POLICY FORUM

ARTIFICIAL INTELLIGENCE

The growing influence of industry in AI research

Industry is gaining control over the technology's future

By Nur Ahmed[1,2], Muntasir Wahed[3], Neil C. Thompson[1,2]

For decades, artificial intelligence (AI) research has coexisted in academia and industry, but the balance is tilting toward industry as deep learning, a data-and-compute-driven subfield of AI, has become the leading technology in the field. Industry's AI successes are easy to see on the news, but those headlines are the heralds of a much larger, more systematic shift as industry increasingly dominates the three key ingredients of modern AI research: computing power, large datasets, and highly skilled researchers. This domination of inputs is translating into AI research outcomes: Industry is becoming more influential in academic publications, cutting-edge models, and key benchmarks. And although these industry investments will benefit consumers, the accompanying research dominance should be a worry for policy-makers around the world because it means that public interest alternatives for important AI tools may become increasingly scarce.

INDUSTRY'S INPUT DOMINANCE

Industry has long had better access to large, economically valuable datasets (*1*) because their operations naturally produce data as they interact with large numbers of users and devices. For example, in 2020, WhatsApp users sent roughly 100 billion messages per day. Thus, it is unsurprising that most large data centers are owned and operated by industry [see supplementary materials (SM)]. In this article, we show that industry's dominance extends beyond data to the other key inputs of modern AI: talent and computing power.

Demand for AI talent has grown much more quickly than supply over the past decade (see SM), generating increased competition for AI talent. Across two different measures of talent, we see that industry is winning this contest. Data on North American universities (where we are able to get the best data) show that computer science PhD graduates specializing in AI are going to industry in unprecedented numbers (see the first figure). In 2004, only 21% of AI PhDs went to industry, but by 2020, almost 70% were. For comparison, this share of PhDs entering industry is already higher than in many areas of science and will likely soon pass the average across all areas of engineering (see SM). Computer science research faculty who specialize in AI have also been hired away from universities to work in industry. This hiring has risen eightfold since 2006, far faster than the overall increase in computer science research faculty (see the first figure). Between the PhD students and faculty leaving for industry, academic institutions are struggling to keep talent (*2*). This concern is not limited to US universities. In the UK, Abhinay Muthoo, Dean of Warwick University's King's Cross campus, said, "The top tech firms are sucking the juice from the universities" (*3*).

The computing power being used by academia and industry also shows a growing divide. In image classification, the computing power being used by industry is larger and has grown more rapidly than that used by academia or by industry-academia collaborations (see the first figure). Here, we proxy for the computing power used in a model with the number of parameters—both because the number of parameters is one of the key determinants of the computing power needed and because the deep learning scaling law literature has shown strong relationships between them. In 2021, industry models were 29 times bigger, on average, than academic models, highlighting the vast difference in computing power available to the two groups. This is not just a difference in approach but a shortfall in computing available to academics. For example, data from Canada's National Advanced Research Computing Platform reveals that academic demand for graphics processing units (GPUs; the most common chips used in AI) on their platform has increased 25-fold since 2013 (see SM), but supply has only been able to meet 20% of this demand in recent years.

Industry's ability to hire talent and harness greater computing power likely arises because of differences in spending. Although investments in AI have gone up substantially in both the public and private sectors, industry's investments are larger and growing faster (see SM). We compare industry with the major source of public-interest AI research: governments, which both fund their own research and are a key source of academic funding. In 2021, nondefense US government agencies allocated US$1.5 billion on AI. In that same year, the European Commission planned to spend €1 billion (US$1.2 billion). By contrast, globally, industry spent more than US$340 billion on AI in 2021, vastly outpacing public investment. As one example, in 2019 Google's parent company Alphabet spent US$1.5 billion on its subsidiary DeepMind, which is just one piece of its AI investment. In Europe, the disparity is smaller but is still present; AI Watch estimates that "the private and public sector account for 67% and 33% of the EU AI investments respectively" (*4*) (see SM). For comparison, in recent decades, research funding in the pharmaceutical industry has been split roughly evenly between the private sector and governments or nonprofits (see SM). An example of the scale of funding needed to pursue AI research comes from OpenAI, which began as a not-for-profit with the claim to be "unconstrained by a need to generate financial return" and aiming to "benefit humanity as a whole" (*5*). Four years later, OpenAI changed its status to a "capped for-profit organization" and announced that the change would allow them "to rapidly increase our investments in compute and talent" (*6*).

