




QILAI SHEN/BLOOMBERG/GETTY

Attendees at the 2023 World Artificial Intelligence Conference in Shanghai, China.

How to track the economic impact of public investments in AI

By Julia Lane, Jason Owen Smith & Bruce A. Weinberg

National statistics systems should recognize the researchers whose ideas drive artificial-intelligence applications, not just machines and factory outputs.

Government spending on artificial intelligence (AI) is surging worldwide. In the United States, for example, the federal government invested more than US\$3 billion in the 2023 fiscal year and an influential US taskforce – the National Artificial Intelligence Research Resource (NAIRR) – recommended channelling at least

\$2.6 billion more to public-funded research over an initial six-year period¹. The private sector is pumping even more into AI research, spending hundreds of billions of dollars each year². The stakes are high.

Why is AI research a priority for public funding? Governments are betting on investments in innovative emerging industries such as AI as a means to transform their economies and generate sustained job growth. But with limited public resources, it's crucial that these bets are well placed – and informed by data and evidence. That is the only way to maximize the return on public AI investments and steer the trajectory of AI towards serving the public.

However, quantifying spending in frontier areas of research and innovation – let alone the return on such spending – is notoriously difficult. Most national and state statistics systems are ill-equipped to track how investments in AI work their way through the economy because the companies and individuals who are driving the deployment of emerging AI tools are dispersed across a variety of

conventional industrial sectors.

The existing statistical classification framework, the North American Industry Classification System (NAICS), was modified in 2022 to add a single category for AI activities: AI research and development laboratories (see go.nature.com/4ayvk5a). In February, Adam Leonard, the chief analytics officer at the Texas Workforce Commission in Austin, applied the new NAICS classification to Texas data and found a mere 298 AI research and development firms employing just 1,021 workers in total³. The real workforce involved in AI-related activities, meanwhile, is likely to be much larger and spread across multiple industry sectors, ranging from hospitality and health care to oil exploration.

Similar challenges relating to the quantification of research spending and estimating the size of the current workforce plague other emerging industries, such as robotics and electric mobility. Indeed, some scholars have postulated that about four-fifths of the economies of some advanced countries can

now be characterized as ‘hard to measure’ (ref. 4). This is a serious concern, because governments can’t manage what they can’t measure. And measurement is particularly crucial in emerging and dynamic areas, in which policy action is most needed.

Here, we outline a way to describe where AI ideas are being used and how they spread – by analysing 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 take cutting-edge ‘AI know-how’ with them. By meshing existing university administrative data with state employment records, we offer a mechanism through which to draw quantifiable inferences about the value of AI research.

A pilot implementation of this system is being developed in the state of Ohio by the University of Michigan’s Institute for Research on Innovation and Science (IRIS), which is based in Ann Arbor; the Ohio State University in Columbus; and the US Social Science Research Council in New York City (see go.nature.com/3vdf5us). It offers a template for governments and policymakers 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 emerging technologies beyond AI.

How to track ideas

Conventional economic accounting is ill-suited for a research-led field such as AI. At this early stage of the technology’s evolution, what constitutes AI-related employment is uncertain. Stanford University’s One Hundred Year Study on Artificial Intelligence (AI100), which aims to convene a study panel once every five years to analyse the effect of AI on society⁵, has noted that “AI can also be defined by what AI researchers do”.

In other words, any attempt to describe the economy-wide impact of public investments

in AI would involve identifying the people at the heart of these investments. It is people who generate ideas, launch start-ups and influence the next generation of innovators through academic and professional networks⁶. In emerging industries, where ideas matter a lot, people are the main value-creating unit – not machines or office floor space.

In the United States, a data system already exists to identify the people who benefit from federal research grants. Proposed more than a decade ago as a vehicle to bring more transparency and accountability into government funding of science, UMETRICS, hosted at IRIS, captures comprehensive information on more than 580,000 grants.

“People who are employed in AI industries tend to earn more on average than those who are not.”

