





The Steep Cost of Capture

🔹 Meredith Whittaker, New York University

This is a perilous moment. Private computational systems marketed as artificial intelligence (AI) are threading through our public life and institutions, concentrating industrial power, compounding marginalization, and quietly shaping access to resources and information.

In considering how to tackle this onslaught of industrial AI, we must first recognize that the “advances” in AI celebrated over the past decade were not due to fundamental scientific breakthroughs in AI techniques. They were and are primarily the product of significantly concentrated data and compute resources that reside in the hands of a few large tech corporations. Modern AI is fundamentally dependent on corporate resources and business practices, and our increasing reliance on such AI cedes inordinate power over our lives and institutions to a handful of tech firms. It also gives these firms significant influence over both the direction of AI development and the academic institutions wishing

to research it. Meaning that tech firms are startlingly well positioned to shape what we do—and do not—know about AI and the business behind it, at the same time that their AI products are working to shape our lives and institutions.

Examining the history of the U.S. military’s influence over scientific research during the Cold War, we see parallels to the tech industry’s current influence over AI. This history also offers alarming examples of the way in which U.S. military dominance worked to shape academic knowledge production, and to punish those who dissented.

Today, the tech industry is facing mounting regulatory pressure, and is increasing its efforts to create tech-positive narratives and to silence and sideline critics in much the same way the U.S. military and its allies did in the past. Taken as a whole, we see that the tech industry’s dominance in AI research and knowledge production puts critical researchers and advocates within, and beyond, academia in a treacherous

Insights

- Big tech’s control over AI resources made universities and other institutions dependent on these companies, creating a web of conflicted relationships that threaten academic freedom and our ability to understand and regulate these corporate technologies.
- To ensure independent and rigorous research and advocacy capable of understanding and checking these technologies, and the companies behind them, we need to organize, within tech and within the university.

IMAGE BY TANOR / SHUTTERSTOCK.COM

position. This threatens to deprive frontline communities, policymakers, and the public of vital knowledge about the costs and consequences of AI and the industry responsible for it—right at the time that this work is most needed.

Reviewing the extent of large tech firms' current influence over AI and AI research, it's helpful to begin with a brief history of the current turn to AI. Given that the AI field is nearly 70 years old and has gone through several "AI winters," why did AI get big in the past decade? And what are we talking about when we talk about AI? Answering these questions highlights the mutability of the term *AI*. It also focuses our attention on the centrality of concentrated corporate resources to the current AI boom, and how monopolistic control of these resources gave a handful of tech companies the authority to (re)define the AI field, while enclosing knowledge about AI systems behind corporate secrecy.

In 2012, a Toronto-based research team created an algorithm named AlexNet that won the ImageNet Large Scale Visual Recognition Challenge. This marked a key moment in recent AI history and was a big deal in the tech industry. It demonstrated that supervised machine learning was shockingly effective at predictive pattern recognition when trained using significant computational power and massive amounts of labeled data [1]. The AlexNet algorithm relied on machine learning techniques that were nearly two decades old. But it was not the algorithm that was a breakthrough: It was what the algorithm could do when matched with large-scale data and computational resources.

AlexNet mapped a path forward for large tech companies seeking to cement and expand their power. The resources on which AlexNet's success was dependent were those large tech firms already controlled: vast computational infrastructure, massive amounts of data (and systems in place to process and store it), entrenched market reach that ensured persistent data collection, and the capital to hire and retain scarce talent. Yoshua Bengio, one of the forerunners of AI research, put it simply: "The [computing] power, the expertise, the data are all concentrated in the hands of a few companies" [2].

The year 2012 showed the commercial potential of supervised machine learning, and the power of the term *AI* as a marketing hook. Tech companies quickly (re)branded machine learning and other data-dependent approaches as AI, framing them as the product of breakthrough scientific innovation. Companies acquired labs and start-ups, and worked to pitch AI as a multitool of efficiency and precision, suitable for nearly any purpose across countless domains. When we say AI is everywhere, this is why.

The rhetoric and capital flowing from these firms served to redefine the AI research field, flooding it with funding and focusing the field's attention on data- and compute-intensive techniques and research questions. University labs and start-ups that wanted to develop and study AI found themselves requiring access to costly cloud-compute environments

operated by big tech firms and scrambling for access to data, a dynamic that has only intensified since 2012. J. Nathan Matias, a Cornell professor and leader of the Citizens and Technology Lab, points to the extent of this dependency when he remarks that "some fields couldn't exist without close industry ties" [3].