THE INCREASING DOMINANCE OF INDUSTRY IN AI RESEARCH

Industry's dominance of AI inputs is now manifesting in an increasing prominence in AI outcomes as well—in particular, in publishing, in creating the largest models, and in beating key benchmarks. Research papers with one or more industry co-authors grew from 22% of the presentations at leading AI conferences in 2000 to 38% in 2020 (see the second figure). Alternate definitions of what constitutes an industry paper yield substantially similar results (see SM). Industry's dominance is even more apparent in the largest AI models (*7*) and in benchmark performance. Industry's share of the biggest AI models has gone from 11% in 2010 to 96% in 2021 [see the second figure; data are from (*8*)]. We use model size as a proxy for the capabilities of large AI models, as is common in the literature. Model size is also often used as a proxy for computing power (see the first fig-

[1]Sloan School of Management, Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. [2]Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA. [3]Department of Computer Science, Virginia Tech, VA, USA. Email: nurahmed@mit.edu; neil_t@mit.edu

ure). This dual usage reflects how important compute is for predicting the performance of deep learning systems (*9*).

We investigate when academia, industry, or academia-industry collaborations led performance on AI benchmarks (see the second figure). When looking across these six benchmarks in image recognition, sentiment analysis, language modeling, semantic segmentation, object detection, and machine translation—as well as 14 more that cover areas such as robotics and common sense reasoning (see SM)—industry alone or in collaboration with universities had the leading model 62% of the time before 2017. Since 2020, that share has risen to 91% of the time. For example, sentiment analysis can be used to understand the emotional tone of written work. Until 2017, academia led this benchmark 77% of the time. But since 2020, industry alone or in collaboration has led 100% of the time. So whether measured by building state-of-the-art AI models (as measured by either size or benchmark performance) or by publishing in leading research outlets, our analysis shows industry's increasing prominence in AI outputs.

AI research inputs

(Top left) Percentage of US artificial intelligence (AI) PhDs hired by industry. (Top right) Growth of US university AI research faculty hired by industry, with a reference line for the total size of computer science research faculty. (Bottom) The total number of model parameters (a rough proxy for compute) for image recognition on ImageNet (see supplementary materials).

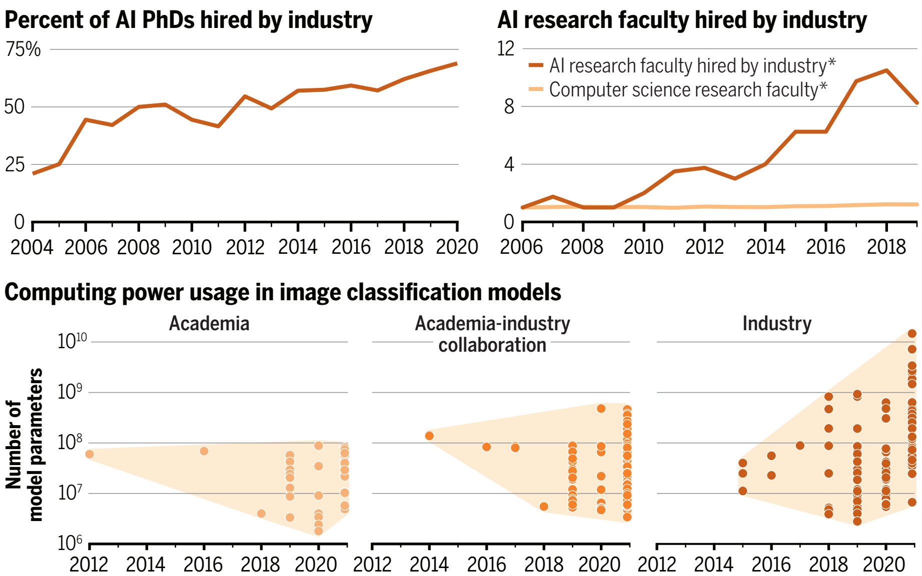

*Data normalized to 2006.

