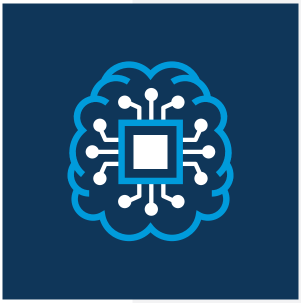

INTEGRATED SUMMARY:

ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) has the potential to substantially increase productivity, output, employment, and scientific discovery across the US economy, but the invention/diffusion process is still in early stages and not all firms, regions, demographics, or scientific fields are benefiting.

Type of critical technology assessment Emerging technology, high economic and security impact

Lead performers Lee Branstetter, Erik Brynjolfsson, Thema Monroe-White, Dewey Murdick, Dashun Wang

Program management Compare different datasets held by different performers to overcome sample and data limitations

Methods Large language models, machine learning, surveys, descriptive statistics, econometrics (causal analyses)

Data Publications, patents, Bureau of Labor Statistics Survey, US Census data

Criticality dimensions measured Economic well-being (S&T competitiveness, productivity, jobs), societal well-being (participation)

Challenges for future critical technology assessment Inadequate availability of and access to timely data—including from private sources—available to top analysts, given the rapid rate of change of the technology; sharing of data and algorithms; broader geographic and demographic participation; demographic impacts of algorithm bias

Additional contributors: Hwijeen Ahn, Catherine Aiken, Sarah Bana, Sagar Baviskar, Krisztina Eleki, Jacob Feldgoise, Jian Gao, Bishu Giri, Emma Herrerra, Eduard Hovy, Christie Ko, Luke Koslosky, J. Frank Li, Nestor Maslej, Tanvi Murke, Maria Ryskina, Shubham Shastri, Sebastian Steffen, Nikolas Zolas

ARTIFICIAL INTELLIGENCE

Different data point in *same direction*
(complementing weaknesses)

