Issue Brief

AI Faculty Shortages

Are U.S. Universities Meeting the Growing Demand for AI Skills?

Authors
Remco Zwetsloot
Jack Corrigan

CSET CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

July 2022

Table of Contents

Executive Summary ... 2
Introduction ... 3
Do Teaching Capacity Gaps Exist? ... 5
- Enrollments Surge ... 7
- Access Restrictions ... 9
- Concerns Among Academics ... 11
- Summary ... 12

The Causes of Teaching Capacity Gaps ... 13
- Hypothesis #1: Private-Sector Poaching of Professors ... 13
- Hypothesis #2: Limited Inflow from PhD Graduates ... 15
- Hypothesis #3: Universities Not Responding to Booming Enrollment ... 18
- Summary ... 21

The Costs of Teaching Capacity Gaps ... 22
- Consequences for Students ... 22
- Consequences for Research ... 24
- Consequences for National Security and the Economy ... 25

Policy Issues and Options ... 26
- Increasing Funding for Education and Research ... 26
- Encouraging Companies to Support University Teaching ... 29
- Facilitating Access to Government Data and Computational Power ... 30
- Other Policy Areas ... 31

Conclusion ... 33
Acknowledgments ... 34
Appendix A: Open Questions and Future Research ... 35
Appendix B: CRA Taulbee Survey ... 38
Appendix C: A Brief History of CS Enrollment Booms and Teaching Capacity ... 39
Appendix D: Data Tables ... 41
Endnotes ... 46

Executive Summary

Universities play an indispensable role in developing artificial intelligence talent, but mounting evidence suggests that computer science departments across the United States do not have enough faculty to meet the growing demand for AI education. The goal of this paper is to help policymakers better assess the scale, causes, and consequences of these "teaching capacity gaps" in AI, and to present policy levers that could help increase the capacity of universities to train the next generation of AI specialists.

While it is difficult to measure the potential mismatch between the supply of instructors and the demand for AI education, available evidence suggests there is indeed a gap. Over the last decade, the increase in computer science enrollments has far outpaced the growth in computer science faculty, who are responsible for much of the AI instruction at U.S. universities. Universities have started restricting access to CS courses, and academic leaders have publicly lamented the difficulty of recruiting and retaining qualified AI faculty.

These teaching capacity gaps can have detrimental effects on students' educational experience, change the quantity and trajectory of academic research, and hamper the country's efforts to build a robust AI workforce. Despite its wide-ranging repercussions, the lack of AI-teaching capacity at U.S. universities has received relatively little attention from policymakers and analysts. While we cannot say for certain that teaching capacity gaps exist, we consider this evidence to be sufficiently strong to warrant further investigation and policy attention.

When teaching capacity gaps are discussed, policymakers often attribute them to a shortage of AI faculty caused by a combination of tech companies "poaching" professors and recent PhD graduates losing interest in academic careers. This argument is based on the implicit notion that AI experts are choosing to forgo academic jobs for better-paying, less-restrictive careers in industry.

While these experts may have correctly identified teaching capacity gaps at U.S. universities, available data suggests they may have misattributed the root cause. Little evidence suggests the outflow of AI experts from academia to industry has distorted the job market, or even that industry hiring is uniformly negative. Furthermore, while a greater share of PhD graduates is indeed flocking to industry, the available data shows that many are still interested in pursuing academic careers, and the total number of PhD graduates who enter academia each year has remained relatively constant. However, we did find evidence that universities have not created enough new faculty positions to accommodate students' growing interest in the field. These findings suggest that universities would be able to close their teaching capacity gaps if they created more faculty positions, though budgetary constraints may limit their ability to do so.

Understanding the specific factors creating teaching capacity gaps is critical, as different causes demand different policy solutions. In pursuing such measures, federal agencies must consider whether their interventions are intended to increase research capacity or teaching capacity. While both research and teaching are integral to universities' missions, they can at times be at odds with one another. If targeted incorrectly, policies meant to increase teaching capacity can in reality exacerbate the problem, and potentially create AI faculty shortages. Still, there are a variety of measures policymakers can explore to close teaching capacity gaps, such as creating federal grants to facilitate faculty hiring, incentivizing industry to support university education, and expanding access to government data and computing resources.

Introduction

In many ways, universities are the powerhouses that fuel the AI innovation ecosystem. They produce novel research that pushes the boundaries of the field, and they generate the steady stream of developers, engineers, and entrepreneurs necessary to put those discoveries into practice. These two products of academia—talent and research—are critical for maintaining a globally competitive tech industry and strengthening the U.S. national security apparatus.

Increasing production of AI talent has become an especially salient goal as policymakers work to secure the country's leadership in the ever-more-competitive field. The number of graduates a university can produce is closely related to its faculty size: schools that employ more computer science professors can likely offer more AI training to their students.[*1] Conversely, students interested in computer science programs might be excluded from enrolling if universities do not have enough faculty. While countless national policies and strategy documents seek to expand the U.S.

---

[*] In this paper, we use computer science education as a proxy for AI education. Many AI courses are taught in computer science departments, and AI specialists account for a growing portion of overall CS faculty and PhD graduates: "Artificial Intelligence Index Report 2021" (Stanford University, 2021), pg. 113, https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf#page=113. For more, see Appendix B.

talent pipeline and research ecosystem, most do not recognize faculty supply as a potential bottleneck for growth in emerging fields such as AI.[2]

Today, evidence suggests that the production of AI talent in the United States is being constrained by a lack of university faculty. (In this paper, we refer to situations in which demand for AI-related education exceeds the supply of qualified instructors and university courses as “teaching capacity gaps.”) Enrollments in CS programs nationwide have surged over the past decade while CS departments have expanded at a much slower rate. As we will discuss in the following section, leaders across academia now regularly discuss the challenges of hiring enough computer science professors to meet students’ growing interest in the field, and some universities have begun limiting students’ access to CS courses. (Throughout this paper, we use data on the broader pool of CS faculty to highlight trends in the narrower field of AI.[3]) In other words, U.S. universities could be graduating more students with AI skills if they had more faculty to train them.

While there is no consensus on what causes teaching capacity gaps, many experts attribute them to a shortage of qualified or interested candidates to fill available CS faculty positions. To explain these shortages, policymakers and tech industry leaders tend to focus on one particular trend: companies poaching AI professors and PhD graduates from universities. Faculty, their argument goes, are lured into industry by high salaries and greater research opportunities, vacating posts that universities then struggle to fill because newly graduated PhD holders would also prefer to work in the private sector.

This argument is gaining traction among tech leaders and policymakers. In a 2017 interview, Oren Etzioni, CEO of the Allen Institute of Artificial Intelligence, remarked “there is a giant sucking sound of academics going into industry.”[4] The National Security Commission on Artificial Intelligence (NSCAI) echoed this idea in its final report, warning that “brain drain from academic institutions to the private sector threatens to hollow out the foundations of the United States’ advantage in basic AI research: its universities.”[5] The Congressional Research Service has also called attention to industry’s poaching of AI faculty, noting this practice of “eating the seed corn” may reduce the country’s capacity to train the next generation of AI experts.[6]

While blaming teaching capacity gaps on industry poaching may make for an appealing narrative, we find the story is not so simple. While there is some evidence that industry poaching does happen, it does not appear to be the primary factor constraining the supply of qualified AI professors.

The goal of this paper is to help policymakers better assess the scale, causes, and consequences of teaching capacity gaps in AI, and the policy levers that could increase universities' capacity to train the next generation of AI specialists. We begin by examining whether AI faculty shortages exist in the U.S. higher education system using data from the Computing Research Association (CRA); we then evaluate the merits of three different explanations for the lack of qualified AI instructors in academia, including both faculty shortages and alternative hypotheses; we go on to discuss the impacts of insufficient teaching capacity on students, AI research, and U.S. economic and national security; finally, we consider what policymakers can do to close these gaps, and we highlight areas for additional research.

Do Teaching Capacity Gaps Exist?

In practice, measuring the AI faculty supply and student demand for AI education is not a straightforward task.[7]

Artificial intelligence is not a neatly defined field, and AI-related concepts and techniques are taught across many classes and departments.[*] As such, even counting the number of current and potential AI faculty can be difficult.[8] Box 1 outlines some of the challenges researchers face when attempting to quantify the supply of AI faculty. Additionally, universities often do not systematically track and publish data on student interest and teaching capacity in computer science, which complicates efforts to determine how much demand is going unmet. Appendix A discusses outstanding policy-relevant questions that would benefit from follow-on research. Despite this lack of data, we identified a variety of indirect evidence that suggests universities are struggling to meet students' growing demand for AI education.

[*] For example, machine learning, the most popular subfield of AI today, is a standard part of the curriculum across computer science, statistics, bioinformatics, economics, and many other natural and social science fields.

Box 1: Difficulties Defining and Measuring "AI Faculty"

Counting the number of AI faculty in the United States is a challenge. On the definitional side, faculty at universities can fulfill several different roles and responsibilities, including:

- Teaching introductory courses to new students and non-majors
- Teaching intermediate-level courses to majors and master's students
- Designing and teaching new courses, and overseeing the curriculum
- Supervising and mentoring PhD students, and running research groups

These different roles require different backgrounds and skills. For example, teaching an introductory course requires less experience than supervising PhD students. Moreover, because AI is an interdisciplinary field, relevant courses are taught across different departments. This makes it difficult to define who counts as "AI faculty" (*current* supply) or who has sufficient skill and experience to be hired as AI faculty (*potential* supply).

When it comes to measurement (i.e., counting the number of AI faculty in the United States), these definitional problems are compounded by a lack of data. We are not aware of a national registry of courses taught or course instructors, which makes measuring teaching activity difficult. Given this dearth of data, we use data on computer science faculty to understand trends in the narrower field of AI.

There is better data on scientific publications by individual faculty, but these figures mainly measure research output and PhD-supervision capacity, not necessarily teaching skills or capacity. Table 1 shows the number of full-time, tenure-track faculty who have published papers at several highly regarded AI conferences. Because of the shortcomings discussed above, we recommend interpreting these numbers as lower-bound estimates of overall faculty counts. However, the numbers are likely to be a reasonable approximation of the number of professors who can supervise AI PhD students in the United States. Based on the available data and adding some room for uncertainty, we estimate there are currently at least 900 tenure-track professors at U.S. universities capable of supervising PhD-level AI research.

Table 1: Number of Full-Time, Tenure-Track CS Faculty at U.S. Universities by Publication Activity

Full-time tenure-track CS faculty with...	Number of CS faculty
At least one AI paper since 1970	2,785
Five or more AI papers since 1970	1,388
At least one AI paper since 2016	2,027
Five or more AI papers since 2016	904

Source: CSRankings.org, accessed November 2021.
Note: CSRankings defines faculty as "full-time, tenure-track faculty who can solely advise a PhD student in computer science."

Of course, full-time professors on the tenure track are not the only instructors capable of educating students in AI. Full-time teaching faculty and part-time adjunct professors—who are not eligible for tenure—are increasingly common across academia, allowing universities to increase their teaching capacity at the undergraduate and master's levels with more flexibility. As we discuss in the following section, the number of full-time teaching professors in U.S. computer science departments tripled between 2011 and 2020, mirroring the increase in bachelor's and master's enrollments. While there is limited data on part-time faculty who teach computer science specifically, the Government Accountability Office estimates between 32 and 50 percent of all U.S. university faculty are employed on a part-time basis, while tenure-track professors account for less than 30 percent of the academic workforce.[9]

Enrollments Surge

Perhaps the best information on CS students and faculty comes from the CRA's annual Taulbee Survey, which collects data from more than 140 computer science departments across the United States (see Appendix B for more details).

Figure 1 shows the total number of enrolled students and faculty in the computer science departments that responded to the Taulbee survey each year between 2011 and 2020. During that period, the total number of students enrolled in computer science programs roughly tripled, from 60,661 to 182,262. Meanwhile, the total number of faculty working in those departments grew just 43 percent, from 4,363 to 6,230.

The spike in enrollments was driven primarily by increases in bachelor's and master's students, whose ranks grew 249 percent and 201 percent, respectively. The growth in faculty numbers was largely a product of schools increasing their number of non-tenure-track teaching faculty.[10] Between 2011 and 2020, the number of non-tenure-track teaching faculty working at surveyed departments grew roughly 200 percent, while the number of tenure track faculty rose just 27 percent. As shown in Figure 1, non-tenure-track teaching instructors represent a relatively small share of the overall faculty pool.[11] Though expanding their ranks relieves some of the pressure of the enrollment boom among bachelor's and master's students, it is unlikely that these additional instructors have entirely met the increased demand for AI education. Between 2011 and 2020, the aggregate student-to-faculty ratio across surveyed CS departments rose from roughly 14-to-1 to 29-to-1.

Figure 1: Enrollment and Faculty Numbers for CS Departments in the Taulbee Survey, 2011–2020

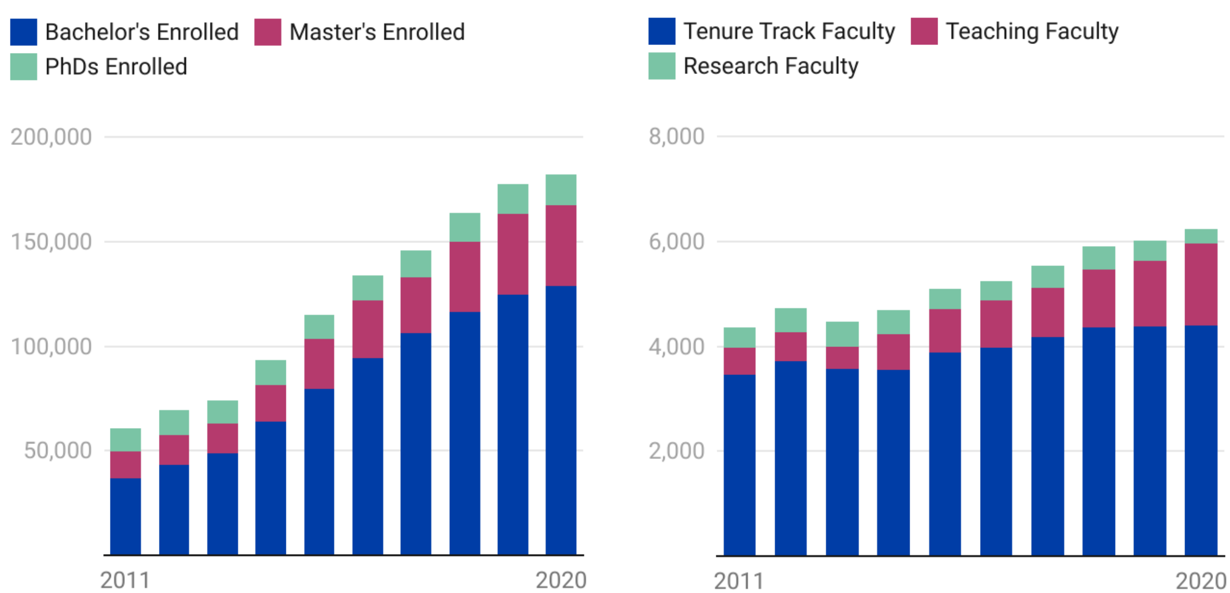

Source: CRA Taulbee Survey (see Table 2 in Appendix D for more detailed data).