This doesn't mean that researchers within these domains are compromised. Neither does it mean that there aren't research directions that can elude such dependencies. It does mean, however, that the questions and incentives that animate the field are not always individual researchers' to decide. And that the terms of the field—including which questions

are deemed worth answering, and which answers will result in grants, awards, and tenure—are inordinately shaped by the corporate turn to resource-intensive AI, and the tech-industry incentives propelling it.

A recent move by Stanford illustrates this dynamic. In August 2021, the university announced the new Center for Research on Foundation Models (CRFM), whose launch was accompanied by a 100-plus author report that characterizes these models as a "paradigm shift" in AI significant enough to justify a new and costly research center [4]. Throughout the report, foundation models are framed as inevitable, cutting edge, and the product of scientific progress. What are foundation models? You'll be forgiven for not knowing. The name was coined by Stanford for its report and CRFM launch materials, rebranding what were previously known as large language models (LLMs). LLMs—think GPT-3 and BERT, among others—are some of the most data- and compute-intensive techniques in AI, and thus among the most industry-captured. They've also garnered a lot of recent media attention, and been subject to sustained

critique focused on significant bias, environmental costs, and concentrated power [5].

Beyond simply valorizing industry-captured techniques as cutting edge, Stanford's rebranding works to distance LLMs from this legacy of criticism. And while the report acknowledges that "research on building foundation models themselves has occurred almost exclusively in industry," it frames questions of concentrated power not as issues that should make us reconsider reliance on these technologies, but rather as problems that can be solved by easing gatekeeping restrictions so that institutions like Stanford also get a piece: "[I]ndustry ultimately makes concrete decisions about how foundation models will be deployed, but we should also lean on academia, with its disciplinary diversity and non-commercial incentives" [4].

Efforts to expand access to AI research also follow this pattern, taking data- and compute-intensive forms of AI as a given and focusing solely on how to get more people access to these concentrated resources. By examining one proposed "solution" to this narrowly framed problem, we come face-to-face with the extent of industry capture.

Tech firms are startlingly well positioned to shape what we do—and do not—know about AI and the business behind it, at the same time that their AI products are working to shape our lives and institutions.

In March 2020, the National Security Commission on Artificial Intelligence (NSCAI), chaired by former Google CEO Eric Schmidt and helmed by other tech executives, recommended that the U.S. government fund what it termed a national AI research infrastructure, in the name of “democratizing” access to AI research. This recommendation was picked up in the 2021 National Defense Authorization Act (NDAA), which mandates the creation of “a system that provides researchers and students across scientific fields and disciplines with access to compute resources, co-located with publicly available, artificial intelligence–ready government and non-government data sets” [6]. Following the NDAA directive, the White House Office of Science and Technology Policy and the National Science Foundation recently launched the National AI Research Resource (NAIRR), appointing a task force to architect its policies and implementation.

Why would a conflicted government body populated by tech executives recommend “democratizing” access to the infrastructures at the heart of their concentrated power? Because this proposal wouldn’t actually reduce that power. Indeed, if implemented, it would almost certainly entrench and expand large tech firms’ power and reach. Big tech’s domination over the infrastructure of AI research and development extends beyond providing “neutral platforms.” These companies control the tooling, development environments, languages, and software that define the AI research process—they make the water in which AI research swims. Even if it were desirable (which, given AI’s harms and flaws, must be open to question), there is no plausible scenario in which a national research infrastructure could be meaningfully constructed outside of the current tech-industry ecosystem. Doing so would require rolling a new platform, developing software, and habituating tens of thousands of researchers to new tooling and interfaces while hiring the thousands of site-reliability engineers, software developers, quality assurance testers, and support personnel necessary to maintain such a large and expensive system in perpetuity.

In practice, then, these proposals to “democratize” access to AI research infrastructures amount to calls to subsidize tech giants further by licensing familiar infrastructure from these firms in ways that allow them to continue defining the terms and conditions of AI and AI research. All while centers like Stanford’s new CRFM are poised to further entrench such dominance by presenting industry-dependent AI techniques as the cutting edge of AI research.

