Building better intelligence about investments in artificial intelligence

Governments are pouring billions of dollars into AI research – here's a way to track impact.

By Julia I. Lane, Jason Owen Smith and Bruce A. Weinberg

Government spending on artificial intelligence (AI) has surged across the world. The US federal government spent over $3 billion in 2022 and an influential taskforce—the National AI Research Resource (NAIRR)—recommended spending at least $2.6 billion more on public-funded research over the next five years (Office of Science and Technology Policy, 2023). While these figures might seem small in comparison to the resources that the private sector is pumping into AI research (Leech et al., 2024), the stakes are high.

In many parts of the world, governments are betting on investments in innovative new industries as a means to transform their economies and generate sustained job growth. But public resources are limited. So, these bets must be well placed, and informed by data and evidence.

Quantifying the return on research investments is notoriously difficult, especially in newly emerging economic sectors. In most countries, national statistics systems are not even equipped to reliably track spending—forget economic impact—in frontier areas of research and innovation. The estimates on the level of public investment in AI research cited above are merely the best available estimates, not figures drawn from a well-defined measurement process.

To illustrate the extent of the current gap in knowledge, consider the state of Texas which in February 2024 applied a new industry code classification for AI, recently developed by the federal statistical system, to find how many firms in Texas were associated with that classification. The state found a mere 298 firms, which collectively employ 1,021 workers (Leonard, 2024). Similar definitional problems also arise in the case of other cutting-edge industries, such as robotics or electric mobility. Indeed, some scholars have postulated that about four-fifths of the economy of some advanced countries can now be characterized as "hard to measure" (Coyle, 2024).

Here, we propose a novel way to describe and analyze where AI ideas are being used and how they spread—by tracing the people and academic communities involved in AI research. When an individual transitions from a government-funded research lab to a private sector company, they carry cutting-edge "AI know-how" with them. Meshing existing university administrative data with state employment records allows several quantifiable inferences about the value of AI research to be drawn from these academia-to-industry migrations.

A pilot implementation of this system is being developed in the state of Ohio. It offers a template for governments and policy makers all over the world. Importantly, the metrics discussed below offer a way to measure the economic impact of scientific research in general, with implications for critical and emerging technologies that go far beyond AI.

How to track ideas

Traditional economic accounting is ill-suited for a research-led field like AI. At this early stage of the technology's evolution, what constitutes AI, let alone AI-related employment, is uncertain. The Stanford University-based One Hundred Year Study on Artificial Intelligence (AI100), which aims to convene a study panel once every five years to analyze the impact of AI on society, has noted that "AI can also be defined by what AI researchers do"(Stone et al., 2022).

In other words, any attempt to describe the economy-wide impact of public investments in AI would involve tracing the people at the heart of these investments. It is people who pollinate ideas, launch startups and influence the next generation of innovators via academic and professional networks(Lane, 2023). In newly emerging industries, where ideas matter a lot, people are the chief value creating unit—not machines or office floor space.

Luckily, in the US, a data system already exists to trace the people impacted by federal research grants. Proposed over a decade ago as a vehicle to document and explain the results of government funding of science, UMETRICS, hosted at the University of Michigan's Institute for Research on Innovation & Science, captures timely and comprehensive information on more than 535,000 grants. The funds tied to these grants support 864,000 employees—including students and research assistants—and 970,000 vendors who supply equipment and technological aids. In the context of AI research, vendors provide highly critical tech inputs, such as the graphics processing units (GPUs) needed to run large language models and the semiconductors needed to fight the chip wars (Miller, 2022). Collectively, the expenditures UMETRICS records from more than 80 different campuses represent about 41% of the US government's R&D spending at universities(Nicholls & Owen-Smith, 2022-04-20).

The subset of the researchers on these campuses supported AI research grants can be identified by cross-referencing against a the authors who present at major

AI conferences. This ‘seed set’ of grant recipients would have direct relationships with larger networks of collaborators (including students at all levels) and vendors”. Government funding enables the work of all these individuals and helps support the organizations that supply them.

