

Testimony before National Artificial Intelligence Advisory Committee Hearing Panel II: Advancing U.S. Leadership in AI Research & Development

Catherine Aiken
Center for Security and Emerging Technology

October 12, 2022

Members of the Committee, thank you for holding this hearing and inviting me to speak on the important topic of U.S. leadership in AI R&D. My name is Catherine Aiken, and I am the Director of Data Science at Georgetown University's Center for Security and Emerging Technology. In this role, my day-to-day involves measuring, counting, and analyzing the trove of information we have available on science and technology. I will focus today on how we measure and assess "progress in AI."

To know how we are doing in AI, or any other field or technology, we measure things to determine where we stand, and where resources should be allocated. Measurement also matters because what we count can define the things that matter to us. But reliance on accepted measures can skew or limit our understanding of those very same things.

What we capture, and what we miss, with current AI R&D measures

For example, when assessing progress in AI, we measure papers published and their citations, benchmarks for specific tasks (like reading comprehension or object detection), patented AI technology, and investment flows into AI research and AI companies. These are useful indicators of progress. CSET relies on these metrics often, including through our collaboration on the [2022 AI Index Report](#).¹ However these measures have come to define AI progress, when they are just steps along the path.

What we really care about when we say "progress in AI" is our ability to accelerate AI research **and** its application for economic prosperity and national security; to create value for people and to benefit society. Measuring AI advancements in this broader sense requires expanding what we measure to include how well we convert research into economic, social, and security benefits. If what we really care about is our ability to convert AI research into technology of relevance, we need to (1) better understand the current AI R&D landscape and; (2) measure both research progress and what happens after research publication.

¹ For some examples, see [Chinese Public AI R&D Spending: Provisional Findings](#); [Comparing U.S. and Chinese Contributions to High-Impact AI Research](#); [Measuring AI Development](#); [Patents and Artificial Intelligence](#); [Research Impact, Research Output, and the Role of International Collaboration](#); and [Research Security, Collaboration, and the Changing Map of Global R&D](#).

Laying the groundwork

To better understand the current AI R&D landscape, we can focus on two things. One is standardization in what constitutes AI across agencies, research teams, and companies.² The lack of a standard here limits our ability to track our investments and use cases. For example, the seemingly simple task of counting federal grants for AI research requires bespoke methods to determine which awards count as AI and which do not.

The other is making information on R&D accessible, so we can assess our investments. We lack comprehensive data here, and there is no central repository of information from which to analyze trends in AI R&D. To fill this gap, government agencies should provide information on awarded grants via an AI R&D dashboard and more universities should share information on what research is being funded via existing initiatives like UMETRICS.³

These gaps limit our ability to map the AI R&D landscape and assess how our investments align with our stated goals.⁴ This is a critical first step, because we need to walk before we can run; to measure our current AI R&D before we can assess our progress and leadership.

Measuring the steps toward AI progress

Next we need to think about AI R&D as accelerating research and its conversion to real-world use, measuring both research progress and what happens after research publication. To do this we similarly need to focus on two things. First, bolstering our analysis of existing measures. Second, identifying measures for what happens after research.

U.S.-based researchers produce impactful AI research and we are home to the AI companies that attract the most investments.⁵ These are reassuring assessments based on existing measures, but we need to invest in the data, infrastructure, and analytic capabilities needed to monitor these, and related, trends and make information available in real time.⁶

² For some discussion on the question of defining AI, see [AI Definitions Affect Policymaking](#).

³ The Universities Measuring the Effects of Research on Competitiveness, Innovation and Science (UMETRICS) dataset was created and is maintained by the Institute for Research on Innovation & Science (IRIS). It includes de-identified transaction-level federal and non-federal research grant data from over 80 U.S. universities. See more at Institute for Research on Innovation and Science, "[Creating Trusted Independent Data About the Impact of Research](#)."

⁴ Preliminary CSET analysis of federal funding at a set of U.S. universities finds that AI research is a small fraction of all funded research, and that fraction has remained stagnant from 2014-2019.

⁵ See CSET's [Country Activity Tracker](#) and [Comparing U.S. and Chinese Contributions to High-Impact AI Research](#).

⁶ [Dewey Murdick's Testimony before the Senate Select Committee on Intelligence](#)

How can we measure AI progress in this broader sense? If we take progress to require accessible communal resources and organizations following agreed-upon best practices; a robust and varied workforce, trained to implement AI solutions in a safe way and empowered to modify implementation as needed; and industry and the public sector implementation of shared frameworks that assure accountability and foster trust - we can measure these things.

Beyond publications and funding amounts, we can consider a host of additional measures. Those I suggest exploring include:

- Measure how communal resources are being used to develop and implement research through activity on GitHub.
- Measure AI use and deployment through use case inventories, AI incident reporting, and surveys of organizations' AI adoption.
- Measure implementation of best practices as the number and kind of frameworks used and the operational evaluations undertaken.
- Measure workforce growth with the number of AI certifications, degree programs, and skills development pathways offered and completed, as well as workforce trends in AI job postings and company hiring, and AI talent career preferences and pathways;
- Measure AI governance as the kind and number of auditing and impact assessments in use, and the structure of bodies that decide how, if, and when AI is used.
- Measure how the public is responding and adapting to AI use via surveys, focus groups, and stakeholder engagement.

Moving forward

Measuring AI progress in this broader sense *will require* being open to qualitative data collection, methods to analyze human behavior, connecting varied data sources, and eliciting information from a wide range of stakeholders on a continual basis.

But to lead in AI R&D we must define progress as a demonstrated commitment toward stated goals, and assess our performance at each step - from the resources needed to push the boundaries in research, to the structures that ensure people are safe and accounted for when AI is deployed. Measuring our ability to turn research into real-world solutions—that is the progress that will ensure our leadership in AI.

Thank you, and I look forward to questions and discussion.