



Public Meeting of the
President's Council of Advisors on Science and Technology (PCAST)

April 23, 2024

Meeting Minutes

MEETING PARTICIPANTS

PCAST MEMBERS

- | | | |
|------------------------------|--------------------------|-----------------------|
| 1. Frances Arnold, Co-Chair | 11. Sue Desmond-Hellmann | 21. William Press |
| 2. Arati Prabhakar, Co-Chair | 12. Inez Fung | 22. Jennifer Richeson |
| 3. Maria T. Zuber, Co-Chair | 13. Andrea Goldsmith | 23. Vicki Sato |
| 4. Dan E. Arvizu | 14. Laura H. Greene | 24. Lisa Su |
| 5. Dennis Assanis | 15. Paula Hammond | 25. Kathryn Sullivan |
| 6. John Banovetz | 16. Eric Horvitz | 26. Terence Tao |
| 7. Frances Colón | 17. Joe Kiani | 27. Phil Venables |
| 8. Lisa A. Cooper | 18. Jon Levin | 28. Catherine Woteki |
| 9. John O. Dabiri | 19. Steve Pacala | |
| 10. William Dally | 20. Saul Perlmutter | |

PCAST STAFF

1. Lara Campbell, Executive Director
2. Reba Bandyopadhyay, Deputy Executive Director
3. Bich-Thuy (Twee) Sim, Assistant Director for Transformative Medicine and Health Innovation
4. Melissa A. Edwards, Assistant Deputy Executive Director
5. Kimberly Lawrence, Administrative Specialist

START DATE AND TIME: Tuesday, April 23, 11:00 AM Eastern Time

LOCATION: Livestreamed via Zoom.gov

WELCOME

PCAST Co-chairs: Frances Arnold, Arati Prabhakar, Maria Zuber

The PCAST co-chairs—Frances Arnold, California Institute of Technology; Arati Prabhakar, Science Advisor to the President; and Maria Zuber, Massachusetts Institute of Technology—called the meeting to order.

SESSION: DISCUSSION AND CONSIDERATION FOR APPROVAL OF PCAST REPORT TO THE PRESIDENT ON SUPERCHARGING RESEARCH: HARNESSING ARTIFICIAL INTELLIGENCE (AI) TO MEET GLOBAL CHALLENGES

Prabhakar introduced the session by providing some context about the work the White House is doing on AI, which became a priority for President Biden and Vice President Harris. She noted that in 2023, PCAST met with the President to discuss how AI is transforming research. Over the course of the year, the White House took several actions, culminating in an October 30, 2023 Executive Order on AI. Broadly, the Executive Order discussed actions the Administration could take under existing law to manage AI's risks and set the stage for using AI to benefit the American people. The Executive Order called for PCAST to develop a report on the potential role of AI, especially given recent developments in AI, in research aimed at tackling major societal and global challenges.

TERENCE TAO

Tao said the October 2023 Executive Order called for PCAST to develop a report on the potential role of AI, especially given recent developments in AI, in research aimed at tackling major societal and global challenges. He noted the Executive Order tasked other agencies and groups to examine the broader effects of AI outside of scientific research, such as the effects AI will have on national security, education, and labor markets.

The AI working group, said Tao, found significant synergies between science and AI. AI technologies, for example, are starting to remove many barriers that make scientific research slow, expensive, or restricted to a few experts. Conversely, he said, the fundamental science of AI can mitigate errors, biases, and other potential harms of AI. In addition, existing scientific values, such as validation, reproducibility, openness, and expert supervision, will help ensure developing a culture of responsible use of AI methods in science.

Tao said there is a strong case for using AI tools in science. In fact, AI is already transforming nearly every scientific discipline and every aspect of the scientific workflow. The AI working group developed with eight examples of fields that AI is transforming grouped into three categories representing different ways AI is changing research: identifying candidate solutions to scientific problems, accelerating and enhancing scientific simulations and models, and analyzing new types of data that traditional methods find difficult to handle.