The funds tied to these grants support 985,000 employees – including students and research assistants – and 1.2 million vendors, who supply equipment and technological aids (see go.nature.com/3erv4e). In the context of AI research, vendors provide crucial hardware, such as the graphics processing units (GPUs) needed to run large language models and the semiconductors needed for microchips⁷. Collectively, the expenditures recorded on UMETRICS represent about 41% of the US government’s research and development spending at universities in 2022 (ref. 8).

The subset of the researchers who receive AI-specific research grants can be identified by cross-referencing grant recipients against authors who speak at big AI conferences (see ‘From the laboratory to the labour market – how AI ideas spread’). This ‘seed set’ would have direct relationships with larger networks

of collaborators, including students and vendors. Government funding enables the work of all these individuals.

To illustrate this point, consider the 3,143 principal investigators (PIs) with US National Science Foundation (NSF) grants in the UMETRICS database who have also presented at AI conferences. The transaction information recorded on UMETRICS links these PIs to more than 46,000 other people. Most – about 30,000 – are students and doctoral or post-doctoral trainees. The rest are research staff and faculty collaborators. The money trail links each PI with, on average, 15 other individuals, who are directly supported by federal funds.

Many of these individuals might never publish a paper, file a patent or become a PI themselves. But conducting AI research teaches them about cutting-edge algorithms and the application of these technologies in several fields that 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 these people key employees for companies across many sectors. In other words, these often unrecognized research-funded people are important, under-examined ‘results’ of grant-funded research and are key to identifying currently unmeasurable workforce effects (see go.nature.com/3vflf7u).

The movement 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 data⁹. This linkage – between academia and private-sector employment – is the new data layer that is being analysed in the Ohio pilot.

The employment footprint of these individuals across conventional industry sectors offers a snapshot of the cross-sectoral workforce of the emerging industry of AI. Initial results using a version of this people-based methodology suggest that AI science investments affect more than 36 million US workers employed in industries that span 18 different sectors – from manufacturing to utilities, health care, finance and IT (see go.nature.com/45pjo2c).

Those industries, and many more, are all home to businesses that employ AI researchers. These preliminary data provide an estimate well in excess of conventional metrics, but it is still likely to be an undercount. The second stage of the pilot project will provide more granular information on employer characteristics and job-market dynamics.


These data suggest that people who are employed in AI industries tend to earn more on average than those who are not. The difference in pay between the workers whose previous research experience demonstrates AI know-how and those without such experience who are employed in the same economic sector is deeply informative. Better pay for the former

FROM THE LABORATORY TO THE LABOUR MARKET — HOW AI IDEAS SPREAD

With artificial intelligence (AI) labs attracting millions of dollars in public funding, tracking the impact of this investment on the private sector and the broader economy is a key challenge.


Stage 1: identify

Academics who present at AI conferences can be identified. On the basis of a list of presenters at 21 major AI conferences, a subset of researchers who also received government grants yielded a ‘seed set’ – more than 7,800 principal investigators (PIs).



Stage 2: expand

These PIs interact with students and research staff, and with vendors who supply lab equipment. These interactions were captured through university administrative data.



Stage 3: integrate

Those who leave an academic lab – after a stint as a student or staff researcher – and then seek employment in the private sector can be identified through state employment records.



Comment

could be seen as a quantifiable return on the initial research investments.

A finer understanding of these emerging pay disparities could reveal not just the market premium attached to AI skills but also how this varies across economic sectors, which could influence the design of academic curriculums and government policies. In the pilot study in Ohio, for instance, it will soon be possible to characterize whether firms hiring AI scientists pay higher wages to new employees across the board, and whether the growth rate in earnings at these firms is greater than that at other companies.

The framework discussed can be generalized to other fields of scientific research. The key insight is this: in some fields, people are the main value-creating unit.

Looking beyond bibliometrics


Researchers and scientists must start paying greater attention to how academic research affects the private-sector job market. This is one way to sidestep the endless race to keep producing scientific publications that often go unread. What we measure will determine the outcomes we get.