From industry-sponsored Ph.D. programs to initiatives that place tech-company offices literally in the middle of universities to the National Science Foundation partnering with Amazon to define the parameters of “fairness” in AI and awarding grants to those who meet their positivist criteria [7], we see myriad schemes to draw academia closer to tech companies. These extend to dual-affiliation arrangements, which are common in the AI field and amount

to companies hiring AI professors while allowing them to retain their academic titles and appointments. Dual-affiliated scholars draw a tech company salary, work closely with tech employees, and avail themselves of corporate research infrastructures, all while publishing research under a university imprimatur. Such arrangements help shield companies from accusations that they’re contributing to brain drain by hiring researchers away from universities. They also allow the companies to enlist practitioners in answering questions interesting to tech firms, while creating the appearance of academic disciplines organically and independently invested in those same questions.

That these conflicted arrangements are treated as standard practice is likely related to the clarity with which AI researchers and universities recognize their reliance on large companies and the resources they control. Maja Pantic, a professor of machine learning who works for Samsung and is dual-appointed at Imperial College London, told the *Financial Times* that she “simply couldn’t continue working solely in academia, we don’t have the computing resources, I couldn’t pay people to work for me, and I didn’t have money to create processing power” [2]. She and many others face the choice of either allying with a company, with all the tacit conditions such dependency requires, or being unable to do the kind of work that equals prestige and scholarly success.

The extent of the tech industry’s influence over the AI research domain has parallels in the U.S. military’s dominance over scientific research during the Cold War. Tech firms are drawing from a similar playbook.

Writing in 1946, just after World War II concluded, General Dwight D. Eisenhower drafted a memo titled “Scientific and Technical Resources as Military Assets”

that proposed bringing scientists and researchers more directly into U.S. military planning, arguing that this would allow the military to build trust with scientists, to have a front-row seat to novel scientific developments, and—via funding and collegial proximity—let the U.S. military steer research questions in ways that would ensure scientists are “familiar with our fundamental problems” [8]. Three years after the memo, in 1949, the U.S. obtained evidence that the Soviet Union was testing nuclear weapons. This helped put Eisenhower’s plan into action, catalyzing the creation of research offices and agencies across military branches dedicated to funding and shaping research [9].

Germane to our case is the power that this gave the U.S. military over the direction of scientific research and the institutions that housed it. This influence was applied not only to ensure that academic research was animated by U.S. military questions and concerns but also to punish whistleblowers, chill dissent, and incentivize complacency in the face of overblown claims masked in scientific authority. It is here, in these darker histories, that we confront the steep cost of capture—whether military or industrial—and its

These companies control the tooling, development environments, languages, and software that define the AI research process—they make the water in which AI research swims.

perilous implications for academic freedom and knowledge production capable of holding power to account.

Aldric Saucier was a scientist working for the U.S. Army on then-President Ronald Reagan's controversial Strategic Defense Initiative (SDI). SDI was a massive military initiative that enlisted scientists across the country in its attempt to build a ballistic missile shield. The proposal was fantastical, and many in the research community viewed it as scientifically unfounded and likely to increase the chances of nuclear war. When Saucier reported waste, fraud, and hyperbole within the program, then-Secretary of Defense Dick Cheney oversaw his firing, along with a campaign to publicly discredit his scientific expertise [10]. Outside of military-run research labs, dissenters were also threatened. Scientists across universities organized a boycott of SDI research and funding. In response, Indiana Congressman Dan Burton threatened to cut funding to universities where professors refused SDI-related grants. Meanwhile, university leadership at Lawrence Livermore National Laboratory, long allied with the U.S. military's nuclear visions, worked to push out physicist Hugh DeWitt, who spoke out against the lab's role in exacerbating the arms race. While DeWitt managed to retain his position, he was denied raises and promotions and was excluded from interesting work [11]. Undersecretary of Defense for Research and Engineering Donald Hicks—at the time in charge of Pentagon research contracts—went as far as to publicly intimidate researchers. In an interview, Hicks stated that, while professors were permitted to speak out in a “free country,” they were also “free to keep their mouths shut...I'm also free not to give them money” [12]. *The Wall Street Journal* printed an editorial applauding Hicks.

With the benefit of hindsight, we know that the SDI's critics were largely correct about the flaws and perilous logic driving the program. But their evidence-supported arguments and analysis did not save them from retaliation, financial coercion, and denigration, even within institutions purportedly dedicated to academic freedom.

