biological Interfaces for Lateral gene Exchange PATHways) AI framework will predict multistep MGE transfer pathways across microbial communities, a capability not currently possible with existing bioinformatic or ML approaches. The AI advantage arises from combining higher-order graph models and causality-aware learning to capture history-dependent, multi-hop transfer dynamics, along with an agentic interface that identifies optimal MGEs for targeted cargo delivery while minimizing off-target spread. This tool will decrease the sequencing costs needed for evaluating MGEs for gene circuit delivery and reduce the number of field trials needed to evaluate these technologies. This AI model also provides a quantitative framework for risk assessment of MGE containment and off-target spread in emerging biotechnologies.

Precision Microbiome Engineering in Anaerobic Communities: From Deconstruction to Bioproduction

- **Lead PI:** Michelle O'Malley, University of California–Santa Barbara
- **Collaborating Institutions:** Lawrence Berkeley National Laboratory; University of California–Berkeley; Cerebras

This project leverages AI to accelerate the bioproduction of medium-chain fatty acids (MCFAs) from engineered anaerobic microbiomes that deconstruct lignocellulose (mixed plant waste). MCFAs are strategic platform chemicals for the bioeconomy because they can be directed towards the synthesis of industrial commodities, plastics and advanced materials, and even pharmaceutical and personal care products. However, there are few natural sources of MCFAs. Despite the potential applications, scientists lack an understanding of how environmental conditions, microbial “dark matter,” and microbe–microbe interactions control the metabolism and assembly of anaerobes into stable consortia with predictable function to produce valuable bioproducts like MCFAs.

This project’s approach to closing that gap is enabled by a high-throughput anaerobic biofoundry that trains and deploys AI interventions to realize engineered microbiomes with a predicted profile of MCFAs. In Phase I, the research team will (1) design, build, and test thousands of natural and synthetic consortia that make MCFAs via a unique high-throughput anaerobic biofoundry

to train AI and ML models; (2) construct a “genome-agnostic” ML pipeline that uses probabilistic modeling to learn the critical drivers of anaerobic consortia assembly; and (3) leverage large language models to build genome embeddings that direct genomic edits of consortia members to improve design-build configurations for synthetic microbiomes. By deploying AI-based tools, this project will reveal a fundamental understanding of microbiome design rules, develop a suite of genetic and nongenetic (e.g., environmental) tools for microbiome-based bioproduction control, and transform native consortia into programmable carbon-conversion engines. The AI-based tools enabled by this project’s biofoundry will accelerate the pace of scientific understanding and translation, achieving precision control of microbiome-based MCFAs bioproduction in 6 months—more than 60 times faster than conventional approaches.

PHOCUS: Phage-Host Interaction Programming for Anaerobic Microbiome Control Using AI-Guided Design

- **Lead PI:** Yinyin Ye, University at Buffalo
- **Collaborating Institutions:** Johns Hopkins University; Pacific Northwest National Laboratory

Medium-chain carboxylic acids (MCCAs) are important platform chemicals for a plethora of applications in industry and agriculture, ranging from aviation fuel to feed additives. Anaerobic microbiomes can effectively and energy-efficiently incorporate molecular building blocks into target MCCA products with minimal carbon loss. However, high-yielding MCCA production is largely challenging due to the presence of competitive anaerobes, such as sulfate-reducing bacteria and acetoclastic methanogens, that divert substrate carbon flux toward undesirable products.

The proposed research aims to design and demonstrate clear advantages of an AI-guided framework to advance phage-mediated engineering of anaerobic microbiomes for high-yield MCCA production. Phage-mediated engineering through targeted discovery and cocktail assembly is a promising strategy for selectively inhibiting metabolic pathways that compete with MCCA production. However, the efficacy of this approach is currently hindered by a limited understanding of environmental phage diversity, which often leads to blind experimental designs for phage isolation.

To overcome these constraints, this project proposes leveraging AI to better predict the associations between phages in diverse environmental sources and key MCCA-competing microbes. By shifting from stochastic phage hunting to evidence-based, tailored isolation, this AI-driven framework should increase the precision and scalability required for successful phage applications. The successful completion of this work should provide an immediate, evidence-based foundation for the metabolic pathway engineering of anaerobic microbiomes. Beyond MCFAs, the developed AI framework and curated datasets may also be able to provide an adaptable and scalable solution for engineering diverse microbial systems to produce high-value products.

Biodesign

AI-Driven Design of Gene Expression Programs in Plants Using Foundation Models and High-Throughput DNA Synthesis

- **Lead PI:** Jennifer Brophy, Stanford University
- **Collaborating Institutions:** Lawrence Berkeley National Laboratory; Cold Spring Harbor Laboratory

Plants are central to major national priorities. They produce the biomass that supports domestic biofuels and biobased materials. Certain species can extract cobalt, nickel, and rare earth elements from U.S. soils, providing a domestic route to critical minerals used in batteries and electronics. Progress on these priorities depends on more precise control over how plants grow and what they produce, which in turn depends on controlling the genetic programs that tell each plant cell which genes to switch on or off. Today, designing these genetic programs is slow and unreliable. Fewer than one in 10 designs work as intended, and those that do rarely work across more than one species.

The research team will build an AI foundation model for plant gene regulation in three steps.

- **Step 1:** Bring together the large collection of plant genetic and experimental data that has accumulated in public databases over the past two decades and convert it into a single, consistent resource suitable for modern AI training.
- **Step 2:** Train a foundation model to learn how DNA sequence shapes gene activity across plant species and measure how the model's accuracy improves as data and model size scale up to help the team estimate how much additional data and computation are needed to reach practical design accuracy.
- **Step 3:** Connect the model to a high-throughput laboratory experiment in sorghum, a widely grown feedstock for bioenergy and critical minerals, to synthesize thousands of AI-designed DNA sequences, measure how they perform in plant cells, and feed those measurements back to refine the model. This iterative design-experiment cycle is the mechanism by which the foundation model is adapted to control gene regulation in specific plant species.

Ultimately, this project's foundation model for plant gene regulation will address a long-standing bottleneck in plant engineering: the slow, unreliable design of genetic programs that precisely control gene activity across plant species. Additionally, the trained model, benchmarks, and software interfaces will be released openly through DOE's Genesis Mission ecosystem, where they can support research and development across academia, national laboratories, and industry to design plants for domestic production of valuable chemicals and materials, recover critical minerals from U.S. land, and develop crops that withstand drought and heat.

Generative AI for Genome Design: Toward Predictable Synthetic Genomes

- **Lead PI:** Farren Isaacs, Yale University
- **Collaborating Institutions:** Harvard University; Sandia National Laboratories

Living organisms are endowed with an extraordinary reservoir of biosynthetic and molecular capabilities that remain largely untapped for applications central to DOE's biodesign mission for energy security: synthesizing advanced materials and chemicals, engineering programmable cells for remediation, or extracting rare earth elements. Realizing this potential requires moving beyond incremental genetic modifications toward the rational design of whole genomes, the cell's operating system. While reading and writing DNA have become increasingly routine, intelligently designing whole, built-to-function genomes remains an unsolved challenge of fundamental importance. The field of biodesign lacks high-quality datasets, quantitative design principles, and AI-integrated workflows needed to construct whole genomes and reliably predict their functional outcomes—a critical gap that limits the ability to engineer biology at precision and speed. This proposal addresses that gap by developing SynGen2AI, a generative AI platform that transforms the traditional and protracted (months-to-years) design-build-test-learn cycle into a streamlined (weeks-to-months) design-build paradigm for predictable genome design in DOE-relevant microorganisms. This platform integrates libraries of synthetic genomes with multiomics data to train and validate AI models outputting synthesis-ready, genome-scale designs.

SynGen2AI will be developed and demonstrated in the cyanobacterium *Synechococcus elongatus* UTEX 2973: a fast-growing, photosynthetic, freshwater microbe with an established track record of genetic engineering. This project will pursue three interconnected objectives. First, the research team will build an experimental and computational pipeline that compresses genome design and construction by tenfold, enabling rapid generation of synthetic genomic variants with recoded genomes, enhanced stability, and predictable function. Second, researchers will generate multiomics datasets across libraries of *S. elongatus* UTEX 2973 variants mapped to defined phenotypes, providing the training and benchmarking data needed for robust AI model development. Third, researchers will develop a multi-hybrid neural network and a containerized, end-to-end agentic workflow that operates directly on user-provided genome assemblies with minimal configuration. By accelerating the training and inference speed, this architecture will outperform standard transformers, making million-token, single-nucleotide genome design computationally tractable by at least a factor of two in terms of current codon language models, increasing confidence scores and annotation time by twofold compared to traditional bioinformatic workflows. SynGen2AI (1) learns fundamental rules of genome recoding, codon usage, and gene expression; (2) captures structural features governing genomic instability and mobile genetic elements; and (3) outputs synthesis-ready genome-scale designs.

SynGen2AI will establish a new platform for generative AI-driven design of stable synthetic genomes, with design principles extensible across organisms. By delivering programmable, predictable, and economically viable microorganisms, this work aims to displace tedious strain-engineering campaigns, enhancing the production of aviation and bio-advantaged fuels, chemicals, and materials. Ultimately, this project will establish new frameworks to predict and engineer life at the genome scale. Delivered through cyanobacteria with controlled operating systems, these frameworks will advance DOE's mission in programmable biology and biomanufacturing and help ensure continuous access to secure, high-quality energy resources.

AI-Driven Computer-Ready Synthetic Cells with Quantum Integration

- **Lead PI:** Roseanna Zia, University of Missouri
- **Collaborating Institutions:** University of Minnesota; SLAC National Accelerator Laboratory; Stanford University; J. Craig Venter Institute; Lawrence Livermore National Laboratory; Princeton University; NVIDIA

Synthetic biology is central to the future of advanced biomanufacturing and biosecurity, enabling capabilities such as on-demand molecular-scale fabrication of valuable chemicals and materials. However, current platforms—cell-free systems and engineered natural cells—have reached hard limits in controllability, scalability, and functional scope. Progress toward reliable, programmable biology therefore requires a transition from modifying evolved systems to constructing synthetic cells that are explicitly designed, measurable, and predictable. This project's long-term goal is to establish a comprehensive workflow for producing such designable synthetic cells as the living factories of the future.

A foundational advance toward this goal was the development of minimal cells, such as the model bacterium JCVI-Syn3A, achieved through systematic genome reduction. Such systems demonstrate that simple cells can live but remain governed by evolutionary constraints rather than by engineering intent.

Constructing synthetic cells from the ground up offers a path toward simple controllable systems with defined chemistry and architectures, where function can be explicitly linked to structure. However, current approaches rely on iterative design-build-test-learn (DBTL) cycles that are resource intensive and do not scale to the complexity of whole-cell design. Advancing beyond DBTL requires models that predict cellular behavior prior to construction and are tightly coupled to what experimentalists can make and measure.

Two primary capabilities must therefore be achieved: (1) the ability to design aspects of the fundamental cell apparatus atomistically at whole-cell scale *in silico* and (2) a tightly linked connection between what experimentalists can build and characterize and what computational models can predict. Realizing such a design-build-measure-model workflow requires coordinated advances across academic research, national-scale computing, and industrial AI/high-performance computing infrastructure. Within this framework, a central challenge is the development of whole-cell modeling approaches that provide sufficient predictive accuracy, tractability, and integration with experimentation to support general-purpose synthetic-cell design.

To address this challenge, this project will establish and demonstrate an integrated design-build-measure-model capability for synthetic-cell engineering by carrying benchmark synthetic-cell design problems through a closed-loop workflow. Using synthetic cells as benchmark systems, the research team will design genetic and cell-organizational interventions associated with defined functional behaviors, build those interventions experimentally, measure their effects using multimodal phenotypic and structural readouts, and predict those outcomes in parallel. AI tools—including graph neural networks and related ML approaches—will be integrated throughout the workflow to accelerate multiscale whole-cell modeling; support genome and cell design; and improve analysis of cryogenic electron tomography, visual proteomics, and dynamic localization measurements, ultimately integrating design-build-measure-model demonstrations and identifying where AI confers measurable advantage.

Predicting U.S. Water for Energy Selections

Water availability is essential for expanding production and utilization of energy, as well as national health and security. However, fundamental scientific gaps in understanding terrestrial and atmospheric systems limit the ability to predict water resources, especially on the time scale of weeks to years. AI that is capable of multiscale temporal reasoning could tackle three interrelated grand challenges: cloud physics, surface and subsurface water flows, and the broader hydrologic cycle. AI could (1) improve, accelerate, and couple exascale-class modeling systems through advances in model initialization and (2) develop surrogates trained on DOE's atmospheric and terrestrial observations and laboratory data, at a fraction of the computational cost of existing models. AI-based model diagnostics for enhanced analysis could refine a model–observational system aligned with decision-making needs.

Under this S&T challenge, BER is involved in three focus areas:

- Cloud Microphysics and Atmospheric Turbulence
- Water and Energy
- Weeks-to-Years Prediction

Cloud Microphysics and Atmospheric Turbulence

ASPIRE: ARM-Driven Simulation for Physical Insight Research Engine

- **Lead PI:** Ann Almgren, Lawrence Berkeley National Laboratory
- **Collaborating Institutions:** Lawrence Livermore National Laboratory; Montana State University; Boise State University

Nearly all water available for U.S. energy production comes from precipitation governed by atmospheric processes spanning microscopic cloud microphysics, turbulent mixing, and large-scale storm dynamics. Accurate prediction of seasonal and year-to-year water availability is essential for hydropower management, while microphysical processes such as freezing rain and wet snow directly threaten energy infrastructure reliability through atmospheric icing. Advancing understanding of these processes—and translating that understanding into actionable prediction—requires deep, process-level insight into individual precipitation events, not just statistical associations derived from surface observations alone.

DOE's Atmospheric Radiation Measurement (ARM) user facility provides the richest observational resource available for this purpose, continuously capturing the evolving atmospheric state across dozens of simultaneous data streams. The scientific power of ARM lies in this multistream simultaneity: when independent data streams converge on a consistent physical explanation, they enable genuine causal understanding rather than statistical inference alone. Yet ARM datasets are significantly underutilized—meaningful scientific analysis currently requires labor-intensive workflows that produce results only years after

data collection, and opportunities for near-real-time insight to guide data collection during active events are routinely missed.

The ARM-Driven Simulation for Physical Insight Research Engine (ASPIRE) is designed to close this gap. Its vision is a future in which scientists have real-time situational awareness during ARM field campaigns, can diagnose the microphysical and turbulent processes controlling surface precipitation as events unfold, and can systematically identify where current process models require improvement, thereby accelerating both the science of cloud microphysics and the applied prediction of U.S. water for energy.

