Stanford University Human-Centered Artificial Intelligence

STANFORD DIGITAL ECONOMY LAB

Student Policy Brief
Winner of the Stanford Emerging Tech Writing Competition
September 2023

Expanding AI Adoption Can Help Create Jobs

Aniket Baksy (PhD '23)
Avi Gupta (MS '23)

The Challenge: Slow Adoption of AI Technologies

Artificial intelligence (AI), including machine learning, natural language processing, and computer vision, has demonstrated immense potential to streamline business processes, improve efficiency, reduce costs, and enhance customer experiences. Despite the increased accessibility to cutting-edge AI provided by open-source technologies such as TensorFlow and PyTorch (and more recently ChatGPT and DALL·E 2), AI adoption remains low among U.S. companies, and skews toward large and well-established firms. As of 2018, the estimated share of U.S. companies that have adopted AI[1] range from about 3.2% to 5.8%. These firms account for up to 12.6% of US workers employed at firms using AI.

Key Takeaways

The concentration of AI adoption in large firms contributes to their market dominance and potentially increases economic inequality.

Expanding AI adoption by reducing adoption costs can unleash AI's productivity-enhancing effects for small businesses, creating medium-skill jobs, enabling entrepreneurship, and promoting diversity in the AI ecosystem.

The fact that AI adoption is concentrated in large firms contributes to high adoption costs. High adoption costs are largely responsible for the slow adoption of AI among small businesses. We focus on two main drivers of these costs:

- AI technologies are expensive to customize to specific business needs and require large intangible investments.

- AI technologies require data, which is expensive to collect, securely store, and analyze.

We propose making AI deployment easier for small businesses by creating a novel clearinghouse-like data licensing and computational resource infrastructure.

- This would democratize the use of AI by enabling easy access to computational resources and government-held datasets.

The skewed pattern of AI adoption toward large firms is likely driven by two advantages they typically enjoy:

- The financial ability to cover upfront costs associated with setting up and deploying AI systems, including access to computational resources and the infrastructure for data collection, storage, and analysis.
- Access to high-quality data, an essential input for training AI systems, which large firms often generate in-house as a byproduct of their regular operations.

As a result, the concentration of AI adoption in leading firms presents several policy challenges:

- Concentrated AI adoption confines AI's productivity-enhancing gains to a subset of already large and productive firms. Inter-firm inequality, in turn, has adverse financial implications for workers, as demonstrated by the well-documented earnings gap between workers at large firms and smaller firms.[2] Large firms also tend to struggle with racial and gender diversity in their workforces, suggesting that the concentration of AI in leading firms could reinforce these trends.[3]
- As AI becomes more advanced, firms will need to incorporate AI into their business practices to remain competitive. If barriers to adoption remain high, expansion will be concentrated among already large firms. Subsequent economic gains by large adopting firms will occur at the expense of non-adopting firms, whose closure or contraction will create employment losses.[4]
- The high adoption costs associated with AI technologies reduces the ability of new firms, especially those starting as small and medium enterprises (SMEs), to compete with established incumbent firms. Consequently, this may reduce business dynamism, where innovative products and services that would have made it to market are bottlenecked out. SMEs play an important role in job creation, innovation, and wealth accumulation for lower and middle-income households. Therefore, it is particularly important to level the playing field with respect to the use of potentially transformative new technologies, including AI.

Concentrated AI adoption confines AI's productivity-enhancing gains to a subset of already large and productive firms.

Proposal and a Theory of Change

In this brief, we offer a solution to reducing AI adoption costs for small businesses. We propose creating a centralized infrastructure that will democratize access to data, pre-trained models, and computational resources. This "clearinghouse" model will make it easier for SMEs, as well as the general public, to customize powerful AI models that fit their needs. Our approach to encouraging AI adoption will create jobs through the following channels:

Direct Effects on Businesses: Allowing businesses to easily access AI technologies will substantially reduce their costs, enabling them to expand and hire more workers.

Indirect Effects: When technologies are used by a wide variety of agents—be they workers operating technologies or firms incorporating them into existing business units— novel applications and extensions of these technologies are discovered, with associated productivity gains.[5] This is particularly true for "general purpose technologies" that are at early stages of adoption and development—as we argue AI technologies are. These additional benefits from technology adoption will lead to the creation of novel tasks and new jobs associated with them, contributing to employment growth.[6]

Effects on Ancillary Employment Generation: The rising adoption of AI technologies will boost demand for services related to data collection and management, including ensuring that data is compliant with legal standards for ownership, stored in a sufficiently secure manner, and compliant with privacy standards. These tasks are particularly suited for middle- and low-skilled workers as a pathway to more advanced jobs in data science and AI more broadly, making them ideal vehicles for retraining workers most likely to be displaced by the adoption of these technologies for careers in a rapidly expanding sector.