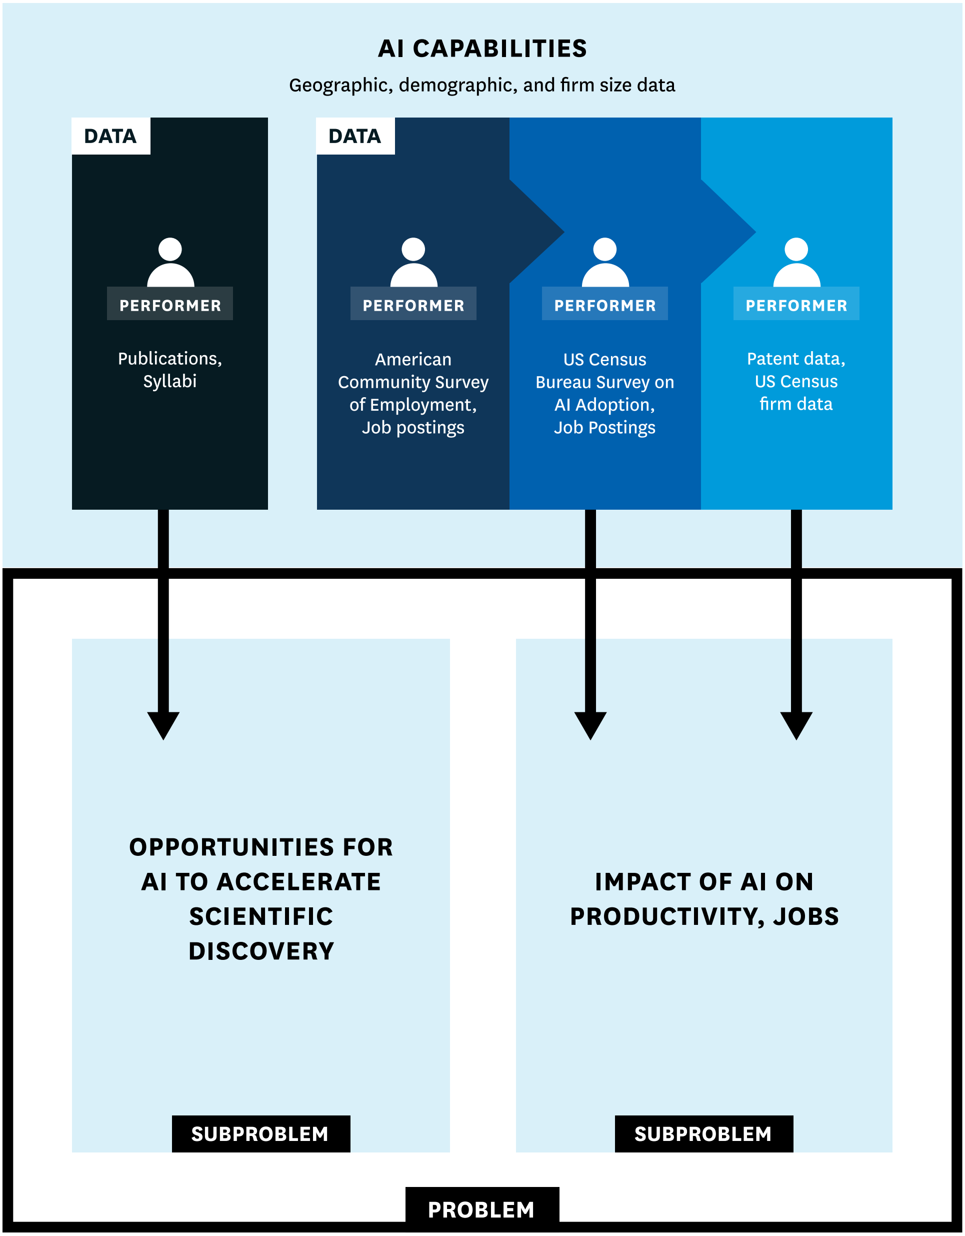

ARTIFICIAL INTELLIGENCE

FINDING: New large-sample survey data indicate that AI adoption is limited to larger, more technologically sophisticated firms and concentrated in a handful of "superstar" cities.

RECOMMENDATIONS: Improve measurement by examining indirect AI adoption through digital services. Expand the ranks of AI workers with the skills needed to work at the disciplinary frontier in AI, through both immigration and support of advanced education of domestic students, to reduce one of the major constraints to AI adoption by smaller enterprises.

FINDING: New firm-level data suggest that AI inventions lead to substantially more rapid growth in the inventing firm's productivity, output, and employment.

RECOMMENDATION: Support basic research and graduate education in AI-related fields while improving methods for measuring AI innovation at the firm level. Create a National AI Research Resource (NAIRR) to provide greater access to the computational resources and datasets for academics, nonprofit researchers, and startups from diverse backgrounds.

FINDINGS: Analysis of US employment and job posting data finds that occupations with AI-relevant knowledge, skills, and abilities represented about 9% of US employment in 2019 and are projected to grow twice as fast as all US occupations. AI occupation supply and demand are also geographically concentrated in several metropolitan areas, including some that are located outside of known "tech hubs."

RECOMMENDATIONS: Authorize funding to staff AI office and workforce support initiatives, such as by increasing staffing at the National Artificial Intelligence Initiative Office for Education and Training; develop a federal framework of technical and nontechnical AI work roles and competencies; and establish federal grant programs for AI industry-academia partnerships, AI-related degree and nondegree programs at community colleges and minority-serving institutions, and equipment at AI labs and related facilities.

FINDING: AI is impacting scientific research, but not all fields and scholars are benefiting from this shift, and the teaching of AI is lagging behind.

RECOMMENDATIONS: Expand the AI-related professoriate immediately by broadening opportunities for foreign graduates of related US PhD programs to remain in the United States; redesign university curriculum to teach more AI skills and facilitate cross-department collaborations with AI experts; and increase funding for female and underrepresented groups to pursue graduate study in AI-related fields.

FINDING: Underinvested and underrepresented segments of the US population are not being engaged in AI in ways that would maximize innovation or national interests, and they experience more stress when pursuing STEM fields.

RECOMMENDATIONS: Targeted programs are needed to increase representation in STEM of diverse identities not only to more fully leverage talent but also to mitigate harms perpetuated by biased AI systems. To uncover inequalities related to AI-powered technology, future work will need to study who is producing algorithms, in what kinds of organizations, for whom, and what data are used in the algorithms.

