

Testimony of Rick Stevens  
Associate Laboratory Director for Computing, Environment, and Life Sciences  
Argonne National Laboratory

Senate Energy & Natural Resources Committee  
Hearing to examine recent advances in artificial intelligence and the Department of Energy's  
role in ensuring U.S. competitiveness and security in emerging technologies

September 7, 2023

Chairman Manchin, Ranking Member Barrasso, and Members of the Committee, thank you for the opportunity to participate in today's discussion about the important role the Department of Energy (DOE) and its national laboratories have in advancing a national, trustworthy, reliable artificial intelligence (AI) capability that will accelerate science, transform innovation, and underpin our future economic and national security.

I am Rick Stevens, the Associate Laboratory Director responsible for Computing, Environment, and Life Sciences research at Argonne National Laboratory and a Professor of Computer Science at the University of Chicago. For over 30 years, my work at Argonne has focused on advancing computer science to tackle grand challenges across many science domains, from fusion energy to medicine. I have helped develop ever-more-powerful generations of supercomputers, culminating in the current generation of exascale supercomputers that are coming online now.

Over the last four years, I have had the fortune to work with my colleagues at the DOE laboratories to organize seven "town hall" meetings on the use of AI to advance science, energy, and national security. Over 1,300 scientists and engineers from all 17 national laboratories, more than 30 universities, and dozens of companies participated. In these meetings we collectively thought broadly and aggressively about how advanced AI systems might be developed over the next ten years to accelerate scientific research, energy technology development, and improve national security.

What has become clear as a result of these town hall meetings and the progress we have made in AI research at the national labs, combined with the speed with which AI tools like ChatGPT are being adopted by the public, is that *AI will completely transform science and society*.

AI is in the same class of transformative technologies as electricity and computers. Both were revolutionary, although they weren't necessarily recognized as such at the beginning. And the beginning is where we are with AI today.

Like all transformative technologies, there are both immense opportunities and risks associated with widespread use of AI. I believe it is imperative that the United States lead the world in the development of advanced AI systems for scientific and national security applications, and that an effort on the scale of the Manhattan Project will be needed for us to do so. I appreciate the

opportunity today to highlight the opportunities and risks associated with AI, and outline how the national labs and DOE are uniquely positioned to build an advanced AI capability for the nation.

### **The opportunities of AI for science and national security—and the risks**

AI is a tool for the mind. An amplifier of thought, a super-smart assistant, and perhaps someday soon, a powerful concierge for nearly all aspects of life and work. Unlike the internet, which enables humans to find information, communicate with each other, and transact business, AI is much more.

Trained on ever-larger ensembles of human and machine-generated knowledge, AI is capable of integrating and synthesizing knowledge and carrying out self-directed sequences of steps to solve problems in response to human prompts or questions.

Generative AI large-language models (LLMs)—such as Open AI’s ChatGPT, Google’s Bard, Anthropic’s Claude2 and other generative systems for images, music, speech, and computer code—have made AI widely available to the public. These new AI interfaces are immensely popular. ChatGPT gained 1 million users in five days and 100 million in two months, making it one of the fastest technologies ever to be adopted. We are seeing the beginning of a vast transformation of how people will work with, and interact with, computers and the internet.

Imagine a world where every person has 24/7 access to an incredibly smart assistant that has been trained on the sum of human knowledge. Your AI assistant will know a hundred human languages, all recorded history and literature, the bible in Hebrew and Greek, non-classified and non-proprietary science and engineering, political theory, military tactics and strategy, popular culture, the locations and relationships between companies, and much more.

You will no longer have to search the internet to find information and read through the resulting pages to synthesize answers. With AI assistants, you pose a query (much like we talk or text today), and the AI assistant does the searching and synthesizing. It writes the paper, answers the questions in your homework, plans your vacations, writes the employee evaluation, summarizes the financial market, or writes your computer code. In the future, we can even envision an AI that’s perceptive enough to ask questions before you realize you have them, just by listening or watching.

