

Genesis Mission Frameworks for AI-Accelerated National Breakthroughs



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

Report from the U.S. Department of Energy Office of Science Advisory Committee

Genesis Mission Frameworks for AI-Accelerated National Breakthroughs

Report from the SCAC Subcommittee on the Genesis Mission

Genesis Mission Subcommittee

Suresh Garimella, Chair
University of Arizona

Makoto Gonokami
RIKEN

Nadya Mason
University of Chicago

Pushmeet Kohli, Vice Chair
Google DeepMind

Michael Mahoney
University of California–Berkeley
International Computer Science Institute
Lawrence Berkeley National Laboratory

David Siegel
Two Sigma

Drew Endy
Stanford University

David Spergel
Simons Foundation
Flatiron Institute

Designated Federal Officer

Harriet Kung
U.S. Department of Energy Office of Science

Coordinating Editor

Ali Douraghy
Google DeepMind

About SCAC

The Office of Science Advisory Committee (SCAC) provides advice to the U.S. Department of Energy's (DOE) Office of Science (SC) on complex scientific and technical issues that arise in the planning, management, and implementation of SC research programs. SCAC's advice includes establishing research and facility priorities; determining program balance among disciplines; and identifying opportunities for interlaboratory collaboration, program integration, and industrial participation. SCAC primarily includes representatives of universities, DOE national laboratories, and industry.

Acknowledgments

SCAC's Genesis Mission subcommittee appreciates the insightful feedback and discussion from subject matter experts consulted throughout development of this report: Paul Adams, Frank Alexander, Rommie Amaro, Tomás Díaz de la Rubia, Carrie Eckert, Justin English, Allison Bennett Irion, Warren Kibbe, Egemen Kolemen, Karl Mueller, Aindrila Mukhopadhyay, Sam Posen, Arvind Ramanathan, Robert Schlberg, Prashant Singh, Elizabeth Villa, Michael Wang, Anne White, and Roseanna Zia.

The subcommittee also acknowledges Ali Douraghy for guiding the report development process across three writing teams. Editorial support and slide preparation was provided by Oak Ridge National Laboratory's Biological and Environmental Research Information System team. Team members are Chloe Freeman, Jessica P. Johnson, Jackie Kerr, Stacey Wysong McCray, Marissa Mills, Natalie Parks, and Holly Haun.

The subcommittee is grateful for the expert guidance, coordination, and support provided by Todd Anderson, Michael Cooke, Matthew Lanctot, and Andrew Schwartz from DOE's Office of Science.

Suggested Citation

SCAC. 2026. *Genesis Mission Frameworks for AI-Accelerated National Breakthroughs: Report from the SCAC Subcommittee on the Genesis Mission*. DOE/SC-2031. Office of Science Advisory Committee. <https://doi.org/10.2172/3387411>.

Genesis Mission Frameworks for AI-Accelerated National Breakthroughs

Report from the SCAC Subcommittee on the Genesis Mission

September 2026



U.S. DEPARTMENT
of **ENERGY**

Office of
Science

Charge Letter



The Under Secretary for Science

Washington, DC 20585

March 26, 2026

Dr. Persis Drell
Chair
Office of Science Advisory Committee

Dear Dr. Drell:

The Genesis Mission was launched as a dedicated, coordinated national effort to create a new age of artificial intelligence (AI)-accelerated innovation and discovery that will help solve the most challenging problems of this century to the benefit of the nation.¹ This will include developing AI methods that will dramatically increase research productivity and impact, accelerate scientific discovery, strengthen national security, and help secure energy dominance. In addition, the Genesis Mission will build an integrated AI platform to harness Federal scientific datasets, train scientific foundation models and create AI agents, automate research workflows, and support acceleration of scientific breakthroughs.

I am asking the Office of Science (SC) Advisory Committee (SCAC) to form a subcommittee to further study how best to prioritize resources across SC to accelerate the application of AI for transformative scientific and engineering discovery focused on pressing national science and technology challenges and to create a decadal roadmap articulating near, mid, and long-term priorities for SC to achieve the Genesis Mission. The roadmap will be based on a rigorous assessment of the current landscape, identifying critical opportunities and gaps and suggesting milestone-based investments spanning facilities, research, and workforce development. The subcommittee's work should consider opportunities to create federal, state, university, philanthropic, and private sector partnerships leveraging the Department of Energy's (DOE) capabilities, private sector advances, and other agency resources to maximize impact.


The subcommittee should address the following questions in their deliberations:

- What should be included in the Genesis platform, and what is the refresh plan?
- What factors should be considered to maintain a world-leading National Science & Technology challenges portfolio that best capitalizes on the Genesis Platform?
- How should DOE engage with federal, state, university, philanthropic, and private sectors to cultivate and grow AI-ready workforce for the Nation?

¹ <https://genesis.energy.gov/>

The subcommittee shall provide an assessment of the current landscape and a decadal roadmap to the SCAC for deliberation, by July 2026, based on the subcommittee's expertise and inclusive and broad stakeholder input. Further charges will follow as the Genesis Mission moves forward. I anticipate the subcommittee's work may involve making recommendations about infrastructure and facilities needs relevant for the Genesis Mission, and therefore I ask that the subcommittee coordinate closely with the SCAC or any subcommittees thereof (i.e., the Facilities Subcommittee) that are charged with making recommendations about the prioritization of SC's facilities needs for the next decade.

Sincerely,

A handwritten signature in black ink that reads "Darío Gil". The signature is written in a cursive, slightly stylized font.

Darío Gil
Under Secretary for Science

Contents

Chapter 1: Catalyzing a New Age of Scientific Discovery	1
Securing U.S Leadership in Discovery and Innovation	2
Empowering Scientists To Achieve Transformational Breakthroughs	2
Mobilizing DOE Capabilities and the Genesis Mission Platform for a New Scientific Infrastructure	3
Marshaling the Expertise of Industry, Interagency, and International Partners	5
Inspiring and Developing the Next-Generation Workforce	6
Leveraging AI for the Public Good	6
Chapter 2: Harnessing Biology at Digital Speed	7
Vision for a Biological Campaign	8
Sectors of National Impact	9
What Must Be Built	11
Pillar 1: AI-Enabled Measurement	11
Pillar 2: AI-Enabled Modeling and Prediction	13
Pillar 3: AI-Enabled Manufacturing and Design	15
A Compounding National Advantage	16
Chapter 3: AI-Accelerated Fusion Energy	17
Vision, Urgency, and Why Now	18
Pillar 1: Measurement and Data—Turning Decades of Data into Discovery	19
Pillar 2: Modeling, Prediction, and Validation—AI as Engineering Multiplier	21
Pillar 3: Making, Manufacturing, and Delivering—The Physical Path to the Grid	22
The Investment Imperative	25
Chapter 4: Magnet Sovereignty	26
The Bottom Line	27
Why This Matters Now	28
What Must Be Built	29
Pillar 1: AI-Accelerated Materials Discovery	29
Pillar 2: AI-Driven Manufacturing	30
Pillar 3: AI-Enabled Domestic Materials Recovery	30
The AI Advantage	31
Magnet Sovereignty Implementation Path	33
What Success Looks Like	33
Broader Scientific Impact	33
Appendix A: Office of Science Advisory Committee Members	34
Appendix B: References	35
Appendix C: Acronyms and Abbreviations	37

Chapter 1

Catalyzing a New Age of Scientific Discovery

DOE's Genesis Mission Brings Together Human Ingenuity and the Power of AI to Meet the Greatest Challenges of Our Day

Suresh Garimella, Subcommittee Chair
University of Arizona

Catalyzing a New Age of Scientific Discovery

DOE's Genesis Mission Brings Together Human Ingenuity and the Power of AI to Meet the Greatest Challenges of Our Day

Throughout history, a small number of foundational technologies have fundamentally transformed what humanity could accomplish. Electricity powered the modern world. The semiconductor launched the information age. The internet connected billions of people. Today, we stand at the threshold of the next great technological transition—one with the potential to be even more consequential than those that came before. Artificial intelligence (AI) is poised to become the defining enabling technology of our era. Developed and deployed responsibly, it has the potential to shape the next generation of scientific discovery, national security and competitiveness, innovation and economic growth, and human prosperity as a force for good.

Securing U.S. Leadership in Discovery and Innovation

The Genesis Mission (The White House 2025) is founded on a core premise: If we can accelerate the rate at which humanity discovers, tests, validates, and deploys new knowledge, we can solve some of society's most consequential challenges decades sooner—and at far lower cost—than would otherwise be possible.

The central question is not whether AI will transform science and engineering. It will. The question is whether the United States will lead that transformation or cede it to strategic competitors. The Genesis Mission is a national strategy to ensure American leadership by building the scientific AI infrastructure, computing capacity, data resources, and partnerships needed to accelerate discovery across disciplines.

The urgency of this effort is clear. Nations such as China are investing aggressively in AI infrastructure, talent, and industrial deployment. Maintaining a competitive edge requires sustained national capability, not isolated breakthroughs. The U.S. Department of Energy (DOE), with its network of world-leading laboratories, state-of-the-art scientific user facilities, university partners, computing expertise, and mission-driven culture of innovation, is uniquely positioned to lead this effort through the Genesis Mission. The time to act is now.

Empowering Scientists To Achieve Transformational Breakthroughs

Artificial intelligence is not a replacement for scientists, engineers, or innovators. Like the digital

The Genesis Mission is a commitment to accelerating discovery; strengthening American competitiveness; powering the economy, securing our energy future, and building a future in which science and technology serve the public good. The question is not whether we should act. We must. The question is whether we will act with sufficient ambition, urgency, and partnership to harness the full potential of this moment.

computer before it, AI is a powerful force-multiplier and collaborative tool. When guided by human ingenuity, AI amplifies researchers' capabilities and empowers them to address complex problems with greater speed, precision, and creativity in a self-reinforcing cycle. AI improves hypothesis generation, makes new discoveries possible (particularly in data-rich fields), enables real-time experimental control, and supports instrumentation in laboratories that expedite experimentation and validation. It compresses the timeline from basic discovery to deployment, reduces development cycles, lowers costs, and increases the productivity of every federal and private research dollar invested. These benefits extend across critical fields such as materials discovery, grid reliability, manufacturing optimization, advanced semiconductors, Earth and environmental systems modeling, quantum information science, biotechnology, and fusion energy, and they will only deepen as our capabilities mature. For example, the 2024 Nobel Prize in Chemistry recognized the dramatic impact made possible when massive protein structure datasets, such as the Protein Data Bank, coupled with new AI algorithms, enabled scientists to develop a highly accurate tool for predicting protein structures where no experimental data previously existed. AI will not eliminate the need for experiments, facilities or human creativity, yet in field after field, it can substantially shorten the path to breakthrough by analyzing vast datasets to rapidly predict outcomes and generate novel hypotheses.

Mobilizing DOE Capabilities and the Genesis Mission Platform for a New Scientific Infrastructure

Realizing this vision demands a new national scientific infrastructure that integrates computing, data, experimentation, and AI into a unified ecosystem for discovery. Agile, purpose-built, AI-optimized hardware and platforms are needed in addition to the leadership-class and exascale computing facilities already in place at DOE national laboratories. Future scientific leadership will depend on a balanced ecosystem that combines hardware; advanced software, including open-source AI models and industry frontier models; specialized computing architectures; trusted data resources; and distributed capabilities accessible to universities, national laboratories, startups, and industry. The Genesis Mission provides the national strategy to build this ecosystem.