AI is essential to this vision for three reasons. First, the volume and diversity of ARM data streams—spanning ground-, air-, and space-based sensors across radar, lidar, radiometer, and *in situ* modalities—exceed what any human analyst can synthesize in real time. Agentic AI enables automated multistream fusion and consistency analysis that elevates output from pattern recognition to physically grounded causal inference. Second, the iterative coupling of observational analysis with process model simulation—running, interpreting, and refining ensembles of atmospheric simulations in response to what observations reveal—is computationally and logistically intractable without autonomous agent coordination. ASPIRE's agentic reflexion loop compresses this cycle from weeks to hours. Third, AI-enabled sensitivity analysis and uncertainty quantification allow ASPIRE to distinguish parametric uncertainty from structural model deficiency, directly identifying where new physics representations or enhanced observations are most needed and providing a principled basis for testing next-generation ML-based microphysics emulators. Together, these capabilities transform ARM data from an underutilized archive into a continuously actionable scientific resource for U.S. energy applications.

Latent Representation Learning for Understanding Ice Microphysical and Precipitation Processes Constrained by Radar and *In Situ* Observations

- **Lead PI:** Christine Chiu, Colorado State University
- **Collaborating Institution:** Pacific Northwest National Laboratory

Ice-containing clouds play a central role in Earth's energy balance and water cycle, generating a substantial share of global rain events from frozen precipitation particles that melt before reaching the surface. Yet they remain among the least understood atmospheric systems because ice crystal growth depends on both local conditions and a particle's growth history, which observations rarely capture directly. Observations also sample only portions of clouds, providing limited connections between cloud processes aloft and the precipitation reaching the surface. These fundamental limitations motivate a new perspective that starts from surface precipitation and works backward to identify the minimum information needed to constrain precipitation formation.

To implement this outcome-oriented strategy, this project leverages recent AI advances together with comprehensive observations from DOE's ARM user facility. AI is foundational to this effort because it can extract compact, physically meaningful latent representations from high-dimensional radar and *in situ* observations that cannot be fully uncovered using traditional analysis methods. These physically interpretable representations provide a new way to connect cloud structure with microphysical states and processes, offering a pathway to understanding how microphysical processes emerge at observable scales.

The primary deliverable will be a first-of-its-kind, ARM-data-driven foundation model that provides reusable representations for a broad range of atmospheric science applications without retraining the core model. This framework can be progressively enriched by incorporating additional measurements from diverse ARM instruments, integrating these observations into a unified system to address emerging scientific challenges. Combined with the scale and breadth of ARM observations, this AI-driven approach establishes a flexible and extensible foundation that expands scientific capabilities, supports energy and water resources applications, and enhances accessibility and usability of ARM datasets for AI-driven analysis.

From Molecules to Precipitation: Machine Learning Bridges Aerosol Microphysics, Atmospheric Turbulence, and Cloud Formation

- **Lead PI:** Celia Faiola, University of California–Irvine
- **Collaborating Institution:** Pacific Northwest National Laboratory

Scientific and applied energy vision. Accurately predicting precipitation location, timing, and intensity is essential for managing U.S. energy infrastructure, including hydropower reservoir inflow forecasting, flood resilience and risk assessment

for thermoelectric plants, and grid resilience planning under severe convection. Current models cannot explicitly represent the subgrid chain of processes linking individual aerosol particles to cloud droplet activation, precipitation initiation, and storm intensity. This gap is especially pronounced for aerosol–cloud interactions, where particle viscosity, phase state, and multiphase chemistry strongly influence cloud condensation nuclei activity, cloud formation, and precipitation. This project addresses that gap by developing the first particle-resolved, physics-informed ML surrogate that will make aerosol microphysics computationally feasible within cloud-resolving models. Phase I of the project will demonstrate an AI advantage in aerosol–cloud microphysics workflows using an ARM site in Alabama as the primary observational benchmark. Contrasting test scenarios will be targeted for developing and testing the surrogate model, including springtime new particle formation and summer organic aerosol condensational growth. Phase II will deploy the validated surrogate within WRF-Chem (Weather Research and Forecasting with Chemistry) for regional-scale precipitation simulations, accelerating scientific discovery and improving predictive tools for energy system planning and resilience.

AI-enhanced scientific and technical workflows. State-of-the-art aerosol descriptive models, such as the Kinetic Multilayer model of Gas-Particle Interactions (KM-GAP) coupled with the 2D Volatility Basis Set with peroxy radical chemistry (r2D-VBS), are prohibitively expensive for per-particle, per-timestep integration in large-scale atmospheric models. AI makes this problem tractable. The proposed project will advance the Genesis Mission by developing AI-accelerated scientific workflows to address the computational challenge of simulating detailed subgrid processes at large scales. The data curation and model training pipeline involves developing a physics-constrained ML model—a Physics-Informed Feedforward Neural Network (PIFNN) surrogate—trained on KM-GAP/r2D-VBS simulations to replicate the full complexity of diffusion-limited mass transport and multiphase chemistry related to organic aerosol formation. The computational cost will be compatible with integration into the particle- and turbulence-resolved Explicit Mixing Parcel Model with Chemistry (EMPM-Chem). By integrating all findings (e.g., best practices, models, and data) with the Genesis Mission Consortium, this workflow pipeline enables transferable AI capability to other regions and larger-scale model applications, such as WRF-Chem and the Energy Exascale Earth System Model (E3SM). By embedding explicit aerosol kinetic processes into particle- and turbulence-resolved simulations at a computationally feasible cost, the AI surrogate will reveal how kinetic aerosol effects shift precipitation onset and storm intensity in southeastern U.S. convective systems, transforming the fidelity of precipitation prediction for energy-relevant applications.

Physics-Informed Machine Learning of Cloud Microphysics for High-Resolution Earth System Modeling with Observationally Constrained Online Training*

- **Lead PI:** Yan Feng, Argonne National Laboratory
- **Collaborating Institutions:** National Center for Atmospheric Research; Aeolus

Accurate prediction of precipitation patterns and extremes is essential for energy sector resilience, yet Earth system models (ESMs) remain limited by persistent biases arising from the parameterization bottleneck in cloud microphysics and turbulence. Microphysical processes, spanning from submicron to centimeter scales, are strongly modulated by small-scale turbulence that existing bulk schemes in ESMs cannot resolve. While high-fidelity simulations such as direct numerical simulations (DNS) offer greater physical realism, their computational cost precludes direct use in ESMs. Existing ML approaches face a structural impasse: “offline” models suffer from numerical instability, while “online” models trained on coarse reanalysis data fail to capture underlying microphysical processes.

To address these challenges, this project proposes PRECISION (Physics-Informed Machine Learning of Cloud Microphysics for High-Resolution Earth System Modeling with Observationally Constrained Online Training), a hybrid ML framework designed to bridge process-level realism and global-scale tractability. The framework comprises three stages: (1) DNS-based pretraining, where deep neural networks (NNs) learn residual corrections to bulk schemes and reduced-order representations of microphysics; (2) neural large eddy simulation (LES) development, integrating NN-based microphysics into a graphics processing unit-first, end-to-end differentiable atmospheric model; and (3) observation-constrained online learning, using DOE ARM data to refine predictions across cloud regimes and reduce systematic biases. Full differentiability enables gradient-based nudging toward real-world observations, creating a “Neural LES” that learns unresolved microphysics–turbulence interactions directly from observations.

Phase I will demonstrate a Neural LES framework. Phase II will enable online training and scale the framework to a neural meso-scale model suitable for kilometer-scale ESM simulations. Key AI innovations include process-resolving learning from DNS, adaptive and stable online training that mitigates distribution shift, and cross-scale generalization within a differentiable framework that transfers consistently from DNS/LES to mesoscale and ESMs. This project represents a paradigm shift, leveraging AI to unify a resolution-spanning hierarchy from DNS (millimeter-scale) to LES (meter-scale) and mesoscale systems, ultimately delivering reliable predictions of the water cycle and extreme weather for the energy sector on the subseasonal to seasonal to decadal scales.

Bridging Resolution Disparities in the Genesis of Mixed-Phase Clouds (BRIDGE)

- **Lead PI:** Anna Gannet Hallar, University of Utah
- **Collaborating Institutions:** University of Texas–El Paso; Pacific Northwest National Laboratory

Mixed-phase clouds are a major source (around 50%) of Earth’s precipitation, yet their formation remains difficult to predict because current models cannot adequately represent the first formation of ice in clouds. Ice initiation depends on ice-nucleating particles (INPs), but existing parameterizations are poorly constrained by aerosol composition, vertical structure, and atmospheric regime. This project will develop a physics-informed AI framework to infer aerosol composition from long-term observations and predict INP activity from vertically resolved aerosol measurements, with the goal of improving mixed-phase cloud prediction and reducing uncertainty in precipitation and water-cycle models.

The proposed research addresses a fundamental scale gap between aerosol chemistry at the nanometer scale and cloud glaciation at kilometer scales. The research team will combine long-term surface aerosol measurements, vertically resolved aerosol composition profiles, and INP observations from DOE user facilities (ARM and the Environmental Molecular Sciences Laboratory) and past field campaign data to build a two-stage AI framework. First, a multimodal inverse model will reconstruct size-resolved aerosol composition from long-term chemical, physical, optical, and meteorological observations. Second, a hybrid physics-informed model based on water-activity freezing theory will use reconstructed and directly measured vertical composition profiles, together with boundary-layer properties, to predict temperature-dependent INP activity under well-mixed and decoupled conditions.

This project leverages a unique observational foundation, including more than 500 vertically resolved aerosol composition profiles spanning millions of particles collected across the continental United States. These data provide direct constraints on aerosol composition at the single-particle level to evaluate how particle source, mixing state, altitude, and thermodynamic conditions influence ice formation. The proposed AI framework adds a major advantage by enabling the integration of heterogeneous multiscale observations that are not easily combined through conventional parameterizations. This integration will produce AI-ready datasets, uncertainty-aware prediction workflows, and physics-informed models that can be incorporated into cloud-resolving and Earth system models.

This project will improve prediction of precipitation and water availability, enhance understanding of aerosol–cloud interactions, and produce new AI tools for DOE’s Water for Energy mission and Genesis Mission priorities. By linking observation, physics, and AI, the project will advance predictive understanding of mixed-phase clouds and support national priorities related to rainfall prediction and water availability.

* BER is lead; co-funded with the DOE Advanced Scientific Computing Research (ASCR) program

Advancing Multimodal and Multidimensional AI To Discover Connections Between Cloud Microphysics and Precipitation

- **Lead PI:** Gabriel Isaacman-VanWertz, Virginia Polytechnic Institute and State University
- **Collaborating Institutions:** Argonne National Laboratory; Loyola University Maryland

Predicting exactly when, where, and how hard it will rain is critical for protecting national infrastructure from extreme weather and for estimating water availability for agriculture and energy production. However, precipitation modeling remains one of the most difficult challenges for models of Earth's weather and environmental systems. Most models struggle to predict the timing and location of heavy storms, the size and number of raindrops, and the intensity of rainfall. This difficulty stems from the complexity of the interactions inside clouds that lead to rainfall, such as the sizes of droplets inside clouds and how they collide and interact with each other. DOE's ARM user facility has gathered decades of highly detailed data on clouds, rain, and the processes that drive them, but the volume and complexity of these measurements exceed the capability for human-driven analysis.

To convert these large observational datasets into predictive power, this project introduces a unique approach by training an AI that is "multimodal," or able to simultaneously analyze numerical data (e.g., precise temperature and humidity readings) and visual representations of that data (e.g., graphs and radar images). Just as a human scientist looks at a weather map or a graph to spot shapes and patterns, the AI will use computer vision to recognize large-scale patterns and at the same time use precise measurements to quantify the relationships within these patterns. In this project, data from some ARM sites with the longest-term data records in the United States will be used to build multimodal AI models that can bridge fine-scale details with large-scale pattern recognition. Furthermore, the AI model's discoveries will be compared against current DOE models to identify failure points in the models to facilitate future improvement in precipitation forecasting.

By coupling the rich, long-term datasets available from DOE with this state-of-the-art AI-driven analytical approach, the team will extract new understanding of the processes inside clouds to better predict precipitation and pave the way for safer, better prepared communities. The new methods and next-generation multimodal AI model will not only unlock new information about clouds and precipitation but will also have broad applications for analysis of complex or large-scale data across science and engineering. In addition, by building a team across academic and national laboratory institutions that combines leaders in atmospheric science with experts in advanced AI techniques and ML, this research lays a foundation for future efforts to tackle the most difficult and pressing issues in environmental and atmospheric science with the most cutting-edge analytic methods.

Cloud Microphysics Multiscale Modeling Moonshot (CM4): Creating the Next Generation of Uncertainty-Aware Cloud Models by Leveraging Multifidelity AI and DOE Observations

- **Lead PI:** Kara Lamb, Columbia University
- **Collaborating Institutions:** Pacific Northwest National Laboratory; National Center for Atmospheric Research

The United States' energy infrastructure, agriculture, and many industries depend critically upon water. Weather and Earth system models, including those used by DOE, have long struggled to accurately represent the details of cloud and precipitation processes, resulting in biases in overall patterns of precipitation intensity and frequency. A major source of uncertainty is cloud microphysics, which includes the processes by which droplets grow and form rain and their interactions with atmospheric turbulence.

To address this challenge, the Cloud Microphysics Multiscale Modeling Moonshot (CM4) project will develop a multiscale framework combining physics-based modeling, ML, and agentic AI to improve understanding of how interactions between cloud droplets and turbulence drive rain formation. CM4 will develop turbulence-aware cloud microphysical parameterizations by machine learning physically constrained process rates from detailed simulations informed by laboratory studies and in-cloud observations. The research team will use agentic AI to rapidly convert an existing model of cloud turbulence and dynamics into a form that can be efficiently tested and improved using observations. The model will incorporate the project's improved description of how turbulence affects the collision and merging of cloud droplets, helping the research team quantify uncertainty and improve the representation of rain formation in weather and Earth system models.

All datasets, models, and workflows will be shared through the American Science Cloud and integrated with the Transformational AI Models Consortium's AI-driven code development effort to support reproducibility and generalizability to other DOE scientific applications.

A Spatial Generative Bayesian Computation Framework for Inferring Turbulence–Microphysics Interactions

- **Lead PI:** Pulong Ma, Iowa State University
- **Collaborating Institution:** Argonne National Laboratory

Water availability is essential to U.S. energy production and national security, yet scientific gaps in understanding of atmospheric systems limit the ability to predict precipitation. Complex nonlinear, multiscale interactions governing cloud microphysics and atmospheric turbulence are central to the formation, phase, and intensity of precipitation. These interactions are poorly understood and hence inadequately represented in Earth system models. Traditional numerical simulation-based approaches to

represent these interactions suffer from reliance on parameterizations, artificial separation of hydrometeors, inability to resolve air motions across scales, and computational bottlenecks. This project aims to recover latent microphysical and dynamical states from real Doppler spectra observations primarily collected at DOE's ARM user facility observatories, thus reformulating this problem as an inverse problem. Solving this inverse problem via optimal estimation or existing ML methods is computationally prohibitive, as it cannot scale with high data volume.