AI is beginning to transform some science and technology fields. AI systems are already being used to explore ways to improve the control of nuclear fusion reactors, increasing stable plasma burn times, and dynamically adjusting controls to avoid plasma disruptions. AI is also being used to accelerate scientific simulations by augmenting and/or replacing the traditional numerical mathematics approaches, in some cases providing massive speedups on existing computer hardware. Early efforts have achieved factors of 100 or more for problems such as molecular electronic structure and weather prediction, but many more areas are ripe for exploration. AI models are already writing code and will soon excel at it, even for quantum computers.

There are enormous opportunities for AI to accelerate discovery in the basic sciences, from new materials for energy applications such as better batteries that require less rare earth minerals, to new types of polymers that are optimal for each application yet can be reprocessed and recycled without loss of performance.

AI coupled with robotics can be used to improve throughput for experimental science many orders of magnitude, leading some to imagine future AI driven “science factories” that could be used to solve problems in many domains, from developing new drugs for infectious diseases and cancer to exploring new types of semiconductors.

AI can also address key challenges in software development, from improving the productivity of programmers for scientific computing, to optimizing performance of codes for our supercomputers. This is important to DOE and the national labs since we maintain over 10 billion lines of code in our laboratories, and it will also be critical for the computing industry.

Trained and tuned AI systems will be able to help us design better and more energy-efficient chips for future supercomputers and optimize the hundreds of processing steps used to manufacture chips in US-based semiconductor factories. AI-based control systems will improve the reliability and stability of the electrical grid, by adjusting operating controls quickly and precisely to balance increasing diversity of energy sources and changes in demand.

AI will improve overall scientific and technical productivity, in some cases speeding science discovery up by 10 times or more. And it will also change the nature of how science is conducted, with AI assistants playing important roles as sounding boards for new ideas and as planners and doers for experiments.

On the cautionary side, AI will pose new risks and threats. As an amplifier of human skills and capabilities, AI has the ability to greatly speed scientific discovery and improve productivity, and it can be used for good and bad purposes. To date, most AI efforts undertaken by industry and other countries have been generally visible, and the scientific community becomes aware quickly when new models are created. In the future, some AI models will be built by large groups working together largely in the open, while others will be built in secret. A small group working in secret with sufficiently powerful AI tools could develop a novel chemical, biological, or cyber threat. We will need to transform how we manage the risks posed by bad actors using the same AI tools we are using to improve science and advance society.

We also need to recognize the fierce international competition that is already underway to build the most capable AI systems. Powerful AI models are already emerging around the globe, including Falcon from Abu Dhabi. In China, hard numbers for public investment in AI are difficult to obtain, but they are estimated at \$2-\$9 billion a year now, up to \$26 billion by 2026. My other colleagues on the panel could expand upon the competition with China in much more detail. Whoever leads the world in AI will lead in science, innovation, energy, and security, and other countries are already positioning themselves for the competition.

The time is now for our country to take the lead in this race, and we believe DOE and the national laboratories have an integral role to play.

### **Why build specialized AI systems for science, energy, and security?**

One could ask why the government or DOE need to develop their own AI capability. Why not just use commercial AI systems like GPT-4 to do the same work?

I believe that researchers should be using commercial AI systems as much as possible, when they are the best tool and when the open nature of their use is not a problem. As commercial AI companies develop and deploy enterprise versions that have more protections of user data, these models should be adopted for use in appropriate settings.

However, we need to go beyond commercial systems in these important ways.