DOE is uniquely positioned to lead the transformation of the national scientific infrastructure to meet the AI needs of the nation through the Genesis Mission. The agency has decades of experience operating the world's most powerful scientific computing systems, extensive experimental capabilities and scientific user facilities, and expertise in integrating large-scale science and engineering. It also brings deep domain expertise in physics, chemistry, biology, and materials science, which no commercial entity or single university can replicate at scale. Through its national

laboratories, university and industry partnerships, and mission-driven culture of innovation, DOE possesses the capabilities necessary to create the next generation of computational infrastructure that will underpin American leadership in science and technology.

The American Science Cloud

The power of advanced AI depends on high-quality data and specialized hardware. The backbone of the Genesis Mission is the American Science Cloud, which will act as the digital fabric linking DOE's computing, data, and experimental facilities, allowing researchers to access AI-ready datasets and train scalable foundation models. An essential driver of future breakthroughs will be the ability to generate and access unprecedented volumes of high-quality, trusted, and interoperable data. This will require dedicated data-generation campaigns and investing in data collection, stewardship, standards, curation, sharing frameworks, and secure repositories that allow data to be combined across disciplines and organizations. Simultaneously, we must build machines that are bigger, smarter, and more agile to support the computational demands of training foundation models and deploying AI at the scale needed to serve the nation's scientific enterprise. Through a series of innovative federal-industry partnerships as part of the Genesis Mission, DOE is building and will operate those AI machines at its national laboratories. To serve the whole-of-government vision of the Genesis Mission, these capabilities should be designed and built with interagency collaboration in mind to ensure data sharing and other capabilities work smoothly across interagency activities.

End-to-End Technology Innovation

The future power of advanced AI demands end-to-end innovation, or modernization across all technology layers, including algorithms, computing architectures, software, data governance, and energy infrastructure. The Genesis Mission platform is designed to meet this need. Unlike commercial AI platforms optimized for general-purpose tasks, this infrastructure must be purpose-built to support the rigorous requirements of scientific research, including physics-informed modeling, uncertainty quantification, and validation

against experimental data. The Genesis Mission platform will also support the development of new classes of AI algorithms that can reason across vast, multimodal data streams and interact seamlessly with physical systems.

To accelerate this innovation, the platform will engage key international partners in the codesign and codevelopment of advanced software frameworks, specialized AI architectures, and interoperable hybrid computing interfaces, enabling collective computational assets to be leveraged efficiently across borderless scientific frontiers. Together, these tools will enable the scientific community to leverage AI to accelerate the end-to-end process of science, from hypothesis generation through computational modeling, experimental and observational implementation, and data analysis—all while facilitating enhanced scientific reproducibility.

Digital Twins

Using the tools developed as part of the Genesis Mission, researchers will be able to combine AI foundation models, advanced supercomputers, and federated data to provide an operating system for modern scientific discovery by replacing traditional design-build-test cycles with AI-driven digital twins and rapidly accelerating the scientific discovery process. Together, these capabilities represent a foundational national asset for advancing scientific discovery, economic competitiveness, energy security, and societal well-being.

The DOE Office of Science is also well positioned to convene and coordinate multidisciplinary teams spanning national laboratories, universities, industry, and international partners to tackle a focused set of high-impact science and technology challenges as part of the Genesis Mission. These ambitious, mission-driven efforts in areas such as fusion energy, resilient energy systems, materials design, advanced manufacturing, quantum information science, and national security can serve as proving grounds for next-generation scientific AI and digital twin capabilities while delivering transformative advances that demonstrate the tangible societal benefits of AI-enabled science.




Fig. 1.1. Maximizing AI Impact Through Transformative Collaboration. The Genesis Mission unites the U.S. Department of Energy, its national laboratories, other federal science agencies, industry, academia, and the private sector—both nationally and internationally—to harness AI for breakthroughs that multiply advances in energy, discovery science, and national security. [FFRDCs: Federally Funded Research and Development Centers]

Marshaling the Expertise of Industry, Interagency, and International Partners

The scale of this challenge makes clear the need for new modes of partnership across the federal government, DOE national laboratories, academia, and the private sector (see Fig. 1.1, this page). With DOE as the lead, the Genesis Mission is a whole-of-government initiative that will unify the expertise, scientific datasets, computing capabilities, and research

facilities of all federal science agencies. The Genesis Mission provides a new opportunity to strengthen existing partnerships and establish new ones between DOE and other agencies to benefit the American people, building on a legacy that includes major advances in human health—including the Human Genome Project—national security, space exploration, and scientific discovery. These interagency collaborations will be essential in realizing the Genesis vision and supporting a workforce equipped to leverage AI for science.

But public investment alone is insufficient. Achieving these ambitions will also require strong partnerships with industry to provide access to the advanced computing infrastructure, specialized hardware, cloud resources, and AI technologies needed to operate at unprecedented scale. Public-private partnerships, including efforts such as the Genesis Mission Consortium, can accelerate innovation while ensuring that the nation's scientific community has access to frontier models, AI for science tools, and the computational resources necessary to address increasingly complex challenges.

Federal leadership remains essential, but philanthropic organizations and private foundations have a critical, catalytic role: funding early-stage research, enabling cross-disciplinary collaboration, supporting workforce development, and advancing public-interest applications of AI that may fall outside traditional funding timelines. Private-sector leaders are equally essential to translating laboratory breakthroughs into real-world industrial implementation. Cooperation with trusted allies and democratic partner nations must be deepened through joint research initiatives, technical standards development, talent exchanges, and coordinated infrastructure planning within a framework that preserves transparency, security, and open scientific inquiry. To this end, the Genesis Mission will provide a collaborative foundation for transparent, reciprocal governance of shared data and codeveloped computational resources, accelerating scientific breakthroughs while preserving security and open scientific inquiry.

Inspiring and Developing the Next-Generation Workforce

Long-term leadership in AI-driven science and engineering ultimately depends on people, not just hardware and software. We must build a durable talent pipeline that expands educational and professional pathways at every tier: inspiring curiosity during K–12 STEM education, training the next wave of researchers through undergraduate programs and graduate fellowships, and creating postdoctoral and interdisciplinary research opportunities that combine computational

expertise with deep domain knowledge. Preparing this next generation of scientists, engineers, technicians, and educators is the foundation upon which everything else rests.

Leveraging AI for the Public Good

Finally, advancing this technology requires meeting public expectations with transparency and accountability. Progress must deliver science and technology breakthroughs that benefit the American people. As with past scientific transformations enabled by new technologies, the United States can lead in the ethical use of AI and in setting standards that promote scientific openness and discovery while respecting individual rights and data privacy. Strengthening a trusted open-science community is essential to achieving national objectives.

With this framework, the Genesis Mission offers an opportunity to integrate domain expertise with computer science innovation to simultaneously advance discovery to improve energy independence, strengthen domestic manufacturing, support high-skilled jobs, and deliver breakthroughs that benefit all of society.

In response to a March 2026 charge from Darío Gil, DOE's Under Secretary for Science (see Charge Letter, p. ii), this report outlines a bold vision for how DOE can prioritize its efforts to achieve such breakthroughs. The Genesis Mission subcommittee identified three opportunity areas that exemplify the potential of AI to transform the U.S. scientific enterprise: biology, fusion, and high-performance magnets. These opportunities align well with issues of national importance, DOE's existing strengths, and the vision of the Genesis Mission.

Ultimately, the Genesis Mission is a commitment to accelerating discovery, strengthening American competitiveness, powering the economy, securing our energy future, and building a future in which science and technology serve the public good. The question is not whether we should act. We must. The question is whether we will act with sufficient ambition, urgency, and partnership to harness the full potential of this moment.

Chapter 2

Harnessing Biology at Digital Speed

A Campaign for AI-Enabled Biology

Authors

Pushmeet Kohli (Lead)
Google DeepMind

Drew Endy
Stanford University

Lara Jehi
Cleveland Clinic

Additional Contributors

Ali Douraghy
Google DeepMind

Paul Adams
Lawrence Berkeley National Laboratory

Rommie Amaro
University of California–San Diego

Roseanna Zia
University of Missouri

Carrie Eckert
Oak Ridge National Laboratory

Aindila Mukhopadhyay
Lawrence Berkeley National Laboratory

Elizabeth Villa
University of California–San Diego

Warren Kibbe
National Cancer Institute

Arvind Ramanathan
Argonne National Laboratory

Robert Schlager
Illumina

Justin English
University of Washington

Harnessing Biology at Digital Speed

A Campaign for AI-Enabled Biology

Vision for a Biological Campaign

From energy independence to supply chain resilience, biology is central to U.S. physical security and economic prosperity. Advancing biology in support of energy and manufacturing leadership also pays dividends for human health, biodefense, and environmental stewardship.

The DOE Genesis Mission's biology campaign proposes to overcome a foundational challenge that has limited all of biology and biotechnology: accurately predicting the functions and behaviors of living systems from their genetic code alone.

For more than three decades, the ability to sequence DNA has outpaced the ability to interpret it, limiting efforts to understand and engineer biological systems on demand (U.S. DOE 2021). Across sequenced genomes, an estimated 30% or more of protein-coding genes still lack a known function (Jeffrey 2023), and the gap is even wider for the complex organisms most vital to national energy and security missions.

Over the next decade, the biology campaign will make partnering with biology as routine and dependable as engineering with silicon or steel, giving the nation's

researchers and industry the power to measure biology at scale, to model it well enough to predict what a change will do, and to manufacture it on demand. That capability turns biology from a slow, manual craft into a national engine for powering and feeding the country, answering biological threats in days rather than years, and curing disease faster. Biology-by-design, from molecules to cells, will be one powerful demonstration of this capability, but will also be extended to organisms and systems to fully realize the applications of this technology.

To deliver these benefits across energy, agriculture, advanced therapeutics, and national defense, biology must move from an observational science to a predictive, engineerable one, practiced at digital speed and amplified by AI. Making biology designable will require fast, experimentally grounded models that bridge molecular physics and cellular function. Until now, that goal has been constrained by slow measurement, the prohibitive computational cost of simulating biology at meaningful scales, and the trial-and-error involved in building and scaling living systems.

When biology becomes reliably measurable, predictable, and engineerable, solving problems in living

systems becomes routine rather than rare. The nation gains the ability to design a therapy, a crop, a material, or a defense and to know that it will work before building it.

Biology is becoming one of the world's largest industries (McKinsey Global Institute 2020). The global bioeconomy is worth about \$4 trillion today and could reach \$30 trillion by 2050 (UNEP-CCC 2025). The biology campaign sets a goal to lift the U.S. share from a quarter to a third, turning a \$1 trillion annual position today (NASEM 2020) into \$10 trillion by 2050. Another return on investment is reduced cost, amounting to a roughly tenfold drop in the price and time to engineer and optimize a target biological process (see Fig. 2.1, p. 10). An improvement of this magnitude is not unprecedented. The cost of sequencing a genome fell more than a millionfold in just two decades, outpacing even Moore's Law for computer chips (Wetterstrand 2022). The biology campaign aims to drive a similar scale of improvement for deploying biology, doing for engineering living systems what sequencing did for reading them. Meeting these goals would also help deliver the Genesis Mission's central aim of doubling the nation's scientific productivity within a decade.