To address these challenges, this project will develop a neural simulation-based inference (SBI) framework via amortized inference to infer turbulence–microphysics interactions under a wide range of weather conditions. The proposed framework will train AI models offline with an upfront cost and enable instantaneous posterior inference for new observations, thus making it scale well for large data volume and high-dimensional data and allowing real-time inference. The project includes an interdisciplinary team with researchers from atmospheric science, remote sensing science, ML, and statistics. This project will leverage the team's expertise to make synergistic advances in these areas and focus on addressing the problem of real-time microphysical retrievals at scale. This framework will be scaled to a community retrieval tool across the full ARM observation network. The trained models and datasets resulting from the proposed framework will be contributed to the American Science Cloud and the Transformational AI Models Consortium for reuse across the Genesis Mission ecosystem, with the framework designed for plug-and-play deployment by other teams.

Upscaling Particle-Resolved Aerosol–Cloud Microphysics with Generative Modeling

- **Lead PI:** Lekha Patel, Sandia National Laboratories
- **Collaborating Institutions:** University of Illinois Urbana-Champaign; Pacific Northwest National Laboratory

Reliable predictions of clouds and precipitation are important for managing water resources and energy systems, but these processes emerge from complex interactions among atmospheric particles, temperature and humidity, and cloud dynamics that remain difficult to represent accurately in large-scale atmospheric models. This project tests whether physics-constrained ML can learn compact aerosol representations that preserve the particle properties most important for cloud droplet formation more effectively than the simplified summaries used today. This project will couple a variational autoencoder with a neural stochastic differential equation, trained on high-fidelity particle-resolved simulations driven by realistic cloud conditions from DOE's ARM user facility, to predict warm-cloud activation and early droplet growth. Trained models, benchmark datasets, and computational workflows will be shared through the Genesis Mission's American Science Cloud platform, with the software released open source for use across the initiative. If successful, this feasibility study will establish the foundation for embedding AI-derived aerosol physics in DOE atmospheric models, improving predictions of clouds and precipitation that underpin water resource and energy-system planning.

EARL: Earth-Atmosphere Agentic Research and Learning—Physics-Constrained AI Closure Development for Cloud Microphysics and Turbulence

- **Lead PI:** Kyle Pressel, Pacific Northwest National Laboratory
- **Collaborating Institutions:** Michigan Technological University; University of Notre Dame

This project advances cloud microphysics and precipitation modeling on two fronts. First, it will develop a physics-constrained ML closure for unresolved subgrid-scale (SGS) turbulence effects on cloud-droplet size distribution (DSD) broadening, formulated within two- and three-moment predicted particle properties (P3) microphysics. The closure is trained using direct numerical simulation with super-droplet microphysics benchmarks for Pi chamber-like conditions and matched P3 box-model simulations, then implemented in the Energy Research and Forecasting (ERF) model. Second, codeveloped with the closure, EARL (Earth-Atmosphere Agentic Research and Learning) provides a human-supervised workflow for ERF closure development and agentic experimentation to explain why a given closure succeeds or fails. Together, they deliver a machine-learned representation of the turbulence–microphysics coupling that drives cloud-droplet growth toward rain and a workflow that is extensible across model components.

DOE's Water for Energy mission depends on accurate forecasts of runoff, reservoir inflow, and cooling-water availability, all of which trace back to precipitation. Two barriers stand in the way. First, existing bulk microphysics schemes omit SGS turbulence-driven DSD broadening that drives droplets to sizes at which collision-coalescence takes over; the result is delayed drizzle onset and errors in rainfall timing and amount. Second, closure development and model experimentation are serial and manual, limiting the range and rigor of candidate evaluation and the pace of growth in process-level understanding.

This project pursues two linked objectives, both demonstrating AI advantage. Under AI-for-physics, the research team will develop an ML closure for SGS supersaturation effects on cloud condensational growth in two- and three-moment P3, formulated as a learned residual correction to analytic grid-mean tendencies. Candidate families include physics-constrained neural networks, symbolic regression, and hybrid parametric-plus-learned residual closures. AI advantage is shown by improved DSD-broadening skill and stable ERF integrations. Under AI-for-workflow, the team will build EARL as a portable, human-supervised agentic workflow that assembles data, generates and screens closures, and stages build-test-run-score cycles, and then expands its role to orchestrate model experimentation and hypothesis generation. AI advantage is quantified by efficiency gains in closure development and experimentation. Combined AI advantage: EARL's agentic experimentation converts faster closure development into greater process-level understanding, delivering an advantage that neither front achieves alone.

The primary deliverable is a validated SGS condensational-growth closure deployed in the project's P3 implementation

in ERF, reducing DSD-broadening errors that propagate into water- for-energy predictions. Because P3 is also the microphysics scheme in the Simple Cloud-Resolving Energy Exascale Earth System (E3SM) Atmosphere Model (SCREAM) and other E3SM configurations, the closure offers a direct transfer path into DOE's global model ecosystem. EARL extends the project's reach further as a portable, human-supervised agentic workflow reusable across ERF model development and experimentation, positioned for Genesis Mission coordination, American Science Cloud readiness, and broader reuse.

An Automated, Multimodal-AI-Enabled Cloud Chamber for Constraining Cloud Microphysical Processes in Earth System Models

- **Lead PI:** Arthur Sedlacek, Brookhaven National Laboratory
- **Collaborating Institutions:** Michigan Technological University; Columbia University; Stony Brook University; Quantum Computing, Inc.

Achieving a transformative understanding of cloud microphysics–turbulence interactions and their role in precipitation formation requires a fundamentally new, AI-centered experimental paradigm. This project proposes a discovery-driven embodied AI framework for controlling next-generation facilities, such as the DOE-envisioned user facility DRACO (DRizzle, Aerosol, Cloud Observatory), for accelerating scientific discovery that directly supports the nation's water and energy security. At its core, this framework addresses the high-dimensional, strongly coupled control challenges inherent with large-scale convective cloud chambers that are fundamentally intractable using existing methods.

Embedding AI into the control loop transforms the cloud chamber from a passive experimental platform into an active, programmable system capable of targeting, achieving, and sustaining specific microphysical states on demand, thereby enabling systematic, repeatable, and precisely controlled investigations of cloud processes. Traditional control and experimental strategies applied to existing cloud chambers are insufficient for navigating the nonlinear, multiscale interplay among cloud turbulence, microphysics, and precipitation. A discovery-driven embodied AI framework not only enables control of various cloud regimes but also fundamentally redefines the experimental process for scientific discovery. In this framework, agentic large language models support high-level reasoning, planning, and coordination of autonomous experimental workflows. Agent decisions are informed by detailed-physics modeling (e.g., large eddy simulations), providing physics-based constraints, then passed to a continuous control model responsible for multimodal sensor fusion and low-level hardware control. The system continuously senses the chamber state and determines optimal control actions to achieve and maintain target cloud conditions, while adapting control policies via reinforcement learning. Collectively, this framework transforms how experiments are designed, executed, and interpreted. Crucially, the system is designed not only to reach predefined states but also to autonomously explore a high-dimensional parameter space,

identify emergent phenomena, and optimize system control to evaluate scientific hypotheses.

An AI-enabled cloud chamber will generate unique, high-fidelity measurements of coupled aerosol–cloud–precipitation processes, providing critical datasets for integration into the American Science Cloud framework and supporting the Transformational AI Models Consortium mission through the development and evaluation of foundation models for atmospheric science. These capabilities will advance fundamental understanding of precipitation formation and improve predictive models. To realize this vision, this project proposes a phased development pathway for a discovery-driven embodied AI cloud chamber. Phase I will leverage the 1 m³ convection cloud chamber at Brookhaven National Laboratory to demonstrate proof-of-concept autonomous control, in which AI dynamically adjusts boundary forcings to achieve targeted cloud states (e.g., liquid water content). Phase II will extend this framework to a three-level configuration, a factor of 30 times more interactions relative to the Pi chamber. Together, these efforts will position the large-scale DRACO chamber to enable transformative studies of cloud microphysics–turbulence interactions and their role in precipitation.

AI-Accelerated Exploration of Droplet Collision-Coalescence Using Observation-Constrained, Multiscale Modeling of a Turbulent-Convection Cloud Chamber

- **Lead PI:** Raymond Shaw, Michigan Technological University
- **Collaborating Institutions:** Brookhaven National Laboratory; Brigham Young University; University of Utah; National Center for Atmospheric Research; Pacific Northwest National Laboratory

Accurately predicting the onset of precipitation, specifically the transition from condensation-dominated droplet growth to collision and coalescence, remains a fundamental challenge in atmospheric science. In turbulent clouds, the evolution of droplet size distributions is driven by a complex interplay of aerosol–cloud microphysics and stochastic turbulent fluctuations. Because resolving the governing mathematical equations for these multiscale interactions is computationally prohibitive for the large-scale Earth system models used in energy and weather decision-making, precipitation formation is currently represented by highly uncertain parameterizations. This project addresses this critical gap by using AI to transform complex data streams from laboratory experiments and multiscale simulations into computationally efficient, interpretable laws of cloud microphysics.

The primary objective of this Phase I effort is to implement a physics-constrained, AI-enabled framework that discovers the underlying principles governing the nonlinear microphysics of precipitation. By overcoming the limitations of sparse experimental data, this project aims to identify key unobservable macroscopic parameters controlling droplet growth. Additionally, the research team seeks to develop a physics-informed reduced-order model (ROM) capable of replacing slow, full-order

fluid dynamics equations for active exploration and real-time simulation. Ultimately, this effort will set the stage for cloud chamber experiments to efficiently connect to broader Earth system modeling communities, including the Transformational AI Models Consortium.

To accelerate the discovery process, the research team is building an AI-ready data foundation that standardizes outputs from a hierarchy of detailed-physics computational models alongside real-world observations from the Michigan Technological University's Pi Convection-Cloud Chamber. Utilizing this integrated data, the project advances through two parallel, physics-constrained AI workflows.

Parameter Identification. The first workflow employs gradient-based optimization coupled with a data assimilation framework to infer hidden parameters controlling droplet growth from sparse experimental data. By embedding established physical equations into an automatic differentiation framework, algorithms will identify the supersaturation variance governing turbulent droplet broadening and pinpoint the exact conditions where systems transition to a collision-dominated precipitation regime.

Forward Prediction via Operator Inference (OpInf). To bypass the computational bottleneck of solving full-order droplet dynamics equations, the second workflow utilizes OpInf to learn a low-dimensional dynamical system directly from the multiscale data. This ROM enforces the parametric differential structure of the underlying physics, utilizing linear operators to capture condensation dynamics, and quadratic operators to capture the nonlinear pairwise interactions of droplet collisions that trigger rapid drizzle growth.

By integrating physical parameter extraction with structure-preserving surrogate models, this project establishes AI workflows that quantifiably reduce model parameter variance. This approach is particularly important for existing and future laboratory cloud chamber science, which often operates in a "scarce-data regime" in which traditional deep learning models fail due to a lack of training data. Embedding physical structure compensates for this data sparsity, yielding accurate models that will significantly enhance the reliability of large-scale atmospheric predictions.

Scalable Twin for Intelligent Turbulence and Cloud Heuristics (STITCH)

- **Lead PI:** Yunyan Zhang, Lawrence Livermore National Laboratory
- **Collaborating Institutions:** Argonne National Laboratory; University of Oklahoma; Pacific Northwest National Laboratory; NVIDIA

Over the past 5 years, a new generation of Earth system models has rapidly advanced, including the development and application of global storm-resolving models (GSRMs)—such as DOE's Simple Cloud-Resolving Energy Exascale Earth System (E3SM) Atmosphere Model (SCREAM)—for global simulations of deep convective processes and their multiscale interactions. However, GSRMs

still face significant predictability challenges due to uncertainties in representing small-scale cloud microphysics and atmospheric turbulence. Biases in clouds and precipitation can lead to errors in surface water and energy budgets.

AI and ML offer strong potential to improve these processes and reduce model biases and uncertainties by:

- Learning from high-fidelity models, such as large eddy simulations (LES), to emulate complex fine-scale physical evolution.
- Emulating model physics to accelerate ensemble prediction and model calibration.
- Integrating observational data with model physics to support physics-informed model development and data assimilation.

Phase I of Scalable Twin for Intelligent Turbulence and Cloud Heuristics (STITCH) proposes to:

- Build a testbed to implement and evaluate process emulators for turbulence and warm-rain microphysics against benchmark LES data and observations by DOE's ARM facility.
- Create an emulator of regionally refined SCREAM (RRM-SCREAM) for simulating clouds and precipitation around the ARM sites.
- Develop a workflow to assess the impact of initial conditions on RRM-SCREAM simulations through data assimilation of ARM observations.

In Phase I, researchers will generate LES training data to develop process emulators using a range of neural operator approaches. The team will also generate RRM-SCREAM training data to build a cloud and precipitation emulator for regions around ARM sites. In parallel, the team will generate initial conditions by assimilating ARM observations, including scanning radar reflectivity and sounding array measurements, and evaluate the impact of those initial conditions on RRM-SCREAM simulations.

If successful, these Phase I efforts will establish the foundation for creating an ARM data cube by integrating RRM-SCREAM and ARM observations. This capability will support the development of an ARM digital twin for discovering new physics and further improving the representation of cloud microphysics, turbulence, and their emulators. Together, these capabilities will form a positive feedback loop that incorporates new observations, model advances, and ML methods to continually improve the digital twin and model predictability of water and energy at regional and global scales.

Water and Energy

AI-Guided Model–Experiment Framework for Predicting Water Availability in Energy Systems Using Multifidelity Watershed Modeling

- **Lead PI:** Xingyuan Chen, Pacific Northwest National Laboratory
- **Collaborating Institution:** University of Illinois Urbana-Champaign

Water availability in watersheds is governed by coupled interactions among atmospheric forcing, land-surface processes, and subsurface hydrology across multiple spatial and temporal scales. Predicting these interactions remains a fundamental challenge because nonlinear feedbacks among precipitation, evapotranspiration, groundwater storage, and stream–aquifer exchange introduce significant uncertainty that is difficult to quantify and interpret. This uncertainty directly impacts energy-relevant decisions related to hydropower generation, thermoelectric cooling, reservoir operations, and drought resilience.

This effort proposes an AI-guided model–experiment (ModEx) framework integrating data, mechanistic modeling, and structured reasoning to improve prediction and reduce uncertainty. The framework uses the Advanced Terrestrial Simulator (ATS) as the mechanistic prediction engine, SciLink as the semantic and executable layer for data integration and model orchestration, and the Agentic Discovery and Exploration Platform for Tools (ADEPT) as the reasoning layer that (1) translates water–energy questions into model-informed hypotheses, (2) guides informative simulations and observational comparisons, and (3) interprets outcomes under uncertainty and ambiguity. Because ATS is computationally expensive at the ensemble sizes and spatial resolutions needed for uncertainty-aware prediction, the team will also use multifidelity and surrogate models to enable rapid scenario exploration while preserving physical realism.

