MERCATUS CENTER
George Mason University

POLICY BRIEF

Building America's Applied AI Workforce: Strategies for Effective Technology Diffusion

Dean W. Ball

September 2024

American businesses and research universities currently lead the world in artificial intelligence (AI) innovations. According to Stanford University's most recent AI Index, US firms and researchers produced 61 cutting-edge AI models in 2023, compared to 15 from Chinese organizations.[1] However, innovations matter little if they are not utilized. The adoption of a new technology by a nation's businesses, individuals, government, and civil society is called technology diffusion. If the United States wants to retain its lead in AI and technology more broadly, it will have to lead not just in innovation but in diffusion.

The US currently produces some of the world's finest AI researchers and is home, by one estimate, to eight of the top ten computer science universities in the world.[2] Doctoral students at Stanford, Massachusetts Institute of Technology, and elsewhere regularly make impactful advancements to AI model architectures, training processes, and similar foundational enhancements. Those researchers pursue careers in academia or in the AI industry, often earning compensation well into the seven figures.

These kinds of researchers produce invaluable innovations, but in general, they do not drive *adoption* of their innovations. They may invent a new kind of model architecture that significantly improves the performance of AI at various tasks, but they are not likely to be interested in deploying those models in, say, a midsize civil engineering company in Alabama. The former describes the profile of a traditional AI researcher, while the latter describes a worker with applied AI skills.

In order to lead in AI diffusion, the US will need to cultivate an applied AI workforce—people who can integrate AI into businesses of all sizes. This is a very different skill set and worker profile from that of the PhD-level AI researcher, and thus will require a different set of policies and other measures than the traditional approaches to strengthening the AI workforce.

Diffusion itself requires quite a bit of innovation to get right. The invention of language models such as ChatGPT is one kind of innovation; a different set of innovations altogether is required to maximize the potential of these language models within businesses small and large. This latter kind of innovation is difficult to track because it is decentralized and incremental. It does not grab headlines, and it is hard to characterize as a single "moonshot," yet AI diffusion is just as, if not more, important than AI innovation.

AI has the potential to accelerate scientific discovery and boost economic productivity through many sectors of industry in every state. To this end, this report will recommend two pathways through which state and local governments can help cultivate an applied AI workforce:

1. Create and support a range of education and workforce training opportunities in applied AI
2. Improve data collection to track AI diffusion

Technology diffusion takes time. When factories first electrified, they did so using the same physical configuration they used with steam power: The power source was at one place within the factory, and machines were all located "downstream" of that source, organized roughly in descending order of their power demands. It took decades for Henry Ford to realize that electrification meant power could be distributed *throughout* the factory and that the placement of machines could be organized according to their logical role in the production chain rather than their power needs. This, in turn, led to the assembly line.[3]

It is quite likely that AI diffusion will follow a similar pattern: AI will be "bolted on" to existing products and business processes. It will take time and skill to adapt those things *to* AI, enabling capabilities that were fundamentally impossible without AI. Finding those unique use cases is a market-oriented search and discovery process. Equipping firms with the skilled labor they need to go about that search process will accelerate the discovery of AI's full value.

Building an Applied AI Workforce

Computer science is one of the most popular undergraduate majors in the United States, and AI is one of the most lucrative professions within the software industry. This talent pipeline produces a wide variety of researchers, engineers, and other similar professionals, but this system does not give students the knowledge or incentives they need to work in applied AI. Rather, it is highly theoretical. This is especially true of undergraduate and graduate programs concentrated on machine learning: The focus tends to be on why and how AI works, rather than how to use it practically. As one might expect, this education is highly technical, multidisciplinary, and takes years to master, often culminating in a PhD.

While this system is essential for US leadership in AI innovation, a different educational approach is needed to create a rich pool of talent in applied AI. As the name suggests, applied AI should be focused on practical, "here and now," applications of AI. Programs geared toward the AI workforce should give students a broad grounding in the theory, history, and potential future of AI research—enough for them to follow the field as informed observers, but not necessarily enough for them to actively make contributions to the bleeding edge.

These programs should also emphasize speed and seek to complement existing educational pipelines and curricula rather than creating new specializations. It is unlikely, for example, that non-AI firms will want to employ a "Chief AI Officer" or a similar position in the long term. Instead, many firms are likely to add AI-related responsibilities and functions to existing roles. This could mean many different things for different employees and companies. For example, an IT leader at a firm may want to pursue a one-year, part-time "Applied AI" program at a local college, whereas a senior executive may opt for a higher-level certificate program that they take online on their own time.