While this disparity between enrollments and department sizes is not necessarily proof that teaching capacity gaps exist, it does show that the increases in student demand for CS education has not been met with a proportional increase in faculty supply. Universities may be able to accommodate the growing demand by increasing class sizes and adopting educational technology that allows each faculty member to serve more students. However, doing so increases faculty workloads and can make academic positions less appealing (see Appendix C for a discussion of historical trends in CS enrollments and teaching capacity). Many universities have also started implementing measures that restrict students' access to CS courses, which suggests there is a limit to how much class sizes can grow.

One major limitation of using enrollment data to analyze student interest in a particular field is that figures show only the students who are admitted into computer science programs—they do not reflect students who want to enroll in computer science programs but are unable due to a lack of teaching capacity or other reasons. Measuring the extent of this *unmet demand* is difficult due to the lack of available data on oversubscribed courses and admissions rates. However, there is evidence to suggest that computer science courses are becoming more difficult for students to access.

Access Restrictions

As the demand for computer science education has spiked, many schools have responded with measures that restrict students' access to computer science courses.

More than 40 percent of computer science departments that participated in a 2018 CRA survey said they had already raised—or were planning to raise—admissions requirements for prospective CS majors (see Figure 2).[*] A majority of respondents reported plans to limit enrollments in popular CS courses (62 percent) and restrict access to upper-level classes to CS majors and minors (60 percent). Such measures limit the overall availability of computer science education and may dissuade, or altogether prohibit, students interested in the field from pursuing a CS major or enrolling in classes. The survey also found 61 percent of CS departments eliminated small-enrollment courses, which may affect the quality of students' education.[†]

---

[*] The questionnaire was administered in 2018 as a part of the CRA's annual Taublee survey. The survey was sent to 161 PhD-granting computer science departments, and about 80 percent responded to questions about actions taken in response to growing enrollments.

[†] Students' educational quality may also be affected by the type of faculty who teach their classes (e.g., tenure-track, full-time teaching, part-time adjunct), though the direction of that effect is unclear. Some studies find students learn better from non-tenure eligible faculty while others found the reverse to be true, and still others found the evidence to be mixed. For more information, see: David N. Figlio, Morton O. Shapiro, and Kevin B. Soter, "Are Tenure Track Professors Better Teachers?" (National Bureau of Economic Research, September 2013), https://www.nber.org/system/files/working_papers/w19406/w19406.pdf; Bettinger, E.P., & Long, B. (2006). The increasing use of Adjunct Instructors at public institutions: Are we hurting students?. In R.Ehrenberg (Ed.), *Assessing Public Higher Education at the Start of the 21st Century*; Ronald G. Ehrenberg & Liang Zhang, 2005. "Do Tenured and Tenure-Track Faculty Matter?," Journal of Human

Figure 2: Measures Implemented or under Consideration at PhD-Granting CS Departments

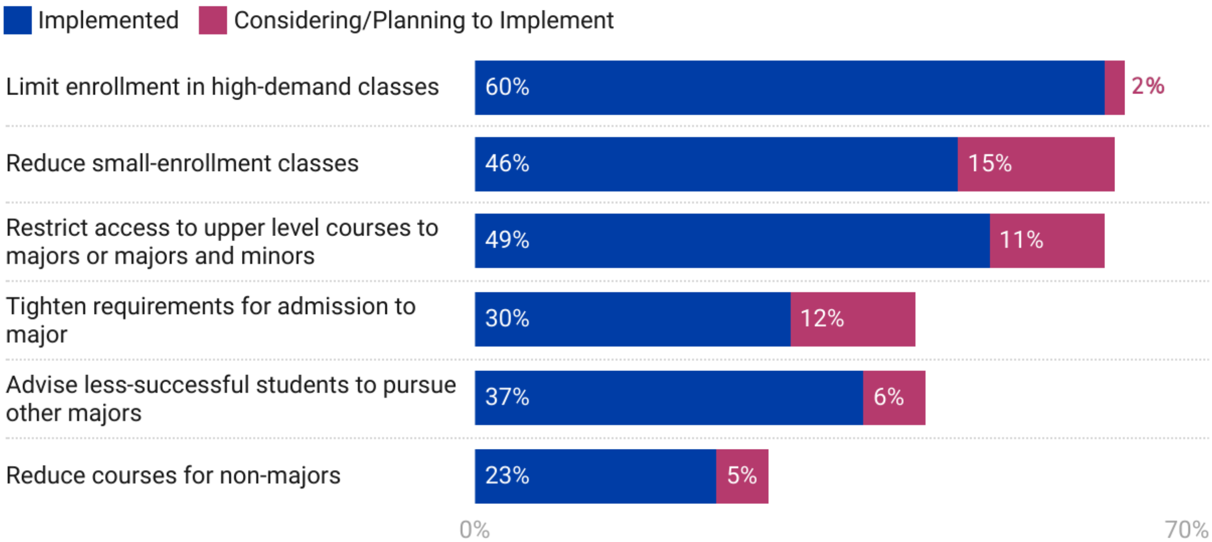

Source: Computing Research Association (2018).

Problems related to quantity are particularly pressing for universities without PhD students or other part-time teaching pools to draw from (e.g., community colleges or liberal arts colleges).[12] For instance, to cope with growing enrollment numbers, Swarthmore College implemented a lottery to select students for computer science courses, restricted the number of classes computer science majors could take, and reduced requirements for the major from nine courses to eight.[13] Students at Haverford, Bryn Mawr, Harvey Mudd, and Pomona colleges have publicly protested similar enrollment restrictions.[14] Even students at large institutions like the University of Texas at Austin report having trouble getting seats in computer science classes.[15]

These examples suggest that at least some student demand for computer science education is going unmet.[16] Researchers at Stanford University's Institute for Human-Centered Artificial Intelligence arrived at the same conclusion after surveying several universities for a 2018 report on the global AI landscape: "nearly every school noted that enrollment, particularly in recent years, is a function of [faculty] supply, rather than student demand."[17]

Resources, University of Wisconsin Press, vol. 40(3), https://ideas.repec.org/a/uwp/jhriss/v40y2005i2p647-659.html.

Concerns Among Academics

Many computer science departments have also explicitly said that they lack the capacity to accommodate all the students who want to study computer science.

In the same 2018 CRA survey, more than 90 percent of CS departments said they lacked enough faculty to meet the current need (see Figure 3). The survey also found that existing CS faculty across the country are being stretched increasingly thin, with mounting workloads and less time for research. More than 70 percent also reported difficulty providing enough seats in required courses and electives for computer science majors, another indicator of unmet student demand.

Figure 3: Issues Reported by PhD-Granting CS Departments in the United States

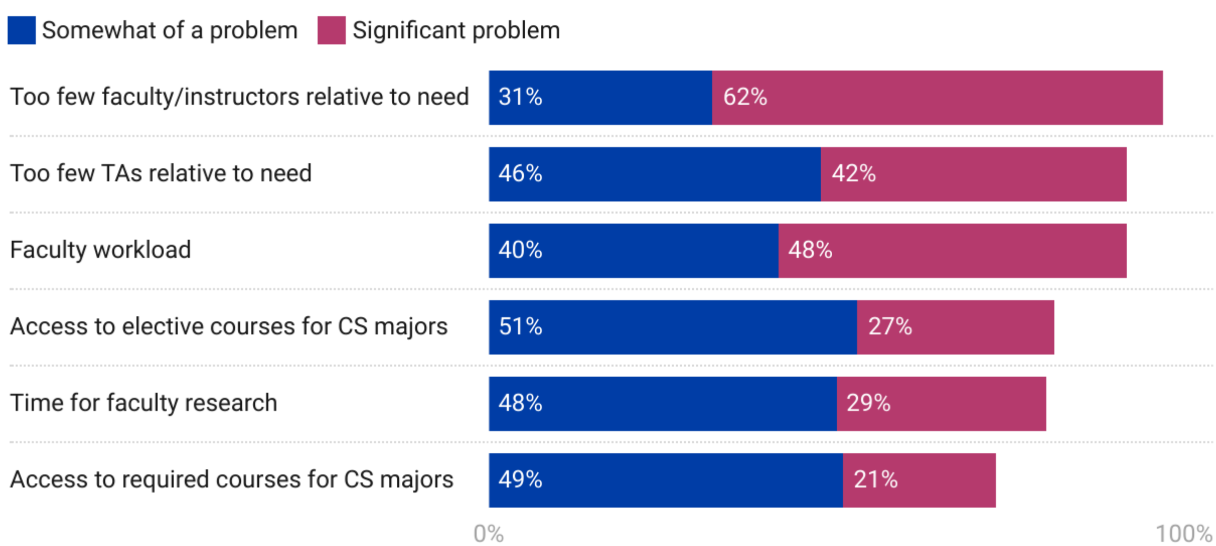

Source: Computing Research Association (2018).

Discussions of teaching capacity gaps are common within academic and computer science circles, focusing mostly on the difficulty of recruiting and retaining faculty. AI professors and researchers frequently write about the challenges of hiring faculty, citing their personal experience, partial statistics, and difficulties they see at university departments.[18] In a Times Higher Education-Microsoft survey of 111 AI researchers and university leaders, the vast majority of respondents said it was either "difficult" (48 percent) or "very difficult" (41 percent) to recruit and retain the academic staff needed to teach and research AI.[19] In 2019, the Computing Community Consortium and the Association for the Advancement of Artificial Intelligence summarized the situation succinctly: "All U.S. universities are looking for AI faculty—and struggling to hire, particularly at senior levels and in areas relevant to industry needs."[20] Several extensive studies have also suggested that computer science departments have struggled to recruit and retain faculty at different points in the past (see Appendix C).[21]

This anecdotal evidence suggests that the mismatch between student demand for AI education and universities' teaching capacity is a product of faculty shortages. However, some experts are skeptical of these claims. Research suggests it is rare for a field to experience faculty shortages, even within STEM disciplines.[22] Moreover, policy measures meant to counter shortages in certain fields have sometimes done more harm than good, with rapid injections of resources followed by stagnant growth and frozen job markets.[23]

Summary

While quantifying student interest in AI and universities' capacity to provide AI instruction remains a challenge, survey data and anecdotal evidence both suggest that demand for CS education does exceed the supply of available CS faculty. Over the last decade, enrollments in undergraduate and master's computer science programs have far outpaced the growth of computer science departments, universities have started restricting access to CS courses, and surveys have shown academic leaders fear they lack the teaching capacity to meet student demand. Without more detailed information from universities, we cannot definitively conclude that teaching capacity gaps exist. However, we find the available evidence to be sufficiently strong to warrant further investigation and policy attention.

It is critical that policymakers understand the factors creating teaching capacity gaps before they begin addressing the problem, as different root causes will require different policy solutions.

The Causes of Teaching Capacity Gaps

In this section, we consider three hypotheses to explain the current teaching capacity gaps at U.S. universities. The first is the "poaching" hypothesis, which is based on the idea that faculty are leaving academia for industry, and universities are unable to fill their vacancies. A second possible explanation is that not enough AI PhD graduates are interested in pursuing academic careers. These two hypotheses imply there is a shortage of AI faculty in the labor market.

By contrast, a third hypothesis is that universities have not changed their hiring practices to accommodate students' growing demand for AI training. In other words, this explanation implies there is a shortage of positions rather than people.

These three hypotheses are all closely related. As the world becomes more digitized, the market value of AI experts has increased across virtually every sector of the economy. While industry has increased salaries commensurate with the growing demand, academia has not been able to respond as nimbly given the tight budgetary and institutional constraints that universities face.

We found little evidence to suggest the outflow of AI faculty from academia to industry has increased in recent years (hypothesis 1), and, though a larger share of new PhD graduates is indeed taking jobs in industry, survey data does not indicate that they are disinterested in academic careers (hypothesis 2). However, we did find evidence that suggests universities have not increased the number of computer science faculty positions in line with the growing demand for AI-related education (hypothesis 3). This constraint may result in increased faculty workloads, which in turn make private sector research positions look more appealing to PhD graduates and existing faculty.

Hypothesis #1: Private-Sector Poaching of Professors

Many policymakers and industry leaders attribute universities' lack of AI-teaching capacity to faculty shortages, which are driven by the private sector—especially large tech companies like Alphabet, Meta, and Microsoft—poaching professors from academia. In their 2019 AI Index report, researchers at Stanford University spoke of "an unprecedented brain drain of AI professors from academia to industry."[24] University administrators worry the private sector is "eating the seed corn," which will make it harder to prepare the next generation of researchers.[25] Even industry leaders have acknowledged their role in pulling talent away from academia.[26]

Most discussions of faculty joining the private sector rely on a small number of prominent examples. Perhaps the most famous case involved Uber hiring a team of 40 researchers away from Carnegie Mellon University's National Robotics Engineering Center in the span of a few months.[27] Another example is Turing Award winner Yann LeCun, who became head of Facebook AI Research in 2013 while maintaining his professorship at New York University.[28] While such "dual-affiliation" arrangements are becoming more common, the exact division of labor varies on a case-by-case basis.[29] Some professors only spend 10 to 20 percent of their time working for a company, whereas others go into industry almost full-time. New dual-affiliation arrangements tend to favor the latter.[30]

Figure 4: Reasons Cited by University Departments for Faculty Departures, 2003–2020

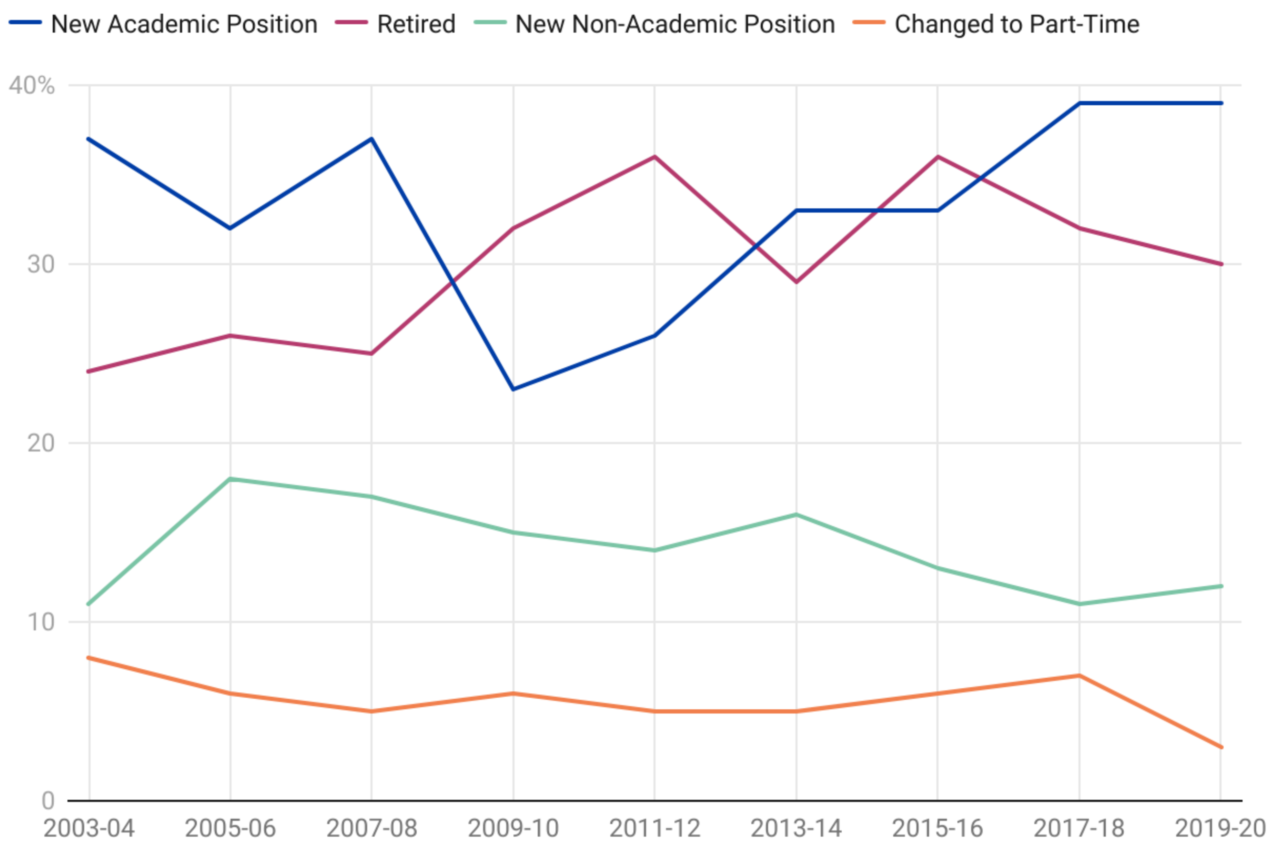

Source: CRA Taulbee Survey (see Table 3 in Appendix D for more detailed data).