Why DOE?

Turning biology into a fast, dependable engineering discipline takes national-scale infrastructure that no single laboratory or company can assemble or sustain, along with the industrial-scale team science to run it. That is DOE's core strength. Its 17 national laboratories already run the country's most powerful exascale computing, its brightest synchrotron and neutron sources, and premier genomics and molecular facilities such as the DOE Joint Genome Institute (JGI) and Environmental Molecular Sciences Laboratory (EMSL). Rebuilt around AI, DOE's new biology mission becomes the backbone of a life-science, digital-biology platform the rest of the country designs and builds upon. This effort can turn the Biological and Environmental Research (BER) program's Bioenergy Research Centers into hubs of predictive biodesign, and through the Genesis Mission, run measurement, modeling, and manufacturing as one national effort. This is a role only DOE can fill, and one the nation now needs it to undertake with speed and scale.

Sectors of National Impact

The campaign's new abilities to measure, model, and manufacture biology at scale do more than improve vital economic sectors; they redefine what is possible. In every arena, from manufacturing, agriculture, and energy to human health and biodefense, biodesigned breakthroughs are concrete and within reach.

Biomanufacturing. A secure domestic supply of biologically derived advanced therapeutics, refined chemicals, and novel materials depends on scaling the full read–annotate–write pipeline used to first understand organisms and then design new and useful functions. This involves sequencing genomes, determining gene functions, and finally building new gene sequences to manufacture bioproducts on an industrial scale. Genomic data feeds AI-driven digital twins that tune production from molecule to reactor, yielding a predictable, full-scale system that maximizes volumetric productivity at minimum cost. Built as a distributed national network, it can place high-value manufacturing in rural communities and make the domestic supply chain resilient to disruption.

Example National Breakthrough Targets:

- An engineered microbe turns carbon dioxide into jet fuel that is cost-competitive with the conventional product, and the industrial-scale production run meets manufacturing specifications on the first batch.
- A U.S. plant makes a top-10 generic drug ingredient, currently sourced from China, at commercial scale and online in 18 months, not a decade.

Agriculture and Bioenergy. Using predictive design to improve photosynthesis would sharply increase the chemical energy crops capture from sunlight, potentially yielding terawatts of new bioenergy (Ort et al. 2015). The same design capability yields resilient, non-food feedstock crops tuned to biorefinery outputs, and food crops that maintain high yields under severe heat, drought, and emerging plant disease.

Example National Breakthrough Targets:

- Soybeans with redesigned photosynthesis show a 50% field-yield gain, the first engineered break past a limit evolution set long ago.




Fig. 2.1. AI Stands to Accelerate U.S. Bioeconomic Growth and Global Competitiveness.

The National Return, in Three Numbers: The biology campaign targets a tenfold-larger U.S. bioeconomy (\$1 trillion today to \$10 trillion by 2050), a tenfold drop in the cost and time to solve a problem using biology, and a doubling of the nation's scientific productivity (a central aim of the Genesis Mission). The precedent is real: The cost of sequencing a genome fell more than a millionfold in two decades, outpacing Moore's Law (Wetterstrand 2022). **What's at Stake:** The United States holds about a quarter of today's \$4 trillion to \$5 trillion global bioeconomy; That market could grow to \$30 trillion by 2050 (UNEP-CCC 2025). Without sustained investment, the U.S. share falls to 15–20% (NASEM 2020), or less than 10% if biomanufacturing also moves overseas the way drug-ingredient manufacturing did (Duman 2025). This AI-enabled biology campaign aims to increase the pace of U.S. bioeconomic growth, targeting a one-third share of the global bioeconomy by 2050. The potential loss to the U.S. bioeconomy with no sustained investment is estimated at \$5 trillion to \$7 trillion a year (McKinsey Global Institute 2020).

- Engineered sorghum feeds a Midwest biorefinery making aviation fuel, maintaining high yield through a drought that halves unimproved crops.

Human Health. The same predictive models DOE builds for its bioenergy apply directly to human disease, making therapeutic discovery a high-value return on infrastructure the United States is already building. DOE's exascale computing, synchrotron light sources, neutron facilities, and structural biology infrastructure form a national platform for resolving the atomic structure and dynamics of disease-relevant proteins long considered undruggable by conventional small molecules (Hopkins and Groom 2002). These proteins resist standard therapeutics because they lack stable binding pockets and instead shift through transient conformational states.

Example National Breakthrough Targets:

- Drugs are designed *in silico* for *Myc* genes involved in several cancers, and long considered undruggable, from target to candidate in nine months.
- All ~20,000 human genes get a validated function, closing a 25-year gap and revealing hundreds of new drug targets at once.

Biodefense. The same measurement and manufacturing infrastructure doubles as a defense system. Interweaving biosecurity and public health monitoring into the nation's bioenergy and biomanufacturing ecosystem creates a passive threat detection network. Routine environmental and biological sampling becomes an engine for novel gene discovery, feeding a distributed biosurveillance capability that can produce targeted countermeasures the moment an engineered or natural threat emerges.

Example National Breakthrough Targets:

- Given a novel engineered pathogen, scientists can sequence, characterize, and match it to a designed antiviral drug within 48 hours of the first case.
- A biosurveillance network catches and treats against an unknown pathogen in wastewater before anyone is harmed.

Competitive Urgency

AI is enabling rapid progress in biology; China has designated biomanufacturing as a core national

industry, capturing a market share of \$157 billion as of 2025, producing over 70% of the world's output of biofermentation products on roughly \$4 billion in annual state investment (Sedzro 2026). In June 2026, the National Security Commission on Emerging Biotechnology warned that China is already executing the very strategies (NSCEB 2025, 2026) the Commission had urged the previous year that the United States should adopt (NSCEB 2025). This is the mechanism by which the American share declines: not a sudden loss, but the steady migration of production capacity, scale-up expertise, and trained workforce to wherever the infrastructure is built first. The first nation to harness AI to interpret and design biological functions at digital speed will set the global paradigm for the crops, fuels, materials, and medicines that underpin industrial leadership.

What Must Be Built

The biology campaign establishes three interdependent capabilities, or pillars: (1) AI-Enabled Measurement, (2) AI-Enabled Modeling and Prediction, and (3) AI-Enabled Manufacturing and Design. Together, these pillars compress the cycle between biological hypothesis and verified physical production to digital speed. Rather than a distributed portfolio of isolated projects, this campaign envisions a continuous, closed-loop national AI-powered engine for biology, a single Measure → Model → Manufacture loop in which each result becomes the starting point for the next.

Pillar 1: AI-Enabled Measurement

Measurement is where the loop begins. This pillar envisions an overhaul of how biological data is acquired, shifting from destructive snapshots that capture a single data layer at a single instant to a live, multiomic view of systems left intact. Conventional techniques tear cells apart, destroying the multilayered spatial and temporal relationships that govern cellular function. Conventional sequencing does the same to DNA, fragmenting it into millions of pieces and obscuring the long-range genomic relationships that govern biological behavior. Capturing these relationships means observing biological systems in their native states and at digital speed—from molecules to whole cells, whole organisms, and communities.

Table 2.1 AI-Enabled Measurement Implementation Path

Biology Target Horizon	High-Throughput Systemic Multiomics	Visual Proteomics and Structural Dynamics
2028 Within Reach	First complete functional gene annotation of a single-celled microbe.	First real-time, tag-free 3D imaging of molecular machines inside a living cell.
2030 Uplift Capability	Automated, AI-driven multiomics that slash sequencing time and cost tenfold, ingesting multimodal instrument outputs into a connected data lakehouse.	Active-learning AI that steers non-destructive imaging such as cryo-ET, quantum sensors, and adaptive-optics light-sheet microscopy in real time, tracking thousands of molecular conformations inside cells and 4D dynamics across cells and tissues.
2035 Transformative Outcome	Sequence-to-function at speed, mapping the full multiomic blueprint of any organism within 24 hours of discovery.	Live instrument streams build maps spanning near-atomistic structure to 4D cell, tissue, and organoid dynamics under physiological conditions.

DOE has a long legacy of building this kind of measurement infrastructure whenever classical biology hits a wall. It launched the first federal genome initiative in 1986 (U.S. DOE 1987) and built early tools like the gene-finding algorithm GRAIL (Uberbacher and Mural 1991). Under the Genesis Mission, that expertise underpins two measurement capabilities (see Table 2.1, this page).

High-Throughput Systemic Multiomics. The first capability targets the data-scarcity gap in genomics, proteomics, and metabolomics. Across DOE user facilities, automated liquid-handling systems and active-learning AI pipelines would sharply cut the time and cost of multiomic profiling. Automated throughput would stream genomic, transcriptomic, epigenetic, proteomic, and metabolomic profiles straight into a lakehouse of AI-ready data linking BER repositories and systems such as JGI and the Systems Biology Knowledgebase (KBase). Sample processing becomes a routine data-ingestion pipeline spanning health, bio-economic, ecological, and national security domains, continuously feeding U.S. models for predictive biology and biomanufacturing, and enabling the detection and characterization of novel biological threats from environmental surveillance data. The same throughput can close biology’s largest knowledge gap: No organism

yet has a complete map of every gene function. By presenting organisms with an ensemble of physical and environmental conditions and reading out the resulting phenotypes, the campaign turns functional annotation from case-by-case discovery into systematic measurement. Because each annotation is a prediction, it feeds a close-the-loop workflow, tested by synthesis and re-measurement and sharpened with every cycle.

In-Cell Visual Proteomics and Structural Dynamics. The second capability, visual proteomics, images the physical mechanics of life inside intact cells without requiring chemical tags that distort or destroy them. Generative AI would interpret the raw, noisy, *in situ* data streams from methods such as cryo-electron tomography (cryo-ET) directly, identifying, segmenting, and classifying molecular structures in real time. These molecular-scale methods pair with emerging technologies such as quantum sensing and quantum-enhanced microscopy to map intracellular magnetic fields, localized temperatures, and chemical kinetics beyond the limits of classical optics. At the mesoscale, adaptive-optics lattice light-sheet microscopy (Liu et al. 2018) would follow whole cells, tissues, and organoids in three dimensions over time, giving Pillar 2 the dynamic, cell-to-tissue measurements its multiscale models require.

AI is transforming biology from a collection of specialized fields into an integrated science of living systems. These capabilities are organism agnostic. The same platforms that read microbes and crops, core BER program priorities, extend to animal and human systems, making human biology not a departure from DOE's mission but its most demanding testbed, required to lead the era of AI-powered biological innovation.

Pillar 2: AI-Enabled Modeling and Prediction

Modeling is the loop's second stage, turning the data streams Pillar 1 produces into predictive, physically grounded simulators. Under the Genesis Mission, the biology campaign will expand two modeling capabilities (see Table 2.2, this page).

Dynamic Molecular Mechanics Modeling. Modeling must shift from describing biology to predicting how molecules move and work, not just as static objects, to comprise entire living cells. Tools like AlphaFold2 already predict a protein's folded structure from its sequence with near-experimental accuracy (see sidebar "AI Breaks Barriers in Understanding and Designing Proteins" p. 14; Jumper et al. 2021), a milestone

showing AI can learn deep rules of molecular biology. But such models stop at a single static snapshot of largely natural proteins; they do not capture how a molecule moves, how it behaves once engineered beyond evolutionary precedent, or the physics of how molecules bind and react.