That a handful of large tech firms currently have similar leverage in relation to AI research should alarm us, especially given mounting evidence of tech's desire to shape a positive narrative in response to growing regulatory and public pressure, alongside the industry's clear willingness to silence and punish critics. Examples abound, from Facebook's revoking data access to NYU researchers examining the company's role in the January 6 insurrection; to Google instructing in-house researchers to “strike a positive tone” in their findings [13], while marshaling external “academic allies” to raise questions about regulatory intervention [14]; to Amazon's directing specious attacks against young Black researchers who revealed racist logics in their products, while retaliating against workers who organized against the company's climate harms. Google also fired Timnit Gebru, after demanding she and her coauthors remove their names

from a paper critical of LLMs that are core to Google's product road map, and that Stanford has recently rebranded and revalorized. The list goes on, providing a good barometer for where these companies draw the line—research and dissent that threatens growth and revenue.

In addition to punishing dissent and denigrating research they find threatening, tech companies are working to co-opt and neutralize critique. They do this in part by funding and elevating their weakest critics, often institutions and coalitions that focus on so-called AI ethics, and frame issues of tech power and dominance as abstract governance questions that take the tech industry's current form as a

In addition to punishing dissent and denigrating research they find threatening, tech companies are working to co-opt and neutralize critique.

given and AI's proliferation as inevitable.

In parallel, tech firms also champion technocratic remedies such as “AI bias bounties” and fairness fixits that stage tech-enabled discrimination as a problem of bad code and “buggy” engineering [15]. Such approaches make great PR. They also serve to cast elite engineers as the arbiters of “bias,” while structurally excluding scholars and advocates who don't have computer science training, but whose focus on the racialized power asymmetries and political economy of AI are essential for understanding and addressing AI harms.

All of this is happening against a backdrop in which academic institutions, increasingly run like businesses in search of large investors, find it hard to ignore the financial and reputational advantages that tech partnerships and funding bring. This dynamic is compounded by the increasing precariousness of academic jobs, in which fewer and fewer academic workers have the job security or union solidarity necessary to safely contest policies that might compromise academic freedom. This gives tech companies increasing leverage not simply over research they fund directly but also over decisions about which work is

included and excluded at the university overall.

Nor can we ignore the attack currently underway against work illuminating structural racism and inequality. Far-right think tanks and Republican apparatchiks are pressuring educational institutions to eliminate pedagogy and research that centers racial justice, which they bucket sloppily under the term *critical race theory*. This attack on intellectual freedom matters for many reasons. Scholarship and movement practice attentive to racial capitalism and structural racism has provided many of the methods and frameworks core to critical work engaging the social implications of tech. It has helped focus tech critique beyond shallow notions of bias to examinations of the ways in which these technologies replicate patterns of racial marginalization and concentrate power in the hands of those with the scarce and expensive resources to develop and deploy AI. This line of critique has already powerfully influenced public discourse and the global regulatory agenda in ways that tech firms are actively resisting.

So what is the path forward? To begin, scholars, advocates, and policymakers who produce and rely on tech-critical work

must confront and name the dynamic of tech capture, co-optation, and compromise head-on, and soon. This means incorporating reflexive critiques of the conditions and of knowledge creation, and the compromises and trade-offs faced by knowledge workers over whom interested institutions have power. Given the politics of collegial proximity that inform academic prestige networks while working to blur the lines between academic and industry workers, this is certain to be uncomfortable. But naming these dynamics is the only way to address them and to stage questions that allow us to envision and demand alternative futures.

This is exactly the kind of intervention that is currently imperiled by the industry capture of AI research. So, how to support such critical work and protect those doing it both within and beyond the academy?

Here we turn to the central role of organized workers, those in tech who have made inroads across the industry over the past five years and academic workers organizing in an environment where the myth of individual genius serves to uphold inequity, even as the labor market crumbles. Academic workers' struggle against the precariousness of the profession is also a struggle for academic freedom. Stable career opportunities and more democratic control over the university would help tip the balance of influence away from wealthy donors and large industry patrons. Organized tech workers, for their part, have a role to play in checking the power of their industry from within, fighting for more control over the work they do, and working to curb the influence of their employers in academia and beyond. In this capacity, we could imagine organized researchers and scientists demanding a redirection of the lavish U.S. congressional appropriations currently underwriting the national AI research infrastructure, using their expertise and standing to instead demand endowments in support of truly public and accessible universities, workers' schools, and programs that integrate communities with lived experience into the pantheon of learners and experts interrogating tech [16]. Of course, given that the tech industry not only writes policy via congressionally appointed councils like NSCAI but also spends more than big oil and big tobacco on lobbying, it is clear any such intervention will require serious organized struggle.