However, while these cases indicate some level of poaching does exist, other evidence suggests this outflow is not the sole cause of teaching capacity gaps. In its annual Taulbee survey, the CRA found no significant increase in the amount of industry "poaching" over the last two decades. Since 2003, the most common reasons cited for faculty departures have been retirement or job changes within academia, as shown in

Figure 4. The proportion of professors leaving their posts for industry has hovered between 11 and 18 percent each year for nearly two decades. Additionally, CRA has not recorded a significant uptick in the number of professors changing from full- to part-time, which would reflect a rise in “dual-affiliation” roles.

Moreover, the flow of AI experts is not a one-way street. While some faculty do indeed leave for industry, there is also a steady stream of industry professionals entering academia. In a separate 2020 CRA survey of 105 university computer science departments, roughly 32 percent of respondents reported having faculty hired away *by* industry while about 38 percent said they had hired new faculty members *from* industry.[31] CRA found roughly one in eight people hired for faculty and academic research roles in 2018 and 2019 came from the private sector.[32] While, on the whole, it appears more AI experts leave academia for industry, this reverse flow somewhat mitigates the net effects of poaching.

It is also worth noting that universities’ relationships with the private sector extend beyond the zero-sum competition for AI talent. Industry also provides funding for university research, shares data and computing resources, and helps transition research findings into practice. Two-thirds of the departments included in the CRA survey report receiving significant research funding from industry.[33] Furthermore, the movement of AI experts between industry and academia helps new ideas develop and spread, potentially leading to future innovations.

In summary, the evidence and consequences of companies poaching professors from academia are mixed. The extent to which AI faculty departures harm the universities they leave remains unclear, and it is unlikely the relatively constant rate of industry “poaching” over the last two decades is entirely to blame for any potential shortages of AI faculty. Though this explanation receives the vast majority of people’s attention, it seems to be only a small slice of the problem.

Hypothesis #2: Limited Inflow from PhD Graduates

A second hypothesis for teaching capacity gaps is that there are not enough PhD holders who want to work at universities. This explanation suggests the *inflow* of new PhD graduates into academia is insufficient, unlike the “poaching” hypothesis, which suggests the *outflow* of academics into industry is excessive. However, both hypotheses rely on the same general premise: industry offers higher salaries and better access to resources than academia, and professors and PhD holders are responding to those market forces. These two explanations both imply that insufficient teaching capacity is caused by faculty shortages.

Academics in the field perceive the lack of inflow as a challenge. At a National Science Board panel on AI, the dean of Georgia Tech's College of Computing stated "we have to actually admit this is a real problem . . . We've had a 30 percent reduction in computer science PhDs who have gone into academia. They're all going into industry."[34] Some computer science departments surveyed by CRA in 2020 criticized industry for hiring away PhD students before they completed their degree.[35]

According to the National Science Foundation's Survey of Earned Doctorates, the share of PhD graduates with AI skills who pursue careers in industry is indeed growing.[36] Between 1995 and 2001, about five computer science PhD graduates entered industry for every three who took jobs in academia. However, since the dot-com crash of the early 2000s, the share of graduates going into the private sector has grown steadily (see Figure 5).[37] By 2020, there were more than three PhD holders going into industry for every one who took an academic position. Sources that collect data specifically on AI PhD graduates suggest the trends in general computer science are representative of those in the narrower field of artificial intelligence.[38]

Figure 5: Number of CS PhD Graduates by Post-graduation Employment Sector, 1995–2020

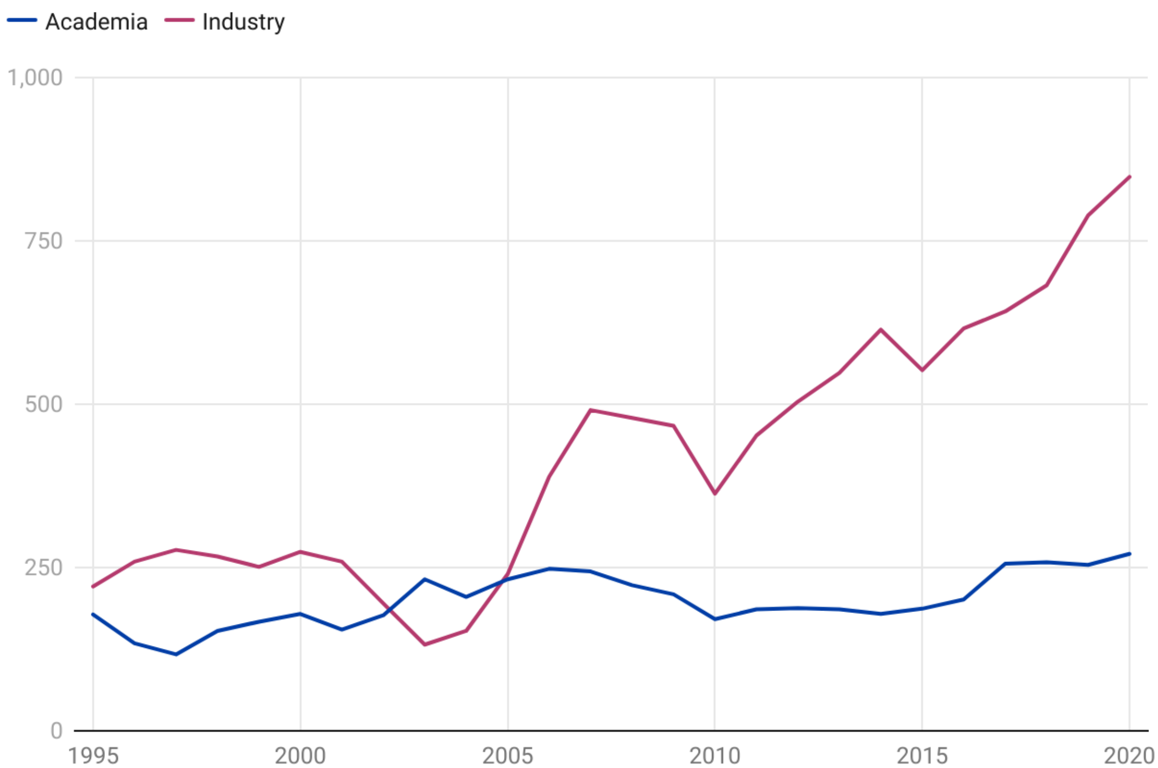

Source: NSF Survey of Earned Doctorates, 1995–2020 (see Table 4 in Appendix D for more detailed data).

However, while the share of graduates who take university jobs has fallen, the total number of PhD recipients who go into academia each year has remained relatively constant. Since 2000, the number of newly minted CS PhD graduates who take academic jobs after graduating has hovered between about 170 and 260 each year. This finding does not necessarily speak to PhD holders' broader interest in academia, but it may suggest that the number of university jobs available to newly minted CS PhD graduates has not changed dramatically over the last two decades.

While the NSF data describes the career *paths* of computer science PhD graduates, it does not necessarily reflect their career *preferences*. Survey data shows that many CS PhD recipients remain interested in academic careers, even as the share of graduates who enter academia declines. Figure 6 shows data from 398 AI-focused PhD students who responded to CRA's Data Buddies survey in 2020. Respondents rated industry research as the most attractive career path (90 percent were somewhat or very interested), but research-oriented academic careers did not rank far behind (71 percent were somewhat or very interested).[39] About 65 percent of respondents also expressed interest in research positions at government labs and agencies. Non-research positions in academia (teaching professors), industry, and government were rated as significantly less attractive, though non-research academic jobs were considered the most appealing of the three categories. These responses suggest that for AI PhD students, the ability to conduct research is a more important factor in choosing a job than the setting of the job itself.

Other studies have echoed these findings. A CSET survey of 254 AI PhD graduates found they viewed careers in academia as slightly more favorable than those in industry. Approximately 74 percent of respondents said they were "extremely" or "somewhat" likely to consider academic careers. The figure was 68 percent for jobs at large companies and 44 percent for roles at small companies. Academic careers received significantly more "extremely likely" ratings than either industry career path.[40]

Figure 6: Career Interests among CS PhD Students Specializing in AI

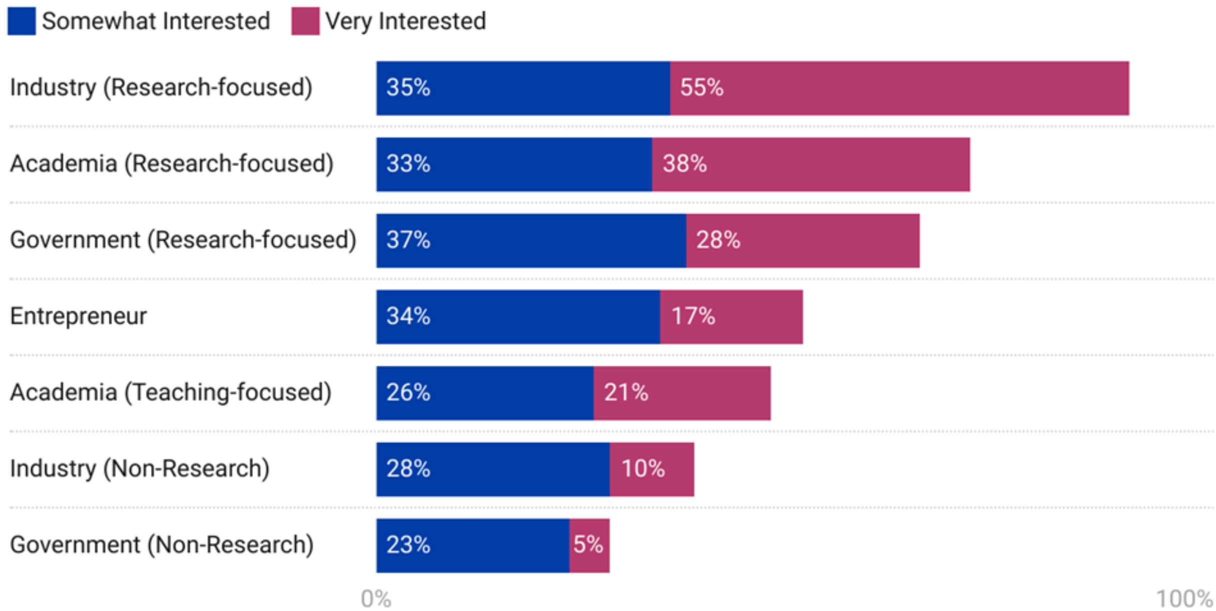

Source: CRA Data Buddies Survey (2020).

In summary, the share of PhD graduates going into university computer science departments has indeed fallen over the last 15 years, but the decline does not correspond with waning interest in academia. A majority of AI PhD recipients continue to view research-oriented professorships as an attractive career path even as they gravitate toward the private sector. Additionally, the total number of graduates who go into academia each year has remained fairly consistent for the last two decades, even as the relative share has declined. While this finding does not speak to PhD graduates' overall interest in academia, it may suggest that universities have not substantially increased the number of positions available to AI experts in recent years (hypothesis 3).

Hypothesis #3: Universities' Slow Response to Booming Enrollment

A third explanation for teaching capacity gaps is that universities, whether by choice or by circumstance, have not created enough new faculty positions to accommodate the increased demand for computer science education. In other words, this is not a problem caused by too few people, but rather by too few available positions.