Three advances are ripe now: extending models from static structures to dynamics, bringing the same predictive power to RNA and mRNA, and grounding predictions in physics so they apply to engineered molecules that evolution never sampled. Physics-based quantum computing offers a longer-term path to that grounding, one that links the campaign to DOE's quantum mission. Blending AI for speed with physics for reach, and checking AI outputs against DOE's simulators before any design is built, makes modeling a reliable design tool. The results are direct: industrial enzymes that work the first time; medical countermeasures on demand; and faster, more affordable drug development.

Whole-Cell Causal Simulation. The second challenge is modeling at scale, moving from single molecules to whole cells. Deep-learning models can bridge kinetic, atomistic, and crowded-cell scales into one

Table 2.2 AI-Enabled Modeling and Prediction Implementation Path

Biology Target Horizon	Dynamic Molecular Mechanics Modeling	Whole-Cell Causal Simulation
2028 Within Reach	First AI model to predict a protein's motion between functional states, validated against experimentation.	First real-time, physics-grounded running model of a single microbial cell.
2030 Uplift Capability	Physics-constrained predictive foundation models that run inference-time search against physical simulators to map a molecule's shifting structural shapes in seconds.	A multiscale virtual-cell prototype that unifies kinetic, atomistic, and mesoscale contexts with calibrated uncertainty, and that acts as a self-correcting discovery engine to replace first-pass wet-laboratory screenings.
2035 Transformative Outcome	A unified model that maps conformational ensembles and cellular mechanics in seconds, under mature scaling laws linking compute to capability gains.	Autonomous digital twins that design synthetic cells in real time with high confidence and experimental constraints, optimizing whole-cell biological behavior before any physical build.

AI Breaks Barriers in Understanding and Designing Proteins

In 2024, the Nobel Prize in Chemistry recognized two AI-driven breakthroughs in biology (NobelPrize.org 2024). David Baker was honored for using computational protein design to create a protein called Top7 in 2003, the first protein with a structure unlike any found in nature. His group has since designed proteins that serve as medicines, vaccines, and molecular sensors. Demis Hassabis and John Jumper were honored for AlphaFold2, an AI model that largely solved a 50-year-old challenge: predicting a protein's three-dimensional shape from its amino-acid sequence. AlphaFold2 has now predicted structures for over 200 million proteins, almost every protein known to science. Together, the power to predict protein structure and to design new proteins from scratch shows how AI can break seemingly impenetrable barriers.

DOE resources helped make this possible. AlphaFold2 was trained on the Protein Data Bank, the global archive of experimentally determined protein structures that DOE supports in part, alongside the vast public sequence databases DOE has long helped build, including through its Joint Genome Institute. Because protein structure informs protein function, these tools are already driving applications across biotechnology, manufacturing, energy, agriculture, and medicine.




Fig. 2.2. Protein Structures Predicted by AlphaFold2. The pictured structures were previously unknown to science: **(left)** part of a large molecular assembly that forms a pore in the nuclear membrane of human cells, **(middle)** a natural enzyme that breaks down plastic, and **(right)** a bacterial enzyme that confers antibiotic resistance. Such predictions can now guide new solutions in human health, waste management, and pathogen control. [©Terezia Kovalova/The Royal Swedish Academy of Sciences]

causal simulator, with in-cell measurements from Pillar 1 entering as physical constraints, so researchers can design a cell's core machinery *in silico* before building it. Because the simulator reports calibrated confidence with every prediction, it becomes a self-correcting partner that proposes experiments worth running rather than serving as a passive

reference. The result is a virtual cell that replaces first-pass wet-laboratory screening, compressing months of trial-and-error experimentation into hours of modeling.

Holding the field to a real measure of progress is just as important. Decades ago, the protein-structure

community created the Critical Assessment of Structure Prediction (CASP), a shared benchmark that scored blind predictions against experimentally solved structures no one had seen. By turning the question of “how good is a model” into a defined metric, CASP made it possible to know when real progress had been achieved, and it is where AlphaFold2 proved its accuracy in 2020. Biology modeling now needs equivalent benchmarks for molecular dynamics, RNA structure and dynamics, and whole-cell modeling; DOE is uniquely placed to convene these fields and define the needed metrics.

Pillar 3: AI-Enabled Manufacturing and Design

Manufacturing closes the loop, converting designs into physical outputs and products whose results flow back to sharpen the next cycle. Under the Genesis Mission, the biology campaign will expand two primary biomanufacturing and biodesign capabilities (see Table 2.3, this page).

Next Generation Biofoundries. Manufacturing must move from empirical, trial-and-error bioengineering to automated, predictable, and scalable

biological foundries. Historically, scaling a successful bench process up to industrial volume has foundered on the fragility of traditional laboratory strains, which behave unpredictably and lose resilience at larger scales. The goal is a network of autonomous, end-to-end design–build–test–learn (DBTL) loops that push production beyond standard laboratory strains into robust industrial organisms, meaning industrial microbes naturally equipped to survive harsh processing conditions. In these closed loops, generative AI uncovers the constraints for transformation and selection of functional genetic parts, designs the metabolic pathways, controls robotic foundries to validate genetic parts and build the cells, and automates analytics to continuously tune performance. Rather than treating cellular survival as an afterthought, the framework uses AI to design the metabolic pathway and the cell’s resilience at the same time. For example, the process can design a stronger cell membrane and add molecular efflux pumps that clear toxic bioproducts before they kill the host.

Predictive Scale-Up. A second objective is to close the long-standing scale-up gap by treating bioreactor physics and cellular mechanics as one integrated computational problem. That means building a predictive

Table 2.3 AI-Enabled Manufacturing and Design Implementation Path

Biology Target Horizon	Next-Generation Biofoundries	Predictive Scale-Up
2028 Within Reach	First self-driving design–build–test–learn (DBTL) cycle in a robust industrial microbe, with researchers directing rather than operating it.	First complex bioproduct taken from bench to 100,000+ liters on the first attempt, including downstream processing.
2030 Uplift Capability	Automated DBTL loops that engineer robust industrial hosts, using AI-designed efflux pumps and membrane changes to improve tolerance to toxic products.	Predictive models that combine multiomics data with bioreactor physics to foresee scale-up bottlenecks up to 500,000+ liters, eliminating trial-and-error runs.
2035 Transformative Outcome	A national network of self-driving foundries that carry out turnkey biological assembly from high-level operational prompts.	First-attempt scaling from <i>in silico</i> design straight to full and distributed scale beyond 1 million liters, a tenfold gain in titer, rate, yield, and scale metrics for novel materials and chemicals.

framework that accounts for titer (amount of product), rate (speed of production), yield (efficiency of conversion), and scale (performance across volume changes)—the TRY&S metrics—from the very first design step. Performance is engineered into the process from the beginning rather than discovered during scale-up. Cross-modal foundation models pair systems biology data with real-time bioreactor physics to predict how physical forces such as fluid shear, oxygen gradients, and temperature fluctuations will affect cellular performance in full industrial-scale vats.

This predictive capacity enables scaling of complex bioproducts on the first attempt, compressing the translation timeline from bench to industry from years to days. This framework spans the whole path from bench to pilot to full commercial and distributed scales. It also treats downstream processing, the purification and finishing steps where many bioprocesses fail, as part of the initial design. Data from pilot and full-scale runs feed back to the design stage, so the biology is optimized for real-world conditions, not

just the laboratory. The same models target volumetric productivity, the grams per liter per hour that decide whether a product is worth making at all. The framework also introduces turnkey workflows that automatically match a target molecule to its optimal metabolic pathway, replacing bespoke, custom engineering with standardized, programmable design rules.

A Compounding National Advantage

By 2035, this AI-enabled biology campaign can place the United States at the center of a predictive bioeconomy, designing crops, fuels, materials, medicines, and biomedical countermeasures *in silico* and building them on the first attempt. The first nation to industrialize biology at digital speed will shape this century's largest emerging economy. The window to lead is now, and the infrastructure to seize it must be built deliberately, at national scale, and under one mission with the computational and scientific power DOE brings.

Chapter 3

AI-Accelerated Fusion Energy

Delivering Abundant Fusion Energy by 2030–2035

Authors

David Spergel (Lead)
Simons Foundation, Flatiron Institute

Tammy Ma
Lawrence Livermore National Laboratory

Mark Papermaster
AMD

Derek Sutherland
Realta Fusion

Additional Contributors

Tomás Díaz de la Rubia
University of Arizona

Egemen Kolemen
Princeton University

Sam Posen
Fermi National Accelerator Laboratory

Anne White
Massachusetts Institute of Technology

AI-Accelerated Fusion Energy

Delivering Abundant Fusion Energy by 2030–2035

Vision, Urgency, and Why Now

Fusion is the power source of the sun. Humanity has spent decades learning to replicate it on Earth. AI is a force multiplier for the investment in fusion, speeding the development of the complex physics modeling and needed digital engineering of fusion systems.

The vision is concrete and near-term: a U.S. fusion pilot plant delivering net electricity to the grid by 2030–2035 with significant scientific and engineering milestones demonstrated within the next one to two years. This is not a distant aspiration. It is an achievable national objective if the United States moves with the urgency the moment demands.

Why fusion, why now? Energy demand is surging, driven by the electrification of transportation, industry, and the explosive growth of AI data centers. Fusion offers a uniquely compelling answer:

- Fusion fuels are derived from abundant natural resources, produce no carbon emissions, and have no risk of meltdown or long-lived radioactive waste.
- On a per-reaction basis, fusion delivers one million times the energy density of chemical fuels.

- Whoever brings fusion to commercial scale first will anchor the energy infrastructure of the AI-driven economy and secure a durable strategic advantage over competitors.

The urgency is geopolitical as much as scientific. China is outspending the United States roughly two-to-one on public fusion research and has capitalized its national fusion company at over \$2 billion (Hiller and Hua 2024). The United States still leads overall, but the margin is narrowing. The window for decisive investment is open, and it will not remain so indefinitely.

Recent breakthroughs reset the timeline. In December 2022, the National Ignition Facility (NIF) at the U.S. Department of Energy’s (DOE) Lawrence Livermore National Laboratory (LLNL) achieved fusion ignition and energy gain for the first time in history, producing more energy out of the target than the lasers put in. That milestone has since been repeated ten times with a record yield of 8.6 megajoules in April 2025, which was more than four times the laser energy input (LLNL 2026a). Alongside NIF’s advances, a wave of private companies—pursuing tokamaks, stellarators, magnetic mirrors, inertial confinement, and

other approaches—has attracted over \$14 billion in cumulative investment (Fusion Industry Association 2026). The landscape of fusion R&D has changed radically in the last decade.

Why DOE and the Genesis Mission? DOE remains the essential foundation beneath this commercial wave. DOE's national laboratories lead the world in high energy density and nonlinear plasma physics, which no individual private company can replicate. The DOE *Fusion Science and Technology (FS&T) Roadmap* (U.S. DOE 2026) was developed with input from over 800 scientists and engineers across more than 15 private companies, 10 national laboratories, and 72 universities. It charts a Build–Innovate–Grow strategy to deliver public infrastructure to support private sector scale-up by the 2030s. Critically, the Genesis Mission, launched by executive order in November 2025, places AI-accelerated fusion among the nation's explicit science and technology priorities. The executive order established a sovereign national platform of supercomputers, scientific data, and foundation models in service of this goal (The White House 2025). Fusion is not only one of 26 challenges the Genesis Mission addresses, it is also among the most strategic.