Research Questions

What are the most effective ways to measure the implications of innovations in artificial intelligence for prosperity, jobs, and equity? What is the potential for AI to drive advances in scientific research? Which firms adopt AI-related technologies and what are the effects of adoption? What does the US AI workforce look like and how can it be leveraged and expanded?

Motivation/Framing

After decades of incremental progress, artificial intelligence (AI) has made impressive strides over the past 15 years, prompting talk of a 4th industrial revolution. However, US aggregate productivity growth remains stuck at historically low levels, holding down growth in living standards, geopolitical power, and fiscal sustainability. Will AI live up to its promise, generating an industrial revolution that raises productivity growth?

Impacts on aggregate productivity of past technological revolutions have taken decades to emerge because of the slow processes of complementary innovation and technology adoption required for a new "general purpose technology" to work its way into the entire economy. A definitive assessment of the impacts of AI is years away, but preliminary evidence can be obtained by exploring the impacts of AI invention and adoption on the inventing and adopting firms, which are likely to be in the vanguard of any AI revolution. To this end, our research has developed new methods for identifying and measuring AI invention and adoption at the firm level—something official government datasets have historically not captured. We have also developed new methods for identifying AI-related scientific publications.

Methods and Sources of Data

We developed new methods for measuring AI invention and adoption at the firm level; for analyzing their impacts on firm output, employment, and productivity; and for identifying AI impacts on scientific research.

Our CMU team developed machine learning algorithms that parse the text of US Patent and Trademark Office patents to identify those that are AI-related. These algorithms also provide a univariate measure of the AI-intensiveness of each patent, allowing us to experiment with various thresholds of "AI-ness." Through a partnership with the US Census Bureau,[1] we link these patents to US firms that create the inventions these patents protect, using the bureau's carefully developed "crosswalk" that links patent owners to US firms. Because both patent data and Census surveys are regularly updated, they can be used to track the impact of AI invention on inventing firms in future years.

Our Stanford team worked with the Census Bureau over several years to create, implement, and refine a survey of AI use and adoption by US enterprises. This provides badly needed visibility into the degree to which, and the processes by which, American firms have adopted AI technologies created by other firms. The labor-intensive nature and expense of these surveys mean they cannot be conducted often, and data access is limited. Nevertheless, the data on adoption provide a useful window through which to observe the impacts of AI on output, employment, and productivity, and one that complements the window provided by our data on AI invention. Over time, the Census Bureau will conduct further surveys, generating a rich panel dimension to the data that will enable continuing statistical analysis of the impacts of AI adoption on firm-level outcomes.

Wang's team used natural language processing techniques and comprehensive data from Microsoft Academic Graph, the Open Syllabus Project, and the Survey of Doctorate Recipients to estimate AI effects on the nature, composition, and impact of scientific research.

[1] Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The DRB codes for this project are DRB-B0027-CED-20190205, CBDRB-FY19-414, CBDRB-FY20-105, CBDRB-FY22-182, and CBDRB-FY22-CES007-004.

Georgetown University's Center for Security and Emerging Technology (CSET) created a series of maps that compares its measure of AI employment, Stanford's measure of AI job postings, and CMU's measure of AI invention (Gehlhaus and Rahkovsky 2021). The CSET team defined the AI workforce by linking the skills and competencies necessary to design, develop, and deploy AI systems to 54 occupations as defined by the Department of Labor. Both technical and nontechnical occupations are needed to develop safe and effective AI systems. The team analyzed data from the Census Bureau's American Community Survey, occupational employment projections from the Bureau of Labor Statistics, and job posting data from Burning Glass (now Lightcast) and LinkedIn Insights (**box 4-2**).

Integrative Findings

AI INVENTION RAISES OUTPUT, PRODUCTIVITY, AND EMPLOYMENT

The CMU team's algorithms identified significant numbers of AI patents since the 1990s, although the early numbers are dwarfed by the scale of AI invention in the 2010s. This long panel dimension to our data makes it possible to compare the productivity growth of AI-inventing firms to that of other firms—a dimension of comparison that economists refer to as the extensive margin. We can also observe how the same firm's output and productivity vary as it invents additional AI-related technologies, a dimension of comparison we refer to as the intensive margin. We see evidence that AI invention boosts firm output per employee by 15–27%, value added by 10–23%, and total factor productivity by 6–8%. These are economically large effects, and they are all statistically significant. While it is not possible to confirm that these effects are causal, tracking firms over time provides a degree of leverage around the possibility that both AI invention and productivity increases are driven by some omitted third variable. Despite concern that AI adoption might lead to significant declines in employment, our results suggest that AI invention leads to growth in employment, although our data do not identify gains or losses for particular types of jobs.