As discussed in the previous section, the surge in CS enrollments has far outpaced the increase in department size. If AI faculty shortages are to blame—as the two previous hypotheses suggest—we would expect to see universities more frequently try and fail to hire CS professors. However, as shown in Figure 7, the success rate of faculty searches has remained consistent for more than a decade. Since 2008, computer science departments have generally filled about 70 to 80 percent of open tenure-track faculty positions and 80 to 90 percent of non-tenure-track teaching positions within a year. The "industry-poaching" and "insufficient-PhD-inflow" hypotheses imply this success rate should have started declining after companies increased investment in AI R&D around 2012.[*]

Figure 7: Success Rate of CS Faculty Searches, 2006–2020

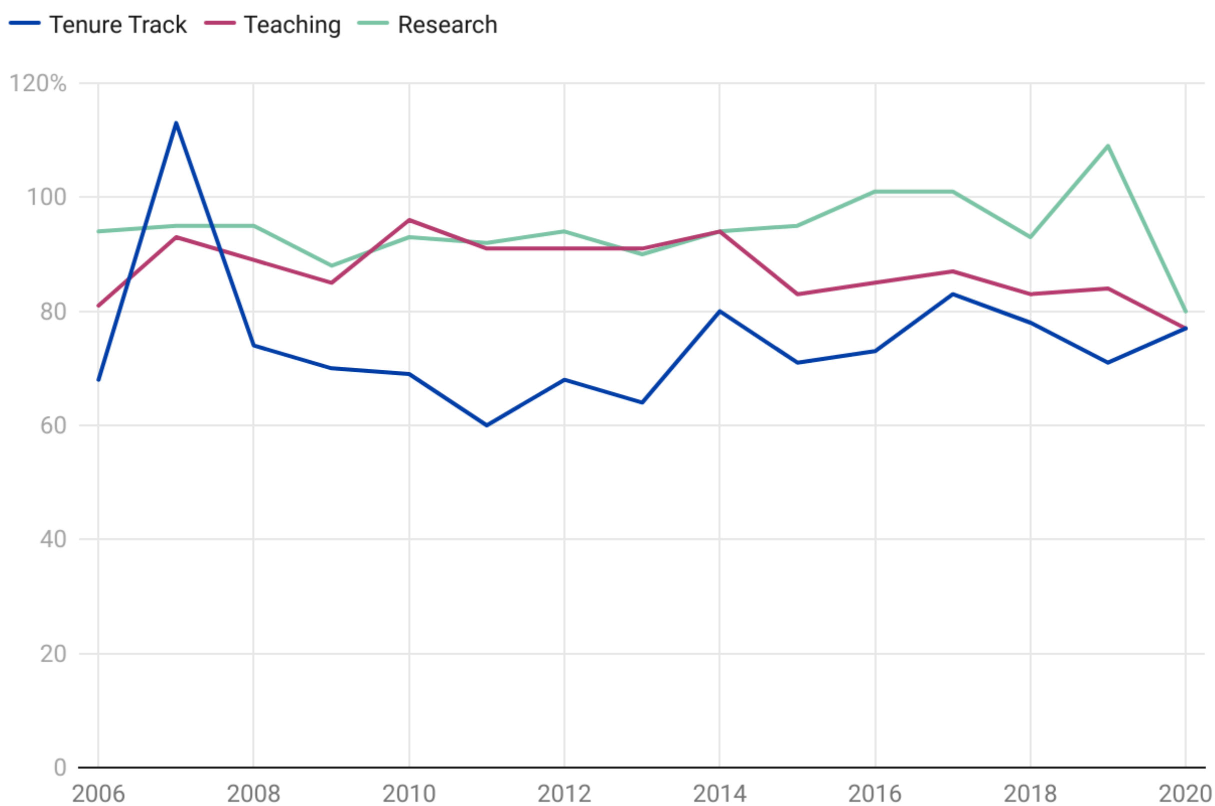

Source: CRA Taulbee Survey (see Table 5 in Appendix D for more detailed data).

If it has not become more difficult for universities to hire faculty, then the most plausible explanation for the relatively sluggish growth in faculty numbers is that universities have not done enough hiring. This is not necessarily an intentional decision on the part of universities—as discussed previously, many academic leaders seem to want to grow

[*] Our analysis of faculty-search success rates does not account for the sources of newly hired faculty. If two universities each hire faculty away from the other, the faculty search success rate would be 100 percent, even though the overall pool of faculty did not change. Thus, success rates may offer a more positive picture of the AI faculty labor market than is warranted. As shown in Figure 4, this sort of faculty "churn" is fairly common.

their CS departments. However, their ability to do so may be constrained by the rigid budgets of their institutions.

One extensive study of computer science education found that computer science students and professors cost universities more than those in the social sciences or humanities.[41] This leaves universities with two options if they wish to increase their computer science offerings: 1) take on the additional costs of increased CS hiring; or 2) reduce funding for faculty and classes in other fields. For institutions with rigid budgets, both options are likely to face substantial internal resistance. “The result,” the study concludes, “is an educational system that responds slowly to student demand.” This problem has been compounded by the relatively slow growth in the federal government’s computer science R&D budget, which directly and indirectly provides universities with funds to hire new faculty.[42]

Indeed, university administrators readily acknowledge the disparity between the demand for computer science classes and the supply of instructors to teach them. In the words of Don Fussell, chairman of the computer science department at the University of Texas at Austin, “as enrollments keep growing without bounds, it's very hard for anyone to keep their faculty size growing as fast as the demand for computer science majors.”[43]

Available data suggests they are trying. As noted previously, the 140-plus computer science departments that participate in the Taulbee survey roughly tripled their number of full-time teaching faculty between 2011 and 2020. Universities also appear to be making more tenure-track positions available as well. A study by Craig E. Wills, head of the Worcester Polytechnic Institute’s computer science department, found the number of tenure-track faculty positions advertised on two major academic job boards roughly doubled between 2014 and 2021.[44] This finding is consistent with CRA data on attempted and successful faculty hires.[45]

Wills’ work also suggests that searches for AI faculty are even more successful than for other types of computer science specialties. In 2019, roughly 18 percent of open faculty positions at surveyed universities were related to AI, data mining, or machine learning, yet about 26 percent of total faculty hires were within those specialties.[46] He found a similar trend in 2017 and 2018.[47] Though imperfect, these findings suggest universities do not face an outsized struggle when trying to recruit AI professors.[*]

Summary

We found little evidence to suggest that industry poaching of AI faculty (Hypothesis 1) or waning interest in academic careers among new PhD graduates (Hypothesis 2) are driving of teaching capacity gaps at U.S. universities. Neither the outflow of professors from academia for industry nor the success rate of CS faculty searches has changed substantially over the last decade. While more CS PhD graduates are indeed taking jobs in industry, their interest in academic careers remains high, and the total number of graduates going into academia each year has not changed significantly in the last two decades.

Though we cannot definitively attribute the cause of teaching capacity gaps, available evidence suggests that universities have not increased the number of computer science faculty positions in line with the growing demand for AI-related education (Hypothesis 3). Data on faculty numbers and job searches suggests universities are trying to grow their computer science departments, though their ability to do so may be limited by budgetary constraints.

Industry poaching of faculty and increased inflow of PhD graduates into the private sector could contribute to teaching capacity gaps in some specific cases, and their underlying causes should be addressed. However, the attention paid to these alternative explanations—especially industry poaching—appears to be disproportionate. To address teaching capacity gaps, policymakers should take the unique economics of academia into account and put forward solutions that target the funding constraints facing U.S. universities.

---

[*] It's also possible that universities have been slow to respond to growing CS enrollments due to the historical volatility in enrollment numbers. Interest in computer science has ebbed and flowed significantly over the last 50 years (see Appendix C), so it may be rational for schools to hedge against potential future downturns in enrollments. A department that is equipped to handle the current demand for CS training may find itself with a surplus of faculty should such a downturn occur, and, given the rigid protections for tenure-track positions, it would be difficult for schools to offload unneeded faculty.

The Costs of Teaching Capacity Gaps

Teaching capacity gaps threaten the two core missions of universities: training students and conducting research. In doing so, they could also have serious downstream consequences for U.S. economic and national security.

Consequences for Students

Teaching capacity gaps can hurt students' education in multiple ways, and different categories of students are likely to be affected differently.

While insufficient teaching capacity can make it harder for students at all degree levels to access computer science courses, the problem is particularly pronounced for PhD students, whose education entails intensive mentorship from senior researchers (typically tenure-track faculty).[48] As shown in Figure 8, undergraduate and master's enrollments in CS spiked in recent years, a trend likely supported (at least somewhat) by the growing number of "teaching" faculty. However, PhD enrollments have grown more slowly, at a rate that closely follows the increase in the number of tenure-track (likely research-oriented) faculty. Anecdotal evidence suggests enrollments in CS PhD programs would be much higher if more slots were available—today, many top PhD programs reportedly reject applicants who do not already have at least one publication in a top journal, a very high bar for students to meet.[49] This relatively sluggish growth in the number of PhD graduates threatens to exacerbate teaching capacity gaps by further constraining the future supply of qualified faculty.

It is worth reiterating that the growth in bachelor's enrollments (92,019) and master's enrollments (25,727) far exceeded the increase in teaching faculty numbers (1,043), even though the trends look similar in percentage terms.

Figure 8: Change in Average CS Enrollments and Faculty Size, 2011–2020

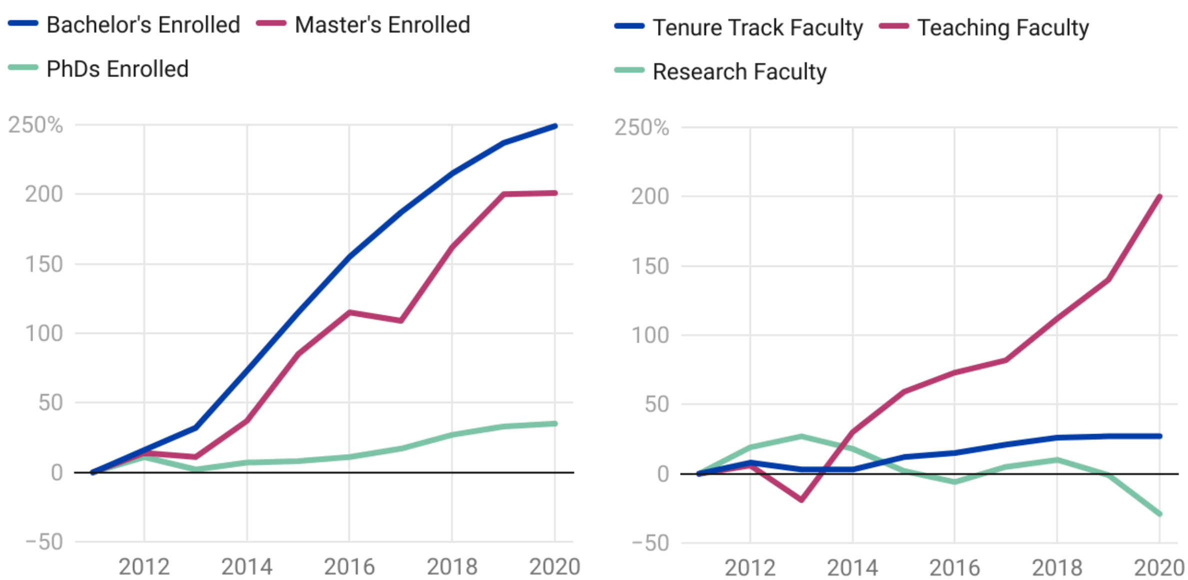

Source: CRA Taulbee Survey (see Table 2 in Appendix D for more detailed data).

Experts note that increased competition for slots and lower instructional quality also hurts diversity by favoring those with pre-existing computer science experience, who are disproportionately white and male. Though computer science courses have proliferated across U.S. high schools in recent years, white students are still more likely to have access to CS classes than those from underrepresented groups.[50] Among students who took the computer science AP exam in 2020, male students outnumbered female students roughly two-to-one.[51] If university enrollments are restricted based on prior class access or test scores, computer science departments—and the tech industry more broadly—will continue to be plagued by a lack of representation.

There is also evidence that faculty departures can have a direct negative impact on student entrepreneurship. A 2020 study by researchers Michael Gofman and Zhao Jin linked the departures of tenured AI professors to a reduction in the number of startups founded by students from the universities where they used to work. The effect was particularly pronounced among PhD students, and it grew when the professor was replaced by untenured faculty or professors from lower-ranked institutions.[52]

Consequences for Research

Teaching capacity gaps also affect both the quantity and type of research being done. In terms of quantity, overburdened faculty likely conduct less research because they must devote more time to their instructional responsibilities. However, if more current faculty and PhD recipients are going into industry, the absolute volume of research may not decrease, but simply shift from academia to industry.

There is no consensus on how shifting AI research from academia to the private sector affects the type of research being done. Some experts have expressed concern that private sector–driven research is profit-driven and therefore more incremental and applied, making fundamental advances less likely.[53] They also worry that private sector AI research could become more closed over time, slowing down progress in the field.[54]

Others believe academia and industry serve two distinct but complementary functions within the country's R&D ecosystem—universities specialize in basic research and experimentation, and then industry transfers breakthroughs from the lab to the real world. In 2019, National Science Foundation Director France Córdova celebrated the interchange of academic and industry researchers, and encouraged more professors to consider doing stints in the private sector.[55] Eric Horvitz, the chief scientific officer at Microsoft, says the reverse flow is just as important, "We should really promote this idea that industry-experienced, industry-hardened scholars head back to academia for a portion of their careers to teach and train."[56]

The potential benefits of talent flow between academia and industry include research becoming more tied to real-world problems, faster technology transfer and commercialization, and greater private sector investment in universities.[57] Moreover, whereas research done in the private sector is generally more applied, there are several companies that also invest heavily in more fundamental AI research (e.g., Alphabet).

In short, while an overall lower supply of senior AI researchers will clearly decrease R&D activity, there is no clear consensus on the net effect of reallocating existing research resources between academia and industry.

Consequences for National Security and the Economy

Teaching capacity gaps limit the amount of talent flowing into the U.S. AI workforce, which in turn negatively impacts economic and national security. Research has shown innovation is partly a function of the absolute number of researchers in a particular field, and the act of generating new ideas is becoming more labor intensive.[58] Less talent therefore means less innovation. The effect is particularly pronounced among PhD-level researchers, who tend to lead the discovery of new AI concepts and techniques, and whose numbers are closely tied to the available supply of faculty supervisors.

Another effect of a stunted workforce pipeline is the offshoring of AI companies. Today, many of America's largest tech companies are establishing research labs abroad to tap into global talent pools.[59] Those companies still house a majority of their AI R&D activity within U.S. borders, but that may change if the country fails to educate enough people with relevant AI skills.[60] Should tech companies shift more of their R&D activities abroad, the U.S. government would lose ready access to AI expertise and the risk of tech transfer would increase. The offshoring of industry R&D would also limit the growth of high-tech jobs in the United States.

The combined effects of these and other consequences of teaching capacity gaps could put the United States at a serious competitive disadvantage. This is especially concerning in light of other countries' investments in strengthening academic AI training. The Canadian government has recruited hundreds of researchers from around the world to university-based AI institutes, where they supervise graduate students as adjunct faculty.[61] In the Netherlands, the government has partnered with major Dutch companies to finance dozens of new positions for AI faculty and researchers.[62]

If the U.S. government fails to increase investment in academic AI research, professors may start drifting toward more generous countries, which could lead to faculty shortages and exacerbate teaching capacity gaps. Experts have linked limited funding availability to increased competition between U.S. universities and universities in other countries, particularly China.[63] Should this outflow increase, the United States risks losing one of its greatest national security assets—the ingenuity and robustness of its academic research institutions—to its biggest geopolitical competitor.

Given the outsized negative impacts teaching capacity gaps can have on the U.S. talent pipeline, research ecosystem, economy, and national security apparatus, we believe policymakers should begin exploring measures to address the issue.

Policy Issues and Options

The three actors that could address the mismatch between student demand for AI education and universities' AI-teaching capacity are academia, industry, and government. To complement ongoing efforts by other organizations that focus on universities and industry,[64] this paper focuses on the role that government can play in addressing these gaps. Federal policy is our primary focus. Because there are many remaining gaps in our knowledge (summarized in Appendix A), we discuss broad categories of measures and the likely pros and cons of different policy approaches. We hope future work will build on this research to make more targeted recommendations.

Increasing Funding for Education and Research

In using education and research funding to address teaching capacity gaps, policymakers have a range of options available to them. Some of these options are institutional (i.e., funding would be given to universities), whereas others are individual (i.e., funding would be given to researchers). Some focus on increasing teaching capacity, whereas others focus on increasing research capacity—two goals that can at times be in tension (see Box 2). The best funding portfolio is one that spreads resources across these different levels and goals.