To illustrate this point, consider the 3,143 principal investigators (PIs) on U.S. National Science Foundation (NSF) grants in the UMETRICS database who have also presented at major AI conferences. The information on science spending in UMETRICS links these PIs to over 46,000 other people. Most, about 30,000 are students and trainees. The rest are research staff and faculty collaborators. The money trail associates each PI with, on average, 15 other individuals, who are directly supported by federal grant funds.

Many of them may never publish a paper, file a patent, or become a PI themselves. But conducting AI research teaches them about cutting-edge algorithms and their application through the lens of nearly every field the NSF supports. It gives them access to specialized professional networks. It makes them both competitive for and interested in AI jobs. All these factors make them key employees for companies across many sectors seeking to develop or apply AI. In other words, these often invisible research funded people are both important, underexamined “products” of grant funded research and a way to identify currently unmeasurable workforce effects(Romer, 2019).

The flow of these trainees and staff through to the wider economy, and the transmission of their ideas, is captured when they get jobs in the private sector. Their earnings and employment are recorded in state administrative data(Zolas et al., 2015). This last-leg linkage—between academia and private sector employment—is the novel data layer currently being piloted in Ohio. The employment footprint of these individuals across traditional industrial sectors—ranging from healthcare to retail—offers a snapshot of the cross-sectoral workforce of the emerging industry of AI. Initial results using a version of this people-based methodology with public workforce data suggest that AI science investments affect more than 36 million American workers employed in industries that span 18 different “old school” sectors from manufacturing, to utilities, health care, finance, and IT. Those industries, and many more, are all home to businesses that employ AI researchers with the skills needed to develop and apply cutting edge technologies for many different purposes. These preliminary data provide a replicable estimate well in excess of conventional estimates, but still likely to be a considerable undercount. The second stage of the project will provide more detailed estimates and measures of job and employer characteristics that cannot be calculated from public sources.

These data also suggest that people employed in "AI industries" tend to make more on average than those who are not, The difference in pay between the workers whose prior research experience demonstrates "AI know-how" and others employed in the same economic sector is deeply informative. Better pay for the former could be seen as a quantifiable 'return' on initial research investments that may dwarf economic returns to particular, published discoveries. Such a pay disparity can reveal not just the market premium attached to AI skills but also show how this varies across economic sectors, which could influence the design of academic curriculums and government policies. In the pilot study in Ohio, it will be possible to characterize whether firms hiring AI scientists pay higher wages to new hires across the board, and whether the growth rate in earnings is greater than in other firms. State agencies can work with their local universities, as in Ohio, to pioneer new ways of describing local AI labor market dynamics.

The framework discussed above can be generalized to other fields of scientific research. The key insight is this: in some fields, scientists moving into firms are the chief value creating unit, not machines, or office space.

Looking beyond bibliometrics

Researchers and scientists must pay greater attention to how academic research impacts the private sector job market. This is one way to side-step the endless arms race to keep producing scientific publications that often go unread. What we measure will determine the outcomes we get. By looking beyond documents and citations to focus on more tangible routes to impact, like the careers of grant-funded students, new avenues for dialogue can be opened up with elected officials about the need and benefits of increased investments in scientific research.

Enough has been written on why tracking the value of academic output purely on the basis on publications is flawed. Women, for instance, are less likely to be credited for their academic contributions in published output, which impacts their career prospects [9]. Importantly, merely tracking publication frequency does not meaningfully engage with the central question that of interest to most governments: How can universities and industry work together to ensure that new discoveries are effectively translated into use by a workforce with the skills to implement and benefit from them?

We should not and cannot accept the status quo. The disruption caused by AI, and its anticipated impact on the economy, has forced the hand of many governments to “do something”. But the response should not just be to spend taxpayer money on research and expect miracles to happen. It should be to understand how science works and build a data infrastructure that is designed to meaningfully measure progress.

This vision can be achieved. The final NAIRR report, which was submitted to President Joe Biden and the US Congress in January 2023, recommended the people-centered evaluation approach we describe here(Office of Science and Technology Policy, 2023). It recommended the use of the type of data systems sketched here, which match rich—though restricted—workforce data with detailed bibliometric and university information. The results could change both how we measure the impact of science investments and how universities, industries, and governments partnerships to ensure that these essential public investments realize their full potential.