DOE has invested billions in world-class physics codes, exascale high-performance computing (HPC) systems, and experimental facilities. Yet these resources operate largely in isolation, connected by human intermediaries who manually translate results from one domain to another. The rate-limiting step is not computational power but the time to orchestrate these resources into coherent, intent-driven campaigns.

Now with Genesis and frontier AI systems, there is an opportunity to accelerate fusion development by 10- to 100-fold through ecosystem-wide orchestration—coordinating physics codes, HPC resources, and experimental facilities across institutional boundaries. The transformational capability frontier AI brings to fusion is not just faster simulation or automated workflows, but intent-driven orchestration across the entire fusion ecosystem. This provides the ability to translate high-level scientific objectives into dynamically reconfigured, multi-institutional, multidomain campaigns that adapt in real time to results. This architecture extends beyond fusion to serve as a DOE-wide AI pathfinder.

Progress in the last 5 years, combined with aggressive and sustained investment in AI–fusion convergence, make it realistic to advance the National Academies of Sciences, Engineering, and Medicine's 2021 target for the first U.S. fusion pilot plant from 2035–2040 to 2030–2035 (NASEM 2021). The ambition horizon presented here is exactly that:

- Significant AI-enabled engineering milestones within one to two years.
- Major capability demonstrations by 2030.
- Grid-connected fusion power by 2035.

Pillar 1: Measurement and Data—Turning Decades of Data into Discovery

More than 100 experimental fusion machines have operated worldwide over the past six decades. The data they generated—including plasma diagnostics, actuator and sensor logs, spectroscopy, physics-model outputs, and images—represent an irreplaceable scientific record. The vast majority of it sits idle, siloed in facility-specific formats, never analyzed in aggregate. This is one of the most consequential untapped resources in the entirety of energy science. The immediate opportunity is to make this data accessible, linkable, and interoperable within a one-year horizon.

DOE's AI-fusion Digital Convergence Platform (DCP), a central component of the Genesis Mission and the *FS&T Roadmap* (U.S. DOE 2026), is designed to be a single national capability that integrates novel algorithms in high-performance computing codes, foundation models for plasma and materials science, neural networks, surrogate models, and digital twins (see Fig. 3.1, p. 20). The Fusion Energy Data Ecosystem and Repository (FEDER) will provide the data-stewardship infrastructure, built around FAIR (Findable, Accessible, Interoperable, and Reusable) principles, to make cross-facility learning possible for the first time.

Big opportunities abound. Princeton Plasma Physics Laboratory's (PPPL) 1990s-era Tokamak Fusion Test Reactor (TFTR) was the only deuterium-tritium-burning tokamak ever operated in the United States (Greenwald 2022). Rebuilding it today would cost




Fig. 3.1. The Integrative Genesis Mission Digital Convergence Platform. This AI-fusion convergence flywheel illustrates the feedback loop between data (simulation and theoretical data coupled with experimental data from more than 100 fusion facilities) and AI foundation models, digital twins, real-time plasma control, materials discovery, and power plant design.

billions of dollars and take years. Its data, however, can be recovered, linked to physics models, and embedded in a digital twin, thereby reclaiming most of its scientific value virtually and at a fraction of the cost. That is not a hypothetical future capability; it is a near-term deliverable (see Table 3.1, p. 21).

Measurement innovation is equally urgent. DOE’s Fusion Energy Science Advisory Committee’s *Basic Research Needs for Measurement Innovation Report* (FESAC 2025) identified critical gaps in both magnetic and inertial confinement fusion. For a future fusion power plant, diagnostics must withstand high radiation environments and function reliably over long operational pulses—requirements that go far beyond anything deployed in today’s experimental devices. For

inertial fusion, diagnostics must operate at high repetition rates, approaching 10 shots per second while surviving intense neutron and X-ray fluxes.

AI changes what is achievable. Multimodal AI models can reconstruct missing measurements from existing sensors, predict how plasma temperature and density will evolve, and effectively expand the diagnostic capability of a machine without adding hardware. A recent result from PPPL demonstrated this, inferring full plasma state from a subset of available diagnostics with striking fidelity (Jalalvand et al. 2025). This class of capability will be critical for operating future power plants with fewer, cheaper sensors while maintaining the measurement precision needed for safety and performance.

Table 3.1 Measurement and Data Implementation Path

Fusion Target Horizon	Milestones
Year 1	- Establish interoperability across major U.S. fusion facility datasets. - Begin training cross-facility foundation models on unified data infrastructure via the Genesis Mission Digital Convergence Platform and American Science Cloud.
Year 2	- Deploy AI-driven diagnostic reconstruction and predictive measurement tools on DIII-D National Fusion Facility and NIF. - Validate against held-out experimental data to assess accuracy.
2030	- Make an operational universal data-based simulator available to the full public–private fusion ecosystem. - Develop verification, validation, and uncertainty quantification (VVUQ) capabilities so they are mature enough to support credible extrapolation to fusion pilot plant conditions.
2035	Qualify radiation-hardened, AI-integrated diagnostic suites for first-generation power plant deployment.

Pillar 2: Modeling, Prediction, and Validation—AI as Engineering Multiplier

The deepest technical challenge in fusion engineering is multiscale physics. A fusion plasma spans six orders of magnitude in both space and time—from reactor-scale magnetic geometry down to the kinetic microstructures that drive turbulence and energy loss. No single simulation captures all scales simultaneously. First-principles codes that resolve the smallest scales require supercomputer runs lasting weeks. This central bottleneck has slowed the design cycle for fusion devices. AI breaks it open (see Table 3.2, p. 22).

AI can speed supercomputer calculations. The convergence of AI with fusion simulation is already producing historically significant results. LLNL researchers have developed AI surrogate models of turbulence and transport at the plasma edge that accelerate calculations by eight orders of magnitude—compressing computations that previously required hours into fractions of a millisecond (Li et al. 2025). This is not merely a convenience; it is the enabling technology for real-time control. A safety-critical plasma calculation that takes

tens of seconds cannot be used to prevent an instability that develops on millisecond timescales, but the same calculation running in under a millisecond can.

AI provides an innovation for controlling plasma instabilities. AI-enabled plasma control represents perhaps the most dramatic near-term breakthrough. A fusion plasma is held in place by magnetic fields and is prone to violent instabilities—tearing modes, edge-localized modes, and Rayleigh-Taylor fingers—that can develop and destroy confinement faster than any human can react. The challenge has been known for decades. The solution has only recently become achievable.

Deep reinforcement learning controllers have now stabilized live tokamak plasmas autonomously, commanding full magnetic coil arrays in real time (Degrave et al. 2022). More significantly, AI models have demonstrated the ability to forecast tearing-mode instabilities up to 300 milliseconds before onset and then autonomously steer the plasma away from disruption entirely, without human intervention. This result has been demonstrated across multiple tokamaks that the model was not trained on, establishing a degree of generalizability that is essential for deployment

Table 3.2 Modeling, Prediction, and Validation Implementation Path

Fusion Target Horizon	Milestones
Year 1	• Make AI disruption-avoidance controllers operational on DIII-D and KSTAR (Korea Superconducting Tokamak Advanced Research). • Validate cross-machine transfer. • Integrate into real-time plasma control systems at target facilities.
Year 2	• Train first fusion foundation model on unified cross-facility dataset. • Deploy the model as interactive, dynamically reconfigurable shot-design tool for experimental campaigns.
2030	• Universal digital twin operational, encompassing data from U.S. and allied international facilities, used to down-select among highest-potential fusion confinement approaches and guide infrastructure investment priorities.
2035	• Validated AI simulation suite with verification, validation, and uncertainty quantification (VVUQ) certification used as primary design tool for fusion plant engineering.

on future devices (see Fig. 3.2, p. 23; Seo et al. 2024; Jalalvand et al. 2025). Physics-informed deep learning has also been used to predict fusion ignition conditions at NIF with high accuracy (Spears et al. 2025), demonstrating that AI can close the loop between simulation and experiment even in the extreme conditions of inertial confinement fusion.

Universal digital twins provide value, speed progress. The universal digital twin is the strategic infrastructure that ties these capabilities together. The concept is straightforward and powerful: train an AI foundation model on the combined data of multiple fusion testbed facilities (experiment logs, physics-model outputs, diagnostic images, and actuator records) to produce a high-fidelity virtual simulator that any physicist or engineer can query. PPPL, in partnership with NVIDIA and IBM, is building this in the STELLAR-AI initiative (Simulation, Technology, and Experiment Leveraging Learning-Accelerated Research), an AI-optimized, fusion-centric supercomputing cluster designed to train the first generation of fusion foundation models. The prototype TokaMind multimodal transformer foundation model, trained on

tokamak plasma dynamics, represents the leading edge of a rapidly advancing field (Boschi et al. 2026).

The value of a universal digital twin compounds across the full pipeline. Fusion power plants will operate in plasma and materials regimes that no current experimental facility can reach—meaning every design decision for a first-generation plant involves extrapolation. A data-grounded digital twin with quantified uncertainty transforms that extrapolation from expert judgment into validated computational inference. The LLNL team has demonstrated a multi-agent AI system for integrated systems engineering design of inertial fusion energy power plants—taking high-level design goals through simulation, machine learning model optimization and validation, to computer-aided design—compressing what previously required months of expert iteration into days (LLNL 2026b).

Pillar 3: Making, Manufacturing, and Delivering—The Physical Path to the Grid

The scientific and computational breakthroughs described above are only as valuable as the physical




Fig. 3.2. AI Tearing-Avoidance Control at DIII-D National Fusion Facility. Researchers at the DIII-D National Fusion Facility tested a breakthrough method, which provides real-time plasma monitoring during fusion reactions and responsive adjustment of the magnetic confinement fields to avoid plasma tearing instabilities, overcoming a major hurdle to the tokamak approach for fusion energy. This method can optimize plasma containment and stability in response to existing conditions in real time. The algorithm is more flexible than traditional plasma control methods, eliminating the need for manual tuning and allowing quicker iteration on designs and scenarios. [Used under a Creative Commons Attribution 4.0 International License (CC BY 4.0) from Seo, J., et al. 2024. “Avoiding Fusion Plasma Tearing Instability with Deep Reinforcement Learning,” *Nature* **626**, 746–51. DOI:10.1038/s41586-024-07024-9. © Nature]

infrastructure that translates them into hardware. DOE’s Build–Innovate–Grow strategy charts a concrete infrastructure pathway across eight streams critical for a fusion power plant: plasma confinement

and performance, enabling technologies, exhaust and high-heat-flux testing, nuclear-effects testing, remote maintenance, blanket development and testing, fuel cycle development and testing, and high-performance

Table 3.3 Making, Manufacturing, and Delivering Implementation Path

Fusion Target Horizon	Milestones
Year 1	• Begin MPEX operations at ORNL. • Engage allied-nation blanket and tritium test facilities under TBDP strategy.
Year 2	• Finalize Nuclear Regulatory Commission fusion licensing framework. • Begin first SPARC (tokamak) plasma campaigns. • Make AI-driven materials-discovery pipeline operational for structural alloys and plasma-facing components.
2030	• Close Fusion Prototypic Neutron Source R&D gaps. • Complete IB-FCTF design. • Anticipate multiple private sector demonstration platforms (SPARC, Infinity One, Eos, and others) will be operating. • Employ universal digital twins to support power plant design decisions.
2035	• Launch first U.S. fusion pilot plant delivering net electricity. • Make domestic fusion supply chain operational. • Define path to nth-of-a-kind commercial deployment.

computing with advanced manufacturing and AI. The near-term horizon is ambitious and specific (see Table 3.3, this page).