An applied AI worker may not exclusively do AI-related work. They may be an IT director at a medium-sized business. They may be a small-business owner. They may even be a PhD-level researcher in other fields. For example, while many biotech researchers have robust skills in lab work and traditional computational approaches to biology, they may have little to no training in machine learning. There are now hundreds of free, online biology-related AI models that can assist in different aspects of biotech research: protein structure prediction, nucleic synthetic sequence prediction, and even modeling interactions between different biomolecules. Many mid-career biotech research professionals are aware that these new tools exist, but they do not concretely know how to take advantage of them.

The following comparison between traditional AI research talent and applied AI talent highlights the need for distinct educational approaches.

Traditional AI research skills

- Advanced mathematical and statistical knowledge in fields such as linear algebra, probability theory, and information theory
- Deep understanding of core machine-learning concepts such as gradient descent, backpropagation, loss functions, and related ideas
- A scientific mindset: creating and iterating on experimental design rather than engineering products
- Expert-level programming skills in languages such as Python and more complex tools such as NVIDIA's CUDA (Compute Unified Device Architecture)
- A PhD, often (though not always) required in a quantitative field such as statistics, math, physics, or computer science

Applied AI skills

- Ability to use application programming interfaces (APIs) provided by OpenAI, Anthropic, Google, and other model providers in production settings
- Skills to fine-tune and otherwise adapt those models on company data, which may further necessitate skills such as data cleaning and preprocessing
- An intuitive grasp of the capabilities and limitations of advanced AI models
- Deep knowledge of an existing company or industry's business processes, and ability to creatively imagine how AI may be used to automate, accelerate, or otherwise augment those processes
- Understanding of the privacy and security implications of using different kinds of AI models, including the differential implications of open- and closed-source models
- Ability to use cloud-computing tools provided by firms like Amazon Web Services, Microsoft Azure, Google Cloud, and others
- Basic experience with programming in languages such as Python

Put more simply: Applied AI workers need an intuitive grasp for what kind of AI advancements—say, a new, more capable model—could benefit their goals and know how to wield those advancements to their advantage, whereas AI researchers actually *create* the advancements.

Expanding educational opportunities

While computer-science students in the US focus on theoretical subjects rather than the practical application of AI to real-world problems, the majority of China's more than 2,300 AI-focused undergraduate programs are focused on applied AI. Nearly half of China's applied AI programs are devoted to applications of AI in the "world of atoms"—fields like construction and manufacturing—rather than the "world of bits" (the digital economy). Even if the US maintains its lead in frontier AI research, it could easily fall behind in applying that research to the practical problems facing US industry.

Indeed, this pattern of superior technology innovation in one country but superior adoption in another has played out in prior periods of technological transformation. Jeffrey Ding, a scholar at George Washington University, has described this dynamic between the United States and Europe in the Second Industrial Revolution. In the 19th century, Europe was home to the world's leading universities, and US students often traveled there to complete their education in scientific and technical fields. European universities and researchers accounted for a significant share of leading papers and prize-winning contributions to various fields. Ding argues, however, that the US led in the adoption of these cutting-edge breakthroughs in industry and other practical endeavors.[4]

It is not hard to see the present-day parallels between the US and China—except this time, the US leads in research and innovation, and China could lead in adoption. What follows are suggestions for how US policies and educational programs can support the creation of an applied US AI talent pipeline that is diverse and flexible enough to meet the wide range of potential outcomes that AI promises.

1. **Support applied AI courses and programs at community colleges and public universities**
    a. State legislatures should allocate funds to develop and staff applied AI course offerings at higher education institutions.
    b. When possible, these programs should be designed and executed in collaboration with industry partners. This can include everything from guest lectures by representatives of local firms that have had success with AI adoption to curricular partnerships with large-scale AI firms such as Microsoft, Google, and Amazon.
2. **Promote AI-related electives and extracurriculars within high schools**
    a. Lawmakers should incentivize school districts to add elective classes focused on AI—not necessarily on applied AI—to their curricula. The goal of these programs would be to give students a broad-based understanding of AI and teach them about the field's history, recent developments, and potential future.
3. **Facilitate increasing adoption of online, AI-related credential programs offered by companies such as Coursera**
    a. Lawmakers should ensure that existing law presents no or minimal barriers to micro-credential courses.
4. **Set reasonable hiring requirements for government-based AI roles**
    a. Just as companies need to adopt AI, so, too, do government agencies. Because AI is still a nascent field, imposing high barriers to entry in AI jobs, for example requiring many years of experience in AI, often can be counterproductive. While some experience is essential, being new to the field is not necessarily a disadvantage. Companies should be willing to hire candidates who can quickly learn new AI-related skills, not just those who already have substantial hands-on experience. Even at frontier AI labs, sometimes relative newcomers to the field make valuable, differentiated contributions. The same will almost surely be true within state government agencies.