The work we are doing is scalable along many dimensions. The data infrastructure is adaptable—in that it draws on administrative records used for human resource management and tax purposes. Such data are typically engineered to meet a small number of standard accounting procedures. Thus code to collect, integrate, and analyse them could be replicated and reused across many organizations working in potentially many sectors and technologies.

Similar data are available internationally and can be applied to innovation-based economies globally. The approach can also be scaled to other emerging technology domains. Scalability is possible because the fundamental building block – using people’s careers to track economic impact – applies equally to all technologies.

While the potential of this approach is clear, several challenges do exist. Change is hard. Policy makers have, to date, settled for counts of publications and patents to draw inferences on how public funds are being utilised. New approaches and databases generate new insights but also require considerable groundwork and a change in mindsets.

Confidentiality issues need to be addressed. Privacy-preserving features are critical in any system anchored on people’s careers (National Institute of Standards and Technology, 2023). There is also the possibility that new metrics could be biased or gamed (Manheim & Garrabrant, 2018). Focusing initially on economic impact can distort the organization of science. But current arrangements are clearly inadequate, and we must make a beginning somewhere. In general, economic outcomes may also prove harder to game than bibliometric outcomes, and economic impact is increasingly a strategic goal.

None of these challenges is insurmountable. The 29 nations that came together in late-2023 to sign on to the Bletchley Park declaration—a commitment to safely and responsibly develop AI—showed that there is determination and political will to take effective policy action on AI. The formation of the UK's AI Safety Institute took less than a year once the initial idea was mooted. An international AI Jobs and Economy Monitor, built on a sound empirical framework like the one described here, could be formed on a similar time scale. We must start now.

Julia Lane is a Professor at New York University, served on the National AI Research Resources Taskforce, and works on the measurement and evaluation of investments in scientific infrastructures. Jason Owen-Smith is the Executive Director of the Institute for Research on Innovation & Science at University of Michigan where he works on the sociology of science and technology using computational and network analytic methods. Bruce Weinberg is Eric Byron Fix-Monda Professor of Economics and Public Affairs at Ohio State University and an expert on the economics of science and innovation, disparities in the labor markets, and the effects of new technologies on labor markets.

References

Coyle, D. (2024). *White Paper: Measuring AI* Workshop on New Approaches to Characterize Industries, Stanford University. https://www.digeconevents.com/people-new

Lane, J. (2023). *The Industry of Ideas: Measuring How Artificial Intelligence Changes Labor Markets* (American Enterprise Institute, Issue. https://www.aei.org/research-products/report/the-industry-of-ideas-measuring-how-artificial-intelligence-changes-labor-markets/

Leech, G., Garfinkel, S., Yagudin, M., Briand, A., & Zhuravlev, A. (2024). Ten Hard Problems in Artificial Intelligence We Must Get Right. *Arxiv:2402.04464*. https://doi.org/https://doi.org/10.48550/arXiv.2402.04464

Leonard, A. (2024). *Outside of the Box Use of Administrative and Wage Data in Texas* Workshop on New Approaches to Characterize Industries, Stanford University. https://www.digeconevents.com/people-new

Manheim, D., & Garrabrant, S. (2018). Categorizing variants of Goodhart's Law. *arXiv preprint arXiv:1803.04585*.

Miller, C. (2022). *Chip war: the fight for the world's most critical technology*. Simon and Schuster.

National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Retrieved from https://doi.org/10.6028/NIST.AI.100-1

Nicholls, N., Ku, R., Brown, C., , & Owen-Smith, J. (2022-04-20). *Summary Documentation for the IRIS UMETRICS 2022 Data Release*. doi: 10.21987/df2a-ha30

Office of Science and Technology Policy. (2023). *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem*. Retrieved from https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf

Romer, P. (2019). The Deep Structure of Economic Growth. https://paulromer.net/deep_structure_growth/

Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., & Kraus, S. (2022). Artificial intelligence and life in 2030: the one hundred year study on artificial intelligence. *arXiv preprint arXiv:2211.06318*.

Zolas, N., Goldschlag, N., Jarmin, R., Stephan, P., Smith, J. O.-., Rosen, R. F., Allen, B. M., Weinberg, B. A., & Lane, J. I. (2015). Wrapping it up in a person: Examining employment and earnings outcomes for Ph. D. recipients. *Science*, *350*(6266), 1367-1371.