Addressing the challenge of developing materials that can tolerate conditions in a fusion power plant. The most economically significant unresolved challenge in fusion engineering is materials. Fusion neutrons, which are produced in abundance by a burning plasma, cause damage to structural materials at a rate and character unlike anything in fission experience. The exceptional materials degradation caused by fusion neutrons is, as the *FS&T Roadmap* states, “one of the single largest factors limiting the economics and safety of fusion energy.”

Two knowledge gaps dominate:

- Can existing material classes survive fusion-relevant irradiation damage levels?
- What is required for an engineering materials database sufficient for power plant design?

DOE is pursuing a Fusion Prototypic Neutron Source (FPNS) as a high-priority facility with aggressive R&D on deuterium-lithium stripping, spallation, and true-fusion neutron source approaches. The Material Plasma Exposure eXperiment (MPEX) at Oak Ridge National Laboratory (ORNL), which is coming online within two to three years, will enable testing of plasma-material interactions under prototypic conditions that go beyond anything currently available in the United States.

Physics-based, multiscale computational modeling can now simulate material evolution under complex loading, irradiation, and tritium exposure histories, identifying candidate materials and failure modes before a single gram of material is irradiated. AI-guided materials discovery—learning process–property performance relationships from existing experimental and computational databases—is compressing the decades-long development cycle for new fusion-grade structural alloys and plasma-facing components into years.

Tritium fuel cycle and blanket technology. This represents another foundational gap on the path to

a self-sustaining fusion power plant. A fusion power plant must breed its own tritium fuel—lithium in a surrounding blanket absorbs fusion neutrons and produces tritium, which is then extracted, processed, and fed back to the plasma. No fusion blanket has yet been built and validated at scale. DOE’s Tritium-Blanket Development Platform (TBDP) strategy coordinates with allied-nation facilities including the LIBRTI Programme in the U.K., UNITY-1 in Japan, UNITY-2 in Canada, and CHIMERA in the U.K. to immediately begin closing low technology-readiness-level gaps. The long-term target is an Integrated Blanket and Fuel Cycle Test Facility (IB-FCTF) to be delivered within a 5- to 10-year window.

Private sector scale-up and supply chain. The DOE *FS&T Roadmap* explicitly identifies supply chain development as a critical public responsibility. Fusion power plants will require radiation-tolerant internal components manufactured at scale—high-temperature refractory metal-based structures, advanced laser optical components, gyrotron heating systems, pellet injection systems, and tritium blanket assemblies. None of these supply chains exists today in forms adequate for commercial deployment. DOE public-private partnership programs cosponsor construction of experimental capabilities with the private sector, state and local governments, philanthropies, and international partners. Eight companies are currently in the DOE Milestone Program including Commonwealth Fusion Systems, Xcimer Energy, Tokamak Energy, Type One Energy, Thea Energy, Zap Energy, Realta Fusion, and Focused Energy. Each is

developing technology roadmaps toward early-stage fusion power plant designs.

The Investment Imperative

The convergence of proven physics, an explosion of private investment, new AI capabilities, and renewed federal commitment has created a genuinely historic moment for fusion energy. The path to fusion power, once defined by fundamental scientific uncertainty, is now a set of solvable engineering challenges. The question is whether the United States moves with the speed and investment the moment demands or cedes the energy foundation of the next century to a competitor already spending twice as much.

The DOE *FS&T Roadmap* provides the fusion infrastructure strategy. The Genesis Mission provides the accelerating AI platform. Winning this race requires investing simultaneously in the essential physical fusion infrastructure and in the digital AI platform that acts as a force multiplier, transforming data into predictive power. The talent exists, distributed across national laboratories, universities, and a dynamic private sector ecosystem that has already demonstrated its capacity to move fast. What is required now is sustained, multiyear investment in the AI-fusion convergence infrastructure—digital twins, foundation models, data platforms, physical testbeds, and intent-driven orchestration—to compress a decade of engineering into 5 years. The dividends are abundant. Always-on power is the energy substrate of American competitiveness for the century ahead.

Chapter 4

Magnet Sovereignty

A 5-Year AI Mission To Rebuild U.S. Permanent-Magnet Capability

Authors

Nadya Mason (Lead)
University of Chicago

Makoto Gonokami
RIKEN

Eric Isaacs
Research Corporation for Science Advancement

Michael Mahoney
University of California–Berkeley;
International Computer Science Institute;
Lawrence Berkeley National Laboratory

David Siegel
Two Sigma

Additional Contributors

Kathryn Moler
Stanford University

Frank Alexander
Argonne National Laboratory

Allison Bennett Irion
Argonne National Laboratory

Karl Mueller
Ames National Laboratory

Prashant Singh
Ames National Laboratory

Michael Wang
Argonne National Laboratory

Magnet Sovereignty

A 5-Year AI Mission To Rebuild U.S. Permanent-Magnet Capability

The Bottom Line

Critical materials are the new oil—and the United States dangerously depends on a single foreign supplier. Permanent magnets made of critical rare-earth materials are essential to nearly every F-35 fighter jet engine, every electric- and hybrid-vehicle motor, every cell phone, every wind turbine, and every advanced medical scanner. China dominates the entire rare-earth material supply chain—from mining, refining, processing, and alloying to final magnet manufacturing (U.S. DOE 2022). When China restricts exports—as it has done and can do again at any time—the United States has no credible domestic or allied alternative. The four rare-earth elements (REEs) at the heart of the most powerful commercial magnets (neodymium, praseodymium, dysprosium, and terbium) have no ready substitutes, are produced in geopolitically concentrated deposits, and require processing capabilities the United States has largely allowed to atrophy.

“Magnet sovereignty”—or securing a domestic supply of critical magnetic materials and manufacturing capabilities—is an urgent challenge of our day. The goal is not to eliminate rare earths but to build multiple credible domestic pathways at once:

- Magnet compositions that use REEs more efficiently or tolerate domestic lower-purity feedstocks.

- Materials that reduce or eliminate the most supply-constrained elements.
- Recovery of rare earths already inside U.S. borders.
- The physical and institutional infrastructure to convert validated compositions into manufacturable U.S. products.

The Genesis Mission provides the opportunity and national strategy to achieve ambitious goals, including magnet sovereignty. To that end, the subcommittee suggests setting a 5-year target to identify, synthesize, and pilot high-performance permanent-magnet compositions codesigned for domestic manufacturability and end-of-life recyclability. AI and computational resources up to the zettascale are what make that speed achievable. The subcommittee proposes the Genesis Mission use AI to compress the materials-discovery and supply-chain-diversification cycles from decades to years through self-running laboratories for the fastest measurements, foundation models for dramatically improved analysis, and supply chain digital twins for more effective manufacturing and distribution. Realizing this vision requires not just discovery but crosscutting machine learning (ML) research as an integrative capability, purchasing commitments that make first-of-a-kind domestic production financeable, and a deliberate focus on building new supply




Fig. 4.1. Magnet Sovereignty Active-Learning Loop. AI enables integration of previously siloed measurement data; multidisciplinary experiments; and material sourcing, processing, and manufacturing. The iterative learning and insight sharing within this feedback loop holds promise for advancing the production of high-performance magnets from domestic feedstocks.

chains rather than feeding AI into existing high-concentration mining pipelines.

Why This Matters Now

The vulnerability is structural, not incidental. Even when rare-earth ore is mined domestically or by trusted partners, the separation, refining, and magnet-manufacturing steps that determine value can still depend on a single overseas pipeline. A disruption at that choke point ripples through defense procurement, grid modernization, transportation, and advanced manufacturing simultaneously. Conventional programs to rebuild domestic capacity operate on decade-long timescales. That pace is not fast enough.

Today, decisions about mineral resources, materials discovery, scientific ML methodology, manufacturing,

and supply chain investment are made from disconnected data, siloed expertise, and models that do not talk to each other. No one is optimizing the full system from ore to deployed product. For example, a geologist does not know what a separations chemist needs, a chemist does not know what a magnet manufacturer needs in terms of material impurity tolerance and scale-up, and a materials scientist does not have models that can draw strength from relevant data in very different domains.

AI makes integration possible for the first time: a single framework connecting mining, separations chemistry, materials science, process engineering, product design, life cycle sustainability, and investment strategy into one active-learning loop (see Fig. 4.1, this page). Data-first decisions can be made about not just which materials perform best, but which can actually be manufactured at scale from domestic feedstocks

and at acceptable environmental and economic cost. The aim is not a single better magnet. The aim is a domestic learning system that lets the United States repeatedly discover, make, qualify, recover, and improve magnets so that no single foreign choke point can again hold its industrial base hostage.

What Must Be Built

Achieving magnet sovereignty requires integrating three pillars: (1) AI-accelerated materials discovery, (2) AI-driven manufacturing, and (3) AI-enabled domestic materials recovery. This integration must operate in a continuous feedback loop where new materials need viable manufacturing processes to reach commercial scale, materials reclamation outcomes depend on alignment with those manufacturing developments, and insights from both feed back into materials research. At every stage, human experts must remain in the loop as active drivers of judgment, error correction, and capability development that AI alone cannot supply.

Pillar 1: AI-Accelerated Materials Discovery

The goal of the first pillar is to discover magnet compositions that free the United States from dependence on the most supply-constrained rare earths (Coe 2020). These materials should approach or exceed the performance of the best magnets available today and activate recycling pathways that allow the recovery and re-use of resources already captured in our domestic supply chain. Crucial to the discovery effort are three activities:

- Identifying materials that substantially reduce or eliminate dysprosium and terbium, or utilizing earth-abundant elements with viable domestic processing or chemistry manipulation pathways.
- Engineering neodymium magnet variants to tolerate the impurity profiles of domestic feedstocks.
- Improving AI-driven purification of domestic rare-earth feedstocks from ore, tailings, or recycled magnets.

A high-risk, high-reward opportunity is the exchange-spring magnet—a 30-year aspirational goal coupling a high-magnetization soft phase to a hard phase

that resists demagnetization (Kneller and Hawig 1991). The barrier is not finding the composition; it is manufacturing the required nanoscale two-phase microstructure at bulk scale. AI offers a path to simultaneously co-optimize anisotropy, magnetization, microstructure, and manufacturability, directing equal effort at the manufacturing barrier and composition search.

AI-accelerated materials discovery requires a shared magnetics data foundation, linking composition, impurity profiles, synthesis conditions, microstructure, and measured magnetic performance so that AI models learn from the full materials pathway rather than composition alone. Feedback-driven active learning lets experimental characterization iteratively refine computational sampling, so that each physical experiment sharpens the model rather than confirms what is already known.