Any increases in education and research funding must also be sustainable. In the past, fields that saw rapid injections of funding experienced significant shocks when investment tapered off. When the National Institutes of Health saw its budget double between 1998 and 2003, researchers flocked to the field of biomedical science. However, after the NIH budget plateaued, many of those same researchers lost their grant funding and saw their long-term career prospects dry up, and established researchers became less likely to propose high-risk but potentially high-payoff projects.[65] Studying the consequences of historical changes in research funding would yield valuable lessons for policymakers seeking to increase AI-teaching capacity.

Box 2: Funding Instruction versus Funding Research

Research and teaching are both integral components of universities' missions. However, these two goals can be at odds with one another. In the words of Northwestern University President Morton Schapiro and Professor David Figlio, "institutions that pride themselves not just on their teaching output but also on the scholarly contributions of their faculty face a multi-tasking problem of the first order."

In the same way, policies aimed at increasing research capacity could unintentionally worsen the teaching capacity problem. One way in which policymakers' desire for more student enrollment and for more R&D can be at odds is when research grant funding is used by professors to "buy out" of their teaching responsibilities. Universities have different policies with respect to how little teaching professors are allowed to do, but, in extreme cases, professors with big grants can go years without doing any undergraduate or master's teaching (they will generally continue or even increase PhD and postdoctoral supervision as part of their research projects).

We recommend that policymakers keep both teaching and research in mind, and anticipate that there may be trade-offs between the two. Research, the focus of R&D funding, is often the more prominent priority, but it should not receive exclusive attention at the expense of training the next generation of workers and teachers. As we discuss below, policymakers may also want to consider creating programs aimed specifically at funding teaching activities.

Institutional funding options

Expand funding programs for faculty hiring. One major reason universities experience persistent teaching capacity gaps in AI and CS programs is that they often lack the money to increase the size of their departments.[66] By offering funding specifically for hiring permanent faculty in fields like AI, the federal government can allow universities to more rapidly increase their research and teaching capacity. There is precedent for such programs. In recent years, the NSF has created grants to increase the number of faculty in other in-demand fields, such as quantum computing and space science.[67]

One factor policymakers must consider when creating hiring grants is what category of faculty they should target. Some faculty devote the majority of their time to research, some focus on instructing students, and others juggle both research and teaching responsibilities. Thus, faculty-hiring grants must be structured differently depending on the goal they are intended to increase: universities' teaching or research capacity.

Support AI education at other organizations. While traditional universities produce the majority of the AI talent in the United States, they are not the only place where students can build AI expertise.[68] Community colleges, massive open online courses (MOOCs), and other programs can provide individuals with the skills necessary for successful careers in AI.[69] Federally funded AI research institutes—such as those proposed by the NSF—already offer education and research opportunities for advanced students, and policymakers can ensure their training mission remains a priority as these organizations evolve.[70] By devoting resources to non-university entities to teach AI skills, the federal government can meet some of the growing demand for AI education and provide alternative career pathways for underserved students.

Individual funding options

Increase AI research funding. Another factor that may reduce universities' AI-teaching capacity is the intensifying competition for federal R&D funds. In its final report, the NSCAI stated that "federal funding that has not kept pace with the growth of the field has led to low grant application success rates and amplified time spent on the bureaucracy of pursuing and completing projects," which has in turn pushed more AI experts to industry.[71]

If grant-application burdens are indeed making universities less attractive for faculty, then increasing the amount and duration of AI research funding should have positive effects on academic recruitment and retention. Based on testimony it received, the NSCAI has concluded that "the NSF's budget for basic AI research would need to double simply to cover only the most highly qualified proposals it receives through its rigorous peer-review process."[72] Many other federal agencies also fund AI research, including the NIH, National Institute of Standards and Technology (NIST), and the departments of Defense, Energy, and Agriculture.[73]

Reducing the application burden could take several forms: First, and most obviously, these funding agencies could receive increased budgets for AI research, as has already been happening to some extent.[74] Second, the efficiency of grant applications, administration, and evaluations could be improved in a variety of ways. One option would be to standardize required application materials across different funding programs and agencies, creating something similar to the Common Application that many universities use for undergraduate admissions.[75]

Some have suggested that, in order to prevent academic brain drain into industry, agencies should have grant programs for AI researchers "that are conditional on remaining in academia for a fixed period of time."[76] We consider this "stick" approach less desirable than a "carrot" approach of increasing the availability and attractiveness of academic jobs. First, if faculty are staying in academia because they feel forced to, they are likely to be less motivated and effective teachers and researchers. Second, preventing the movement of faculty into industry could mean foregoing the productive benefits of cross-fertilization between academia and the private sector. Expanding the pool of R&D funds—and making those funds easier to access—would be a more effective approach.

Add a teaching option to "scholarship for service" graduate fellowships. Several government scholarships, such as CyberCorps and SMART, require recipients to complete a term in public service in exchange for the grant.[77] This requirement is typically fulfilled working for the agency that granted the scholarship. But in fields with a shortage of instructors, it may make sense for some percentage of recipients to fulfill their service requirements as teachers when positions are available but unfilled. This would benefit both society at large, and government agencies specifically, by maintaining the country's broader talent pipeline.

Evaluate and expand faculty reentry grants. Teaching capacity gaps could also be addressed by incentivizing academics who left for industry to return to teaching posts. As noted in the prior section, CS departments already do some amount of hiring from industry, illustrating that such transitions are feasible. The NSF runs at least one such fellowship program aimed at drawing former academics back into tenure-track roles.[78] The existing fellowship focuses on researchers in the fields of chemistry, bioengineering, environmental, and transport systems, but similar programs could be used to expand the ranks of faculty in AI and CS more broadly. Such programs are also a good way to encourage cross-pollination of ideas between industry and academia. If evaluations show such programs to be effective, their scope and number could be expanded.

Encouraging Companies to Support University Teaching

Companies are concerned about teaching capacity gaps in AI, even as they play some role in causing them. Realizing its need for a robust talent pipeline, the private sector already funds university teaching and research in a variety of ways. Since 2015, Google has awarded dozens of grants to academics studying ways to scale and increase the accessibility of computer science education, and DeepMind has endowed several faculty chairs at universities in Canada and the United Kingdom.[79] After it hired 40 of the 150 researchers at CMU's National Robotics Engineering Center, Uber gave the center a $5.5 million donation for student and faculty fellowships.[80] In a 2020 CRA survey, nearly two-thirds of computer science department chairs reported receiving "significant research funding" from industry.[81]

With the proper incentives, policymakers could leverage government resources to encourage more such private sector investment in university capacity.

Public-private matching grants. Policymakers could make certain types of funding conditional on matching investments from companies, as other countries have started doing. For example, in the Netherlands, government agencies are partnering with large players in the AI industry, such as Philips and ASML, to create 50 new AI professorships at Dutch universities.[82] In the United Kingdom, 11 companies are helping fund several hundred new graduate student positions to complement simultaneous government investment.[83] We have already seen similar public-private partnerships in the United States focused on supporting federal research: in 2020, the NSF partnered with Amazon, Accenture, Google, and Intel to stand up eight new AI institutes.[84] Incentivizing similar financial support for faculty-hiring programs could help increase teaching capacity.

Support for collaborative curricula. Companies also contribute to universities in non-monetary ways, such as granting access to datasets and computing resources, allowing employees to take part-time or full-time teaching sabbaticals, or providing in-house or public educational programs.[85] Policymakers could explore measures to further incentivize such industry educational investments, for example through tax credits.[86] Companies can also partner with universities to develop educational resources that equip students with industry-relevant, in-demand AI skills, which may evolve over time as the field changes. NIST's National Initiative for Cybersecurity Education, which supplies resources developed in tandem with private companies, could serve as a potential model.[87]

Facilitating Access to Government Data and Computational Power

A common explanation for teaching capacity gaps is that certain types of AI research now require such large amounts of data and computing power that AI researchers prefer working in industry over working in academia. For example, a Computing Research Association report attributes the increase in faculty dual-appointments with companies to the idea that "many important computing research problems require resources at a scale now unavailable to academics."[88]

Others in the field, however, disagree. One professor, who also works at Google part-time, said "I feel the data-and-computing issue is largely a myth."[89] Survey data collected by CSET indicates that access to data and computing resources ranks among the less important factors affecting AI PhD graduates' job choices. And universities were actually not seen by AI PhD graduates as much worse than industry when it comes to data and compute access.[90] However, it could be that this is true only for elite universities—which often have their own supercomputers—but inaccurate for non-elite universities.[91]

Given this disagreement and uncertainty, we recommend further studies into whether and how data and computational partnerships could in fact help increase faculty attraction and retention. If access to government compute and data do turn out to be helpful for universities, policymakers have several options for improving that access. In June 2021, the White House created a federal task force to establish a National Artificial Intelligence Research Resource (NAIRR), a shared platform which would offer data and computing infrastructure to academics and startups.[92]

Other Policy Areas

State-level policies. Much of the funding for public universities comes from state governments. Several state governments have already considered or implemented policies to increase local CS capacity.[93] Such policy options should be explored further. However, state-level policies by themselves will not solve national problems. As the chancellor of North Dakota's university system argues in the case of cyber education, "A national approach is essential. If one state invests substantially in [cyber programs], a brain drain from staff and faculty of institutions in less-well-funded states will likely occur, further exacerbating interstate inequities."[94] The same argument applies to AI education. State-level policies are likely to be most impactful when they involve experiments that can be scaled up to the federal level if they turn out to be successful.

Innovative academic-industry collaborations. Certain solutions to teaching capacity problems require cultural and policy changes within universities and companies, for example when it comes to intellectual property issues in academic-industry partnerships.[95] Even if government may not be the leading player here, policymakers and agencies should continue engaging with academic associations and other stakeholders to see where they can facilitate such progress. For instance, agencies could fund innovative pilot programs for joint-degree programs in computing-related fields—an idea university associations have already expressed some support for—and

Congress could raise awareness by commissioning studies or holding hearings on how the private sector can help grow the country's talent pipeline.[96]

Immigration reform. Around half of current computer science, mathematics, and statistics faculty at U.S. universities were born abroad, as were around two-thirds of PhD students and more than half of all doctorate holders in these fields employed in the United States.[97] Universities' ability to recruit and retain foreign-born faculty is thus essential to the U.S. AI talent pipeline. Compared to other employers, universities face fewer immigration issues due to their exemption from annual H-1B caps. However, this exemption only concerns temporary residency, not permanent residency, and universities' immigration offices still say that they regularly face issues when trying to attract and retain foreign-born faculty. Immigration reform could therefore help address teaching capacity gaps.[98]

Conclusion

University faculty play an indispensable role in developing AI talent, but mounting evidence suggests that institutions across the United States do not have enough faculty to meet students' growing demand for AI education. Many leaders in academia blame these teaching capacity gaps on excessive industry poaching of AI professors and PhD recipients. While these experts may have correctly identified teaching capacity gaps at U.S. universities, available data suggests they may have misattributed the root cause.

We found little evidence to suggest that industry poaching of AI faculty and waning interest in academic careers among new PhD graduates are driving teaching capacity gaps at U.S. universities. Neither the outflow of professors from academia to industry, nor the success rate of CS faculty searches has changed substantially over the last decade. While a larger share of CS PhD graduates is indeed taking jobs in industry, their interest in academic careers remains high, and the total number of graduates going into academia each year has not changed significantly in the last two decades.

Though we cannot definitively attribute the cause of teaching capacity gaps, available evidence suggests that universities have not increased the number of computer science faculty positions in line with the growing demand for AI education. Increasing the number of faculty positions available at U.S. universities may help close teaching capacity gaps. Data on faculty numbers and job searches suggests universities are trying to grow their computer science departments, but their ability to do so may be limited by budgetary constraints.

While industry poaching of faculty and increased inflow of PhD graduates into the private sector may contribute to faculty shortages in some geographic regions and AI subfields, the attention paid to these trends is disproportionate to their relative importance. Like teaching capacity gaps more broadly, faculty shortages are difficult to measure given the complexities of the academic labor market and multidisciplinary nature of artificial intelligence. There remain many open questions regarding the number of qualified AI faculty in the United States, the extent of the unmet demand for AI training, the scale of teaching capacity gaps at different types of institutions, and the potential impacts of policies targeting research and teaching capacity, among other issues. Appendix A discusses these and other areas for additional research in more detail.

Authors

Remco Zwetsloot is a trustee fellow in the International Security Program at the Center for Security and International Studies, and was formerly a research fellow at CSET, where Jack Corrigan is a research analyst.

Acknowledgments

Special thanks to Jacob Strieb. For research and data support, we would like to thank Catherine Aiken, Betsy Bizot (CRA), James Dunham, Rebecca Gelles, Rebecca Kagan, Tianwei Liu, and Burcin Tamer (CRA). For conversations and comments, we would like to thank Andy Bernat, Emily Cook, Obum Ekeke, Melissa Flagg, Alex Friedland, Diana Gehlhaus, Peter Hersha, Jesse Lastunen, Justin Lynch, Igor Mikolic-Torreira, Emily Oehlsen, Stephanie O'Sullivan, Ariel Procaccia, Eric Roberts, and Lynne Weil.

Document Identifier: doi: 10.51593/20190049

Appendix A: Open Questions and Future Research

There are many open questions about teaching capacity gaps that this paper has not fully addressed. Answering these questions will be critical to understanding the scope and size of AI teaching capacity gaps across the United States and crafting targeted policies to address them. We encourage interested researchers to explore these areas, including:

- **The scope of teaching capacity gaps:** How many AI-focused faculty are there in the United States? How many people are there that *could* teach AI classes? How are they distributed across public versus private universities, elite versus non-elite universities, across regions, and so forth? Does teaching capacity differ by AI subfield (e.g., because some subfields are currently more commercially valuable)? Is it possible to predict future trends?

- **The extent of the unmet demand for AI education among students**: How many students are barred from enrolling in computer science programs and courses due to insufficient teaching capacity? How does the impact of insufficient teaching capacity differ among undergraduate, master's, and PhD students?

- **The prevalence of dual-affiliation arrangements among faculty:** What share of AI faculty maintain simultaneous affiliations in academia and industry? How do those who participate in these arrangements distribute their time and resources? Are these arrangements more prevalent at certain universities or in certain AI subfields?

- **The dynamics of academic poaching:** What are the career trajectories of the AI faculty who spend their careers in academia? How prevalent is poaching within academia? How does this "churn" affect teaching capacity and quality at different types of institutions?

- **The effectiveness of different types of faculty**: Do tenure-track, teaching, and adjunct professors have different impacts on students' educational outcomes? How prevalent are part-time professors within computer science departments? What is the optimal distribution of faculty types for different universities? Can remote instruction, MOOCs, and other education technology be used to increase teaching capacity in bachelor's and master's programs?

- **The potential tradeoffs between research and teaching:** To what extent is research funding currently used to "buy out" of teaching? What policy interventions would simultaneously increase research and teaching capacity? How feasible and desirable is it for universities to have separate research- and teaching-focused faculty tracks?

- **The career trajectories of graduates from top CS programs:** Are graduates from the highest-ranked CS PhD programs more or less likely to pursue careers in academia? Do they congregate at particular universities or companies?