Phase I will build and test this framework to predict late-summer streamflow and groundwater heads relevant to municipal supply reliability and prospective cooling-water demand in the South Branch Kishwaukee Watershed in Illinois, with the Oak Creek Watershed in Washington serving as the technical benchmark and validation watershed. Using an agentic workflow that links data ingestion, simulation, evaluation, reasoning, and human validation in a closed loop, the team plans to quantify the dominant drivers and uncertainties of seasonal water availability, including the roles of forcing, parameters, and model structure. The effort will demonstrate measurable gains in predictive skill, uncertainty attribution, and workflow efficiency relative to a nonagentic baseline.

This project will provide a validated prototype demonstrating AI advantage through improved predictive skill, stronger uncertainty attribution, and faster model–experiment iteration. By tightly coupling execution, mechanistic modeling, and reasoning in a closed-loop framework, this project will advance uncertainty-aware, decision-relevant prediction of water availability for energy systems. This capability can support better planning for

drought, cooling-water reliability, and flood-related infrastructure risk. More broadly, the project will lay the groundwork for scalable AI-enabled water prediction tools with regional and national relevance.

RIVER-AI: Reservoir–Groundwater Interactions for River Flow Variability, Energy, and Resilience with AI

- **Lead PI:** Hong-Yi Li, Lehigh University
- **Collaborating Institutions:** University of Texas–Austin; The City College of New York; Pacific Northwest National Laboratory; Oak Ridge National Laboratory

Rivers supply much of the water that the nation's power plants and hydropower facilities depend on, yet predicting river behavior in heavily managed basins remains a persistent scientific challenge. Reservoirs, groundwater, and river networks interact in ways that physics-based models only partially capture. A key barrier is information itself—reservoir operations are governed as much by operating manuals, policy provisions, and institutional practice as by hydrology, and this unstructured knowledge has not been used by conventional numerical models. Managed river basins are therefore a compelling proving ground for whether modern AI can deliver measurable advantages in science and decision support.

This project will develop RIVER-AI (Reservoir–Groundwater Interactions for River Flow Variability, Energy, and Resilience with AI), a hybrid framework that couples three AI advances with a physics-based regional water model. Large language models will turn reservoir manuals and policy documents into machine-readable operating constraints—bringing context awareness to infrastructure simulation—while a computationally efficient AI emulator preserves groundwater physics critically governing river flows during droughts and floods. By binding these components into a coherent whole, a trustworthy AI workflow will produce uncertainty-aware predictions, flag conditions outside the training distribution, and trace any AI gains to physically meaningful processes rather than black-box correlations. This framework will be built and tested in the Delaware River Basin, a heavily managed, data-rich basin serving Mid-Atlantic water and energy systems.

Phase I will put these ideas to a strict test, with three checkpoints that must be passed. First, the research team will test whether AI that learns from operating documents can anticipate reservoir releases better than conventional rule-based methods while following the requirements written in those documents more faithfully. Second, the team will evaluate whether accounting for both reservoir decisions and groundwater's long memory improves predictions of unusually high and low river flows. Third, the team will determine whether these improvements—if they appear—offer useful guidance for hydropower planning and flood and drought preparedness and whether they hold up in parts of the basin where AI was not trained, paving the way for knowledge and capacity transfer to other regions and energy sectors. If successful, this project will help energy providers, water managers, and communities better anticipate and prepare for droughts and flood risks. Just as importantly, it could show—in a real river basin

that supplies water to millions of people—when AI and physics together outperform either approach alone.

HydroORBIT: A Physics-Informed Multimodal AI Foundation Model for Coupled Surface–Groundwater Prediction in Energy-Critical Regions

- **Lead PI:** Dan Lu, Oak Ridge National Laboratory
- **Collaborating Institutions:** Georgia Institute of Technology; Tennessee Valley Authority; Advanced Micro Devices, Inc.

Reliable prediction of water availability is essential for ensuring secure energy production and effective water resource management. Hydropower generation, thermoelectric power plant cooling, and flood preparedness all depend on understanding how water moves above and below the Earth's surface. However, accurately predicting these interconnected surface and groundwater systems remains challenging because of complex, multiscale physical processes, diverse landscape and subsurface conditions, and limited observations.

This project will develop HydroORBIT, a physics-informed AI foundation model designed to improve predictive understanding of coupled surface–groundwater systems. HydroORBIT will integrate weather data, streamflow measurements, groundwater observations, landscape characteristics, and physics-based hydrologic simulations while incorporating established physical principles that govern water movement. By combining scientific knowledge with advanced AI, this project seeks to produce predictions that are more accurate, computationally efficient, and physically consistent than existing approaches.

This project will demonstrate HydroORBIT in three energy-critical applications: forecasting river inflows for hydropower operations in the Tennessee River Basin, mapping groundwater availability in California, and predicting coastal flooding in Florida. These demonstrations will evaluate whether the new AI framework can provide faster, more transferable, and more reliable predictions than current physics-based and ML methods.

HydroORBIT will establish a scalable scientific AI framework for improving predictive understanding of water systems across diverse regions of the United States. The resulting capability will support more informed decisions for energy production, water resource management, and infrastructure planning while providing reusable AI models, datasets, and workflows that accelerate future scientific discovery through DOE's Genesis ecosystem.

RAIN: Rigorous AI and Information Network for Surface–Subsurface Water System Coupling

- **Lead PI:** Mikhael Seeman, EnviTrace, LLC
- **Collaborating Institutions:** Oak Ridge National Laboratory; Lawrence Berkeley National Laboratory; University of California–Davis

Reliable prediction of water availability remains a critical constraint on U.S. energy production, reservoir operations, and flood

resilience. In systems like the Colorado River Basin, water supply emerges from coupled interactions among snow accumulation and melt, runoff generation, soil moisture, recharge, groundwater storage, baseflow, and reservoir dynamics. These coupled dynamics are nonlinear, heterogeneous, and strongly lagged, limiting forecast usefulness at weekly-to-seasonal-to-decadal scales and directly affecting reservoir release planning, water–energy coordination, and flood risk mitigation.

RAIN will develop a physics-informed, information-theory-guided AI approach to coupled surface–groundwater modeling. The project will integrate multisource observations and Earth system simulations with unsupervised regime discovery, temporal information-theoretic analysis, and a dual-branch temporal transformer. These tools will identify physically meaningful exchange pathways, predictive lags, and memory scales, encode them in the AI model while limiting spurious correlations, and support rigorous benchmarking and transfer across hydrologic regimes and scales.

Phase I will produce an AI-ready, uncertainty-aware dataset for selected watershed-reservoir systems spanning distinct Colorado River Basin watersheds and regimes. Planned accomplishments include interpretable hydroclimatic-regime classifications and lag-aware coupling maps that identify candidate exchange variables, memory scales, and uncertainty structure. The project will deliver a calibrated dual-branch surface–groundwater forecasting model informed by those maps. The team will demonstrate and benchmark the model for weekly reservoir-inflow quantiles and flood-threshold exceedance probabilities, evaluating forecast skill, calibration, transferability, and computational efficiency and producing a Phase II scale-up plan.

The resulting capability will aim to improve water-availability forecasts for reservoir operations, water–energy coordination, and flood resilience while advancing fundamental understanding of complex multiscale, multiphysics Earth systems. To broaden public benefit, visibility, and adoption, project outcomes will be shared through peer-reviewed publications, conference presentations, engagement with water and energy stakeholders, and, where permitted, openly accessible software, benchmark datasets, and reproducible workflows. The proposed effort aims to further expand fundamental understanding of how to predict complex multiscale and multiphysics Earth science processes to empower water management strategies toward U.S. energy resilience.

A Foundational Generative AI Framework To Advance Water–Energy Security

- **Lead PI:** Charuleka Varadharajan, Lawrence Berkeley National Laboratory
- **Collaborating Institutions:** Los Alamos National Laboratory; National Laboratory of the Rockies; Massachusetts Institute of Technology; Google

Accurately predicting water availability—both in rivers and underground—is essential for managing hydropower, thermoelectric, fossil, and geothermal energy production, particularly as droughts and floods become more frequent and severe. Existing

models are limited in their ability to predict surface water and groundwater together at the level of detail needed for energy planning, and they typically do not indicate how confident those predictions are.

This project—led by Lawrence Berkeley National Laboratory (LBNL) in partnership with Los Alamos National Laboratory (LANL), the National Laboratory of the Rockies, the Massachusetts Institute of Technology, and Google—will develop an AI framework to address this gap. The approach combines detailed land-surface information from Google’s AlphaEarth Foundations model with underground properties derived from LANL’s Transparent Earth model to identify the surface and subsurface conditions that control water availability. This information will feed into a generative AI model developed by LBNL that predicts how rivers and groundwater will behave together and indicates the reliability of each prediction. Because this AI approach learns directly from data rather than requiring the detailed physical parameters used in traditional models, it offers a faster and more scalable path to high-resolution water predictions over large areas.

This framework will be developed and tested at a well-studied research watershed in Colorado, then applied across the larger Upper Colorado River Basin. Predictions will be benchmarked against existing streamflow and groundwater estimates. The project will evaluate how additional data and model size improve prediction accuracy and reliability. A complementary analysis will use these results to identify which watersheds are more or less able to withstand drought and flooding.

If successful, this research will lay the groundwork for extending these predictions across the continental United States, giving water managers, utility providers, and energy planners more accurate and detailed information to support energy- and water-resource decisions nationwide. The underlying approach may also offer broader value for related challenges, such as understanding subsurface resources relevant to critical minerals and other DOE mission areas.

Multifidelity AI Foundation Model for Coupled Surface–Groundwater Predictions of Water Availability and Flood Hazards Across CONUS

- **Lead PI:** Eugene Yan, Argonne National Laboratory
- **Collaborating Institutions:** Lawrence Berkeley National Laboratory; Princeton University; University of Arizona; University of Illinois Urbana-Champaign; Georgia Institute of Technology; Cognition AI

Reliable water supplies and flood forecasting are essential for U.S. energy production, infrastructure, agriculture, and communities. Power plants, hydropower facilities, data centers, and irrigation and municipal water systems all depend on accurate information about drought, groundwater, soil moisture, streamflow, and flooding. Current large-scale hydrologic models are often too computationally expensive to provide local detail across

the continental United States, and many do not fully represent human influences such as groundwater extraction.

This project will develop an AI foundation model that learns jointly from real-time observational data and multiple physics-based surface water and groundwater simulations spanning resolutions from tens of meters to roughly one kilometer. By fusing these sources within a single generative modeling framework, the model will produce rapid, uncertainty-aware forecasts of groundwater levels, soil moisture, streamflow, and flood depth under both natural and human-influenced conditions. Research will begin with a focus on Midwest water availability and flood hazards and expand toward the first continental-scale, high-resolution AI-driven hydrologic prediction capability.

The resulting tools will support analyses of drought and flood impacts on energy systems and other water-dependent infrastructure, informing both near-term operational decisions and long-term resilience planning. This work will advance scientific understanding of coupled surface water and groundwater processes across scales and provide actionable tools to strengthen national water–energy security.

Weeks-to-Years Prediction

DL4MCS: Deep Learning Methods To Enhance Subseasonal Predictions of Mesoscale Convective Systems by Physics-Based Systems

- **Lead PI:** Hansi Singh, Planette AI
- **Collaborating Institutions:** Pacific Northwest National Laboratory; University of Wyoming

This project will use AI and advanced Earth system models to make weeks-ahead forecasts of dangerous storm systems more accurate and useful for communities, businesses, and governments across the United States. These storms, called mesoscale convective systems (MCS), produce intense rain, hail, high winds, and sometimes tornadoes and are a major cause of flooding, infrastructure damage, and economic disruption.

Why these storms matter. In many parts of the country, MCS supply a large share of the warm-season rainfall that feeds rivers, reservoirs, and groundwater, directly affecting water available for drinking, irrigation, and energy production. At the same time, they are responsible for more property loss than wildfires, hurricanes, and earthquakes combined, damaging homes, roads, power lines, and critical facilities. Better advance warning of where and when these systems are likely to strike can improve preparedness, reduce losses, and support safer, more efficient use of water and energy resources.

The forecasting gap. Traditional weather forecasts are very good for a few days in advance, but their skill drops sharply beyond about a week, especially for complex storm clusters like MCS. Longer-range subseasonal forecasts (i.e., 7 days to 6 weeks) can capture broad patterns in temperature and moisture but often

miss the fine-scale details that determine whether a community experiences a routine thunderstorm or a damaging multi-hour deluge. This gap leaves water managers, emergency planners, utilities providers, and insurers without the reliable, location-specific guidance needed to plan for high-impact events weeks ahead.

What this project will do. This project will build a new forecasting system that combines the strengths of physics-based weather models with modern AI methods to improve storm predictions on subseasonal timescales. First, the team will use AI to expand existing forecast ensembles, effectively generating many more plausible future scenarios so that rare but dangerous events are better represented. Second, the team will train AI models to blend forecasts with key observations of the ocean and land surface—such as soil moisture and upper-ocean heat content—because these slowly changing conditions strongly influence whether conditions become favorable for MCS.

High-resolution storm detail. Finally, this project will apply advanced generative AI techniques to translate coarse global forecasts into high-resolution maps, detailed enough to represent individual storm systems and their evolution over the central and eastern United States. These methods are trained using state-of-the-art regional simulations and cloud observations, allowing the system to better capture how storm structure, rainfall intensity, and cloud processes respond to large-scale atmospheric patterns. By comparing different plausible representations of cloud physics, this project will also quantify uncertainty in rainfall predictions and identify configurations that best match observations.

Benefits to society and the water–energy system. The end result will be a prototype forecast pipeline that can predict MCS activity over the United States up to several weeks in advance with substantially higher skill than today’s operational models. These improved forecasts can inform reservoir operations, drought and flood preparedness, grid reliability planning, and risk management decisions in sectors such as agriculture, transportation, and insurance. By directly supporting DOE’s goal of predicting U.S. water for energy, this project will help communities and industries build resilience to extreme storms while making smarter use of water and energy.

From Improved Coupled Ocean–Ice States to Improved Water Prediction: A Generative AI Approach

- **Lead PI:** Luke Van Roekel, Los Alamos National Laboratory
- **Collaborating Institutions:** Lawrence Livermore National Laboratory; Hewlett Packard Enterprise; University of Colorado–Boulder

AI-based forecasting systems have dramatically improved predictions of U.S. water availability on timescales from hours to weeks. Yet on subseasonal-to-interannual timescales, forecast skills remain limited due in part to errors in the initial state that grow rapidly. This shortfall matters to DOE: reliable weeks-to-years forecasts of regional water availability are needed to support energy

infrastructure resilience and energy systems under weather risk. Improving predictions requires advances in the initialization of coupled ocean and sea-ice states, where traditional data assimilation (DA) methods (e.g., Kalman Filter) are computationally expensive and struggle with extreme data sparsity. DA-based initial conditions often lead to strong model drifts, rapidly degrading predictions. This project proposes to develop a generative data assimilation (GenDA) method that pairs latent score-based diffusion models with graph neural networks (GNNs). This method will produce physically consistent initial conditions with quantified uncertainty for the coupled ocean and sea-ice system in DOE’s Energy Exascale Earth System Model (E3SM).