By looking beyond publications and citations and focusing on more tangible measures of impact – such as the career trajectory of grant-funded students – dialogue on the need to increase investments in scientific research can be opened up with elected officials.

Enough has been written on why tracking the value of academic output purely on the basis of publications is flawed. Women, for instance, are less likely to be credited for their academic contributions in published content, which affects their career prospects¹⁰. The disruption caused by AI, and its anticipated effect on the economy, has forced 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 accurately 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-centred evaluation approach we describe here¹. It recommended the use of the type of data systems outlined here, which match rich – although restricted – workforce data with detailed bibliometric and university information. The results could change how we measure the impact of science investments.

The work we are doing is scalable to many industries. The data infrastructure is adaptable, because 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. The code to collect, integrate and analyse the data could be replicated and reused



AI technologies were displayed at the 2024 Mobile World Congress in Barcelona, Spain.

across many organizations.

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. This is possible because the fundamental building block – using people's careers to track economic impact – applies equally to all technologies.

Although the potential of this approach is clear, several challenges do exist. Change is hard. Policymakers have, so far, settled for using the numbers of publications and patents to draw inferences on how public funds are being used. Fresh approaches and databases generate insights but also require considerable groundwork and a change in mindsets.

Confidentiality issues need to be addressed. Privacy-preserving features are crucial in any system that uses information about people's careers¹¹. There is also the possibility that new metrics could be biased or manipulated¹². Focusing on economic impact can distort the organization of science away from the pursuit of scientific discovery. But current arrangements are clearly inadequate, and we must make a start somewhere. In general, economic outcomes might be harder to manipulate than bibliometric outcomes, and economic impact is increasingly becoming a goal of national science policies as laid out by governments.

None of these challenges is insurmountable, however. The 29 nations that came together at Bletchley Park near Milton Keynes, UK, in late-2023 to sign the Bletchley Declaration – a commitment to develop AI safely and responsibly – 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 after the initial idea was mooted. An international AI jobs and economy

monitor, built on a sound empirical framework such as the one described here, could be formed on a similar timescale. We must start now.

The authors

Julia Lane is a professor of public policy at New York University, New York. **Jason Owen-Smith** is the executive director of the Institute for Research on Innovation and Science (IRIS) at the University of Michigan, Ann Arbor, Michigan. **Bruce A. Weinberg** is a professor of economics and public affairs at the Ohio State University, Columbus, Ohio.
e-mail: julia.lane@nyu.edu

1. Office of Science and Technology Policy. *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem* (OSTP, 2023).
2. Leech, G., Garfinkel, S., Yagudin, M., Briand, A. & Zhuravlev, A. Preprint at arXiv <https://doi.org/10.48550/arXiv.2402.04464> (2024).
3. Leonard, A. *Workshop on New Approaches to Characterize Industries: Outside of the Box Use of Administrative and Wage Data in Texas* (Stanford Univ., 2024).
4. Coyle, D. *Workshop on New Approaches to Characterize Industries: Measuring AI* (Stanford Univ., 2024).
5. Stone, P. et al. Preprint at arXiv <https://doi.org/10.48550/arXiv.2211.06318> (2022).
6. Lane, J. *The Industry of Ideas: Measuring How Artificial Intelligence Changes Labor Markets* (American Enterprise Institute, 2023).
7. Miller, C. *Chip War: The Fight for the World's Most Critical Technology* (Simon & Schuster, 2022).
8. Chang, W.-Y., Garner, M., Basner, J., Weinberg, B. & Owen-Smith, J. *Harvard Data Sci. Rev.* <https://doi.org/10.1162/99608f92.1e23fb3f> (2022).
9. Zolas, N. et al. *Science* **350**, 1367–1371 (2015).
10. Ross, M. B. et al. *Nature* **608**, 135–145 (2022).
11. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST, 2023).
12. Manheim, D. & Garrabrant, S. Preprint at arXiv <https://doi.org/10.48550/arXiv.1803.04585> (2018).

The authors declare no competing interests.