A future where the U.S. Congress richly supports truly democratic and independent critical work does not appear near on the horizon. But organizing within academia and tech workplaces can also help us protect ourselves and the public interest in the short term by preparing us to stand up for one another in the face of institutional pressure, and developing muscles of care and mutual accountability that let us name dynamics of coercion and capture more safely. This won't be easy; it will require confronting cultures of competition and turf-claiming that scar both tech and academic workplaces. But the stakes are too high, and those aiming to shape what we do (and do not) know about AI and the industry responsible are well organized and supremely well resourced. In short, this is a battle of power, not simply a contest of ideas, and being right without the strategy and solidarity to defend our position will not protect us.

ACKNOWLEDGMENTS

A sincere thank you to Theodora Dryer, Lilly Irani, Nantina

Vgontzas, Sarah Myers West, and Michael Decker for their careful and generous close reading and suggestions, and to Nicole Weber for excellent research assistance. J. Khadijah Abdurahman's rigorous and careful editing was also fundamental in shaping this piece.

ENDNOTES

1. Krizhevsky, A., Sutskever, I., and Hinton, G.E. ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25 (NIPS 2012); <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>
2. Murgia, M. AI academics under pressure to do commercial research. *Financial Times*. Mar. 13, 2019; <https://www.ft.com/content/94e86cd0-44b6-11e9-a965-23d669740bfb>
3. Matias, J.N. Why we need industry-independent research on tech & society. Citizens and Tech Lab. Jan. 2020; <https://citizensandtech.org/2020/01/industry-independent-research/>
4. Bommasani, R. et al. On the opportunities and risks of foundation models. 2021; *arXiv preprint arXiv:2108.07258*
5. Bender, E.M., Gebru, T., McMillan-Major, A., and Shmitchell, S. On the dangers of stochastic parrots: Can language models be too big? *Proc. of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, 2021, 610–623; <https://dl.acm.org/doi/10.1145/3442188.3445922>
6. <https://www.congress.gov/116/bills/hr6395/BILLS-116hr6395enr.pdf>
7. https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505651
8. Melman, S. *Pentagon Capitalism: The Political Economy of War*. McGraw-Hill, New York, 1970.
9. Krinsky, R. Swords and sheepskins: Militarization of higher education in the United States and prospects of its conversion. *Bulletin of Peace Proposals* 19, 1 (1988), 33–51; <http://www.jstor.org/stable/44481371>
10. Lardner, Jr., G. Army accuses SDI critic of falsifying credentials. *Washington Post*. Apr. 14, 1992; <https://www.washingtonpost.com/archive/politics/1992/04/14/army-accuses-sdi-critic-of-falsifying-credentials/13ffe75f-50f8-4654-9027-536c30880c13/>
11. Martin, B. Science: contemporary censorship. In *Censorship: A World Encyclopedia, Vol 4*. D. Jones, ed. Fitzroy Dearborn, London, 2001, 2167–2170; <https://documents.uow.edu.au/~bmartin/pubs/01cescience.html>
12. Hiatt, F. Official seeks like minds in 'Star Wars'. *Washington Post*. May 13, 1986; <https://www.washingtonpost.com/archive/politics/1986/05/13/official-seeks-like-minds-in-star-wars/28ddbdc0-d55f-4cac-a742-b17176e6adcf/>
13. Dave, P. and Dastin, J. Google told its scientists to 'strike a positive tone' in AI research – documents. *Reuters*. Dec. 23, 2020; <https://www.reuters.com/article/us-alphabet-google-research-focus/google-told-its-scientists-to-strike-a-positive-tone-in-ai-research-documents-idUSKBN28X1CB>
14. Satariano, A., and Stevis-Gridneff, M. Big tech turns its lobbyists loose on Europe, alarming regulators. *New York Times*. Dec. 14, 2020; <https://www.nytimes.com/2020/12/14/technology/big-tech-lobbying-europe.html>
15. Vanian, J. Why Microsoft and Twitter are turning to bug bounties to fix their A.I. *Fortune*. Aug. 10, 2021; <https://fortune.com/2021/08/10/why-microsoft-and-twitter-are-turning-to-bug-bounties-to-fix-their-a-i/>
16. I want to thank particularly Lilly Irani and Nandina Vgontzas for their care and intellectual generosity, which helped me map out this set of points.

🔗 **Meredith Whittaker** is the Minderoo Research Professor at NYU and faculty director of the AI Now Institute. Her research focuses on the social implications of AI and the industry responsible for it. Prior to joining NYU, she founded and led Google's Open Research Group. → mw3900@nyu.edu