These policies will foster a wide range of applied AI education opportunities without requiring large amounts of new funding.

Improving Measurement of AI Diffusion

Without good data, it is difficult to determine whether a policy is having its desired effect. And as AI begins to have broader societal effects, understanding how it is being used by businesses, even if only in broad statistical terms, will be useful to policymakers. Because technology diffusion is primarily the result of a series of decentralized decisions made by many different economic actors, it is difficult to accurately measure. However, there are steps state policymakers can take to develop a more refined sense of how AI is being used within their borders.

The first step is to better track the adoption of AI within state government agencies. State governments that have established Chief AI Officers, or individuals similarly responsible for AI adoption, within each state agency are a natural starting place. Governments that have not adopted such requirements for state agencies can place reporting requirements on the agencies' senior information technology (IT) officers. Regardless of the personnel, state agencies can be required to report the following metrics to any statewide AI or technology coordinating bodies, the legislature, and the public:

- IT budget items related to AI deployment
- The dollar amount of contracts with AI companies and cloud computing providers for AI services
- To the extent possible, state agency employee use of AI models (measured, for example, by the number of prompts made to language models deployed for agency use)
- State-funded AI research activities at public universities
- The availability, usage (as measured by declared majors or degrees and certificates conferred), and nature (AI research and engineering vs. applied AI) of AI-related curricula at public higher education institutions, including community colleges

While these metrics largely measure inputs (dollars spent and programs created) versus outputs (productivity enhancements from AI, dollars saved, and others), they still capture a reasonable baseline of AI activity and use within state governments.

Collecting data on AI diffusion within businesses is a larger challenge. Requirements to report fine-grained metrics such as those listed above may impose a burden on small businesses in particular. However, state governments already collect business data via tax documents, payroll censuses, and economic surveys. To the extent possible, policymakers should consider adding simple, yes-or-no questions to these forms. Examples could include: "Is your business using AI services for internal purposes?"; "Is your business using AI services for external (customer-facing) purposes?"; or, "Does your business employ any personnel who spend the majority of their time working on deploying AI?"

Furthermore, state governments can partner with state-based industry and trade associations to conduct more granular surveys of their membership, potentially gathering data on the amount of time businesses spend using AI services, the amount of money they have invested in AI, and their top use cases for AI.

Is Diffusion Necessary for AI?

Some have questioned whether a talent pipeline for AI diffusion is necessary. They say the technology is so intuitive and powerful that it practically diffuses itself. ChatGPT was the fastest-adopted consumer technology in history after its launch (though it was soon surpassed by Meta's social network Threads). Microsoft CTO Kevin Scott said in a May 2024 interview that while typical enterprise software "takes a fairly long time . . . to diffuse throughout the organization," advanced AI systems like the company's Copilot has "the highest level of engagement we've seen in any new Microsoft 365 or Office product, maybe in history."[5]

One of the novel and powerful characteristics of frontier AI systems is their ability to seamlessly integrate into existing workflows and other processes. Because of their general intelligence, they can often adapt to different circumstances. Most observers of the AI field predict that these capabilities will increase considerably in the coming years: Agentic AI models, for example, are expected to be able to take action on behalf of users. These models could write the code and use the tools necessary to integrate themselves into existing business operations. In other words, sufficiently powerful AI is part of the talent pool for AI diffusion because it is capable of integrating itself. However, this does not mean there is no need for an applied AI talent pool. There are a few reasons for this:

1. Not all AI falls into the category of generalist assistants and agents such as OpenAI's ChatGPT and Anthropic's Claude. For example, the most impactful uses of AI in industrial settings involve the application of comparatively simple computer vision models, which allow AI to see and even understand the world, to existing factory processes. AI-based computer vision models that can meaningfully assist in industrial quality assurance by visually scanning for product defects have been available—including for free as open-source models—for more than five years, accounting for the plurality of industrial computer vision use cases in 2022.[6] Though these models may be low-hanging fruit in many cases, their adoption will not be obvious to people who do not have knowledge of and skills in applied AI.
2. Even if AI models can handle some of the technical diffusion tasks that would otherwise be completed by a human, there will still be need for human oversight. AI models will be deployed in settings of kaleidoscopic diversity, with different constraints on energy use, budget, data privacy, and many other things. Ensuring this is done in compliance with applicable law, company policies, and company needs and desires will ultimately need to be handled by a person who understands both the high-level goals and the technical details.
3. When and if generalist assistants like the ones discussed above do enter the market, they are likely to have a profound impact on the nature of many businesses and on productive economic activity more broadly. Contending with this change will, no doubt, be a whole-of-company effort, but that effort will be greatly impoverished without staff who have expertise in AI and empowerment within their organizations.