Unlike traditional DA, the proposed GenDA-GNN system naturally handles sparse observations and nonlinear structures. Training the diffusion process in latent space drastically reduces computational cost relative to other GenDA approaches. Coupling to a GNN backbone allows this system to operate directly on E3SM’s unstructured mesh, thereby preserving important multivariate covariances, preventing interpolation-induced information loss, and allowing for assimilation of multimodal observations without retraining the network. Performance will be evaluated through hindcast experiments using both E3SM and its emulator, targeting improvements in weeks-to-years prediction of U.S. water availability.

Phase I success criteria include (1) reducing ocean and sea-ice initialization cost by at least a factor of 2 relative to conventional spin-up approaches and (2) statistically significant improvements in forecast skill and reliability relative to baseline E3SM large ensembles. To ensure reusability and cross-domain impact within the Genesis Mission, all models, datasets, and workflows will be deployed on the American Science Cloud with integration into the Transformational AI Models Consortium infrastructure, leveraging established frameworks for scalable AI and high-performance computing orchestration. This will enable downstream projects to reuse trained models, observational pipelines, and initialization workflows without reimplementing. Improving coupled ocean and sea-ice initialization and quantifying forecast uncertainty will contribute to the Genesis Mission’s goal of advancing skillful water prediction on weeks-to-years timescales. The resulting capability will enable improved anticipation of hydroclimate variability and extremes, supporting DOE’s mission in energy security, resilience, and water and energy systems management. This project establishes the foundation for Phase II efforts extending to land processes, interpretable AI for predictability, and regional prediction relevant to energy.

Integrated AI-Driven Weeks-to-Years Prediction of Water for Energy in the Tennessee Valley Authority Region

- **Lead PI:** Dali Wang, Oak Ridge National Laboratory
- **Collaborating Institution:** Georgia Institute of Technology

This project will develop an AI-driven Earth system prediction framework for weeks-to-years forecasts of water availability, streamflow, and river temperature to support hydropower and

thermoelectric operations in the Tennessee Valley Authority (TVA) region. Reliable outlooks at these timescales help energy providers plan fuel use, schedule maintenance, manage reservoirs, and reduce outage and efficiency risks for hydropower and power plants that depend on cooling water.

Domestic energy providers need better foresight of hydrological conditions that affect operations and planning. Today, many decisions still rely on historical weather and water patterns that may not capture changing environmental conditions. This project aims to move from retrospective understanding of hydrological variability to operationally relevant forecasts at the watershed- and reservoir-scale, including the likelihood of high-risk conditions that matter for water allocation, thermal compliance, and energy resilience. This project will link Earth system observations, E3SM simulations, and stakeholder-defined decision metrics for TVA use cases. At the project's core is an advanced Earth system foundation model developed at Oak Ridge National Laboratory and designed for DOE leadership-class computing. The model will use an extended information context including large-scale environmental signals, local land-surface features, and longer-term system memory to predict energy-relevant water conditions at flexible lead times.

Rather than treating AI as a black box substitute for physical understanding, the team will use the model as a scientific tool to quantify how much predictability exists, where it comes from, and whether it is useful for energy decisions. The project will evaluate forecast skill and reliability against baselines and examine how predictability transfers between real-world observations and E3SM results. Success will be measured through improved forecast performance for key TVA-relevant water indicators, calibrated probabilistic outlooks, and clear evidence that the approach provides actionable water-for-energy risk information for decision-makers. If successful, this work will create a practical science-to-operations pathway in which an AI-driven Earth system prediction framework and high-performance computing strengthen utility-relevant adaptation under uncertain and shifting water supplies, while advancing environmental prediction science for broader community use.

GENESIS-AR: Generative Models and New Observations for Improved Subseasonal-to-Seasonal Prediction of Atmospheric Rivers

- **Lead PI:** Da Yang, Stanford University
- **Collaborating Institutions:** Pacific Northwest National Laboratory; University Corporation for Atmospheric Research

Atmospheric rivers (ARs) are long, narrow corridors of concentrated water vapor that produce much of the extreme precipitation and water-supply variability along the U.S. West Coast. They are essential to regional water resources, but the strongest events can cause flooding, damage energy infrastructure, disrupt hydropower operations, and contribute to wind-related power outages. Although current forecast models can often predict ARs several days in advance, reliable forecasts 2 to 6 weeks ahead remain

difficult. Better probabilistic forecasts at these lead times could help energy and water managers prepare for hydropower inflow variability and flood hazards.

GENESIS-AR (Generative Models and New Observations for Improved Subseasonal-to-Seasonal Prediction of Atmospheric Rivers) will develop and evaluate a physics-guided AI framework to improve subseasonal AR prediction. The project combines three complementary advances: (1) generative AI to correct forecast errors and represent uncertainty; (2) an efficient AI weather model informed by the physical energetics of moisture transport; and (3) a feasibility study of seismic noise as an additional observational source. AI is central to the proposed workflow because it can learn complex forecast-error distributions, generate large forecast ensembles at low computational cost, and identify patterns in large atmospheric and seismic datasets that are difficult to extract using conventional methods alone.

First, diffusion models will learn the full range of forecast errors—rather than only the average bias—to test whether they improve estimates of the probability of rare and extreme AR events and subseasonal AR forecasts. Second, the project will retrain the ACE2 (Ai2 Climate Emulator version 2) model to incorporate integrated vapor kinetic energy, a measure of the strength and energetics of atmospheric water-vapor transport. This retraining will test whether incorporating physical knowledge improves forecast skill, maintains reliable performance over extended forecasts, and clarifies the processes controlling AR predictability.

Third, the team will test whether microseisms—weak, continuous ground vibrations generated through interactions among the atmosphere, ocean, and solid earth—contain detectable signals of ARs. Phase I will use overlapping seismic observations and satellite-era AR records to evaluate this possibility. If AR signals are detected in microseisms, future work could apply the method to earlier seismic records to provide a longer observational context for rare events and long-term variabilities in the Earth system.

Phase I will develop and evaluate prototype methods, compare them with established forecasting and observational baselines, and quantify where AI provides improvement. Successful capabilities could support DOE's mission by improving information relevant to hydropower planning and energy infrastructure resilience along the U.S. West Coast.

An Agentic AI Framework for Seasonal-to-Interannual U.S. Water Prediction: Coupling Large AI-Based Prediction Models and Wavelet Diffusion with E3SM*

- **Lead PI:** Jinwoong Yoo, University of Maryland
- **Collaborating Institutions:** Oak Ridge National Laboratory; University of Florida; University of Colorado–Boulder; Silurian AI

Extreme weather events—including hurricanes, floods, droughts, and winter storms—increasingly threaten U.S. energy infrastructure, causing billions of dollars in damages and disrupting power

* BER is lead; co-funded with the DOE Advanced Scientific Computing Research (ASCR) program

generation, fuel refining, and water supply for energy production. Predicting water availability with sufficient accuracy and spatial detail months to years in advance remains a fundamental challenge, leaving energy providers unable to adequately plan for and mitigate these risks.

This project develops an agentic AI framework that couples large AI prediction models and wavelet diffusion techniques with E3SM to deliver faster, more detailed, and more reliable seasonal-to-interannual water and flood predictions for U.S. energy planning. This framework integrates five innovations: (1) improved model initialization by assimilating NASA satellite observations of soil moisture and groundwater into E3SM using ML, (2) a large generative AI prediction model that produces probabilistic seasonal forecasts more than 100 times faster than conventional simulations, (3) wavelet diffusion models that sharpen coarse predictions from approximately 100-km resolution down to 3 km and generate flood inundation maps at 1-m resolution, (4) a diagnostic system that verifies AI predictions are physically consistent and explainable, and (5) an autonomous AI agent system that orchestrates the entire prediction workflow with minimal human intervention.

If successful, this framework will transform seasonal water prediction from a slow, manual research exercise into a fast, autonomous capability that helps protect U.S. energy systems from extreme weather. Phase I focuses on two devastating U.S. hurricane events—Harvey (2017) and Helene (2024)—to demonstrate that AI can measurably improve seasonal water prediction accuracy, speed, resolution, and operational readiness. In Phase II, the framework would extend to full continental U.S. coverage with operational energy stakeholders; comprehensive teleconnection and explainability diagnostics; and production-grade autonomous prediction covering drought, winter storms, atmospheric rivers, and wildfires.

Drift-Aware Initialization of Coupled Earth System Models for Scalable Subseasonal-to-Seasonal Prediction Using Agentic AI*

- **Lead PI:** Shixuan Zhang, Pacific Northwest National Laboratory
- **Collaborating Institution:** Texas State University

Reliable prediction of water availability and hydrologic extremes at subseasonal-to-seasonal timescales is critical for U.S. energy systems, including hydropower operations, reservoir management, drought preparedness, and resilient energy-system planning. However, coupled Earth system forecasts can lose skill after initialization because the atmosphere, land, ocean, and sea-ice components may not be dynamically balanced. These imbalances can produce initialization-induced drift, accelerate early forecast error growth, and degrade forecast trajectories for water-cycle processes. Existing data assimilation, reanalysis-based initialization, and conventional drift correction reduce some errors, but they do not directly resolve state-dependent inconsistencies across the initialized coupled system. This leaves a gap between

realized and potential forecast skill for water-for-energy applications. This project will address this gap by developing a drift-aware, agentic AI framework within DOE's E3SM to improve the dynamical consistency of initialized coupled states.

The framework will use a two-stage learn-then-correct strategy. First, a neural operator will learn lead-time-dependent and state-dependent drift as a structured response of the coupled system. Second, a physics-constrained offline reinforcement learning method will generate bounded, physically consistent adjustments to selected initialization variables before forecast integration. Together, these components will enable an AI agent to assess the initialized coupled state, anticipate its likely drift, and select corrective actions under physical constraints. This shifts AI from passive forecast post-processing toward active, constraint-aware initialization control. Multiyear initialized E3SM hindcasts will be employed to (1) characterize drift across seasons and hydrologic regimes, (2) validate the AI framework with independent test data, and (3) deploy it in both standard E3SM and data-assimilation-enabled E3SM Atmosphere Model Data Assimilation Research Testbed (EAM&DART) workflows.

Phase I performance will be compared with uncorrected forecasts and a conventional linear drift-correction baseline to quantify AI advantage. Evaluation will focus on reduced early forecast error growth during weeks 1 and 3, improved forecast skill during weeks 4 and 12, statistically significant gains in anomaly correlation coefficient, reduced forecast error, stable coupled-model behavior, and improved cost and efficiency. Hydrologic evaluation will emphasize precipitation, snow water equivalent, soil moisture, runoff, total water storage anomaly, and river discharge in representative U.S. basins. This project's primary deliverable will be a validated, AI-enhanced E3SM initialization and hindcast workflow integrating the trained drift-prediction and correction framework with standard and EAM&DART-enabled coupled simulations. By reducing initialization-induced drift, this project advances the capability from Technology Readiness Level 3 toward a validated Technology Readiness Level 5 demonstration within E3SM.

WISE: Water Intelligence System for Energy Infrastructures

- **Lead PI:** Qing Zhu, Lawrence Berkeley National Laboratory
- **Collaborating Institutions:** University of California–Berkeley; Sandia National Laboratories; Pacific Northwest National Laboratory; Oak Ridge National Laboratory

As the United States transforms its energy landscape to meet increasing demands, the dependency of energy generation and operation on water resources is a critical vulnerability. In addition, rapidly expanding AI infrastructures (e.g., data centers) require massive, uninterrupted supplies of water for cooling and daily operations. However, water availability is constantly shifting, and existing computer models struggle to predict these changes

* BER is lead; co-funded with the DOE Advanced Scientific Computing Research (ASCR) program

on the weeks-to-years timescales that decision-makers need. To address this challenge, this project introduces the Water Intelligence System for Energy Infrastructures (WISE). WISE is an innovative, AI-powered system designed to forecast water availability and identify potential risks to vital energy infrastructure.

During Phase I, this project will develop the core AI prediction engine, using Texas as a proof-of-concept state. Texas is an ideal testbed because it is a major energy producer; relies heavily on water-intensive cooling technologies; and has faced persistent, long-term drought conditions. By focusing on Texas, the project will demonstrate how advanced AI technology can protect vital infrastructure under real-world environmental stress. This

project will develop a physics-informed hybrid AI approach that ensures the AI's predictions are not only highly accurate but also physically realistic, even when forecasting unprecedented weather events.

Ultimately, this Phase I effort will produce detailed forecasts of water balance (e.g., rainfall, runoff, and soil moisture) across Texas from 2027 through 2031. By overlaying these predictions onto maps of existing power plants, WISE will help resource managers pinpoint potential water-stress hotspots before they cause operational failures. This AI tool will empower stakeholders to make informed planning decisions and build a more resilient energy future.

Securing America's Critical Minerals Supply Selections

America's dependence on foreign supply chains for critical minerals and materials (CMM) threatens national security, economic competitiveness, and the deployment of technologies essential for energy independence. Domestic critical mineral production is expensive, complex, and time-consuming—involving numerous steps to identify, extract, refine, and concentrate from complex heterogeneous sources across critical mineral supply chains.

AI will revolutionize the entire critical minerals supply chain and development of alternative materials by integrating geophysical data, other fundamental science data, process optimization, cost estimation, and economic modeling into one connected system. Solving this challenge demands AI that can reason scientifically, understand complex structure–property relationships, and design alternatives with different compositions. Physics-based AI offers advanced predictive capabilities to identify alternatives and understand processes underlying critical mineral availability, recovery, refinement, and replacement.

Under this S&T challenge, BER is involved in two focus areas:

- Geological Finders/Keepers
- Biological Pathways to CMM

Geological Finders/Keepers

AI-Guided Map for Biomining of Critical Minerals (AIM-BioCM): Integrating Microbial, Hydrological, and Geochemical Data

- **Lead PI:** Kung-Hui Chu, Texas A&M Engineering Experiment Station
- **Collaborating Institutions:** U.S. Geological Survey; Lawrence Berkeley National Laboratory; Emory University; Montana State University; University of Montana

Critical minerals (CM) are vital to U.S. technology, energy, and security, yet the United States heavily relies on foreign imports. At the same time, millions of tons of mining waste across the country contain valuable CM once considered uneconomical to recover. Biomining is increasingly recognized as a promising and innovative option for their extraction. By tapping into the mineral potential locked within mine waste, the United States can strengthen mineral supply chains and reduce reliance on foreign producers while reducing the footprint of mining.

This project's vision is to advance U.S. energy and technology security by establishing a domestic supply of CM through AI-guided biomining of mine waste, minimizing impact and reliance on foreign imports. This project aims to develop a way to recover CM from abandoned mine waste using naturally occurring microorganisms. The overall goal of this project is to develop an AI-powered predictive framework [i.e., AI-Guided Map for Biomining of Critical Minerals (AIM-BioCM)] that integrates microbial, hydrogeologic, and geochemical data from unconventional environments to guide targeted extraction of rare earth elements (REEs) through biomining.