- 1) **Specialized data.** Commercial systems will never be deeply trained on the vast quantities and diverse types of data that are currently being managed by DOE laboratories. DOE has thousands of times more scientific data than the largest AI models are currently trained on. This data comes from experiments and from computer simulations and spans dozens of scientific fields. The scientific community needs AI systems that deeply understand scientific data, scientific processes, and scientific reasoning. DOE needs AI systems that have mission-related datasets. For national security applications, DOE needs models that can operate on that type of data in a classified environment.
- 2) **Trust and reliability.** Many of our use cases for AI (scientific data analysis, mission-critical decisions, operations of complex and life-critical systems) require high levels of accuracy in predictions. They also require reliable models whose training data is known, and whose outputs can be trusted and validated at scale.
- 3) **Security.** Whether for national security applications or sensitive science (such as understanding the risks posed by third-party models) we need to have the most powerful AI models in secure enclaves where they can be developed, used, and evaluated. An important class of uses of advanced AI will be to manage, control, and analyze experiments and large-scale simulations. In these roles the AI will need to be tightly integrated with robotic laboratories and supercomputing infrastructures, often in complex and secure environments.

### **Building the most advanced AI systems is truly a large-scale activity**

The DOE national laboratories were founded out of the Manhattan Project—our nation’s first example of leveraging teams of thousands of scientists and engineers, stationed at centers across the country, to accomplish a Herculean scientific endeavor in a very short period of time.

State-of-the-art AI systems are difficult to build, requiring millions of dollars of computing time and large teams of researchers that only exist in relatively few institutions. Building an advanced AI capability for science, energy, and national security—and doing it before, and better than, other countries—will require a project of similar scale and investment as the Manhattan Project.

World-leading AI research groups operate at the scale of 1,000 researchers and developers working on a small number of models and systems. Even at this scale, with access to sufficient computing power and large-scale data, these teams produce a major model system on an annual or biannual pace. It is not unusual to take months to initially train a model, many more months to refine the model and to evaluate it and improve its safety and alignment, and many more months after that to produce versions for specific purposes. Once models are developed, they must be deployed and refined as they ramp up usage and become integrated into the production environment.

We expect that DOE national labs will be a primary developer of AI models for scientific and national security uses. Some models developed by DOE laboratories will be widely used by the scientific community, perhaps hosted by laboratories or in partnership with companies. Others will be deployed in secure environments. A single foundational science model could be deployed for hundreds or thousands of downstream applications, and derivative models could be produced for even more.

As a trusted neutral party, DOE can also play a major role in development of the science of “AI risk management,” especially when it involves assessing the risks of third-party models using scalable methods (i.e., using AI systems and large-scale computing to automate as much of risk assessment as possible). By partnering with other agencies and leveraging its expertise in scalable computing, DOE can play a major role in developing the technical means for AI risk assessment and determining fitness for purpose. These techniques could be made available to regulators with responsibility for AI regulation and assessment.

### **A possible model for a national AI project for science, energy, and national security**

A new initiative by DOE and the national laboratories could be created to position the United States as the undisputed leader in AI research, development, and deployment for science, energy, and national security.

Through a committed partnership between the DOE’s Office of Science, the DOE National Nuclear Security Administration, and the DOE’s applied energy programs, a coordinated effort could build the world’s most powerful AI systems explicitly targeting accelerating scientific research and technology development. This AI leadership would be enabled by harnessing the DOE’s exascale computing resources and its integrated data and research infrastructure supporting the world’s largest collection of national user facilities and collections of scientific

data. This capability could allow DOE to bring AI research challenges via the largest scientific and mathematics workforce under one organization.

A network of groundbreaking AI research clusters could be established, built on DOE's unparalleled ecosystem of high-performance computing capabilities and research facilities. This network would operate as public-private partnerships, fusing the expertise and perspectives of national labs, academia, and industry leaders. It would also tap into the thousands of scientists, mathematicians, and engineers within the DOE lab complex to tackle multifaceted challenges requiring integrated, multidisciplinary thinking. By bringing together interdisciplinary teams, innovation cycles would be accelerated, ideas rapidly translated into deployable solutions, and insights and models broadly shared across institutional boundaries.