BOX 4-2

Combining Data Sources for a Whole Greater than the Parts

Lee Branstetter

The AI team found that the synergistic combination of multiple datasets can make up for significant flaws in any one dataset. AI-related patents matched to firms and assigned the date of application provide rich, detailed data on AI invention, but all patent data are subject to the problem that not all patents result in real inventions and not all real inventions are patented. Thus patents alone may or may not correspond to economically meaningful innovation. By matching patent data to census firm-level input and output data, one can observe statistically significant and economically meaningful changes in output, employment, and productivity that could be statistically associated with AI-related patenting (invention) in both the intensive and the extensive margin. By bringing in Annual Business Survey data on firm-reported AI adoption and matching these data to the same firms, researchers can simultaneously observe the firms' AI invention and adoption. But these survey data are costly to obtain, and it will be many years before they acquire a time series dimension sufficient for the econometric techniques that are used to make causal inferences from observational data, although patents have a deep time series dimension that offsets this shortcoming in adoption data. Finally, Lightcast/Burning Glass data on AI hiring could be linked to the same census firms, providing yet another dimension to indicate which firms are investing in AI capability. In this way the "holes" or shortcomings in any one data series are partially compensated for by the others, and the combined complementary pictures of AI adoption, use, hiring, and innovation sketched out by different datasets yield a much richer, and likely more accurate, picture of the phenomenon.

In addition to these regression-based results, data on AI-related patenting enable us to examine the distribution of AI invention across geographic boundaries, time, firms, and industries (**figure 4-7**). These results complement the Stanford team's findings that AI adoption also is correlated with growth and increased employment.

AI ADOPTION IN THE UNITED STATES IS CORRELATED WITH SUBSEQUENT GROWTH, BUT ITS INCIDENCE IS HIGHLY UNEVEN ACROSS FIRMS AND GEOGRAPHY

The Stanford team analyzed data from the Census Bureau's 2018 Annual Business Survey of over 850,000 firms to establish a number of stylized facts about early AI adoption in the United States. While less than 6% of firms use any of the AI technologies we measure, adoption is prevalent in firms with the following characteristics: over 5,000 employees; owners who are more educated and experienced with AI, younger, and motivated by aspirations such as bringing new ideas to market or helping the community; early markers of high-growth entrepreneurship, innovation, and growth-oriented strategies; and location in a handful of "superstar" cities. AI use is conditionally correlated with significant later-stage firm growth. In addition, AI job postings are correlated with increases in job postings outside AI. The concentration and growth potential of AI's leading edge portend economic and social impacts far beyond this limited early diffusion, along with a potential "AI divide" if early patterns persist.

We characterize AI adoption patterns at the core-based statistical area (CBSA) level and find significant geographic disparity. We focus on single unit firms to pinpoint the exact location of AI use, then calculate the number of those firms in the CBSA (weighted by employment) and the percentile rank of the CBSA in terms of AI usage rate (lighter colors correspond to higher rankings). We look separately at all single unit firms and young startups. Regions that are well known for pioneering technologies, such as Silicon Valley and the Research Triangle, stand out with high AI intensity. Areas in the Northeast and Midwest have lower AI intensity as a share of the number of firms, as indicated by the size of bubbles. Further discussion of our results is in our working paper, "AI Adoption in America: Who, What, and Where" (McElheran et al. 2021).