Leveraging AI advantages, AIM-BioCM will (1) produce the first genome-resolved functional map of CM-accumulating microbial communities in the Berkeley Pit and Clark Fork River, (2) predict CM-rich hotspots, and (3) deliver a numerical score to assess the feasibility of biomining CM-containing resources in mining-impacted areas. The developed decision tool is expected to be applicable to hundreds of acid mine drainage sites across the country. By turning mining waste into a valuable domestic resource, this research could strengthen U.S. supply of critical minerals, reduce dependence on foreign imports, lessen the impacts of mining, and support the nation's energy security and economic competitiveness.

AI-Driven Probabilistic Source Mapping for Critical Minerals at Continental Scale

- **Lead PI:** David Costello, Kent State University
- **Collaborating Institutions:** U.S. Geological Survey Upper Midwest Water Science Center; Lawrence Berkeley National Laboratory; Virginia Polytechnic Institute and State University

The United States depends on a reliable supply of critical minerals for energy deployment, advanced manufacturing, and national security. Yet domestic discovery is not keeping pace with demand, and conventional exploration methods are too costly to scale across the continent. At the same time, federal and state agencies have already built a nationwide stream-sampling network that records chemical signals from upstream geology. This dataset represents an untapped federal asset for identifying potential mineral resources.

This project will convert that existing water chemistry record into a continental-scale critical mineral discovery tool aligned

with federal directives to inventory mineral deposits on public lands. The approach recognizes that water chemistry in streams reflects both geological sources and biological processes. A sample from a river or stream is a composite of the geologic signal from the upstream watershed. Microbial communities in streams can take up trace elements, which can average out noisy water chemistry through time but alters the concentrations measured downstream for any single water sample. This project aims to understand the geologic, hydrologic, and biological effects on trace metals within the water quality record for accurate mineral assessment.

The research team will develop an AI-enabled framework that integrates hydrologic transport, geochemical behavior, and biological modification of streamwater chemistry. Modern ML methods are required because the dataset spans hundreds of thousands of sampling locations within complex river-network connections, and samples are collected over different seasons and flow conditions. The resulting models will learn how chemical signals move through watersheds and how biological activity modifies them, producing reach-level predictions across the continental United States.

Phase I will deliver a harmonized national dataset; continental-scale prediction layers; models of biological modification; and an initial prospectivity tool for copper, cobalt, and manganese. These products will be designed to integrate with DOE's emerging data platforms and evaluated against known mineral deposits using established performance metrics.

Bringing together expertise from DOE, the U.S. Geological Survey, and universities, this project leverages complementary federal datasets. By transforming routine water-quality monitoring into a strategic mineral-assessment capability, this work will help expand the nation's domestic mineral resource base and inform future exploration decisions.

GEM-AI: Generative Exploration of Minerals Using AI

- **Lead PI:** Artur Dubrawski, Carnegie Mellon University
- **Collaborating Institutions:** Colorado School of Mines; Sandia National Laboratories

Securing a reliable domestic supply of critical minerals is vital for modern technology and national prosperity. Traditional mineral exploration relies on computer models that treat all data equally, leading to overly optimistic predictions in regions where data is sparse or outdated. This project addresses this challenge by developing an adaptable AI framework that knows how reliable its data is at every location. Instead of making uncertain assessments in data-poor areas, Generative Exploration of Minerals using AI (GEM-AI) tracks data quality and coverage to focus its attention on where the data is strongest.

Initial application of this approach focuses on demonstrating its utility at two environmental frontiers to locate copper sulfide deposits across the Pacific Northwest and the Arizona copper belt.

- **Soil microbiology.** This frontier will use large soil metagenome datasets from DOE to investigate how underground mineral deposits alter the microbial communities living in the soil above them to enhance identification of biological fingerprints of hidden minerals.

- **Space-based mapping.** This frontier will use advanced NASA hyperspectral imaging data originally designed to study dust in the atmosphere to spot subtle mineral changes on Earth's surface, even in areas heavily covered by vegetation.

This project brings together geological, satellite, and biological data into a single, cohesive map. By showing geoscientists and explorers where more data is needed to improve the reliability of prospective assessments, this framework actively reduces both the footprint and the financial risk of mineral exploration. The resulting analytic toolkit will run on a federally funded supercomputer, providing the broader scientific community with affordable access to next-generation AI-driven exploration technology.

Physics-Informed AI for Accelerated Discovery of Ion-Adsorption-Type Rare Earth Element Resources in Granitic Systems

- **Lead PI:** Xin Gu, Oak Ridge National Laboratory
- **Collaborating Institutions:** University of Maryland; The Pennsylvania State University; Syracuse University

Rare earth elements (REEs), especially heavy REEs, are essential for energy technologies, advanced manufacturing, and national supply chain security. More than 95% of current heavy REE production comes from ion-adsorption-type deposits formed in weathered granitic regolith in East and Southeast Asia. The southeastern United States shares similar humid subtropical conditions and widespread granitic bedrock, suggesting significant but largely unquantified potential for analogous resources. This project will develop a physics-informed AI framework for accelerated discovery of ion-adsorption-type REE resources in granitic systems through parameter inference, reactive transport model (RTM) emulation, and predictive scaling. By combining mechanistic geoscience with AI-enabled analysis, the vision is to move beyond descriptive case studies toward predictive, uncertainty-aware discovery of REE regolith.

AI is needed because REE enrichment emerges from coupled processes across scales, including mineral weathering, adsorption to secondary minerals, hydrologic transport, and regolith evolution. RTMs provide a mechanistic framework for these processes, but they are difficult to calibrate because adsorption behavior, source terms, and transport controls are only partially constrained by limited field data. Purely data-driven approaches face the complementary problem that observations are sparse, heterogeneous, and weakly connected to the governing process structure. This project addresses that gap by integrating field and laboratory observations with RTM-generated synthetic data in a unified workflow. Phase I will focus on three granitic watersheds in North Carolina and will combine depth-resolved regolith

characterization, adsorption experiments, reduced-complexity RTMs, and AI models for parameter inference, surrogate modeling, and prediction of REE enrichment and regolith structure.

The Phase II-scale breakthrough will be an interpretable, uncertainty-aware discovery platform that can predict REE enrichment potential, enrichment depth, and regolith structure across southeastern U.S. granitic terrains. Such a capability would advance the field by linking mineral-scale reactions and watershed-scale transport to regional prospectivity in a single predictive framework and would establish a generalizable approach for combining physics-based modeling and AI to discover critical mineral resources in complex subsurface systems.

A key 6-month deliverable is demonstration of an initial AI advantage at the watershed scale. By Month 6, the project will deliver (1) quantified pH-dependent adsorption behavior for neodymium and dysprosium, including dissolved organic matter and competing-cation effects; (2) depth-resolved characterization of REE enrichment patterns and major REE-bearing and adsorption-host phases; (3) calibrated 1D RTMs that reproduce first-order observed enrichment behavior; and (4) an initial AI model for regolith-thickness prediction or RTM parameter inference, with predictive skill exceeding simple empirical or geostatistical baselines. These Phase I outcomes will establish a credible basis for regional scaling in Phase II.

AI-Guided Discovery of Dissolution Regimes and Optimal Conditions for *In Situ* Critical Mineral Extraction

- **Lead PI:** Peter Kang, University of Minnesota
- **Collaborating Institutions:** Los Alamos National Laboratory; Colorado School of Mines

Securing domestic supplies of critical minerals and materials (CMMs) is a national priority for energy resilience and supply chain security. This project develops a transformative, AI-enabled framework for predicting and optimizing CMM extraction from fractured rock formations and mine tailings, where fluid flow, interfacial processes, and mineral–fluid geochemistry interact nonlinearly across scales to govern extraction efficiency. Rather than attempting exact spatiotemporal prediction of these inherently complex systems, this project introduces a regime-based framework that identifies predictable classes of behavior—preferential bypass, premature reprecipitation, and uniform reactive access—each leaving characteristic signatures in produced-water chemistry and flow rates. These observable signatures serve as the primary data streams for AI-enabled inference, optimization, and operational control of extraction conditions.

AI is essential for this problem because the governing parameter space—including fracture network structure, pore-scale architecture, fluid chemistry, flow conditions, and mineralogy—is

high-dimensional and nonlinearly coupled, making exhaustive experimental or computational exploration intractable. The proposed framework integrates four mutually reinforcing components: (1) high-throughput microfluidic and rock-on-a-chip experiments providing spatially resolved geochemical and physical ground truth across systematic ranges of mineralogy, fracture geometry, and flow conditions; (2) a multifidelity modeling framework spanning pore-network models and field-scale discrete fracture network (DFN) simulations and using dfnWorks coupled with PFLOTRAN for reactive transport, capturing dissolution, precipitation, and aperture evolution along preferential flow paths; (3) AI surrogate models based on neural operators and latent-space representations that predict the temporal evolution of reactive transport fields and classify dissolution regimes from observable inputs; and (4) a closed-loop active learning framework employing Bayesian uncertainty quantification and multi-objective optimization to adaptively allocate experimental and computational resources toward maximizing mineral recovery while avoiding clogging. Together, these components form an AI-guided workflow linking pore-scale experiments with field-scale prediction.

Phase I will demonstrate the value of this AI-enabled framework by showing it can identify dissolution regimes and estimate CMM recovery more accurately than conventional approaches while substantially improving data efficiency through active learning and enabling reliable knowledge transfer from pore-scale experiments to field-scale systems. These advances will establish the foundation for a predictive, data-efficient, and scalable framework for optimizing critical mineral extraction.

AI-Enabled Geobiological Critical Minerals Characterization and Recovery in Geothermal Systems*

- **Lead PI:** Karen Lloyd, University of Southern California
- **Collaborating Institutions:** University of Tennessee; Woods Hole Oceanographic Institution; Georgia Institute of Technology; University of North Carolina–Chapel Hill; Oak Ridge National Laboratory

Critical minerals and materials (CMMs) are essential to U.S. energy technologies and may occur in economically meaningful concentrations in geothermal and other subsurface fluid systems. However, these systems remain undercharacterized for CMM recovery due to fragmented exploration approaches. This project will develop an interpretable, AI-enabled framework to characterize and prioritize geothermal systems for CMM recovery by integrating geophysical, geochemical, hydrologic, biological, and remote sensing signals across scales ranging from large (e.g., regional fluid flow and geologic structure) to small (e.g., microbial activity and biomineral interactions).

Conventional workflows typically analyze geophysical, geochemical, and microbial datasets independently, limiting the ability to detect cross-scale relationships that govern CMM mobilization, concentration, and recoverability. AI-enabled exploration will

* BER is lead; co-funded with the DOE Basic Energy Sciences (BES) program

integrate these multimodal observations into interpretable models, enabling a consistent ranking of geothermal systems by recovery potential and reducing uncertainty in Phase II site selection.

Proposed Research

This Phase I project tests the hypothesis that geothermal fluids preserve coupled signatures of fluid origin, transport, redox conditions, and biological activity that indicate where CMMs are mobilized, concentrated, and recoverable. Researchers will analyze existing datasets from four tectonically distinct regions—the Yellowstone hotspot track, Central American subduction zone, Hikurangi subduction zone (New Zealand), and Chile subduction zone—to identify transferable indicators of high-priority systems.

Magnetotelluric geophysical data constraining subsurface fluid distribution will be integrated with methane, oxygen, metals, sulfide, and noble gas (helium and neon) isotopes that trace fluid sources, connectivity, and energy gradients. These data will be combined with microbial measurements, including multi-omics, single-cell activity assays (e.g., BONCAT), and isolate-based validation with a known metal-concentrating bacterium, *Exiguobacterium* sp. Microbial features will indicate *in situ* processes for metal mobility and concentration. A key advantage of this approach is that the team has complete multimodal datasets in hand for 10 to 30 thermal springs distributed across each of the four testbeds. This collection offers more interpretation power than single-point measurements at each location. The observed landscape-scale gradients will enable AI-driven ranking of the thermal springs with the highest potential for CMM recoverability in each of the four tectonic regions.

Specifically, this project will construct a harmonized, AI-ready dataset linking these variables to CMM occurrence and develop interpretable multimodal AI models to identify predictive relationships. Models will be benchmarked against conventional single-modality and PI-driven screening approaches to quantify the AI advantage. Phase I will specifically determine which combinations of data types and model architectures are required to produce a measurable advantage before scaling to Phase II field campaigns.

Outcomes and Impact

Phase I will deliver a curated multimodal dataset, an interpretable AI workflow, and a ranked set of geothermal systems prioritized for CMM recovery potential, with uncertainty bounds. These outputs will provide the framework for Phase II field campaigns and recovery-oriented modeling. By replacing fragmented, single-discipline, single-sample screening with an integrated, AI-enabled approach, this project advances DOE priorities in energy security and scientific discovery. The resulting framework will accelerate identification of concealed CMM resources, reduce exploration risk, and establish a scalable strategy for linking subsurface processes to resource recovery in geothermal systems.

* *The DOE Basic Energy Sciences (BES) program is lead; co-funded with BER*

AI-Driven Multiscale Framework for Critical Mineral Discovery and Recovery from Subsurface and Unconventional Sources*

- **Lead PI:** Ghanashyam Hari Neupane, Idaho National Laboratory
- **Collaborating Institutions:** Boise State University; Mississippi State University; Idaho State University

Critical minerals such as cobalt, rare earth elements, and thorium are essential ingredients in electric vehicles, energy systems, and modern electronics. The United States currently relies heavily on foreign sources for these materials, creating vulnerabilities in U.S. energy and technology supply chains. Finding and recovering more domestic sources of these minerals is a national priority, but doing so is difficult: deposits are often hidden deep underground, spread across large and rugged landscapes, or locked up in overlooked sources, like old mine waste piles or groundwater.

This project builds GeoFinderAI, a smart, AI-powered tool that acts like a highly trained geological detective. Instead of relying on any single clue, it combines many different types of information at once—satellite images of the ground surface, measurements of underground rock structures, water and soil chemistry, and even the DNA of microbes living in the soil—to build a complete picture of where critical minerals may be hiding and how they might be recovered. The AI learns the patterns that connect surface clues to what lies underground, allowing it to make predictions far faster and across far larger areas than traditional methods.

The team will test GeoFinderAI across two well-known mineral-rich areas in Idaho: the Lemhi Pass belt, rich in rare earth elements and thorium, and the Blackbird belt, known for cobalt deposits. Over 9 months, the system will be built, trained, and rigorously tested to demonstrate that it can outperform conventional approaches.

All tools and data produced by this project will be made freely available to the broader scientific and mining community, supporting a stronger and more self-sufficient domestic critical minerals industry.

A Novel AI Workflow for Critical Minerals in Heterogeneous Geosystems

- **Lead PI:** Carl Steefel, Lawrence Berkeley National Laboratory
- **Collaborating Institutions:** Purdue University; Stanford University; Montana Technological University; Group Ten; Spruce Oregon Holdings, LLC

This project develops a data-agnostic AI framework to accelerate predictive understanding of the diffusion, transport, and reaction processes controlling critical mineral discovery and recovery from complex, heterogeneous systems, integrating multimodal mineralogical, geochemical, and geophysical data.