- **The impact of international students on domestic-student enrollment:** A 2017 study found that between 1995 and 2005, increases in international-student enrollments in graduate programs were associated with increases in domestic-student enrollments.[99] Has this trend held up over time? Is it broadly applicable to all fields at all types of universities?

- **The effect of regulations on teaching capacity:** Are public universities constrained in addressing teaching capacity gaps by any national- or state-level regulations (e.g., caps on faculty salaries)? Are academia-industry collaborations and industry funding of instructors constrained due to conditions placed on government funding?

- **The impact of resources on career decisions**: How does access to data and computing resources affect the career preferences of AI experts? How can these resources be shared with universities to improve research and teaching capacity?

- **The process of allocating funds within universities:** What factors do universities consider when allocating resources among different departments? What metrics does the federal government use to distribute funds? How can teaching capacity in different fields be quantified, and how should those metrics factor into funding decisions?

Answering these and other questions may involve interviews and surveys with stakeholders (e.g., university administrators, professors, industry scientists); further data collection on AI course offerings, job openings, and university staffing (collected via surveys, web scraping, or some other method); and further analysis of scientometric data on academic and industry research output (as in our preliminary analysis of CSRankings data to estimate U.S. AI faculty counts in Table 1).

Case studies of teaching capacity gaps and faculty shortages in other fields could also help inform AI policy. Historically, faculty shortages are said to have plagued fields related to semiconductors (e.g., electrical engineering) and the internet and consumer computing (e.g., computer science) when those technologies experienced their boom years.[100] Several fields—notably the life sciences in the late 1990s, when the NIH budget nearly doubled within five years—have also gone through periods of rapid growth, with sometimes unintended negative consequences.[101] Studying the consequences of measures (not) taken during these relevant historical periods could yield valuable lessons for increasing teaching capacity today. Lessons learned could also be applied in fields besides AI. Quantum-related disciplines, for example, are also currently thought to be experiencing teaching shortages and industry poaching.[102]

Appendix B: CRA Taulbee Survey

Our analysis leans heavily on data from the Computing Research Association's Taulbee survey, an annual survey that asks North American PhD-granting computer science, computer engineering, and information departments about their students and faculty. This paper uses data specifically from computer science departments at public and private U.S. universities. Not every department employs every type of faculty or educates every type of student, which prevents us from calculating average enrollments and faculty sizes at the department level.

While response rates vary year to year, the sample sizes are consistent enough to highlight broad trends in student enrollment and department sizes. Between 142 and 152 departments responded to the survey each year from 2011 to 2020.[103]

The Taulbee survey does not include data on the subset of CS faculty who specialize in artificial intelligence. However, though the mapping is imperfect, data on the broader pool of CS faculty is useful for understanding employment trends in the narrower discipline of AI. Many AI courses are taught in computer science departments, and AI specialists account for a growing portion of CS faculty overall.[104] Thus, any increases in industry poaching or other trends among AI faculty would likely be reflected in the data on CS faculty.

Appendix C: A Brief History of CS Enrollment Booms and Teaching Capacity

The current surge in computer science enrollments is not without precedent. Over the last half century, we have seen two similar spikes in the number of students pursuing computer science degrees: the first in the late 1970s and early 1980s, and the second in the late 1990s and early 2000s. In both cases, the rapid rise in enrollments was followed by an equally rapid downturn. These two boom-and-bust cycles are depicted in Figure 9.

Figure 9: Number of CS Bachelor's Graduates, 1971–2020

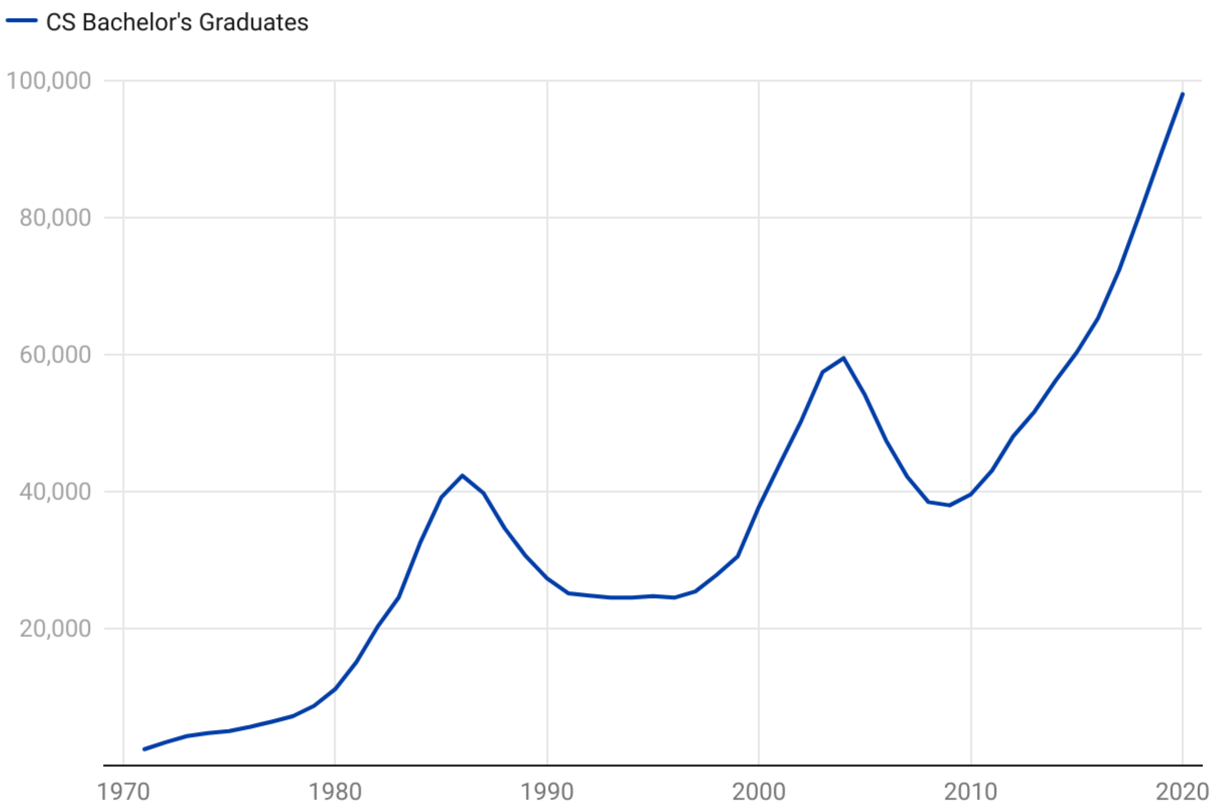

Source: Adapted and updated from Eric Roberts, A History of Computer Science Capacity Challenges (2016).

In a 2016 study, Stanford University computer scientist Eric Roberts attributed the more recent spike to the "dot-com bubble," with students flocking to computer science amid the expansion of the internet and growth of the tech industry and losing interest in the field after the bubble burst in the early 2000s.[105] According to Roberts, faculty shortages *did not* play a significant role during this boom-and-bust cycle.

However, he claimed faculty shortages were to blame for the downturn after the first enrollment boom. Like the second, the first wave of computer science enrollments also coincided with a new technology trend: the rise of personal computers. However, the field was still in its early days and qualified experts were in short supply. Those who opted for academia over industry found themselves overburdened by increasing class sizes and teaching requirements. Despite efforts to increase their teaching capacity, universities could not keep pace with student demand. Many schools ultimately raised the requirements for enrollment, creating an impression among students that computer science was competitive and exclusionary. Their enrollment numbers dropped as a result. “Students in the mid 1980s did not *decide* against majoring in computer science but were instead *prohibited* from doing so by departments that lacked the resources to accommodate them,” Roberts wrote.

The downturn of the 1980s serves as a warning that faculty shortages in computer science, if left unaddressed, can undermine the AI talent pipeline. Overburdened professors, increased class sizes, and restricted enrollments can create an unwelcoming learning environment and dissuade students at all levels from furthering their academic pursuits. As reviewed below (“Access Restrictions”), many of these same measures are being implemented today. In his paper, Roberts drew parallels between the past downturn and our present situation: “the problems we see in computer science education today closely resemble those from the beginning of the 1980s.”[106]

Appendix D: Data Tables

Table 2: Number of Enrollments and Faculty at Computer Science Departments That Responded to the Taulbee Survey

	Enrollments			Faculty		
	Bachelor's	Master's	PhD	Tenure Track	Teaching	Research
2020	129,034	38,501	14,727	4,390	1,564	276
2019	124,850	38,270	14,412	4,384	1,249	382
2018	116,439	33,432	13,802	4,366	1,107	426
2017	106,287	26,755	12,689	4,176	947	408
2016	94,501	27,423	12,109	3,971	903	364
2015	79,737	23,650	11,783	3,880	826	396
2014	64,141	17,488	11,620	3,559	679	455
2013	48,866	14,172	11,065	3,564	421	491
2012	43,105	14,537	12,033	3,725	550	460
2011	37,015	12,774	10,872	3,455	521	387

Source: CRA Taulbee Survey, 2011–2020.

Table 3: Reasons Cited by University Departments for Faculty Departures, 2003–2020[107]

Year	Changed to Part-Time	New Non-Academic Position	New Academic Position	Retired	Other	Total
2020	10	33	113	91	65	312
2019	11	43	139	103	31	327
2018	23	34	126	94	26	303
2017	12	26	85	80	31	234
2016	13	42	89	90	36	270
2015	16	24	77	94	26	237
2014	15	44	86	65	36	246
2013	11	32	74	74	41	232
2012	11	27	62	89	32	221
2011	12	34	52	67	48	213
2010	12	27	46	73	50	208
2009	11	33	46	53	48	191
2008	10	50	97	71	47	275
2007	17	42	103	60	37	259
2006	11	38	74	55	29	207
2005	16	39	61	56	41	213
2004	21	26	87	45	48	227
2003	13	22	74	59	41	209

Source: CRA Taulbee Survey, 2003–2020.

Table 4: Share of U.S. CS PhD Graduates Going into Academia and Industry, 1999–2019

Year	Total CS PhD Graduates	Percent into Academia	Percent into Industry
2020	1,981	27%	65%
2019	1,877	28%	64%
2018	1,721	31%	61%
2017	1,583	31%	61%
2016	1,479	30%	62%
2015	1,364	32%	59%
2014	1,335	33%	59%
2013	1,310	35%	57%
2012	1,225	36%	55%
2011	1,077	36%	51%
2010	1,061	41%	47%
2009	1,232	41%	50%
2008	1,219	40%	53%
2007	1,253	42%	51%
2006	1,190	44%	46%
2005	906	55%	38%

2004	741	62%	29%
2003	725	64%	29%
2002	738	54%	37%
2001	845	46%	45%
2000	938	48%	44%
1999	877	48%	44%
1998	917	44%	45%
1997	850	39%	49%
1996	806	47%	46%
1995	830	55%	37%

Source: National Science Foundation, Survey of Earned Doctorates, 1995–2020.

Table 5: Success Rate of CS Faculty Searches, 2006–2020

	Tenure Track			Teaching			Research		
	Vacant	Filled	Success Rate	Vacant	Filled	Success Rate	Vacant	Filled	Success Rate
2020	412	316	77%	128	99	77%	41	33	80%
2019	437	309	71%	129	108	84%	47	51	109%
2018	440	341	78%	140	116	83%	67	62	93%
2017	434	360	83%	219	191	87%	72	73	101%
2016	418	304	73%	169	143	85%	68	69	101%
2015	346	245	71%	160	133	83%	84	80	95%
2014	282	227	80%	161	152	94%	78	73	94%
2013	275	176	64%	137	125	91%	78	70	90%
2012	322	219	68%	137	125	91%	124	116	94%
2011	204	122	60%	115	105	91%	121	111	92%
2010	172	119	69%	178	171	96%	58	54	93%
2009	199	140	70%	131	112	85%	60	53	88%
2008	411	306	74%	131	117	89%	109	104	95%
2007	245	278	113%	136	127	93%	97	92	95%
2006	278	190	68%	106	86	81%	65	61	94%

Source: CRA Taulbee Survey, 2006–2020.

Endnotes

[1] The 2018 AI Index team surveyed several universities for AI enrollment trends and found that "nearly every school noted that enrollment, particularly in recent years, is a function of supply, rather than student demand." For more information, see: "Artificial Intelligence Index 2018 Annual Report" (Stanford University, 2018), pg. 75, https://hai.stanford.edu/sites/default/files/2020-10/AI_Index_2018_Annual_Report.pdf#page=75.

[2] For example, few programs at agencies like the National Science Foundation are designed to raise the number of faculty in emerging fields, such as AI. Likewise, recent calls for a 21st century version of the 1958 National Defense Education Act tended to focus on increasing student demand and student funding while neglecting the issue of teaching capacity: Tom Donilon, "Trump's Trade War is the Wrong Way to Compete with China," *Foreign Affairs*, June 25, 2019, https://www.foreignaffairs.com/articles/china/2019-06-25/trumps-trade-war-wrong-way-compete-china.

[3] There is little detailed data on the AI faculty workforce, so we use data on the broader pool of CS faculty to understand trends in this narrower field. Many AI courses are taught in computer science departments, and AI specialists account for a growing portion of CS faculty overall (see: "Artificial Intelligence Index Report 2021" (Stanford University, Human-Centered Artificial Intelligence, 2021), pg. 113, https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf#page=113; Craig E. Wills, "Outcomes of Advertised Computer Science Faculty Searches for 2019" (Worcester Polytechnic Institute, October 2019), pg. 10, https://web.cs.wpi.edu/~cew/papers/outcomes19.pdf#page=11.) Thus, any increases in industry poaching or other trends among AI faculty would likely be reflected in the data on CS faculty. However, as discussed in Box 1, the mapping is far from perfect.

[4] Cade Metz, "Tech Giants are Paying Huge Salaries for Scarce A.I. Talent," *New York Times*, October 22, 2017, https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html.

[5] "Final Report: National Security Commission on Artificial Intelligence" (National Security Commission on Artificial Intelligence, March 2021), pg. 188, https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf#page=188.

[6] Laurie A. Harris, "Artificial Intelligence: Background, Selected Issues, and Policy Considerations," *Congressional Research Service*, May 19, 2021, https://crsreports.congress.gov/product/pdf/R/R46795.

[7] For a broader discussion of issues that arise when trying to measure labor shortages, see Michael S. Teitelbaum, *Falling Behind? Boom, Bust, and the Global Race for Scientific Talent* (Princeton: Princeton University Press, 2014), chapter 5.

[8] Potential supply is especially hard to measure. In addition to U.S. faculty who have AI-adjacent skills, other potential sources of supply include scientists and engineers not currently working at U.S. universities (e.g., in the private sector) and scientists and engineers not currently within the United States (e.g., those working in academia abroad).

[9] "Contingent Workforce: Size, Characteristics, Compensation, and Work Experiences of Adjunct and Other Non-Tenure-Track Faculty" (Government Accountability Office, October 2017), pg. 11, https://www.gao.gov/assets/gao-18-49.pdf#page=17.