AI-approximation models make searching vastly larger magnet compositional spaces tractable. One example is HydraGNN, a stand-in for expensive quantum-mechanical (DFT) calculations that speeds up materials screening nearly 1,000-fold (Lupo Pasino et al. 2025). Leveraging AI, cost-penalized optimization algorithms rank candidate materials against supply chain constraints, separability, manufacturability, and performance simultaneously. Additionally, AI agents based on multimodal large language models (LLMs) read scientific literature and patents to generate physically grounded hypotheses and identify unexplored design space.

Genuinely sharing the underlying physics across separations, manufacturing, and recycling enables a shared AI model that works across domains, drawing statistical strength from abundant data in one domain to make predictions in a data-poor one. Ground truth from DOE scientific user facilities—high-field magnetic characterization, synchrotron X-rays, and neutron scattering—provides the atomic-scale data on which reliable AI models depend. Tight coupling between autonomous laboratories, user-facility scheduling, and novel AI methods is essential for closing the discovery loop at speed.

Pillar 2: AI-Driven Manufacturing

U.S. magnet manufacturing today depends on unwritten expertise built up over decades, processes derived from trial and error, and quality control that happens after the fact rather than in real time. Each stage of manufacturing—melting, strip casting, powder production, alignment, pressing, sintering, heat treatment, and coating—can degrade final performance, and these effects compound at scale. The goal of this pillar is to replace that fragile, expertise-dependent system with AI-driven process control: manufacturing lines that adapt in real time to feedstock variability, flag defects before they compound, and accumulate institutional knowledge in models.

AI-informed process control enables tighter compositional control during melting, more consistent phase formation during strip casting, narrower grain-size distributions during jet milling, and more uniform grain-boundary structure across large-format sintered components. These advantages directly address persistent challenges in getting strip casters and related equipment to reliably produce target microstructures. Additional advances are possible by pairing digital-twin models of manufacturing lines with computer vision and nondestructive evaluation for in-line defect detection. Such capabilities provide real-time process control for variable feedstocks rather than the post hoc inspection common today.

The AI-driven manufacturing pillar requires a deliberate mix of modeling approaches rather than a collection of narrow point solutions. For example, domain-specific ML is useful for well-posed problems, while cross-domain scientific foundation models can leverage physics shared across separations, manufacturing, and recycling. Low-precision, hardware-aware computing reduces the energy footprint of model development and enables deployment directly at processing sites. This pillar also requires pilot-scale alloying, powder-production, and magnet-manufacturing lines instrumented for real-time learning so that lines adjust process parameters as data arrive. The objective is a manufacturing knowledgebase that accumulates nationally, improves over time, and transfers to additional facilities.

Pillar 3: AI-Enabled Domestic Materials Recovery

The United States has billions of dollars' worth of rare earths locked in discarded hard drives, decommissioned wind turbines, scrapped military hardware, mine tailings, and industrial waste streams. Almost none of it is recovered. The goal of this pillar is to build AI-enabled recovery of domestic critical materials that treats end-of-life magnets as a domestic ore body, recovering and reintroducing rare earths into U.S. production at a scale that meaningfully reduces import dependence (see Fig. 4.2, p. 31).

The first bottleneck in magnet recycling is not separations chemistry but the inability to rapidly identify, sort, and assess heterogeneous end-of-life product streams at economic scale (Zakeri and Tafaghodi 2025). Computer-vision systems can identify magnet-containing products automatically; predictive models can assess REE composition from limited sensor signals without destructive testing; and AI-guided disassembly can optimize recovery yield per unit time and cost, turning discarded products into a classified feedstock stream. Downstream, self-driving separations laboratories replace sequential trial-and-error with model-directed autonomous experimentation, targeting an order-of-magnitude reduction in experiments needed to reach a target recovery process.

DOE's life cycle analysis model (GREET) shows that even mature supply chains lose half or more of a critical element between mine and finished material, with losses concentrated at identifiable stages (Wang et al. 2026). Applying the same discipline to REE supply chains via AI-enabled digital twins would provide meaningful utilization-efficiency gains.

A further opportunity lies in moving separations chemistry beyond traditional equilibrium paradigms toward models that account for the nonequilibrium kinetics and interfacial transport that actually govern industrial-scale separations. Self-driving laboratories can accelerate this work by autonomously running experiments, collecting data, and testing AI-predicted candidate molecules for solvent extraction, with AI acting as an active partner in directing the next experiment rather than only interpreting the last one. Standardized AI-to-instrument interfaces can connect self-driving laboratories to shared high-performance




Fig. 4.2. Rare-Earth Element Supply Chain. AI has the potential to be a powerful optimization tool for recognizing, unlocking, and recovering previously unavailable rare-earth elements at economic scale from all steps of the supply chain for reintroduction into U.S. production.

computing (HPC) resources across facilities, letting separations efforts at different sites draw on the same models and data rather than duplicating effort. Recovered material must feed domestic alloying and magnet manufacturing—continually and not as a one-off demonstration—which requires pilot-scale recycling capacity and reliable demand for domestically recovered rare earths.

The AI Advantage

Without AI, the search for new magnet compositions is bound by the experiments a team can physically run and the composition space accessible to human intuition. AI breaks both constraints, enabling fast approximation models that make vast composition spaces tractable, active learning that ensures each experiment maximally sharpens the model, and cost-penalized optimization that embeds supply chain and manufacturability constraints directly into candidate ranking. The result is a search that is simultaneously broader,

faster, and more strategically targeted than any conventional program.

ML-based surrogates or quantum computers coupled with HPC may also offer a near-term path to more accurate electronic-structure calculations for magnetically complex systems where standard quantum chemistry methods struggle. In addition, specialized quantum algorithms can refine energy landscapes that inform AI candidate ranking.

In manufacturing, AI replaces tacit, difficult-to-transfer expertise for diverse feedstocks with data-driven process control that adapts in real time. For example, neural digital twins ingest sensor data and update process parameters continuously, computer-vision systems flag defects at machine speed, and ML models trained on diverse feedstocks generalize across input variability.

AI-driven fast approximation models can compress qualification dataset-building from years of physical testing to just months of model-guided prioritization,

all while meeting or exceeding standards with a faster, more transparent, and more predictive evidence base.

In reclaiming critical materials, AI converts heterogeneous end-of-life streams—too complex to process economically by hand—into a data challenge. Computer vision and predictive composition models replace manual inspection, AI-guided disassembly maximizes recovery yield automatically, and self-driving laboratories reach optimized separation protocols in a fraction of the conventional time. Across all three pillars, this is not automation of existing practice, it is replacing slow, post hoc control with fast, adaptive, and predictive control.

Effective coordination is as important as any individual AI capability. The full system spans geology, separations chemistry, materials science, ML, manufacturing, product design, recycling, and economics—no single laboratory, company, or agency can optimize it alone. The mission should operate as an integrated network with shared data standards, protected pathways for proprietary information, common AI-to-instrument interfaces, and clear ownership for moving candidates from discovery to pilot production. AI makes that integration achievable for the first time; realizing it requires institutional commitment to match the technical ambition.

Table 4.1 Magnet Sovereignty Implementation Path

Magnet Target Horizon	Milestones
Year 1	- Stand up a national magnetism data foundation. - Deploy literature- and patent-mining AI agents to map the known design space and identify underexplored compositions. - Connect autonomous laboratories to DOE user facilities. - Launch pilot projects in recycling detection and manufacturing-data capture. - Establish end-user requirements at the outset so the mission optimizes for real products, not scientific novelty alone.
Years 1 to 2	- Demonstrate an impurity-tolerant neodymium magnet composition viable with domestic-grade or recycled feedstock. - Deploy AI-based process-control models on at least one pilot-scale strip-casting or sintering line, demonstrating measurable improvement in grain-size or phase-uniformity consistency. - Deploy AI-based detection and sorting pilots at one or more existing e-waste, mining, or end-of-life processing facilities.
Years 2 to 4	- Integrate discovery, manufacturing, and recycling into closed loops. - Move candidate magnets from gram-scale synthesis to pilot-scale quantities. - Deploy digital twins on at least one pilot strip-casting, powder-production, or sintering line. - Demonstrate a self-driving separations laboratory achieving order-of-magnitude time reduction in experiments-to-target for at least one REE recovery process.
Years 3 to 5	- Validate one or more high-performance magnet systems at pilot scale, including at least one REE-light or REE-free candidate. - Validate a full collection-to-recovered-material pipeline, with recovered REE content feeding a domestic alloying or magnet-manufacturing line. - Publish a replicable playbook for expanding capacity across additional facilities and regions.

Magnet Sovereignty Implementation Path

The roadmap for designing and manufacturing high-performance permanent magnets at pilot scale using domestic feedstocks includes targeted milestones across 5 years. These milestones demonstrate how AI can significantly accelerate solutions to open challenges in materials discovery, manufacturing, and recovery (see Table 4.1, p. 32).

What Success Looks Like

Magnet sovereignty is achievable if treated as a mission: integrated, measurable, AI-enabled, tied to real manufacturing, and sustained until the first qualified domestic magnets enter use.

By the end of 5 years, success should be visible in practical terms. The target is a 50-fold increase in discovery efficiency, a 10-fold reduction in experiments needed, model confidence exceeding 80% before physical validation is prioritized, and greater than 75% autonomy for routine experimental cycles.

The United States should have:

- Pilot-scale production of high-performance permanent magnets designed for domestic feedstocks.
- AI-controlled manufacturing that reduces scrap and improves lot-to-lot consistency.

- A transformative collection-to-recovered-material recycling pipeline.
- Qualification protocols that end users trust.
- A shared data and modeling infrastructure that keeps improving after the first demonstration.

Broader Scientific Impact

The payoff extends well beyond magnets. The same methods—active learning, autonomous experimentation, surrogate modeling, digital twins, and supply chain-aware optimization—apply directly to batteries, semiconductors, catalysts, structural alloys, and many other systems where performance must be balanced against criticality or manufacturability. The orchestration layer connecting self-driving laboratories across institutions establishes shared data standards, model interfaces, and compute access that lower the barrier for the next autonomous-laboratory effort in any domain.

While capturing the supply chain remains strictly domestic, engaging with cross-border open-science initiatives in basic AI and materials modeling ensures that the U.S. research enterprise continuously integrates the world's finest scientific insights, ultimately sharpening our competitive edge.

Magnet sovereignty is a focused, urgent test case; the AI methods and data infrastructure it requires are direct investments in magnetic technology as well as down payments on the broader U.S. AI-for-science research enterprise.