Conventional mining relies on bulk excavation with heavy waste and energy costs. *In situ* leaching (ISL) offers a selective, low-disturbance alternative but lacks robust models for realistic geochemical and physical heterogeneity. This project pairs DOE's CrunchFlow reactive transport software with modern AI to close this gap, applying the approach to three deposit types: platinum group element (PGE)-bearing ultramafic rock (Stillwater, Mont.), nickel laterite (Red Flat, Ore.), and supergene copper sulfide (Butte, Mont.).

Project work spans a curated kinetics library, AI-assisted multi-modal core curation, transformer-based masked autoencoders for heterogeneity mapping, AI-orchestrated automated reactive transport modeling, and a transformer-based digital twin for closed-loop ISL prediction, with Phase I proof-of-concept at the Butte site (architecture designed to extend to electrokinetic ISL in Phase II).

Success metrics include order-of-magnitude speedups in kinetics assembly, core interrogation, and heterogeneity prediction, plus a validated digital twin demonstration.

Phase I establishes the scientific foundation for Phase II expansion across all three sites with greater complexity and heterogeneity.

Biological Pathways to Critical Minerals and Materials

BIND: Biophysics-Informed Learning of Coordination for Metalloprotein Design

- **Lead PI:** Yongqin Jiao, Lawrence Livermore National Laboratory
- **Collaborating Institutions:** Brookhaven National Laboratory; Los Alamos National Laboratory; University of Washington; Columbia University

BIND will develop a new physics-informed AI framework to enable the predictive design of metal-binding proteins for applications in critical mineral recovery, radionuclide management, and environmental monitoring. Metal-binding proteins exhibit high selectivity for specific metal ions, but current engineering approaches rely largely on trial and error and must be repeated for each new metal or application. BIND addresses this challenge by integrating quantum chemistry, structural biology, experimental measurements, and ML to uncover the fundamental coordination chemistry principles that govern how proteins recognize and bind metals. Rather than treating each metal as an independent design problem, the framework will learn transferable design rules that can be applied across related elements and chemical environments. Phase I will focus on rare earth elements as a model system, combining computational modeling with high-throughput experimental validation in a closed-loop design-build-test-learn workflow to rapidly evaluate and improve AI-generated protein designs. By integrating molecular-scale physics with modern AI methods, BIND aims to

significantly reduce the experimental effort required to discover high-performance metal-binding proteins while improving predictive accuracy and generalizability. If successful, the project will establish a broadly applicable framework for AI-enabled metalloprotein engineering, accelerating the development of bio-based technologies for critical mineral recovery, environmental remediation, and advanced manufacturing while advancing DOE priorities in predictive biosystems design, AI-enabled scientific discovery, and fundamental understanding of structure–function relationships in biological systems.

Ultra-High-Throughput Characterization and Predictive Design of CMM Binding Proteins

- **Lead PI:** Benjamin Keitz, University of Texas–Austin
- **Collaborating Institution:** Sandia National Laboratories

Rare earth elements (REEs) and other critical minerals and materials (CMMs) underpin modern energy technologies, but their extraction, purification, and recovery are challenging. Naturally occurring REE-binding proteins, such as lanmodulin (LanM), have demonstrated that biological systems can achieve selectivity among elements with similar ionic radii, but the molecular rules governing this discrimination are poorly understood and the sequence space of functional REE-binding proteins remains largely unexplored. Rational and ML-guided design of high-performance REE-binding proteins has been constrained by laborious protein isolation procedures, limited structural data, the dynamic nature of metal binding in these proteins, and the inability to efficiently traverse the enormous combinatorial sequence space of functional variants. More importantly, current protein structural and language models trained on evolutionary sequence data encode fitness in a natural biological context, not binding selectivity for REE-relevant coordination environments that have no evolutionary precedent in most organisms. This disconnect between evolutionary fitness landscapes (genotype) and application-relevant phenotypes is a fundamental limitation of ML sequence- and structure-based predictions for REE-binding performance and a key barrier to AI advances in this space.

The proposed work directly addresses these challenges by developing an integrated experimental–computational platform that couples ultra-high-throughput protein variant screening with task-specific ML model fine-tuning. The team previously developed a yeast surface display system in which libraries of LanM protein variants are expressed on the cell surface and sorted by fluorescence-activated cell sorting in the presence of REEs. Following the sort, both responsive and unresponsive variants are characterized via next-generation sequencing. The platform enables large-scale, quantitative assessment of REE binding across millions of variants simultaneously and generates large-scale datasets in which variant sequences are annotated with enrichment scores reflecting their binding affinity to specific target elements. Critically, by performing both positive selections in the presence of a target REE and negative counter-selections in the presence of competing elements, the experiments generate labeled datasets that encode both binding affinity and selectivity.

Together, the experimental platform lays the foundation for ML advances in designing new REE- and CMM-binding proteins.

This project will leverage empirical datasets and fine-tune state-of-the-art protein language models (e.g., Evolutionary Scale Model) to predict REE binding directly from sequence and nominate new variants with improved selectivity. Further, model precision will be increased by integrating structural priors via physics-aware deep learning from structural ensemble computer simulations. Enforcing these physical constraints will allow essential cooperative interactions and coordination geometry to assist in nominating high-selectivity variants for further validation. Finally, the project will validate ML-nominated variants against complex REE sources including primary ores, recycling streams, and low-concentration environmental matrices. Feeding these data into models will allow for the prediction of performance for complex feedstocks.

HARVEST-CMM: Hybrid AI-Driven Chelator–Protein Codesign for Robust Extraction and Selective Separation of Critical Mineral Materials

- **Lead PI:** Justin Wilson, University of California–Santa Barbara
- **Collaborating Institution:** Los Alamos National Laboratory

AI is creating new opportunities to accelerate scientific discovery by enabling the rapid design and optimization of complex molecular systems. This project will develop an AI-enabled framework for the codesign of synthetic metal-binding molecules (chelators) and engineered proteins to address one of the fundamental challenges in chemistry: the selective separation of chemically similar metal ions. Critical minerals, including rare earth elements, nickel, cobalt, and related metals, are essential for advanced energy technologies, electronics, and national security. However, many of these elements exhibit nearly identical chemical properties, making their separation from natural resources, recycled materials, and industrial waste streams one of the most difficult problems in chemical processing. Existing separation

technologies often require large numbers of sequential processing steps, consume substantial amounts of energy and chemicals, and provide only modest selectivity for closely related metals. This project introduces a fundamentally new molecular recognition strategy that combines two complementary mechanisms of selectivity. First, AI-guided chelator design will produce molecules that selectively bind target metal ions and undergo metal-dependent structural changes. Second, AI-guided protein design will generate engineered proteins that selectively recognize the resulting metal–chelator complexes. Together, these complementary recognition processes are expected to provide substantially greater molecular selectivity than either synthetic chelators or engineered proteins can achieve independently.

This project will integrate ML, quantum chemical calculations, inorganic chemistry, and *de novo* protein design within an iterative design-build-test-learn workflow. AI models will predict promising chelator structures, identify complementary protein recognition interfaces, and prioritize candidates for experimental evaluation. Newly designed chelators and proteins will be synthesized, expressed, and characterized using structural, spectroscopic, and analytical methods to measure molecular recognition, metal-binding strength, selectivity, and separation performance. Experimental results will be incorporated back into the AI models to continuously improve predictive accuracy while reducing the number of molecules that must be synthesized and tested experimentally. The project will establish fundamental design principles for integrating coordination chemistry with biomolecular recognition and will demonstrate how AI can accelerate molecular discovery by coupling computation with targeted experimental validation. Beyond critical mineral separations, the resulting AI-enabled design framework may be broadly applicable to other challenges involving molecular recognition, including catalyst discovery, sensing, environmental remediation, and advanced materials development. By advancing both AI-driven molecular design and fundamental understanding of selective chemical recognition, this research will lay the foundation for next-generation separation technologies that support secure domestic supplies of energy-relevant materials.

Unleashing Subsurface Strategic Energy Assets Selections

Delivering cost-effective energy from the Earth's subsurface entails the use of heterogeneous reservoirs dominated by fractures. Tools capable of predicting reservoir behavior and the interactions of complex biogeochemical, mechanical, and hydrologic processes are critical to developing innovative, cost-effective extraction of subsurface energy sources, including unconventional oil and gas, geothermal resources, and coalbed methane.

Developing AI capable of reasoning under extreme uncertainty, integrating heterogeneous data types (e.g., seismic, geochemical, biological, and hydrologic), and building predictive models of systems that cannot be directly observed has broad applicability to any domain requiring inference from indirect evidence. For subsurface science, AI that connects molecular-scale mechanisms to field-scale resource availability will transform the field from descriptive to predictive. DOE's vast biological, geochemical, geomechanical, and hydrologic datasets can be combined with high-performance modeling capabilities, laboratory analytics, and a suite of AI technologies (e.g., surrogate models, physics-informed machine learning, and digital twins). This will enhance engineering evaluation and subsurface control during characterization, drilling, stimulation, and production.

Under this S&T challenge, BER is involved in one focus area: Chemical and Hydrologic Transport in the Subsurface.

Chemical and Hydrologic Transport in the Subsurface

MIRACLE-AI: Molecular-to-Reservoir Integrated AI for Coupled Landscape and Earth-System Subsurface Intelligence

- **Lead PI:** Umakant Mishra, Sandia National Laboratories
- **Collaborating Institution:** Lawrence Berkeley National Laboratory

Developing geothermal energy and other subsurface energy resources depends on accurately understanding how water, heat, chemicals, minerals, microbes, and plant roots interact below ground. These processes influence fluid movement, subsurface chemistry, and the long-term performance of geothermal reservoirs. However, current modeling approaches require a trade-off between scientific detail and computational efficiency. Detailed mechanistic models can represent complex biological, chemical, and physical processes but are too computationally demanding for routine forecasting and scenario analysis. Faster, large-scale models are more practical but often omit small-scale interactions near the land surface that can affect deeper subsurface conditions.

The objective of this Phase I project is to develop and demonstrate a fast, scientifically grounded, AI-enabled framework for representing complex shallow-subsurface processes in models used to support geothermal energy forecasting. The project will focus on the rhizosphere, the narrow region surrounding plant roots where water, roots, microbes, minerals, and organic compounds interact. Although this zone is located near the surface, it can influence water movement and chemical conditions that

propagate into deeper subsurface systems. Research will evaluate whether these complex interactions can be represented using a reduced-order AI model that retains physical realism while substantially improving computational efficiency.

This project will integrate existing observations of soils, plant traits, roots, microbes, minerals, and environmental conditions from DOE-supported and publicly available databases with outputs from mechanistic subsurface simulations. Generative AI methods, including diffusion and flow-matching approaches, will be used to create physically plausible synthetic data where observations are limited or spatially imbalanced. These data will support development of a supervised neural-network surrogate capable of rapidly reproducing key outputs from detailed subsurface models. An agentic AI system will guide model development by identifying uncertainty, selecting additional simulations, comparing alternative model structures, and screening predictions for violations of known physical and scientific constraints. The resulting reduced-order model will then be incorporated into an existing large-scale land model to test whether improved representations of root-microbe-mineral interactions enhance prediction accuracy while reducing computational demand.

Project success will be evaluated using quantitative performance targets. The AI surrogate will aim to reproduce mechanistic-model outputs with a coefficient of determination greater than 0.90, reduce prediction errors by at least 30%, and operate at least 10 times faster than current process-based simulations. It will also provide uncertainty estimates and generate predictions that remain consistent with established physical principles. Successful completion of this project will establish the technical foundation for a larger Phase II effort addressing broader geothermal and subsurface energy applications. The anticipated outcomes include faster and more reliable subsurface

forecasting tools, improved representation of biologically and chemically mediated processes, and greater capacity to evaluate alternative resource-development scenarios. By combining AI, mechanistic science, and existing observational resources, this project will support more informed and trustworthy decisions related to geothermal exploration, reservoir management, and long-term energy recovery. Project datasets, models, and workflows will be made openly available through relevant Advanced Scientific Computing Research, modeling, and computational science initiative platforms, supporting reuse and broader application across the Genesis Mission research community.

An AI-Enabled Digital Twin Framework for Coupled Surface–Subsurface Hydrologic and Biogeochemical Simulations in the Northeast U.S. Coastal Region

- **Lead PI:** Wei Ren, University of Connecticut
- **Collaborating Institution:** Pacific Northwest National Laboratory

Water, carbon, nutrients, and dissolved chemicals move between soils, groundwater, rivers, and coastal waters, yet these movements are hard to predict because many of the most important processes happen underground and cannot be seen directly. This project will develop an AI-enabled digital twin of the Northeast U.S. coastal region to better understand and predict how water and chemicals move through connected surface and underground systems. The project will combine field measurements, satellite observations, maps of soils and geology, and computer model results from 2002 to 2025 into a consistent regional dataset.

The project will use advanced computer models, including PFLOTRAN and the Advanced Terrestrial Simulator (ATS), and AI to simulate groundwater flow, chemical transport, and interactions among water, carbon, nitrogen, and related environmental conditions. AI will help make these simulations much faster, fill in gaps where observations are limited, update predictions as new datasets become available, and estimate uncertainty. The main products will be a publicly usable dataset and a prototype AI-assisted digital twin that supports improved prediction for groundwater management, coastal water quality, elemental cycle studies, and subsurface energy applications such as biogeochemical storage and geothermal energy, while strengthening predictive Earth system science relevant to DOE missions.

Transforming Subsurface Prediction with an Autonomous, Multi-Agent AI Workflow for Multiscale Coupled Hydro-Biogeochemical Modeling

- **Lead PI:** Scott Saleska, University of Arizona
- **Collaborating Institution:** Lawrence Berkeley National Laboratory

The goal of this project is to use AI to improve predictions of how water, nutrients, minerals, plants, and microorganisms interact belowground. These interactions are important because

they control soil formation, nutrient availability, organic matter decomposition, groundwater chemistry, and the movement of chemicals through soils and aquifers. Better understanding of these processes is needed for characterizing and predicting subsurface conditions in U.S. landscapes relevant to DOE missions.

This goal will be achieved by integrating two advanced computer models that represent different scales: the EcoSIM model and the PFLOTRAN model. EcoSIM represents whole-ecosystem processes such as plant growth, microbial activity, decomposition, and water and carbon cycling. PFLOTRAN represents smaller-scale chemical reactions and the movement of water and dissolved chemicals through soils and groundwater. Combining these models will allow researchers to represent complex interactions among plants, microbes, minerals, water, and subsurface chemistry. For example, mineral weathering releases nutrients needed by plants, while plants and microbes alter water flow, carbon cycling, and chemical conditions that then influence further weathering and soil development.