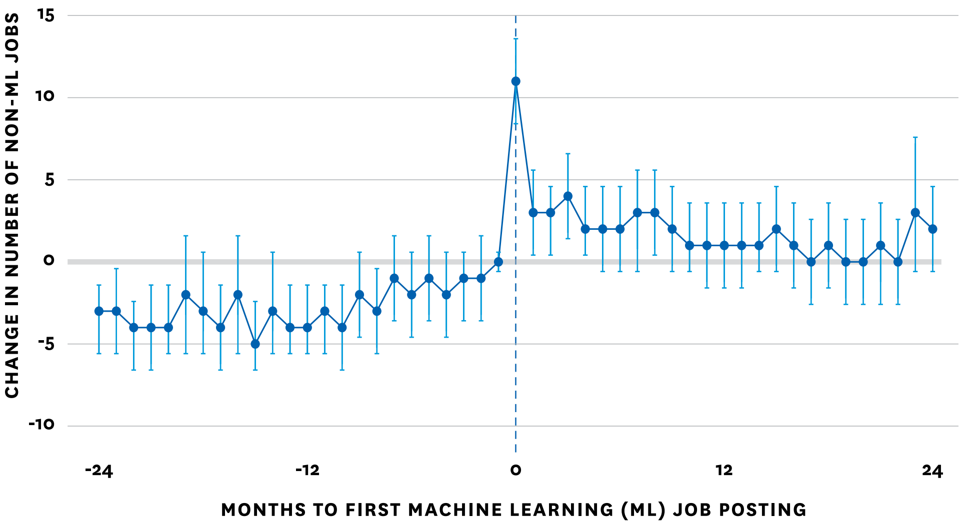

FIGURE 4-7. The Stanford team's work shows that more AI job postings correlate with more non-AI job postings at firms.

AI BENEFITS SCIENTIFIC RESEARCH, BUT BENEFITS ARE UNEVEN ACROSS FIELDS AND CATEGORIES OF RESEARCHERS

The direct and potential impacts of AI on scientific research are analyzed using semantic analysis of AI papers and patents, and scientific papers across fields. Direct impact is measured using the frequency with which words and phrases from AI papers and patents appear in papers in other fields. Potential impact is measured by extracting verb-noun pairs from the titles of AI papers and patents (i.e., what AI can do) and comparing these to verb-noun pairs in the titles of papers across fields (what the field does).

First, the use of AI appears widespread throughout the sciences, growing especially rapidly since 2015, and papers that use AI exhibit a citation impact premium. Second, despite heterogeneity in AI's impact across research areas, almost every discipline has some subfields that benefit substantially from AI innovations. Third, analysis of university course syllabi across 17 disciplines reveals a systematic misalignment between the teaching of AI in higher education and its impact on scientific research (**figure 4-8a**), suggesting that the preparation and supply of AI talent in scientific disciplines is not commensurate with AI research demand. Fourth, rapid advances pose growing knowledge demands on individual scientists, who increasingly rely on collaborators with AI expertise instead of working to push AI applications forward in their disciplines (**figure 4-8b**). Fifth, women and underrepresented minority scientists benefit substantially less from AI advances, which may exacerbate existing inequalities in science.

Options and Tradeoffs for the US Government

AI OFFERS A PROMISING POTENTIAL ROUTE TO FASTER PRODUCTIVITY GROWTH

The most important determinant of growth in future US living standards, economic size, and global power is arguably the country's rate of productivity growth, which has been stuck at low levels since the mid-2000s.