GRAPHIC: ADAPTED FROM AHMED *ET AL* BY K. FRANKLIN/*SCIENCE*

POLICY IMPLICATIONS

Industry's increasing investment in AI has the potential to provide substantial benefits to society through the commercializing of technology. Firms can create better products that benefit consumers [for example, machine translation benefits international trade (*10*)] and can streamline processes that drive down a firm's costs. Industry's investment in AI also produces tools that are valuable to the whole community (such as PyTorch and TensorFlow, which are widely used in academia), hardware that facilitates efficient training of deep-learning models [such as tensor processing units (TPUs)], and publicly accessible pretrained models (such as the Open Pretrained Transformer model by Meta).

At the same time, the concentration of AI in industry is also worrisome. Industry's commercial motives push them to focus on topics that are profit oriented. Often such incentives yield outcomes in line with the public interest, but not always. Were all cutting-edge models from industry, situations would arise when no public-minded alternatives would exist. This possibility raises concerns akin to those about the pharmaceutical industry, where investment disproportionately neglects the needs of lower-income countries (*11*). Recent empirical work finds that "private sector AI researchers tend to specialise in data-hungry and computationally intensive deep learning methods" and that this is at the expense of "research involving other AI methods, research that considers the societal and ethical implications of AI, and applications in sectors like health" (*12*). These questions about the trajectory of AI and who controls it are also important for debates about job replacement and AI-induced inequality. Some researchers are concerned that we may be on a socially suboptimal trajectory (*13*) that focuses more on substituting human labor rather than augmenting human capabilities.

Even with a growing divide between industry and academia, one might imagine that the field could settle into a division of labor similar to that of other disciplines, in which basic research is primarily done in universities, and applied research and development is primarily done by industry. But in AI, such a clear divide does not exist; the same applied models used by industry are often those pushing the boundaries of basic research [a situation akin to what Donald E. Stokes referred to as "Pasteur's Quadrant" because of a similar overlap between applied and basic research in pasteurization (*14*)]. For example, transformers, a type of deep-learning architecture, were developed in 2017 by Google Brain researchers. Not only was this an important step forward in basic research, it was also applied almost immediately in models being used by industry. One benefit of this overlap is that it means that academic work can benefit industry directly (and industry has been supportive of efforts to increase public investment in AI). But this overlap also has a drawback: It means that industry domination of applied work also gives it power to shape the direction of basic research. Given how broadly AI tools could be applied across society, such a situation would hand a small number of technology firms an enormous amount of power over the direction of society. For many around the world, this concern is further heightened because these organizations are "foreign firms" to them. For example, the Future of Life Institute argues that "European companies are not developing general-purpose AI systems and are unlikely to start doing so anytime soon due to their relative competitive disadvantage vis-a-vis American and Chinese players" (*15*).

Even absent public alternatives to industry research, one might imagine that regulation, through auditing or external monitoring of industry AI, could be the solution. For example, in 2018 Joy Buolamwini, an academic, and Timnit Gebru, then a Microsoft employee, documented gender and racial biases in commercial face recognition systems (*16*). Establishing monitoring or auditing requirements (such as those in the Liability Rules for AI in Europe) can help mitigate these types of harms. However, if academics do not have access to industry AI systems, or the resources to develop their own competing models, their ability to interpret industry models or offer public-interest alternatives will be limited. This is both because academics would be unable to build the large models that seem to be needed for cutting-edge performance, but also because some useful capabilities of AI systems seem to be "emergent," meaning that systems only gain these capabilities once they are particularly large (*17*). Some negative characteristics of models also seem to scale with size [for example, toxicity in AI-generated language, and stereotyping (*7*)]. In either case, academics without access to sufficient resources would be unable to meaningfully contribute to these important areas.