Directly coupling these models is computationally challenging because they represent processes occurring over very different spatial and temporal scales. To overcome this limitation, the project will use AI tools to develop a new approach called surrogate-mediated coupling, or SMC. In this approach, a physics-informed neural network is trained using PFLOTRAN simulations to create a faster approximation (called a surrogate) of the full reactive PFLOTRAN model. Once created, the surrogate should greatly reduce computational costs of model simulations while still preserving physical laws such as mass and energy conservation. This AI-based approach thus brings efficiencies that should make it possible, for the first time, to connect detailed subsurface chemistry with ecosystem-scale processes that include capturing all the key complex interactions among plants, microbes, soil minerals, and hydrology.

AI enhances this project in three key ways. First, the research team will use the AI technique of physics-informed ML to create the surrogate model of PFLOTRAN. The surrogate will be trained on simulations of the full model in a way that is also constrained by physical laws. This enhancement can achieve orders-of-magnitude greater computational speed while maintaining physical consistency. Second, AI agents will assist in model ensemble generation and development of the combined EcoSIM–PFLOTRAN, accelerating otherwise time-intensive configuration and integration tasks. Third, a framework called agentic adaptive multitarget calibration (A2MC) will automate important aspects of the process, including calibrating models to different field settings, designing model experiments, executing model runs, and testing hypotheses. Together, these capabilities will allow the modeling system to use multiple types of observations and adapt efficiently to different sites.

Phase I will deliver proof-of-concept at the Landscape Evolution Observatory (LEO) at the University of Arizona's Biosphere 2, a uniquely instrumented hillslope system providing the multiple observations required to rigorously constrain the biogeochemical dynamics of the coupled model. This phase will deliver and

evaluate a functional EcoSIM–PFLOTTRAN SMC framework and A2MC calibration workflow. The goal is to achieve measurable improvements in prediction of subsurface fluxes and states relative to uncoupled baselines.

Anticipated future efforts will extend the validated framework to DOE Watershed Function Science Focus Area sites and other energy-relevant subsurface systems. The goal is to advance a multi-agent AI framework toward an autonomous, self-improving system that would continuously assimilate new observations to refine predictions of subsurface structure, flow, and reactive transport at regional scales across the conterminous United States.

AI-Accelerated Discovery of Multiscale Subsurface Architecture and Coupled Hydro-Biogeochemical Processes

- **Lead PI:** Mohamad Soltanian, University of Cincinnati
- **Collaborating Institutions:** New Mexico State University; Pacific Northwest National Laboratory; Certerra Subsurface Imaging

Solving many of the nation's energy challenges depends on understanding what lies beneath the Earth's surface. Whether recovering critical minerals needed for advanced technologies, developing geothermal energy, protecting groundwater resources, or managing subsurface energy systems, accurate prediction of how water, chemicals, heat, and other fluids move through complex underground formations is essential for designing, operating, and managing these systems. Because the subsurface cannot be observed directly, these predictions are often highly uncertain. This project will develop an AI-enabled framework that uses limited field observations and physics-based models to improve characterization of the subsurface structure. Rather than relying on massive preexisting datasets, the new approach combines geological information, geophysical measurements, monitoring data, and advanced computational modeling, allowing the framework to learn directly from site-specific data. This capability will improve predictions in data-limited environments, which are common across energy applications.

This project will bring together researchers from the University of Cincinnati, Pacific Northwest National Laboratory, New Mexico State University, and Certerra Subsurface Imaging to develop and test this new AI framework. The team will demonstrate how AI can improve understanding of subsurface structure, reduce uncertainty in predictions, and support more informed decision-making for complex underground systems. The technologies developed through this research will support a wide range of DOE priorities, including critical mineral exploration and recovery, geothermal energy, geological carbon storage, groundwater protection, and other subsurface energy applications. By combining AI with physics-based scientific models and field observations, the project will help accelerate scientific discovery while improving the nation's ability to manage energy and environmental resources.

AI-Enabled Subsurface Biogeochemical Modeling*

- **Lead PI:** Benjamin Sulman, Oak Ridge National Laboratory
- **Collaborating Institutions:** University of California–Riverside; Auburn University; Lawrence Berkeley National Laboratory

This project plans to unlock an AI advantage in subsurface hydro-biogeochemical modeling and deliver an order-of-magnitude increase in computational efficiency compared with current approaches while increasing predictive accuracy by incorporating both continuous and discrete observational data.

Spanning from the surface through deeper soils and into the bedrock, the subsurface critical zone (SCZ) represents a crucial area for resource availability that is shaped by complex geochemical and biological interactions and ecohydrological processes. Biogeochemical cycling in the SCZ affects important processes such as organic matter degradation, nutrient cycling, contaminant mobilization, and critical mineral availability. These processes are sensitive to organic–mineral interactions, redox fluctuations, root interactions, and microbial processes. Although DOE supports state-of-the-art subsurface biogeochemical and reactive transport models, current models are computationally expensive to run at scale, lack representation of critical subsurface processes, and have limited capacity to ingest observational data. A major challenge is integrating geographically sparse continuous datasets with temporally sparse discrete sample data of subsurface hydro-biogeochemical processes such as those collected in Critical Zone Observatories and the Critical Zone Network. This project will combine process-based models and observational data to develop an AI-enabled subsurface biogeochemical modeling system that leverages existing DOE-supported interoperability frameworks to integrate seamlessly with land surface and hydrological models, with the following objectives:

- Train and evaluate a generalizable, multimodal AI biogeochemical solver using a combination of process-based model simulations, *in situ* measurements from instrumented soil pits, and laboratory quantifications of whole-profile deep soil biogeochemical properties. This will include continuous measurements of metrics such as moisture, temperature, and redox potential, as well as discrete sampling of soil structure, microbial community structure and function, and organic matter properties.
- Couple the AI biogeochemical solver into DOE transport and land surface models by implementing a flexible, interoperable AI interface to support efficient computations across host models and computer architectures.
- Evaluate the accuracy and computational efficiency of AI-enabled reactive transport simulations across a diverse set of deep soil profiles using multiple ecological and hydrological models, including the Energy Exascale Earth System Model (E3SM) Land Model (ELM) and the Advanced Terrestrial Simulator (ATS).

* BER is lead; co-funded with the DOE Advanced Scientific Computing Research (ASCR) program

This project will lay the groundwork for a next-generation, AI-enhanced modeling paradigm for complex environmental systems. Training AI biogeochemical solvers using a multimodal combination of process-based biogeochemical model simulations, field measurements from deep soil pits and cores, controlled experiments, and microbial sequence data will produce a foundational subsurface biogeochemistry capability that supports rapid, spatially and temporally continuous predictions of subsurface biogeochemistry across the United States and advances Genesis Mission priorities in critical zone and subsurface resource science.

Bridging Genomics and Reactive Transport Models with AI for Subsurface Prediction

- **Lead PI:** Jiaze Wang, University of Maine
- **Collaborating Institutions:** University of Delaware; Argonne National Laboratory

The predictive management of subsurface energy systems, including hydrological resource extraction and enhanced hydrocarbon recovery, requires a detailed understanding of coupled hydrogeochemical and biogeochemical processes. Microbial communities are central to these processes, mediating critical reactions that influence fluid chemistry, mineral stability, and contaminant fate. However, integrating high-dimensional omics data into traditional reactive transport models remains a significant challenge. The complexity of microbial functional potential, coupled with heterogeneous subsurface environments, makes exhaustive physics-based simulations computationally prohibitive. This project assembles a uniquely capable multi-institutional team with expertise in bioinformatics and computational biology, biogeochemistry, reactive transport modeling, computer science and AI to develop an innovative AI-empowered framework that seamlessly bridges high-dimensional multiomics data and reactive transport modeling, with the goal of creating the adaptive and predictive capability needed to manage subsurface energy systems with confidence.

Leveraging DOE's ongoing COMPASS-FME (Coastal Observations, Mechanisms, and Predictions Across Systems and Scales—Field, Measurements, and Experiments) program and its publicly available comprehensive datasets as a testbed, this project will (1) adapt large genome-scale language models to extract structured genome-informed traits from multiomics datasets; (2) embed omics-explicit microbial representations directly into the selected reactive transport model, PFLOTRAN; and (3) train physics-, omics-, and chemistry-informed AI emulators that relieve computational and biogeochemical complexity while preserving physical, biological, and chemical realism, collectively bridging molecular-scale measurements and field-scale biogeochemical prediction. This framework (1) enables genome-informed predictions of system behavior, guiding interventions and operational strategies to optimize energy recovery, minimize impacts, and ensure long-term system stability, and (2) allows researchers and operators to assess the impacts of microbial variability on subsurface energy systems in near-real time, supporting robust decision-making under uncertainty.

This project aims to transform subsurface biogeochemistry from a descriptive to a predictive science by creating an innovative AI-empowered framework that systematically incorporates multiomics data into reactive transport models, enabling mechanistically grounded, field-scale predictions of microbial, mineral, geochemical, and fluid interactions. Predictive subsurface models carry direct significance for applied energy, underpinning critical decisions in energy system siting, contaminant stewardship, carbon storage, and the long-term integrity of subsurface energy assets. By bridging molecular biology and large-scale, real-world simulation, this project lays the foundation for a pioneering microbiome digital twin that advances DOE's mission to unleash strategic subsurface energy assets while providing a generalizable, scalable platform for AI-accelerated Earth system modeling.

Scaling the Grid To Power the American Economy Selection

The electric grid faces reliability challenges and infrastructure limitations as it struggles to accommodate dramatic increases in electricity demand from data centers, manufacturing, and electrification while maintaining affordable power for Americans.

Using deep and reinforcement learning techniques on newly integrated big data sources, AI will reduce uncertainty, improve insights, and speed processes in grid planning, interconnection, operations, and security. This effort aims to enable 20 to 100 times faster decision-making and at least 10% improvement in grid reliability and reduction in electricity cost.

Under this S&T challenge, BER is involved in one focus area: Uncertainty Quantification.

Uncertainty Quantification

AI-Powered Frameworks To Enhance Grid Resilience Under High-Dimensional Uncertainties

- **Lead PI:** Karen Fisher-Vanden, Arizona State University
- **Collaborating Institutions:** The Pennsylvania State University; University of California–San Diego; Lawrence Berkeley National Laboratory; Pacific Northwest National Laboratory

The electric power sector faces unprecedented challenges from the confluence of multiple stressors and uncertainties, including changes in demand, generation technologies, and extreme weather. Traditional capacity expansion models (CEMs) are too computationally intensive to explicitly represent the large number of plausible future scenarios that reflect the many dimensions of uncertainty. A computationally efficient AI emulator of a high-resolution CEM can explore many more future scenarios at much lower computational burden than the CEM. The AI emulator can then select the best representative subset of scenarios for further detailed planning studies.

This project will develop an AI-accelerated scenario selection method for capacity expansion planning to enhance grid resilience. Interpretable Bayesian networks—probabilistic, network-based learning algorithms—will be used to improve traditional scenario selection tools by (1) modeling the dependencies between scenario drivers and generation and transmission mix and (2) reducing the scenario set from thousands of possibilities to a few representative scenarios. A case study of the Northeast United States will demonstrate the AI advantage of this approach. AI benefits will be quantified in terms of reduced computation time to achieve the same quality result (target 60% reduction), higher-quality results within the same computation time (target 50% increase), and increased robustness across repeated trials (target 30% increase).

This project will build on prior successes among the investigators with grid optimization, energy planning, and AI methods to develop a dynamically adaptive AI workflow that identifies the most critical scenarios for maintaining reliability. The proposed team will extend existing partnerships with several grid operators [e.g., the Pennsylvania-New Jersey-Maryland Interconnection (PJM) and the New York Independent System Operator] to make these AI methods implementable by grid operators and maximize their value for grid resilience in practice.

Artificial Intelligence in Fluid Flow for Energy Components and Technologies (AFFECT) Selection

The dynamic behavior of fluids (gases, liquids, plasmas, and multiphase mixtures) dictates the performance, safety, and longevity of many energy technologies. Examples are diverse, including extreme weather impacts on energy infrastructure, reacting flows for combustion, magnetically confined fusion plasmas, cooling of advanced reactors and data centers, resin flows through composites, and fuel storage and transport. Achieving a secure and resilient energy future depends on understanding and optimizing fluid flow, but current modeling and design approaches struggle with the multiscale, multiphysics nature of these fluid systems. This leads to high development costs, suboptimal performance, and slower deployment of critical energy technologies.

AI in Fluid Flow for Energy Components and Technologies (AFFECT) leverages sophisticated AI and ML techniques, including physics-informed neural networks and reinforcement learning, to analyze vast datasets from high-fidelity simulations (e.g., partial differential equation–based computational fluid dynamics), *in situ* and remote sensing observations, and flow experiments in the field and laboratory. This will enable rapid discovery of governing equations and surrogate models for complex fluid phenomena and optimization of control strategies. The topic aims to accelerate the design cycle of energy technologies by five- to 10-fold, reduce operational inefficiencies by 15 to 25%, and significantly enhance safety and reliability.

Under this S&T challenge, BER is involved in one focus area: Physics-Informed AI for Complex Flow Modeling.

Physics-Informed AI for Complex Flow Modeling

A Physically Informed Neural Network for Power-Grid Resilience to Extreme Wind Events

- **Lead PI:** Sean Bailey, University of Kentucky
- **Collaborating Institutions:** Brookhaven National Laboratory; National Center for Atmospheric Research

Damaging winds are a leading cause of power distribution failures, producing widespread outages, costly infrastructure damage, and, in fire-prone regions, elevated wildfire risk. Despite improvements in numerical weather prediction, operational models often cannot resolve the fine-scale wind variability that drives extreme wind loads on individual power line assets. This limits the ability of utility providers and partners to anticipate localized hazard conditions, which can be influenced by terrain, land cover, and the built environment.

This project will develop an AI-enabled, physics-based downscaling capability to predict near-surface wind conditions at resolutions relevant to power line infrastructure. The vision is to advance DOE's science and applied energy initiatives by improving how multiscale atmospheric flow physics are represented and by translating that capability into faster, more reliable tools for extreme wind hazard assessment. Ultimately, this project

aims to strengthen power system resilience by enabling earlier and more accurate identification of dangerous wind states.

The research team will use a conditional, physics-informed ML framework that incorporates atmospheric forcing from operational weather prediction while enforcing governing flow physics. To enhance scientific and technical workflows, AI will reduce the need for repeated, computationally expensive high-resolution simulations by providing a reusable predictive model that can efficiently be adapted to new sites and conditions. This supports rapid evaluation of many weather and site scenarios, enabling decision-relevant analysis that would be impractical with simulation-only approaches.

Model training and evaluation will use observations from ARM's Southern Great Plains (SGP) facility, augmented with high-resolution flow fields generated using FastEddy—the National Center for Atmospheric Research's Large Eddy Simulation model. Performance will be assessed through (1) prediction of extreme wind events within the ARM SGP dataset and (2) transfer to land-use conditions in the Lexington, Ky., region, quantified using wind speed accuracy, gust exceedance prediction skill, and spatial agreement.

Phase I deliverables include a validated downscaling prototype and a quantitative assessment of performance gains over baseline approaches.