[10] Our count of non-tenure-track teaching faculty includes both "Teaching Professors" and "Other Instructors," as defined by the Taulbee survey. According to the 2020 survey, "'teaching professors' on average have more varied responsibilities in teaching, scholarship, service/governance, etc., and higher expectations for visibility outside the unit or the institution. 'Other Instructors' are more focused on teaching introductory or mid-level courses and tend to have shorter contract lengths, though they are still full-time faculty." For more information, see: Stuart Zweben and Betsy Bizot, "Bachelor's and Doctoral Degree Production Growth Continues But New Student Enrollment Shows Decline," *Computing Research Association*, May 2021, https://cra.org/wp-content/uploads/2021/05/2020-CRA-Taulbee-Survey.pdf.

[11] By contrast, the student-to-faculty ratio across the U.S. university system fell from 16-to-1 to 14-to-1 during approximately the same time period: Ricardo Azziz, "COVID-19: Will the pandemic worsen U.S. higher education's excess capacity?" *TIAA Institute*, https://www.tiaainstitute.org/about/news/covid-19-will-pandemic-worsen-us-higher-educations-excess-capacity. Unlike our analysis, these figures reflect changes in the number of part-time faculty.

[12] Colleen Flaherty, "System Crash," *Inside Higher Ed,* May 9, 2018, https://www.insidehighered.com/news/2018/05/09/no-clear-solution-nationwide-shortage-computer-science-professors; Adams Nager and Robert D. Atkinson, "The Case for Improving U.S. Computer Science Education" (Information Technology and Innovation Foundation, May 2016), pg. 20, https://www2.itif.org/2016-computer-science-education.pdf#page=20.

[13] "2018-19 Changes to the CS Major," *Swarthmore College*, https://www.swarthmore.edu/computer-science/2018-19-changes-to-cs-major; "Lottery & Waitlist Information," *Swarthmore College*, https://www.swarthmore.edu/computer-science/lottery-waitlist-information.

[14] Colleen Flaherty, "System Crash," *Inside Higher Ed,* May 9, 2018, https://www.insidehighered.com/news/2018/05/09/no-clear-solution-nationwide-shortage-computer-science-professors.

[15] Natasha Singer, "The Hard Part of Computer Science? Getting Into Class," *The New York Times*, January 24, 2019, https://www.nytimes.com/2019/01/24/technology/computer-science-courses-college.html.

[16] It is important to note that students who are currently excluded from majoring in computer science would not necessarily graduate with CS degrees had they been admitted. The attrition rate for students who enroll in STEM programs is high (see: Peter Arcidiacono, Esteban M. Aucejo, and V. Joseph Hotz, "University Differences in the Graduation of Minorities in STEM Fields: Evidence from California," American Economic Review. Mar 2016, Vol. 106, No. 3: Table 1) It is possible that raising requirements for the CS majors would prevent some students who would otherwise drop out of the program from being admitted in the first place. However, other measures—such as eliminating small-enrollment classes and limiting enrollment in high-demand classes—adversely affect all students.

[17] "Artificial Intelligence Index 2018 Annual Report" (Stanford University, 2018), pg. 75, https://hai.stanford.edu/sites/default/files/2020-10/AI_Index_2018_Annual_Report.pdf#page=75.

[18] Sarah McBride and Ashlee Vance, "Apple, Google, and Facebook are Raiding Animal Research Labs," *Bloomberg*, June 18, 2019, https://www.bloomberg.com/news/features/2019-06-18/apple-google-and-facebook-are-raiding-animal-research-labs; Jonathan Douglas, "University struggles to keep up as AI industry booms," *The Brown Daily Herald*, November 7, 2017, https://www.browndailyherald.com/article/2017/11/university-struggles-to-keep-up-as-ai-industry-booms; Cade Metz, "Facebook Adds AI Labs in Seattle and Pittsburgh, Pressuring Local Universities," *The New York Times*, May 4, 2018, https://www.nytimes.com/2018/05/04/technology/facebook-artificial-intelligence-researchers.html; Tony Peng, "Are Commercial Labs Stealing Academia's AI Thunder?" *Medium*, July 10, 2019, https://medium.com/syncedreview/are-commercial-labs-stealing-academias-ai-thunder-dd51cf4bd8d6; Ariel Procaccia, "Tech Giants, Gorging on AI Professors is Bad for You," *Bloomberg*, January 7, 2019, https://www.bloomberg.com/opinion/articles/2019-01-07/tech-giants-gorging-on-ai-professors-is-bad-for-you.

[19] Rachael Pells, "The THE-Microsoft Survey on AI," *Times Higher Education*, March 28, 2019, https://www.timeshighereducation.com/features/microsoft-survey-ai#survey-answer.

[20] "A 20-Year Community Roadmap for Artificial Intelligence Research in the US," *Computing Community Consortium and Association for the Advancement of Artificial Intelligence*, August 2019, https://cra.org/ccc/wp-content/uploads/sites/2/2019/08/Community-Roadmap-for-AI-Research.pdf.

[21] John Stankovic and William Aspray, "Recruitment and Retention of Faculty in Computer Science and Engineering," *Computing Research Association*, 2003, http://archive2.cra.org/uploads/documents/resources/workforce_history_reports/rrfaculty.pdf; Eric Roberts, "A History of Capacity Challenges in Computer Science" (Stanford University, March 7, 2016), https://cs.stanford.edu/people/eroberts/CSCapacity.pdf; Colleen Flaherty, "System Crash," *Inside Higher Ed,* May 9, 2018, https://www.insidehighered.com/news/2018/05/09/no-clear-solution-nationwide-shortage-computer-science-professors; "Generation CS: Computer Science Undergraduate Enrollments Surge Since 2006" (Computing Research Association, February 2017), https://cra.org/wp-content/uploads/2017/02/Generation-CS.pdf.

[22] Paula Stephan, *How Economics Shapes Science* (Cambridge, MA: Harvard University Press, 2012), chapter 7.

[23] Stephan, *How Economics Shapes Science*; Teitelbaum, *Falling Behind?*

[24] "Artificial Intelligence Index 2019 Annual Report" (Stanford University, 2020), pg. 120, https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf.

[25] Cade Metz, "Facebook Adds AI Labs in Seattle and Pittsburgh, Pressuring Local Universities," *The New York Times*, May 4, 2018, https://www.nytimes.com/2018/05/04/technology/facebook-artificial-intelligence-researchers.html.

[26] "NSCAI Conference - Panel I: Strengthening Our Core: The Way Ahead for American R&D," uploaded by National Security Commission on Artificial Intelligence, November 15, 2019, https://www.youtube.com/watch?v=1x8npd-9Su8.

[27] Mike Ramsey and Douglas MacMillan, "Carnegie Mellon Reels After Uber Lures Away Researchers," *Wall Street Journal*, May 31, 2015, https://www.wsj.com/articles/is-uber-a-friend-or-foe-of-carnegie-mellon-in-robotics-1433084582.

[28] "Courant's LeCun to Lead Facebook's New Artificial Intelligence Group," *New York University,* December 9, 2013, https://www.nyu.edu/about/news-publications/news/2013/december/courants-lecun-to-lead-facebooks-new-artificial-intelligence-group-.html.

[29] Cade Metz, "Tech Giants are Paying Huge Salaries for Scarce A.I. Talent," *New York Times*, October 22, 2017, https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html.

[30] Shwetak Patel, Jennifer Rexford, Benjamin Zorn, and Greg Morrisett, "Evolving Academia/Industry Relations in Computing Research," *Computing Community Consortium*, June 2019, https://cra.org/ccc/wp-content/uploads/sites/2/2019/06/Evolving-AcademiaIndustry-Relations-in-Computing-Research.pdf; Cade Metz, "Facebook Adds AI Labs in Seattle and Pittsburgh, Pressuring Local Universities," *The New York Times*, May 4, 2018, https://www.nytimes.com/2018/05/04/technology/facebook-artificial-intelligence-researchers.html.

[31] Vivek Sarkar et al., "CRA Industry/Academia Committee Report," *Computing Research Association*, May 29, 2020, https://cra.org/wp-content/uploads/2020/07/CRA-Industry_Academia-Committee-Report.pdf.

[32] Stuart Zweben and Betsy Bizot, "2019 Taulbee Survey," *Computing Research Association*, May 2020, https://cra.org/wp-content/uploads/2020/05/2019-Taulbee-Survey.pdf.

[33] Sarkar et al., "CRA Industry/Academia Committee Report."

[34] Paul Basken, "Engineering and Computer Science: Time to Separate?" *PRISM*, September 2018, http://www.asee-prism.org/engineering-and-computer-science-time-to-separate/.

[35] Sarkar et al., "CRA Industry/Academia Committee Report."

[36] Computer science is not the only STEM field seeing this shift. In 2020, more than half of all PhD graduates in chemical, electrical, and industrial engineering also took jobs in industry. The share of industry-bound graduates in each field increased roughly 18 percent since 2010.

[37] The popularity of industry jobs seems to coincide with broader trends in the U.S. economy. We observed a decline in number of PhD graduates who entered the private sector in the years following the dot-com crash of the early 2000s and the Great Recession of the late 2000s.

[38] The CRA Taulbee Survey also finds that roughly 60 percent of AI PhD recipients go into industry and 35 percent into academia, as does CV data on U.S. AI PhD graduates collected by CSET; see Remco Zwetsloot, James Dunham, Zachary Arnold, and Tina Huang, "Keeping Top AI Talent in the United

States" (Center for Security and Emerging Technology, December 2019), https://cset.georgetown.edu/wp-content/uploads/Keeping-Top-AI-Talent-in-the-United-States.pdf.

[39] Career preferences of AI-focused CS PhD graduates generally mirrored those of CS PhD graduates who did not specialize in AI.

[40] Catherine Aiken, James Dunham, and Remco Zwetsloot, "Career Preferences of AI Talent" (Center for Security and Emerging Technology, June 2020), https://cset.georgetown.edu/publication/career-preferences-of-ai-talent/.

[41] Adams Nager and Robert D. Atkinson, "The Case for Improving U.S. Computer Science Education" (Information Technology and Innovation Foundation, May 2016), pg. 20, https://www2.itif.org/2016-computer-science-education.pdf#page=20.

[42] Between 2011 and 2019 (the latest year for which there is data), federal funding for computer science research rose approximately 32 percent, significantly slower than the growth in computer science enrollments. See: National Center for Science and Engineering Statistics (NCSES). 2021. *Federal Funds for Research and Development: Fiscal Years 2019–20*. NSF 21-329. Alexandria, VA: National Science Foundation. Available at https://ncses.nsf.gov/pubs/nsf21329/.

[43] Esther Shein, "The CS Teacher Shortage," *Association for Computing Machinery,* October 2019, https://cacm.acm.org/magazines/2019/10/239667-the-cs-teacher-shortage/fulltext.

[44] Craig E. Wills, "Analysis of Current and Future Computer Science Needs via Advertised Faculty Searches" (Worcester Polytechnic Institute, November 2014), pg. 2, https://web.cs.wpi.edu/~cew/papers/CSareas15.pdf#page=2; Craig E. Wills, "Analysis of Current and Future Computer Science Needs via Advertised Faculty Searches for 2022" (Worcester Polytechnic Institute, December 2021), https://web.cs.wpi.edu/~cew/papers/CSareas22.pdf.

[45] The Taulbee survey showed the number of new tenure track faculty hires rose consistently between 2011 and 2020.

[46] Craig E. Wills, "Outcomes of Advertised Computer Science Faculty Searches for 2019" (Worcester Polytechnic Institute, October 2019), https://web.cs.wpi.edu/~cew/papers/outcomes19.pdf.

[47] Craig E. Wills, "Outcomes of Advertised Computer Science Faculty Searches for 2017" (Worcester Polytechnic Institute, October 2017), https://web.cs.wpi.edu/~cew/papers/outcomes17.pdf; Craig E. Wills, "Outcomes of Advertised Computer Science Faculty Searches for 2018" (Worcester Polytechnic Institute, June 2018), https://web.cs.wpi.edu/~cew/papers/outcomes18.pdf.

[48] Elizabeth Gibney, "AI talent grab sparks excitement and concern," *Nature*, April 26, 2016, https://www.nature.com/articles/532422a.

[49] Alexander Johansen, "How to get into the Stanford Computer Science PhD Program," *Medium*, October 7, 2021, https://medium.com/@alrojo_github/how-to-get-into-the-stanford-computer-science-phd-program-71c8e1169b34; Andreas Madsen, "Becoming an Independent Researcher and getting published in ICLR with spotlight," *Medium*, December 29, 2019, https://andreas-

madsen.medium.com/becoming-an-independent-researcher-and-getting-published-in-iclr-with-spotlight-c93ef0b39b8b.

[50] Katie Hendrickson et al., “2021 State of computer science education: Accelerating action through advocacy” (Code.org, 2021), pg. 21, https://advocacy.code.org/2021_state_of_cs.pdf#page=25.

[51] Hendrickson et al., “2021 State of computer science education: Accelerating action through advocacy,” pg. 17, https://advocacy.code.org/2021_state_of_cs.pdf#page=21.

[52] Michael Gofman and Zhao Jin, “Artificial Intelligence, Education, and Entrepreneurship,” http://gofman.info/AI/AI_GofmanJin.pdf.

[53] Hannah Knowles, “CS research shifting toward industry,” *Stanford Daily*, February 12, 2016, https://stanforddaily.com/2016/02/12/cs-research-shifting-toward-industry/; Kevin McLaughlin, “To Find AI Engineers, Google and Facebook Hire Their Professors,” *The Information*, August 7, 2017, https://www.theinformation.com/articles/to-find-ai-engineers-google-and-facebook-hire-their-professors; Peng, “Are Commercial Labs Stealing Academia’s AI Thunder?”

[54] Elizabeth Gibney, “AI talent grab sparks excitement and concern,” *Nature*, April 26, 2016, https://www.nature.com/articles/532422a.

[55] “NSCAI Conference - Panel I: Strengthening Our Core: The Way Ahead for American R&D,” uploaded by National Security Commission on Artificial Intelligence, November 15, 2019, https://www.youtube.com/watch?v=1x8npd-9Su8.

[56] “NSCAI Conference - Panel I: Strengthening Our Core: The Way Ahead for American R&D,” uploaded by National Security Commission on Artificial Intelligence, November 15, 2019, https://www.youtube.com/watch?v=1x8npd-9Su8.

[57] Patel, Rexford, Zorn, and Morrisett, “Evolving Academia/Industry Relations in Computing Research”; Elizabeth Gibney, “AI talent grab sparks excitement and concern,” *Nature*. For evidence on this from other fields such as biotechnology, see Lynne G. Zucker and Michael R. Darby, “Virtuous circles in science and commerce,” *Papers in Regional Science,* August 23, 2007, https://rsaiconnect.onlinelibrary.wiley.com/doi/full/10.1111/j.1435-5957.2007.00133.x.

[58] Nicholas Bloom, Charles I. Jones, John Van Reenen, and Michael Webb, “Are Ideas Getting Harder to Find?,” *American Economic Review* 110, no. 4 (2020): 1104-1144.