Around the world, this concern about academia's resource disadvantage in AI research is being recognized, and policy responses are beginning to emerge. In the United States, the National AI Research Resource (NAIRR) task force (*18*) has proposed the creation of a public research cloud and public datasets. In Canada, the national Advanced Research Computing platform has been serving the country's academics and has been oversubscribed since its launch almost a decade ago. Chinese authorities have recently approved a "national computing power network system" (*19*) that will enable academics and others to access data and computing power. In Europe, similar initiatives have yet to emerge, although there is a clear recognition of the risk. As French president Emmanuel Macron said, "if you want to manage your own choice of society, your choice of civilization, you have to be able to be an acting part of this AI revolution" (*20*). For many countries, the scale needed for these types of investments may be daunting. In such cases, the key question for policy-makers will be whether they can pool sufficient resources with like-minded collaborators to reach the scale needed to create AI systems that reflect their own priorities.

Computing power is not the only area in which remedies should be offered. Steps must also be taken for the other key inputs to AI. Building public datasets will be important but also a challenge because modern AI training datasets can be billions of documents. Of particular interest should be creating important datasets for which there are no immediate commercial interests. It is also important to provide the resources to keep top AI researchers in academia. For example, the Canada Research Chairs Program (CRCP), which provides salaries and research funds, has proven to be a successful means of attracting and retaining top talent in Canada.

For policy-makers working on this problem, the goal should not be that academia does a particular share of research. Instead, the goal should be to ensure the presence of sufficient capabilities to help audit or monitor industry models or to produce alternative models designed with the public interest in mind. With these capabilities, academics can continue to shape the frontier of modern AI research and benchmark what responsible AI should look like. Without these capabilities, important public interest AI work will be left behind. ■

AI research outputs

(Top) The proportion of papers at leading AI conferences that have at least one industry co-author. (Middle) The fraction of the largest AI models that are from industry (3-year rolling average). (Bottom) Periods when the state-of-the-art model for leading AI benchmarks were from academia, industry, or collaborations (see supplementary materials).