Appendix A

Office of Science Advisory Committee Members

Chair

Persis Drell, *Stanford University*

Members

Scott DeBoer, *Micron Technology*

Drew Endy, *Stanford University*

Cynthia Friend, *The Kavli Foundation*

Suresh Garimella, *University of Arizona*

Makoto Gonokami, *RIKEN*

Anna Grassellino, *Fermi National Accelerator Laboratory*

Supratik Guha, *University of Chicago and Argonne National Laboratory*

Eric Isaacs, *Research Corporation for Science Advancement*

Lara Jehi, *Cleveland Clinic*

Pushmeet Kohli, *Google DeepMind*

Tammy Ma, *Lawrence Livermore National Laboratory*

Michael Mahoney, *University of California–Berkeley, International Computer Science Institute,
and Lawrence Berkeley National Laboratory*

Nadya Mason, *University of Chicago*

Mark Papermaster, *AMD*

Mayly Sanchez, *Florida State University*

Martin Schmidt, *Rensselaer Polytechnic Institute*

Edward Seidel, *University of Wyoming*

David Siegel, *Two Sigma*

David Spergel, *Simons Foundation and Flatiron Institute*

A.N. Sreeram, *Dow*

Derek Sutherland, *Realta Fusion*

Appendix B

References

- Boschi, T., et al. 2026. “TokaMind: A Multi-Modal Transformer Foundation Model for Tokamak Plasma Dynamics,” *ArXiv*. DOI:10.48550/arXiv.2602.15084.
- Coe, J. M. D. 2020. “Perspective and Prospects for Rare Earth Permanent Magnets,” *Engineering* **6**(2), 119–31. DOI:10.1016/j.eng.2018.11.034.
- Degrave, J., et al. 2022. “Magnetic Control of Tokamak Plasmas Through Deep Reinforcement Learning,” *Nature* **602**, 414–19. DOI:10.1038/s41586-021-04301-9.
- Duman, E. 2025. “Global Manufacturing Capacity for Active Pharmaceutical Ingredients Remains Concentrated.” USP Quality Matters. Accessed July 9, 2026. <https://qualitymatters.usp.org/global-manufacturing-capacity-active-pharmaceutical-ingredients-remains-concentrated>.
- FESAC. 2025. *Basic Research Needs for Measurement Innovation Report*. Eds. Delgado-Aparicio, L. F., and S. P. Regan. Fusion Energy Sciences Advisory Committee. <https://www.pppl.gov/basic-research-needs>.
- Fusion Industry Association. 2026. *The Global Fusion Industry in 2026*. <https://www.fusionindustryassociation.org/fusion-industry-reports>.
- Greenwald, J. 2022. “How the Princeton Plasma Physics Lab Contributed to the New World Record in Clean Fusion Energy.” Princeton University. <https://www.princeton.edu/news/2022/02/14/how-princeton-plasma-physics-lab-contributed-new-world-record-clean-fusion-energy>.
- Hiller, J., and S. Hua. 2024. “China Outspends the U.S. on Fusion in the Race for Energy’s Holy Grail.” *Wall Street Journal*. <https://www.wsj.com/world/china/china-us-fusion-race-4452d3be>.
- Hopkins, A. L., and C. R. Groom. 2002. “The Druggable Genome,” *Nature Reviews Drug Discovery* **1**(9), 727–30. DOI:10.1038/nrd892.
- Jalalvand, A., et al. 2025. “Multimodal Super-Resolution: Discovering Hidden Physics and Its Application to Fusion Plasmas,” *Nature Communications* **16**, 8506. DOI:10.1038/s41467-025-63492-1.
- Jeffrey, C. J. 2023. “Current Successes and Remaining Challenges in Protein Function Prediction,” *Frontiers in Bioinformatics* **3**, 1222182. DOI:10.3389/fbinf.2023.1222182.
- Jumper, J., et al. 2021. “Highly Accurate Protein Structure Prediction with AlphaFold,” *Nature* **596**, 583–89. DOI:10.1038/s41586-021-03819-2.
- Kneller, E.F., and R. Hawig. 1991. “The Exchange-Spring Magnet: A New Material Principle for Permanent Magnets,” *IEEE Transactions on Magnetics* **27**(4). DOI:10.1109/20.102931.
- Li, N. M., et al. 2025. “Exploring the Transition from Continuous Turbulence Fluctuations to Bursting ELMs in High SOL Density Regimens,” *Nuclear Fusion* **65**, 076023. DOI:10.1088/1741-4326/ade0d0.
- Liu, T.-L., et al. 2018. “Observing the Cell in its Native State: Imaging Subcellular Dynamics in Multicellular Organisms,” *Science* **360**(6386), eaaq1392. DOI:10.1126/science.aaq1392.
- LLNL. 2026a. “Target Breakthrough Enabled Fusion Record at NIF.” Lawrence Livermore National Laboratory. <https://lasers.llnl.gov/news/target-breakthrough-enabled-fusion-record-nif>.
- LLNL. 2026b. “LLNL Showcases AI-Enabled Science, National Security, and Energy Innovation at AI+ Expo.” Lawrence Livermore National Laboratory. <https://www.llnl.gov/article/54416/llnl-showcases-ai-enabled-science-national-security-energy-innovation-ai-expo>.
- Lupo Pasini, M., et al. 2025. “Scalable Training of Trustworthy and Energy-Efficient Predictive Graph Foundation Models for Atomistic Materials Modeling: A Case Study with HydraGNN,” *The Journal of Supercomputing* **81**, 618. DOI:10.1007/s11227-025-07029-9.
- McKinsey Global Institute. 2020. *The Bio Revolution: Innovations Transforming Economies, Societies, and Our Lives*. https://www.mckinsey.com/~media/mckinsey/industries/life%20sciences/our%20insights/the%20bio%20revolution%20innovations%20transforming%20economies%20societies%20and%20our%20lives/may_2020_mgi_bio_revolution_report.pdf.
- NASEM. 2020. *Safeguarding the Bioeconomy*. National Academies of Sciences, Engineering, and Medicine. The National Academies Press, Washington, DC. DOI:10.17226/25525.
- NASEM. 2021. *Bringing Fusion to the U.S. Grid*. National Academies of Sciences, Engineering, and Medicine. The National Academies Press, Washington, DC. DOI:10.17226/25991.
- NobelPrize.org. 2024. “Nobel Prize in Chemistry 2024.” Accessed July 10, 2026. <http://www.nobelprize.org/prizes/chemistry/2024/press-release>.
- NSCEB. 2025. *Charting the Future of Biotechnology: An Action Plan for American Security and Prosperity*. National Security Commission on Emerging Biotechnology. <https://www.biotech.senate.gov/final-report/chapters>.

- NSCEB. 2026. “U.S. vs. China: Government Action on NSCEB Recommendations.” National Security Commission on Emerging Biotechnology. Accessed July 9, 2026. <https://www.biotech.senate.gov/us-vs-china-govt-action-on-nsceb-recs>.
- Ort, D. R., et al. 2015. “Redesigning Photosynthesis To Sustainably Meet Global Food and Bioenergy Demand,” *Proceedings of the National Academy of Sciences of the USA* **112**(28), 8529–36. DOI:10.1073/pnas.1424031112.
- Sedzro, T. 2026. “Positioning for Growth: Navigating China’s Biomanufacturing Industry.” *China Briefing*. Accessed July 9, 2026. <https://www.china-briefing.com/news/chinas-biomanufacturing-industry-opportunities>.
- Seo, J., et al. 2024. “Avoiding Fusion Plasma Tearing Instability with Deep Reinforcement Learning,” *Nature* **626**, 746–51. DOI:10.1038/s41586-024-07024-9.
- Spears, B. K., et al. 2025. “Predicting Fusion Ignition at the National Ignition Facility with Physics-Informed Deep Learning,” *Science* **389**, 727–31. DOI:10.1126/science.adm8201.
- The White House. 2025. “Launching the Genesis Mission. Executive Order 14363.” <https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission>.
- Uberbacher, E. C., and R. J. Mural. 1991. “Locating Protein-Coding Regions in Human DNA Sequences by a Multiple Sensor–Neural Network Approach,” *Proceedings of the National Academy of Sciences of the USA* **88**(24), 11261–65. DOI:10.1073/pnas.88.24.11261.
- UNEP-CCC. 2025. *Climate Technology Progress Report 2025: Advancing Biobased Technologies in the Bioeconomy*. United Nations Environment Programme Copenhagen Climate Centre. Copenhagen, Denmark. <https://unepccc.org/publications/the-climate-technology-progress-report-2025>.
- U.S. DOE. 1987. *Report on the Human Genome Initiative, Prepared by the Health and Environmental Research Advisory Committee*, DOE/ER-0382. U.S. Department of Energy. <https://www.osti.gov/servlets/purl/6746220>.
- U.S. DOE. 2021. *Biological Systems Science Division Strategic Plan*, DOE/SC-0205. U.S. Department of Energy Office of Science <https://genomicscience.energy.gov/2021bssdstrategicplan>.
- U.S. DOE. 2022. *Rare Earth Permanent Magnets: Supply Chain Deep Dive Assessment*, DOE/OP-0008. U.S. Department of Energy Office of Policy. <https://doi.org/10.2172/1871577>.
- U.S. DOE. 2025. “U.S. Department of Energy Announces Selectees for \$107 Million Fusion Innovation Research Engine Collaboratives, and Progress in Milestone Program Inspired by NASA.” U.S. Department of Energy. <https://www.energy.gov/articles/us-department-energy-announces-selectees-107-million-fusion-innovation-research-engine>.
- U.S. DOE. 2026. *Fusion Science and Technology Roadmap*. U.S. Department of Energy Office of Science. <https://www.energy.gov/documents/fusion-science-and-technology-roadmap>.
- Wang, M., et al. 2026. [Computer software] R&D GREET 2025 Rev.1©. Argonne National Laboratory. <https://doi.org/10.11578/GREET-Net-2025Rev1/dc.20260610.2>.
- Wetterstrand, K. A. 2022. “DNA Sequencing Costs: Data from the NHGRI Genome Sequencing Program (GSP).” National Human Genome Research Institute. Accessed July 9, 2026. <https://www.genome.gov/sequencingcostsdata>.
- Zakeri, A., and L. Tafaghodi. 2025. “A Review of the Current Progress in High-Temperature Recycling Strategies for Recovery of Rare-Earth Elements from Magnet Waste,” *Journal of Sustainable Metallurgy* **11**, 88–113. DOI:10.1007/s40831-025-01010-9.

Appendix C

Acronyms and Abbreviations

AI	artificial intelligence	KBase	DOE Systems Biology Knowledgebase
BER	Biological and Environmental Research program	KSTAR	Korea Superconducting Tokamak Advanced Research
CASP	Critical Assessment of Structure Prediction	LLM	large language model
cryo-ET	cryo-electron tomography	LLNL	Lawrence Livermore National Laboratory
DIII-D	DIII-D National Fusion Facility	ML	machine learning
DBTL	design–build–test–learn	MPEX	Material Plasma Exposure eXperiment
DCP	Digital Convergence Platform	NASEM	National Academies of Sciences, Engineering, and Medicine
DFT	Density Functional Theory	NIF	National Ignition Facility
DOE	U.S. Department of Energy	ORNL	Oak Ridge National Laboratory
EMSL	Environmental Molecular Sciences Laboratory	PPPL	Princeton Plasma Physics Laboratory
FAIR	Findable, Accessible, Interoperable, and Reusable	REE	rare-earth element
FEDER	Fusion Energy Data Ecosystem and Repository	SC	Office of Science
FESAC	Fusion Energy Sciences Advisory Committee	SCAC	Office of Science Advisory Committee
FFRDC	Federally Funded Research and Development Center	STELLAR-AI	Simulation, Technology, and Experiment Leveraging Learning-Accelerated Research
FPNS	Fusion Prototypic Neutron Source	TBDP	Tritium-Blanket Development Platform
FS&T	Fusion Science and Technology	TFTR	Tokamak Fusion Test Reactor
HPC	high-performance computing	TRY&S	titer, rate, yield, and scale
IB-FCTF	Integrated Blanket Fuel Cycle Test Facility	VVUQ	verification, validation, and uncertainty quantification
JGI	DOE Joint Genome Institute		