The first example Tao discussed was how materials scientists have used AI to predict millions of new, and importantly stable, materials that can be automatically synthesized in a laboratory. The hope is that in the future, AI will predict room temperature superconductors and new materials for batteries to power electric vehicles, among other useful applications. Tao noted that fabricating and testing new materials is expensive, so if AI technologies can reduce the number of candidates before finding a successful one, the cost and time savings would be substantial.

For drug discovery, said Tao, AI has succeeded in identifying a promising new antibiotic and led to the discovery of new antimicrobial mechanisms that drug developers can now target. Researchers have also used AI to design vaccines and targeted therapies that block the binding regions of viruses. As AI technologies improve, researchers will use AI to develop custom-tailored therapies for cancer and autoimmune diseases, prolong the vibrancy of cells, and provide a deeper understanding of how antibiotics function.

The third example Tao discussed illustrated how the semiconductor industry is using AI to generate initial circuit designs for computer chips. In the future, AI should be able to perform lower-level chip design tasks, leaving designers to focus on high-level design challenges and improving designer productivity by more than an order of magnitude.

Tao, turning to examples illustrating how AI is accelerating and enhancing scientific simulations and modeling, said that AI is already accelerating aspects of weather forecasting and climate modeling by as much as a thousand-fold. Going forward, AI may accelerate data gathering and assimilation, major bottlenecks in weather forecasting. It may also create whole Earth models that would combine both traditional and AI models with real-time data to simulate climate processes and advance areas such as catastrophe modeling, water management, and greenhouse gas monitoring.

Cell biology, said Tao, is another area in which AI models are beginning to make progress. The AI-powered AlphaFold program has successfully predicted the structure and function of protein, and a future goal is to not just model individual proteins but the entire human cell. An AI-assisted model of the human cell would help unlock the secrets of cell biology and allow testing new therapies *in silico* rather than *in vivo* or *in vitro*.

Tao said AI has accelerated simulation of cosmological hypotheses on the nature of the universe. In the future, AI-assisted modeling may generate new fundamental theories in physics.

Regarding new types of data, Tao said social scientists have used AI-powered large language models to analyze non-quantitative data, such as large-scale behavioral data in social media posts that can help track changes in sentiment, for example. Data-driven social science may enable more effective, responsive, and fair delivery of public services.

Tao said the health and wellness field generates large amounts of unstructured or complex data, and AI-powered analysis of this data is already helping physicians to diagnose cancer earlier and detect potential errors in medical care that endanger patient safety. In the future, AI may make ultra-personalized medicine based on an individual's genetics and medical history more widely available. However, there are non-trivial issues that the field needs to address before AI is deployed widely in medicine, such as how to ensure patient privacy is protected.

LAURA GREENE

Greene discussed the working group's vision of AI-assisted research and development. To gain the greatest net benefit of science adopting AI, the working group envisioned pursuing three broad goals: empower human scientists, use AI tools responsibly, and share AI resources. To achieve these goals, the working group made five recommendations.

AI, said Greene, can empower human scientists in several ways. AI assistants, for example, can complement and empower human scientists rather than replace them. AI tools can process huge data streams; manage tasks such as laboratory work, coding, writing; and identify promising solutions to scientific problems, which will enable human scientists to focus on high-level directions for research. She noted that traditional, non-AI-assisted forms of scientific research will remain an essential component of the scientific landscape by complementing AI-assisted paradigms, training the next generation of scientists, and maintaining intellectual diversity. New ways of collaborating may emerge, such as human scientists directing a network of interlinked AI assistants and enabling large, interdisciplinary, or decentralized projects.

To accomplish this, said Greene, the working group recommended supporting both basic and applied research in AI that involves collaboration across academia, industry, national laboratories, and federal agencies. The National Science Foundation's (NSF) Materials Innovation Platforms, for example, is developing a data sharing infrastructure while employing AI tools as part of the community building with other agencies and industrial partners. Future projects could include collaborations to develop next-generation quantum computing qubits, whole cell modeling, developing whole Earth foundation models, or building high-quality scientific databases in a broad range of disciplines.