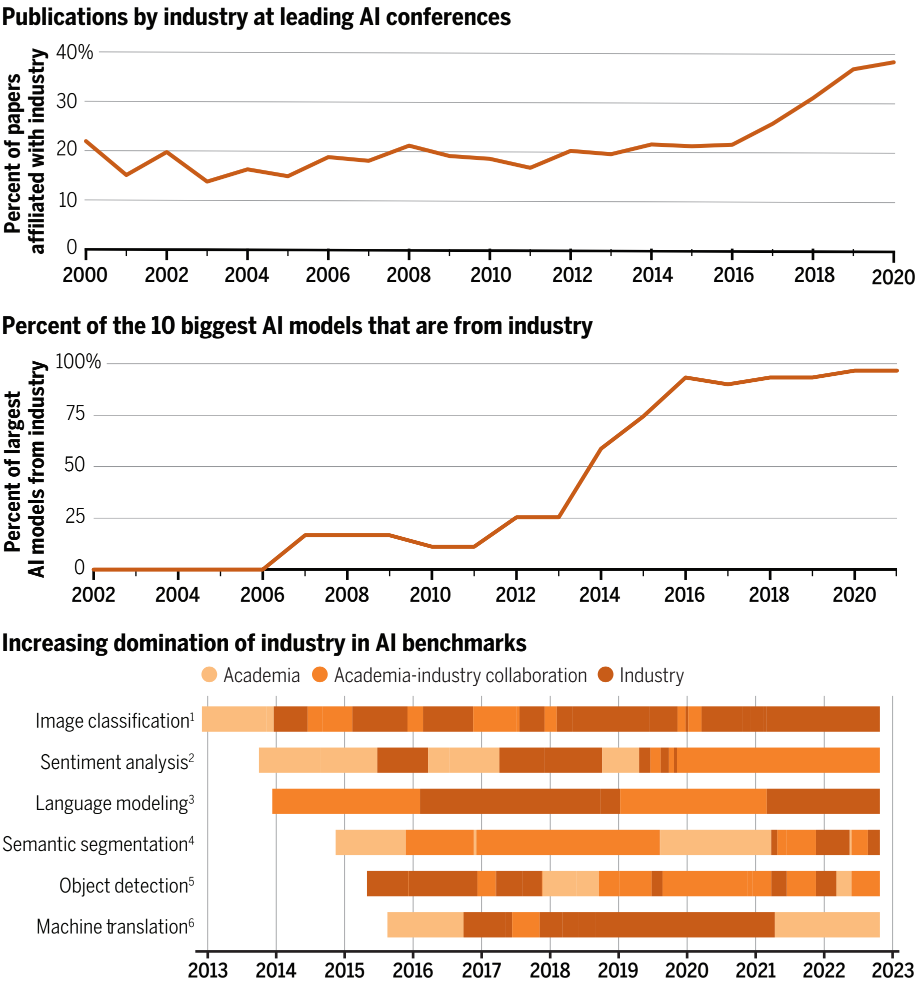

Benchmarks: [1]ImageNet. [2]SST-2. [3]One Billion Word. [4]ADE20K. [5]COCO test-dev. [6]WMT2014.

REFERENCES AND NOTES

1. R. Shokri, V. Shmatikov, in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (ACM, 2015), pp. 1310–1321.
2. R. Jurowetzki, D. Hain, J. Mateos-Garcia, K. Stathoulopoulos, arXiv:2102.01648 [cs.CY] (2021).
3. "UK universities alarmed by poaching of top computer science brains," *Financial Times*, 9 May 2018.
4. T. Evas *et al.*, "AI Watch: Estimating AI investments in the European Union" (Publications Office of the European Union, 2022).
5. G. Brockman, I. Sutskever, OpenAI, "Introducing OpenAI," *OpenAI*, 11 December 2015.
6. G. Brockman, I. Sutskever, OpenAI, "OpenAI LP," *OpenAI*, 11 March 2019.
7. D. Ganguli *et al.*, in *2022 ACM Conference on Fairness, Accountability, and Transparency* (ACM, 2022), pp. 1747–1764.
8. J. Sevilla, L. Heim, M. Hobbhahn, T. Besiroglu, A. Ho, arXiv:2202.05924 [cs.LG] (2022).
9. N. C. Thompson, K. Greenewald, K. Lee, G. F. Manso, arXiv:2007.05558 [cs.LG] (2020).
10. E. Brynjolfsson, X. Hui, M. Liu, *Manage. Sci.* **65**, 5449 (2019).
11. P. Trouiller *et al.*, *Global Health* 267 (2017).
12. J. Klinger, J. Mateos-Garcia, K. A. Stathoulopoulos, *SSRN* 10.2139/ssrn.3698698 (2020).
13. E. Brynjolfsson, *Daedalus* **151**, 272 (2022).
14. D. E. Stokes, *Pasteur's Quadrant: Basic Science and Technological Innovation* (Brookings, 1997).
15. Future of Life Institute, "Emerging non-European monopolies in the global AI market" (Future of Life Institute, 2022); https://bit.ly/3k2ckD9.
16. J. Buolamwini, T. Gebru, in *Conference on Fairness, Accountability and Transparency* (PMLR, 2018), pp. 77–91.
17. J. Wei *et al.*, arXiv:2206.07682 [cs.CL] (2022).
18. D. E. Ho, J. King, R. C. Wald, C. Wan, "Building a national AI research resource," white paper (Stanford University Human-Centered Artificial Intelligence, 2021)
19. CAICT, "White paper on China's computing power development index" (CAICT, 2021); https://cset.georgetown.edu/wp-content/uploads/t0402_compute_white_paper_EN-2.pdf.
20. N. Berggruen, N. Gardels, *The Washington Post*, 27 September 2018.

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.ade2420

10.1126/science.ade2420

GRAPHIC: ADAPTED FROM AHMED *ET AL.* BY K. FRANKLIN/*SCIENCE*