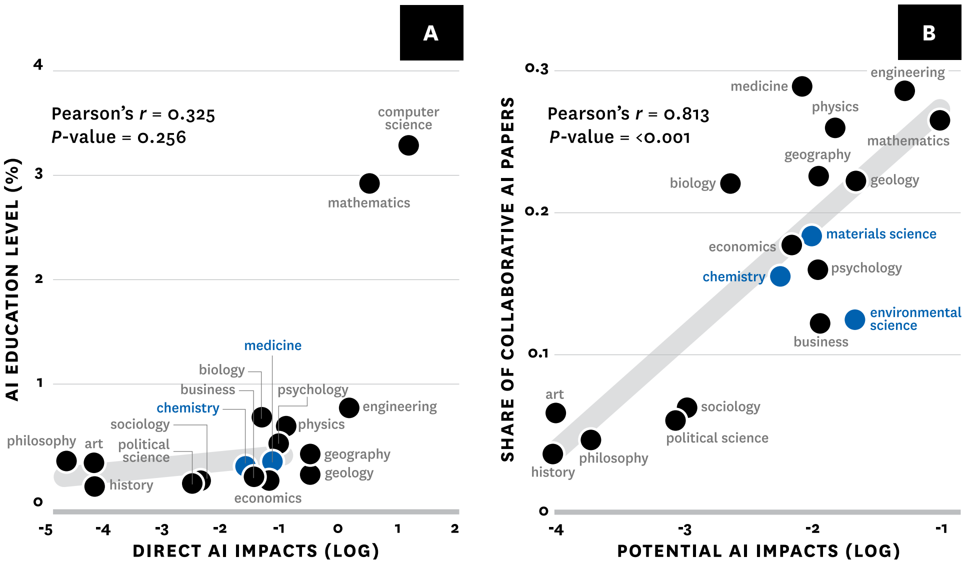

FIGURE 4-8. Estimating the benefits of AI in science. (a) Correlation between AI impact score and AI education levels. (b) Correlation between the share of collaborative AI papers and potential AI impact.

Our results provide grounds for optimism that continued innovation in AI and firm adoption of AI inventions could help spur significant and lasting acceleration in productivity growth. The federal government should seek to support this by (i) continuing to invest in AI-related basic research, (ii) expanding the domestic pipeline for AI talent by supporting graduate education in AI-related disciplines, (iii) taking meaningful steps to increase the number of foreign graduates of US AI-related programs who receive permission to work in the United States, especially in teaching positions at US universities, and (iv) investing in continued efforts to measure the invention and adoption of AI at the firm level. Realizing the potential productivity benefits of AI will also require continued societal attention to issues related to how AI changes the nature of jobs, increasing some kinds of employment while decreasing other opportunities.

AI IS HAVING PROFOUND — BUT UNEQUAL — IMPACTS ON SCIENTIFIC RESEARCH

The pervasive impact of AI across disciplines and its rapid advances pose growing AI knowledge demands on scientists. In particular, the misalignment between AI education and AI's impact on science indicates a critical need to redesign university curricula for teaching more AI skills and/or to facilitate cross-department collaborations with AI experts. Both AI education and collaboration will upskill scientists, and this has implications for preparing next-generation scientists to take full advantage of cutting-edge AI advances in their research. It is also important to recognize that, as AI becomes increasingly capable of performing research tasks, it may create unequal impacts on the research workforce. Our analysis reveals inequalities in AI's benefits for science, with implications for building a diverse, equitable, and inclusive research workforce.

THE US AI WORKFORCE AND PATENTS ARE GEOGRAPHICALLY CONCENTRATED

Figure 4-9 compares CSET's measure of AI employment, Stanford's measure of AI job postings, and CMU's measure of AI invention (patents). There is geographic concentration in AI occupations and skills demand, primarily in Los Angeles, San Francisco, Chicago, New York, and Seattle.

Vision for Future Analytic Work

Not all AI inventions are patented. How do we measure AI invention when patents are not generated? Firms seeking to use AI to either introduce or substantially reengineer products or services need to hire AI experts trained up to the technology frontier. The CMU team is using publication data to identify star AI scientists and the doctoral students and postdocs with whom they coauthor. We then use a mix of publication and social media data to trace the movement of these experts from the academy where they are trained and into firms. Using our link to Census data, we can test the hypothesis that firms acquiring a critical mass of PhD-level AI experts trained by star scientists experience large productivity gains. The CSET team is identifying other subsets of AI talent and mapping their education and career histories. The team is also drawing on novel data to explore trends in the Chinese AI workforce, which can provide important insight and help inform US policy actions.

A better understanding of AI's impact on science may not only help guide AI development, bridging AI advances more closely with scientific research, but also have implications for science and innovation policy. The work by Wang's team takes an initial step in assessing how AI might impact scientific research. As AI research evolves rapidly, there is a critical need for continuous monitoring and updates to estimates of AI's benefits for scientific research.

The team is using large-scale datasets covering about 6 million research grants and resulting publications to study whether funding support for AI research is commensurate with AI's scientific impacts. This analysis may inform funding allocation strategies to better support AI research that may benefit the development of many research fields.