Greene said the second recommendation to empower human scientists is to encourage innovative approaches to integrating AI assistance into scientific workflows. Funding agencies, she said, should recognize the emergence of new workflows and design flexible procedures, metrics, funding models, and problems that encourage strategic experimentation with new AI-assisted ways to organize a scientific project. However, while AI-assisted science can be superior to either unassisted human science or fully automated science, traditional human scientific research must complement these tools.

Regarding the responsible use of AI, Greene said the scientific community should develop a culture of responsible AI use in which AI outputs are verified externally, protections for private data are in place, algorithmic bias is measured and compensated for, and models and data are as transparent, replicable, open, and explainable as possible. A continuous dialog among the physical and social scientists, the humanities, and policy makers should shape this culture. She noted that the development of lightweight models and shared resources will likely reduce the environmental impacts and other costs of AI use.

Greene said the working group recommended adopting principles of responsible, transparent, and trustworthy AI use throughout all stages of the scientific research process. AI tools, she said, could strengthen scientific integrity by detecting algorithmic biases, enabling more efficient replication and validation of scientific findings, and protecting personal privacy and intellectual property rights. Funding agencies should require researchers to develop responsible AI use plans that would assess potential AI-related risks, including aligning with the principles in the blueprint for the AI Bill of Rights and the National Institute of Science and Technology's (NIST) AI Risk Management Framework.

Greene, referring to the goal of developed shared and open resources, said multiple open-source models powerful enough for research applications will emerge avoiding dependence on a few private technology companies. A shared AI infrastructure comprising models, data, standards, and best practices, researchers will avoid redundant efforts, democratize access to AI for research purposes, and significantly reduce the time and resources needed to use AI tools. In addition, as previous PCAST reports have noted, the federal

government should have significant scientific data sets, such as Federal Emergency Management Agency data to improve catastrophe modeling, defense satellite data to improve wildfire prediction, and health care data to make patient safety measures more effective.

The first recommendation to achieve this goal, said Greene, is to expand existing efforts to broadly and equitably share basic AI resources. Toward this end, the working group strongly recommended that Congress fully fund the National AI Research Resource (NAIRR) as an important first step to establishing a broader national AI infrastructure that can compete with private and international efforts. In the longer term, NAIRR could be a stepping stone for more ambitious “moonshot” programs such as creating multimodal models of complex systems such as the Earth and the human cell.

Greene said the working group’s second recommendation regarding shared and open resources is to expand secure access to federal data sets for approved critical research needs, with appropriate protections and safeguards. The working group encouraged the federal government to expand the National Secure Data Service Demonstration and Federal Statistical Research Data Centers, with strengthened enforcement of data sharing mandates for federally funded research. Shared AI infrastructure, said Greene, may enable upgrading raw federal datasets to more valuable highly curated databases to democratize resources and benefit society.

In conclusion, Greene said scientific research will be radically accelerated by a broad spectrum of AI technologies. With the right AI infrastructure in place, this will enable scientists to address urgent global challenges. However, while AI is powerful, it has weaknesses such as algorithmic bias and can require enormous amounts of computational power, energy, and data. There are ways to mitigate these weaknesses and reduce the net resources required, particularly for scientific applications. When used responsibly, said Greene, AI can unlock great advances in science and technology.

After the working group’s co-leads’ presentation, discussion among PCAST members followed.

ARNOLD MODERATED THE Q&A AND DISCUSSION BETWEEN PCAST MEMBERS AND TAO AND GREENE

Arnold asked Tao and Greene how open-source resources can also be made secure. Greene said that is something that needs to be worked out, and this called out specifically in the NAIRR pilot program. William Dally added that open-source models can be cryptographically signed so users know they are working with the model as developed and not an altered model. Tao said there are anonymizing and randomizing technologies available to protect the privacy of the data used to train a model. There are also efforts underway using synthetic data to train models rather than real-world data. This is research that needs to be supported, he said. Arnold noted that shared and open models could mitigate the sustainability issues associated with AI training and model building.