As this data is collected, state governments should create publicly accessible AI diffusion dashboards. Responsibility for creating, hosting, and maintaining these dashboards will be clear for states that have a Chief Data Officer (or similar role) and open data portals. States that do not have these resources in place should consider creating them. If that is not feasible, the responsibility should rest with the most relevant agency, such as the state commerce departments. Deployed with the suggestions highlighted in the previous section, these approaches to measurement should yield a tight feedback loop between policy action and real-world results.

Conclusion

The enormous potential of AI will not be unlocked merely by training increasingly advanced and expensive models. For the United States to realize the benefits of AI, individuals, businesses, and governments will have to use these innovations.

There is no way to determine a priori how to best employ a general-purpose technology. So, the best way to unlock AI's potential is through diverse and various experiments. Some of those experiments will fail. Others, however, will succeed. The key to enabling these experiments is to equip as many workers as possible with the knowledge and skills they need to subject recent innovations to the trials of the real world.

Policymakers should be humble about what public policy can accomplish. The track record on skills-based workforce development policies is mixed. Larry Good and Ed Strong argue in a 2015 Aspen Institute report that this is primarily because US workforce development takes place over a series of fragmented and targeted programs that are not unified by any single policy framework. "We do not believe there is a real workforce development 'system' in the United States. Our national workforce investments are essentially a series of separate domestic policy programs, each designed to serve a specific need or target group . . . We believe attempting to solve workforce issues through programs is fundamentally flawed."[7] This would suggest that yet another narrow program is not the ideal approach for bolstering the applied AI workforce in the US.

Instead, a comprehensive approach, including the suggestions above, with programs involving K–12 schools, community colleges, universities, online microcredential programs, and firms is more likely to yield results. Additionally, addressing public opinion toward AI can help lawmakers pave the way for more engagement with AI. Polls have consistently shown that Chinese citizens are significantly more enthusiastic about AI and other emerging technologies than US citizens are.[8] This is, at least in part, likely related to the media climate in the United States, which has encouraged negative views toward technology in general and AI in particular. Political leaders can work to shape their rhetoric toward AI more conscientiously: Rather than propagating fear about AI and associated technologies, they should encourage widespread adoption and lead by example, integrating AI into government service delivery.

In the long run, some applications of AI may encounter barriers presented by existing state law. State policymakers should seek to find such barriers through reviews of current statutes conducted by state agencies or statewide AI coordinating bodies. Even these barriers, though, can only be fully discovered through real-world experimentation. As these barriers are discovered, state lawmakers should carefully evaluate the need for reform on a case-by-case basis.

AI development itself is a rigorously empirical discipline. It is a field of live science, and our theoretical understanding of how advanced models work often lags behind our empirical observations about what *does* work. Thus, the best AI developers often avoid dogma and seek to advance approaches that are demonstrated to work in practice. In this sense, the practice of AI research itself offers a template for policymakers, businesses, and individuals as they incorporate AI into their work and daily lives. This approach will pay dividends as we collectively navigate the coming industrial revolution.

About the Author

Dean W. Ball is a research fellow at George Mason University's Mercatus Center and author of the Substack *Hyperdimensional*. His work focuses on artificial intelligence, emerging technologies, and the future of governance.

Notes

1. Nestor Maslej et al., "The AI Index 2024 Annual Report," AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2024.
2. CSRankings: Computer Science Rankings (website), accessed August 26, 2024, csrankings.org/#/index?all&world.
3. David Nye, *Electrifying America: Social Meanings of a New Technology, 1880–1940* (Cambridge, MA: The MIT Press, 1992).
4. Jeffrey Ding, "The Diffusion Deficit in Scientific and Technological Power: Reassessing China's Rise," *Review of International Political Economy* 31, no. 1 (2023).
5. Ben Thompson, "Interviews with Microsoft CEO Satya Nadella and CTO Kevin Scott About the AI Platform Shift," *Stratechery*, May 23, 2024.
6. Grand View Research, "Computer Vision Market Size, Share and Trends Analysis Report," 2023.
7. Larry Good and Ed Strong, "Reimagining Workforce Policy in the United States," in *Transforming US Workforce Development Policies for the 21st Century*, ed. Van Horn et al. (Kalamazoo: W.E. Upjohn Institute for Employment Research, 2015).
8. Michelle Faverio and Alec Tyson, "What the Data Says About Americans' Views of Artificial Intelligence," Pew Research Center, November 21, 2023.