AI EMPLOYMENT IN 2019 BY US CORE-BASED STATISTICAL AREA (CBSA)

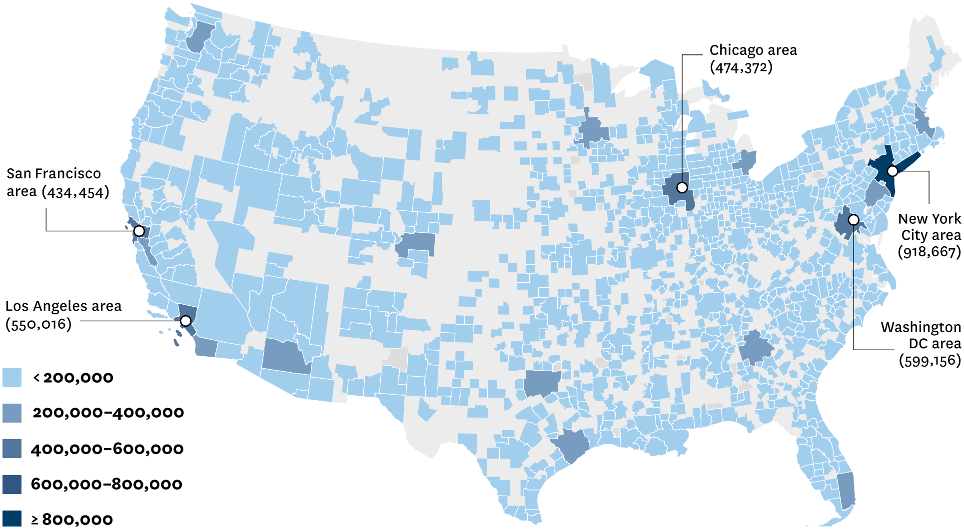

Note: 357, 188 AI employment records did not have location data.
Map: Center for Security and Engineering Technology. Source: American Community Survey

AI-RELATED JOB POSTINGS IN 2019 BY US CORE-BASED STATISTICAL AREA (CBSA)

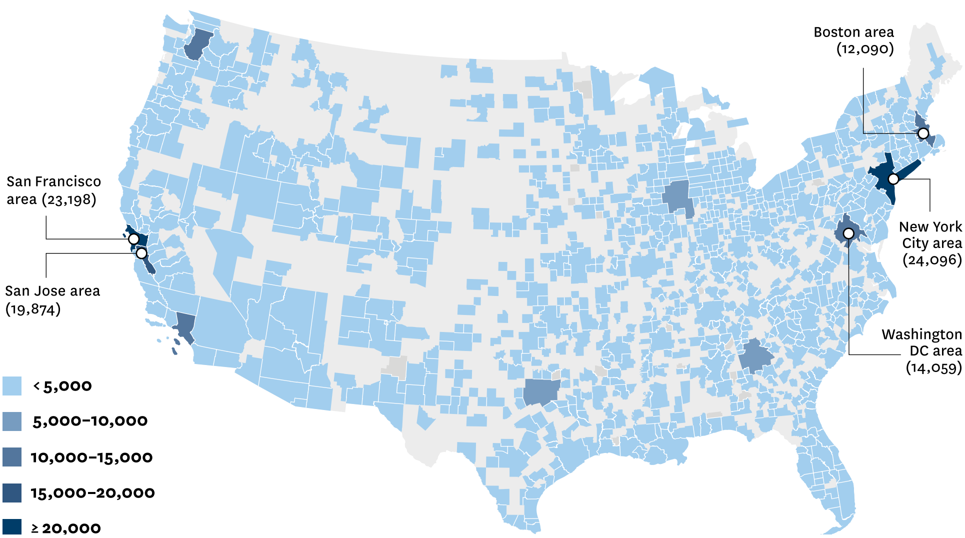

Map: Center for Security and Engineering Technology
Source: Lightcast, Stanford University NNCTA Team

CUMULATIVE AI PATENTS THROUGH 2018 BY US CORE-BASED STATISTICAL AREA (CBSA)