Jennifer Richeson said she was glad to see social science research as a topic ripe for discovery assisted by AI tools and approaches, and she asked if the working group could provide examples of ongoing AI-assisted research relevant to the health, wealth, and well-being of the American public. One example the PCAST social science working group is exploring, responded Jon Levin, is how AI tools can now measure the effect of federal programs and policies on smaller populations, rather than on the whole U.S. population, to enable refining programs so that more Americans benefit from them.

A mundane application, said Levin, is to use AI to make federal programs easier for people to access by synthesizing their data and auto-filling reports such as tax returns and student loan applications. Eric Horvitz said he is seeing the first wave of AI tools that reduce the administrative burden on health care clinicians, giving them more time to interact with their patients. There has also been intensive work on tools to provide insights on bias and fairness of AI algorithms and track over the time the implications and powers of these tools for different groups, such as urban versus rural populations.

Horwitz commented that the responsibility to use these tools in the best way possible and look for possible downsides must rest with scientists, whether the models are open source or not. He added that for biosecurity applications, the Office of Science and Technology Policy, the National Security Council, and NIST will be weighing in on what needs to be done regarding new kinds of screening tools for DNA synthesis to understand when there might be challenges arising.

Zuber asked if the working group had discussed what it means to have a high-quality, trusted dataset with requirements for archiving and peer review. Tao replied there are standards for datasets to be open and replicable and for provenance tracking and noted the value of the NAIRR project, which can set the standards for open datasets. Both Tao and Greene noted the FAIR data standard as an example of what is possible. Greene added that AI tools may address the problem of non-reproducibility in science by examining original datasets for signs that something is incorrect.

William Press said it is important that the working group's report emphasizes that scientists will have to learn how to use AI responsibly. He recounted how he had used a large language model to write code that would analyze a large quantity of public census data and then asked the model how he could check to see if the resulting code might introduce biases into the output. The large language model then wrote code to post-process the output to check for biases. In his view, today's AI tools are ethically neutral in that they are not necessarily biased, and that they can check themselves with no pride of ownership, as a human might have.

Lisa Cooper asked the working group members to identify some challenges in using AI to inform public policies. One issue, said Levin, regards how data used to power facial recognition software will represent different populations. Algorithms might also be optimized during training in a way that might favor different populations. How to make AI-powered models unbiased and fairer was a major topic of discussion for the working group. Greene noted the importance of anonymizing data as part of the responsible use of data.

Horwitz commented on the growing sensitivity regarding the limitations of AI tools and methods, saying that understanding these limitations will be valuable for extending the utility of these tools. One example is using AI methods to determine where there are gaps for specific populations or demographics in a dataset.

Dally said he is excited that AI tools are empowering scientists and engineers to accomplish new tasks and do so with fewer people and computational resources. It is encouraging, he said, that the United States is at the forefront of AI research, but the Nation is in danger of losing that lead because of underinvestment. As the report recommends, immediately authorizing and funding NAIRR is critical. He noted that India recently announced it is investing \$1.2 billion to build a national AI

infrastructure, and that China and Europe are also making large investments in AI. Tao added that NAIRR is functioning as a small pilot program using donated resources.

Andrea Goldsmith asked how the Nation can democratize the AI-enabled acceleration of science for universities or individual researchers that may not have the necessary computational resources. Tao said NAIRR provides computational resources to address that very issue. Horvitz added that the NAIRR pilot program is focused on ensuring that regions that typically do not have a significant amount of computational resources get priority when requesting computational time. NAIRR's first call for proposals, Horvitz noted, is focused on responsible AI applications that address fairness, bias, safety, and reliability.

Goldsmith remarked that writing proposals for computational resources can be challenging, which might interfere with the solving the democratization problem. Tao said a future development that might help would be the creation of scientific foundation models that are already pre-trained and provided on platforms such as NAIRR. These would be much less resource intensive for individual researchers to use. That is a longer-term goal.