[59] Roxanne Heston and Remco Zwetsloot, “Mapping U.S. Multinationals’ Global AI R&D Activity” (Center for Security and Emerging Technology, December 2020), https://cset.georgetown.edu/publication/mapping-u-s-multinationals-global-ai-rd-activity/.

[60] Heston and Zwetsloot, “Mapping U.S. Multinationals’ Global AI R&D Activity.”

[61] “Canada CIFAR AI Chairs,” *CIFAR*, Accessed May 2021, https://cifar.ca/ai/canada-cifar-ai-chairs/; Davide Castelvecchi, “AI pioneer: ‘The dangers of abuse are very real’,” *Nature*, April 4, 2019, https://www.nature.com/articles/d41586-019-00505-2?mod=djemAIPro.

[62] Norbine Schalij, "100 million euros, 50 professors for new Artificial Intelligence institute in Eindhoven: EAISI," *InnovationOrigins.com*, July 2, 2019, https://innovationorigins.com/100-million-euros-50-full-professors-for-new-artificial-intelligence-institute-in-eindhoven-eaisi/; "Five Dutch companies to further boost artificial intelligence in the Netherlands," *Ahold Delhaize,* October 10, 2019, https://www.aholddelhaize.com/en/media/latest/media-releases/five-dutch-companies-to-further-boost-artificial-intelligence-in-the-netherlands/.

[63] Daniel Tenreiro, "Why American Scientists Take Chinese Money," *The National Review*, February 3, 2020, https://www.nationalreview.com/2020/02/charles-lieber-case-why-american-scientists-take-chinese-money/; Ellen Barry and Gina Kolata, "China's Lavish Funds Lured US Scientists. What Did It Get In Return?" *The New York Times*, February 7, 2020, https://www.nytimes.com/2020/02/06/us/chinas-lavish-funds-lured-us-scientists-what-did-it-get-in-return.html.

[64] See, for example, the Computing Research Association's work on academia-industry relations in computing research: Patel, Rexford, Zorn, and Morrisett, "Evolving Academia/Industry Relations in Computing Research"; Sarkar et al., "CRA Industry/Academia Committee Report."

[65] Richard B. Freeman and John Van Reenen, "Be Careful What You Wish For: A Cautionary Tale About Budget Doubling," *Issues in Science and Technology*, Fall 2008, https://issues.org/p_freeman/; Chris Woolston, "Proposed NIH windfall raises hopes—and fears," *Nature*, July 27, 2021, https://www.nature.com/articles/d41586-021-02064-x?proof=t%2Btarget%3D.

[66] Many universities rely on adjunct professors to quickly fill gaps in their faculty, though these faculty generally make less money and have less job security than their full-time counterparts: "Contingent Workforce: Size, Characteristics, Compensation, and Work Experiences of Adjunct and Other Non-Tenure-Track Faculty" (Government Accountability Office, October 2017), pg. 33, https://www.gao.gov/assets/gao-18-49.pdf#page=39; Adams Nager and Robert D. Atkinson, "The Case for Improving U.S. Computer Science Education" (Information Technology and Innovation Foundation, May 2016), pg. 19, https://www2.itif.org/2016-computer-science-education.pdf#page=19. We were unable to study the labor market dynamics of adjunct faculty in computer science departments due to a lack of data.

[67] "NSF invests $9.75 million into growing the academic faculty in quantum computer science and engineering," National Science Foundation, August 4, 2020, https://www.nsf.gov/news/news_summ.jsp?cntn_id=301001&org=CCF; "Program Solicitation: NSF 19-558," National Science Foundation, May 24, 2020, https://www.nsf.gov/pubs/2019/nsf19558/nsf19558.htm. Other countries appear to be more active on this front. For example, professors in Montreal have said the Canadian province will be doubling its number of machine learning faculty due to funds committed in Canada's AI strategy; see Davide Castelvecchi, "AI pioneer: 'The dangers of abuse are very real'," Nature, April 4, 2019, https://www.nature.com/articles/d41586-019-00505-2?mod=djemAIPro.

[68] Diana Gehlhaus and Santiago Mutis, "The U.S. AI Workforce: Understanding the Supply of AI Talent" (Center for Security and Emerging Technology, January 2021) pg. 13, https://cset.georgetown.edu/wp-content/uploads/US-AI-Workforce_Brief-2.pdf#page=14; Diana Gehlhaus and Ilya Rahkovsky, "U.S. AI Workforce: Labor Market Dynamics" (Center for Security and Emerging Technology, April 2021), https://cset.georgetown.edu/publication/u-s-ai-workforce/.

[69] For more information on utilizing community colleges for AI education, see: Luke Koslosky and Diana Gehlhaus, "Training Tomorrow's AI Workforce," *Center for Security and Emerging Technology*, April 2022, https://cset.georgetown.edu/publication/training-tomorrows-ai-workforce/. For more comprehensive recommendations on strengthening the U.S. AI workforce, see: Diana Gehlhaus, Luke Koslosky, Kayla Goode, and Claire Perkins, "U.S. AI Workforce: Policy Recommendations," *Center for Security and Emerging Technology*, October 2021, https://cset.georgetown.edu/publication/u-s-ai-workforce-policy-recommendations/.

[70] "Program Solicitation NSF 22-502," *National Science Foundation*, https://www.nsf.gov/pubs/2022/nsf22502/nsf22502.htm.

[71] "Final Report: National Security Commission on Artificial Intelligence" (National Security Commission on Artificial Intelligence, March 2021), pg. 186, https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf#page=188.

[72] "Interim Report" (National Security Commission on Artificial Intelligence, November 2019) pg. 26, https://www.nscai.gov/wp-content/uploads/2021/01/NSCAI-Interim-Report-for-Congress_201911.pdf#page=31.

[73] "Artificial Intelligence R&D Investments: Fiscal Year 2018 – Fiscal Year 2022," *Networking and Information Technology Research and Development*, https://www.nitrd.gov/apps/itdashboard/AI-RD-Investments/.

[74] "American Artificial Intelligence Initiative: Year One Annual Report" (Executive Office of the President of the United States, February 2020), pg. 5, https://www.nitrd.gov/nitrdgroups/images/c/c1/American-AI-Initiative-One-Year-Annual-Report.pdf#page=11.

[75] See, e.g., Lisa Mosley, Jeremy Forsberg, and David Ngo, "Reducing Administrative Burden in Federal Research Grants to Universities" (IBM Center for The Business of Government, February 2020), https://cohortforresearch.com/wp-content/uploads/2020/02/Reducing-Administrative-Burden-in-Federal-Research-Grants-to-Universities.pdf.

[76] "Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan RFI Responses," *Center for Data Innovation*, October 26, 2018, https://www.nitrd.gov/rfi/ai/2018/AI-RFI-Response-2018-Joshua-New-CDI.pdf.

[77] "CyberCorps: Scholarship for Service," *Office of Personnel Management*, https://www.sfs.opm.gov/; "About SMART," *Department of Defense*, https://www.smartscholarship.org/smart?id=about_smart.

[78] "Program Solicitation NSF 20-586," *National Science* Foundation, https://www.nsf.gov/pubs/2020/nsf20586/nsf20586.htm.

[79] "Computer Science Education Research (CS-ER) Awards," *Google*, https://research.google/static/documents/outreach/featured-research-collaborations/cser/cser-awardees.pdf; Yuqing Li, "Amazon ML Director Begins DeepMind Professorship at Cambridge," *Medium*, September 24, 2019, https://medium.com/syncedreview/amazon-ml-director-begins-deepmind-professorship-at-cambridge-bc8f348d816; Jennifer Pascoe, "Deepening artificial intelligence expertise,"

University of Alberta, January 30, 2018, https://www.ualberta.ca/science/news/2018/january/deepening-artificial-intelligence-expertise.html.

[80] Elizabeth Gibney, "AI talent grab sparks excitement and concern."

[81] Sarkar et al., "CRA Industry/Academia Committee Report."

[82] Norbine Schalij, "100 million euros, 50 professors for new Artificial Intelligence institute in Eindhoven: EAISI," *InnovationOrigins.com*, July 2, 2019, https://innovationorigins.com/100-million-euros-50-full-professors-for-new-artificial-intelligence-institute-in-eindhoven-eaisi/; "'AI helps us squeeze every possible bit of performance our of our machines.'," *High Tech Campus*, April 2, 2021, https://blog.hightechcampus.com/htce/ai-helps-us-to-squeeze-every-possible-bit-of-performance-out-of-our-machines.

[83] "Next generation of artificial intelligence talent to be trained at UK universities," *Gov.UK*, February 21, 2019, https://www.gov.uk/government/news/next-generation-of-artificial-intelligence-talent-to-be-trained-at-uk-universities.

[84] "A strong S&E ecosystem depends on partnerships: Maintaining US leadership in AI," *National Science Foundation,* September 2, 2020, https://beta.nsf.gov/science-matters/strong-se-ecosystem-depends-partnerships-maintaining-us-leadership-ai.

[85] Patel, Rexford, Zorn, and Morrisett, "Evolving Academia/Industry Relations in Computing Research," pg. 8.

[86] Erica York, "Tax Treatment of Worker Training," *Tax Foundation*, March 21, 2019, https://taxfoundation.org/tax-treatment-of-worker-training; Diana Gehlhaus, Luke Koslosky, Kayla Goode, and Claire Perkins, "U.S. AI Workforce: Policy Recommendations" (Center for Security and Emerging Technology, October 2021), https://cset.georgetown.edu/publication/u-s-ai-workforce-policy-recommendations/.

[87] "Education and Training Provider Resources," *National Institute of Standards and Technology*, https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center/education-and-training-provider.

[88] Patel, Rexford, Zorn, and Morrisett, "Evolving Academia/Industry Relations in Computing Research," pg. 8.

[89] Kyle Gorman, "What to do about academic brain drain," *Wellformedness.com*, April 8, 2017, http://www.wellformedness.com/blog/academic-brain-drain/.

[90] Catherine Aiken, James Dunham, and Remco Zwetsloot, "Career Preferences of AI Talent" (Center for Security and Emerging Technology, June 2020), https://cset.georgetown.edu/publication/career-preferences-of-ai-talent/.

[91] Our survey had too few respondents to reliably test whether this is the case or not. CSET is currently fielding another survey to examine the importance of compute resources and other factors for AI researchers at different types of institutions.

[92] While proponents claim the NAIRR will democratize access to AI development tools, critics argue the platform will further entrench the power of Big Tech companies: David Ingram, "Big Tech is pushing a 'national cloud.' Critics say Big Tech would profit from it." *NBC News*, October 21, 2021, https://www.nbcnews.com/tech/tech-news/big-tech-pushing-national-cloud-critics-say-big-tech-profit-rcna2971; "The Biden administration launches the National Artificial Intelligence Research Resource Task Force," *National Science Foundation*, June 10, 2021, https://www.nsf.gov/news/news_summ.jsp?cntn_id=302882&org=NSF.

[93] Adams Nager and Robert D. Atkinson, "The Case for Improving U.S. Computer Science Education" (Information Technology and Innovation Foundation, May 2016), pg. 28, https://www2.itif.org/2016-computer-science-education.pdf#page=28.

[94] Mark Hagerott, "Time for a Digital-Cyber Land Grant System," *Issues in Science and Technology*, Winter 2020, https://issues.org/time-for-a-digital-cyber-land-grant-system/.

[95] For an overview, see Patel, Rexford, Zorn, and Morrisett, "Evolving Academia/Industry Relations in Computing Research."

[96] Patel, Rexford, Zorn, and Morrisett, "Evolving Academia/Industry Relations in Computing Research," pg. 6; "Driving U.S. Competitiveness Through Improved University-Industry Partnerships" (Association of Public and Land-Grant Universities, 2021), https://www.aplu.org/library/driving-us-competitiveness-through-improved-university-industry-partnerships/file.

[97] For current faculty, see "Science and Engineering Indicators 2018: Table 5-17" (National Science Board, 2018), https://www.nsf.gov/statistics/2018/nsb20181/assets/968/tables/at05-17.pdf; for current PhD students, see Zwetsloot, Dunham, Arnold, and Huang, "Keeping Top AI Talent in the United States," pg. 2–3; for overall U.S. employment, see "Science and Engineering Indicators 2019: Figure 3-24" (National Science Board, September 2019) https://ncses.nsf.gov/pubs/nsb20198/immigration-and-the-s-e-workforce#characteristics-of-foreign-born-scientists-and-engineers.

[98] For more background on the U.S. immigration system and how it affects the U.S. AI sector, along with specific immigration reform recommendations, see Zachary Arnold, Roxanne Heston, Remco Zwetsloot, and Tina Huang, "Immigration Policy and the U.S. AI Sector," (Center for Security and Emerging Technology, September 2019), https://cset.georgetown.edu/wp-content/uploads/CSET_Immigration_Policy_and_AI.pdf.

[99] Kevin Shih, "Do international students crowd-out or cross-subsidize Americans in higher education?" *Journal of Public Economics* 156, (December 2017): 170-184.

[100] Elizabeth Gibney, "AI talent grab sparks excitement and concern"; Eric Roberts, "A History of Capacity Challenges in Computer Science" (Stanford University, March 7, 2016), https://cs.stanford.edu/people/eroberts/CSCapacity.pdf; Patrick Coffey, *Cathedrals of Science: The Personalities and Rivalries That Made Modern Chemistry* (Oxford University Press: 2008).

[101] Richard B. Freeman and John Van Reenen, "Be Careful What You Wish For: A Cautionary Tale About Budget Doubling," *Issues in Science and Technology*, Fall 2008, https://issues.org/p_freeman/; Stephan, *How Economics Shapes Science*; Teitelbaum, *Falling Behind?*.

[102] Jacob D. Biamonte, Pavel Dorozhkin, and Igor Zacharov, "Keep quantum computing global and open," *Nature*, September 11, 2019, https://www.nature.com/articles/d41586-019-02675-5; Mark Piesing, "'How can we compete with Google?': the battle to train quantum coders," *The Guardian*, January 15, 2020, https://www.theguardian.com/education/2020/jan/15/how-can-we-compete-with-google-the-battle-to-train-quantum-coders.

[103] The annual response rate hovered between 73 percent (2018) and 80 percent (2012 and 2016).

[104] Eric Roberts, "Outcomes of Advertised Computer Science Faculty Searches 2019" (Worcester Polytechnic Institute, October 2019), pg. 10, https://web.cs.wpi.edu/~cew/papers/outcomes19.pdf#page=11; "Artificial Intelligence Index Report 2021" (Stanford University, 2021), pg. 113, https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf#page=113; We recognize that other academic disciplines, such as engineering, also offer training in AI.

[105] Eric Roberts, "A History of Capacity Challenges in Computer Science" (Stanford University, March 7, 2016), https://cs.stanford.edu/people/eroberts/CSCapacity.pdf; Note: the number of bachelor's degrees awarded in a given year lags roughly two to three years behind enrollment data.

[106] Roberts, "A History of Capacity Challenges in Computer Science."

[107] The "Other" category includes faculty who died or departed for unknown or "other" reasons, as defined by the CRA Taulbee Survey.