A primary research goal for this national effort would be pioneering advanced AI techniques for science and engineering tasks that push beyond human reasoning and learning capabilities. Progress in this direction would enable AI systems to flexibly adapt to novel contexts and tasks by building causal, intuitive understandings of the world. These advanced AI systems would become key assistants and autonomous tools for accelerating scientific research.

The effort would also emphasize developing responsible and trustworthy AI that balances robust performance with safety and security. This involves mathematical methods, software frameworks, datasets, and testing protocols to ensure model fairness, interpretability, reliability, and integrity. AI systems would be designed with transparency, auditability, and human oversight in mind from the start. The program would additionally support fundamental research into AI safety, aiming at avoiding harmful behaviors in future advanced systems.

The effort's scale and visibility could position DOE to lead national and global conversations on managing risks posed by ever-improving AI models. DOE and national lab researchers would collaborate with academic, industry, and government partners to pioneer frameworks and best practices for addressing AI risks to individuals and groups. The effort would work to establish international norms and deterrents for global risks from potential adversarial uses of advanced AI.

The effort would target developing world-class AI-enabling technologies to amplify returns on investment in fundamental AI research. Future AI systems will require advances in parallel computing hardware and software, new algorithms to make training and evaluation more data efficient and advances in data-oriented computing methods.

Science-focused AI systems would accelerate discovery and optimization across disciplines like materials science, particle physics, molecular biology, chemistry, and environmental sciences. Advanced software libraries, execution frameworks, high-performance workflows and cloud-based development environments will lower barriers to applying AI techniques. Specialized AI hardware architectures and low-power processing units will enable deploying high-performance AI affordably at scale.

DOE's unparalleled national lab network and high-performance computing resources will combine to create an integrated data and experimentation ecosystem. By pooling the exabytes of observational data, experimental results, and simulation output datasets into curated repositories with common standards, AI training and discovery would be empowered at a scale exceeding any academic or industrial counterpart. Tight integration of edge devices, high-speed networks, and cloud compute resources will support efficient distributed workflows.

Industrial partnerships and investments in AI hardware research and testbeds—aimed at achieving 100-fold improvements in performance per watt over ten years—would continuously improve AI hardware optimized for performance, scalability, cost efficiency, and sustainability. These solutions would expand access to capable AI resources across public and private sectors.

Partnerships with universities and other agencies would expand, deepen, and improve the diversity of the AI workforce, with a special emphasis on the AI for science and engineering uses and the connection of AI to the high-performance computing and modeling and simulation community.

Finally, this national effort would enable partnerships with other federal agencies through interdisciplinary programs, personnel exchanges, and coordinating strategies. These collaborative efforts will enhance the federal ecosystem's ability to develop and deploy advanced AI systems. This holistic approach would synergize national efforts toward responsible leadership in artificial intelligence.

## **Conclusion**

AI is rapidly becoming the most important tool in the scientific and technical toolbox. It is quite possible that by leveraging sufficiently powerful AI we could make a century's worth of progress in 10 years. I strongly believe that we must explicitly commit to building an advanced AI capability for science, energy, and national security. DOE oversees an unmatched ecosystem of national laboratories and high-performance computing capabilities that together constitute an ideal environment for creating and deploying advanced AI systems.

Whether it is the use of next generation AI models to advance theory in physics, or to synthesize knowledge across scientific literature and databases for a cure for a rare disease, or more specialized AI models exploring novel materials for future microelectronics, three things are clear.

Whoever leads the world in AI for science will lead the world in scientific discovery and will have a head start in the translation of discoveries to products that expand our economy and address modern needs, securing the innovation frontier.

Whoever leads the world in AI for energy will lead the world in developing and deploying next-generation energy technologies such as modular nuclear reactors, super-efficient combustion

systems, new approaches to carbon capture, and new strategies for electrification of the economy, securing the energy and climate frontier.

And finally, whoever leads the world in understanding and mitigating the risks of AI and the use of AI to improve national and global security will determine the landscape in which we and our allies work for the future.

Thank you for your time and careful consideration. I would be happy to answer any questions.