Paula Hammond commented that the more the scientific community can educate a broad set of the U.S. population about AI and create AI tools the broader population can use, the more likely it will be there will be investigators asking important questions about AI and that AI can address. Tao replied that academia needs to change the way it educates students about how to use AI responsibly. This is a place where AI itself could be useful as a learning tool. He also sees the potential for AI tools to allow lay citizen scientists to contribute to scientific projects, something he is seeing already in his field of mathematics.

Arnold noted that model interfaces are now easier for non-experts to navigate, but there is still the need to interrogate a model and understand what it is doing. Tao agreed and said AI tools will need to be integrated with other technologies to address that issue. In mathematics, for example, there is a technology called formal verification that can check whether a mathematical argument produced by a human or AI tool is correct. It is important, he added, to pair the power of AI with the rigor of other verification tools.

When Arnold asked how verification would work in with materials science, Greene pointed out that humans must be involved at every step of the discovery process. One thing she has been look at is whether it would be possible to interview researchers who have had hunches about and successes with producing new superconducting or thermoelectric materials, for example, and then use AI to go through the literature to try to identify new materials. She noted, though, that nobody in her field of correlated electron materials believe that DeepMind discoveries will help identify new materials. Tao added that for materials, laboratory experiments are the final arbiter.

Hammond mentioned there is an opportunity for high schools and two-year learning institutions to use AI tools and engage a larger population of students, though as Arnold noted, that will require teachers trained to use AI tools. Dennis Assanis agreed with importance of provide AI-related education to all members of society. At the same time, he said, while the power of AI is vast, it is important to acknowledge the challenges it poses to current systems of knowledge production and dissemination and the reward systems that support these activities. There is the challenge, for

example, of attribution and verification, which will require guidance on citational practices, data preservation, and research integrity for field beyond science and technology and the agencies that support them. This is not just a problem for NSF, the National Institutes of Health, the Department of Energy, and the Department of Defense to address, but also for the National Endowment for the Humanities, the Institute for Museum and Library Science, and others.

Saul Perlmutter said there is a real opportunity to step back and remember the need to teach how to think clearly when doing scientific analysis or in one's day-to-day life. The tools available today have the potential to make people armed but dangerous as far as preparing what looks like real analysis but is answering the wrong question or falling into different kinds of standard mental traps it takes the typical graduate student years to learn how to be watching out for when doing their research. Educators have started to develop courses that sets out a problem that students will solve with AI and then gets them to see the traps they have fallen into by using AI tools.

Perlmutter pointed out that AI is good at translating jargon and terminology, which creates the opportunity for people outside of the sciences to work with scientific ideas. It also opens the door for researchers in one area to work better with researchers in other areas of science.

Horwitz called for AI model creators and AI researchers to get feedback from the scientists using their tools about their needs and requirements regarding scientific integrity. Scientific research, he said, needs well-calibrated confidence in inferences and attribution. For example, it is important to know where information comes from that generates syntheses and inferences.

With the discussion concluded, PCAST voted unanimously to accept the report. Arnold said she believed this report would be useful for charting the future of AI in the sciences and research.

PUBLIC COMMENT

Brien A. Seeley provided two minutes of public comments.

CLOSING COMMENTS

Zuber and Prabhakar thanked the working group for producing a useful and impactful report, and Arnold commented she is excited about the future but there is still work to do to make sure AI serves society. Zuber then adjourned the public meeting.

PUBLIC MEETING ADJOURNED: 12:00 PM Eastern Time

I hereby certify that, to the best of my knowledge, the foregoing minutes are accurate and complete.

Frances Arnold, Ph.D.
Co-Chair
President's Council of Advisors on Science and Technology

Arati Prabhakar, Ph.D.

Co-Chair

President's Council of Advisors on Science and Technology

Maria Zuber, Ph.D.

Co-Chair

President's Council of Advisors on Science and Technology