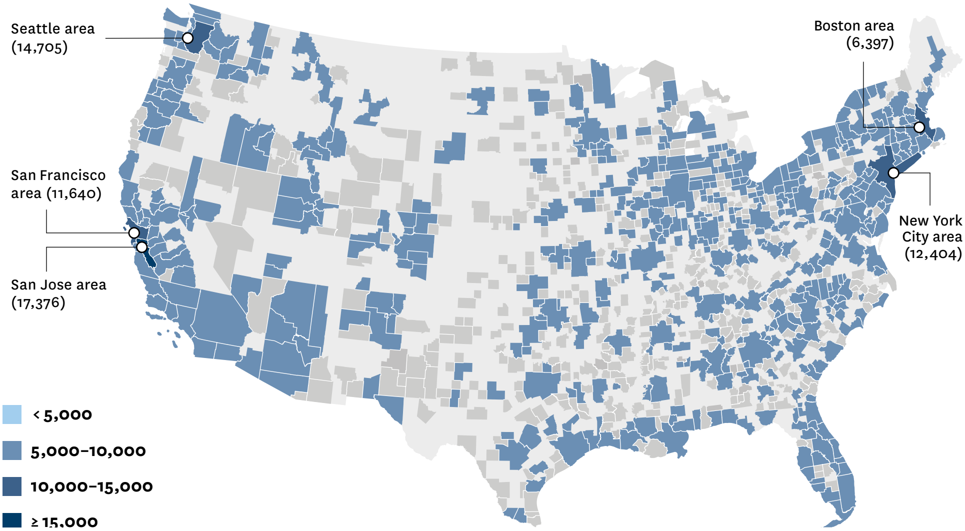

Map: Center for Security and Engineering Technology
Source: USPTO, Carnegie Mellon University NNCTA Team

FIGURE 4-9. Comparison of US core-based statistical area capabilities in AI according to different NNCTA teams' measures: employment, job postings, and cumulative patents.

Going forward, important questions about AI, equity, and labor need to be addressed. AI-powered technologies may affect different communities differently; in particular, racial, educational, and immigration status disparities in paid work may be exacerbated with AI and automation. Evidence also suggests that the United States is failing to leverage substantial STEM and AI talent (e.g., Black, Indigenous, Latinx, rural communities, and women of all races). Addressing these issues will require systematic development and collection of metrics that capture how AI impacts different types of jobs and different types of workers.

Increased representation in STEM of diverse intersectional identities (e.g., race, gender, among others) is necessary to mitigate harms perpetuated by biased AI systems. To understand how inequalities relate to AI-powered technology, research on AI should consider who is developing AI, based on what knowledge, in what kinds of organizations, and for whom and what uses. To that end, Hoffman et al. (2022) articulate five critical questions: (1) What do data mean? Problems occur when AI system designers and users fail to see that neutral-seeming data (e.g., criminal record, ZIP codes, location of hospitals) also reveal socially significant inequalities (e.g., class, gender, race, segregation, racist policing practices). (2) What are myths about AI? A myth that AI accomplishes human-level tasks without human intervention can make it more difficult to observe how social actors are shaping where, why, and by what means AI is used in practice. (3) How do interlocking structures of inequality influence AI systems? Intersectional analyses can show which human actors and values drive AI development and identify harms from AI systems across age, race, ethnicity, gender, sexuality, and class. An intersectional approach can also help everyone imagine new futures in which benefits (and harms) are distributed more equally. (4) Where is labor to support AI going unnoticed? Firms that provide seemingly futuristic AI capabilities often outsource or offshore the necessary work of contract laborers who engage in a range of small tasks that help ensure automated systems' accuracy and efficiency, labor hidden behind platform interfaces. (5) What more just AI futures can be imagined? Problems in AI development are not inevitable. Research must be used to create more equitable knowledge production contexts for this critical technology.

Potential Broader Lessons for Critical Technology Assessment

In principle, the methods applied to measure AI innovation and adoption and their effects on inventing firms could be adapted to other critical technologies. Machine learning algorithms could be used to parse patent documents and identify those associated with other critical technologies, showing the distribution of inventive activity across geography, time, and firms. Then the Census Bureau's patent-assignee-to-enterprise crosswalk could be used to connect the patents to the inventing firms. This would enable researchers to (i) estimate the impact of the critical technology on inventing firm output, productivity, and employment; and (ii) place invention in the targeted critical technology, and its effects, in the larger context of US aggregate innovation and productivity growth. Expanding and regularly conducting Census Bureau surveys to examine adoption of other critical technologies would enable the government to assess the impact of adoption on firm outcomes such as output, employment, and productivity.