Policy Proposal: The Clearinghouse

Our policy proposal seeks to encourage AI adoption by SMEs by combining three pillars: access to data, access to pretrained models, and access to computational power.

Today, massive amounts of data are publicly available both on private repositories and via the federal government. We propose centralizing existing data resources to release large public datasets to private industry. To ensure interoperability, we propose the creation of standards for model input. This will reduce some of the high upfront costs associated with developing datasets for AI deployment.

To allow SMEs to leverage these datasets, we propose creating an easy-to-use web "clearinghouse" where companies can upload models (or choose from pretrained models) to be trained on data licensed from other companies and/or publicly available data. The computation would be performed using cloud infrastructure available through the platform.[7] After training, the user would receive optimized coefficients for their model.

This approach enables firms of all sizes to unlock the benefits of large-scale data and computing infrastructure while minimizing risks to privacy and intellectual property. The clearinghouse platform would interface directly with low/no-code AI tools to simplify model development and incorporate cutting-edge, publicly available, pretrained AI models.[8] The clearinghouse would expose sufficient metadata[9] to allow users to optimize their model hyperparameters, but would not permit access to the underlying data.[10] The clearinghouse model can expand access to proprietary data because a trusted intermediary allows companies to train models on others' data without actually accessing it. Privacy is enhanced because datasets are not shared with end users, and therefore cannot be cross-referenced to expose personally identifiable information.

Our policy builds on some of the proposals suggested in the Interim Report by the National AI Research Resource (NAIRR) Task Force. Specifically, recommendation 4-1 from the NAIRR proposes that "the NAIRR should coordinate a network of trusted data and compute providers and hosts for a robust, transparent, and responsible data ecosystem." Our proposed clearinghouse model builds on and synergizes with this vision. The NAIRR report also contains recommendations for privacy, safety, and security measures that could be

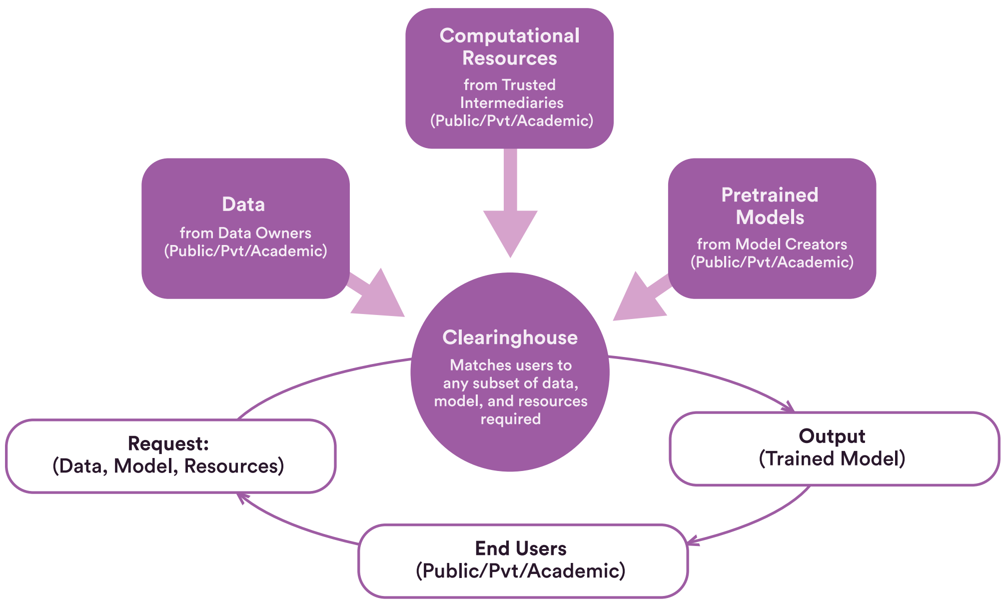

Figure 1: The Clearinghouse model for AI model deployment

End users, typically businesses or academics, submit a request consisting of any combination of a dataset, a model design, and computational resources. The clearinghouse matches the end user to data owners, model creator willing to license their pretrained models and trusted intermediaries willing to rent their computational resources. Model coefficients are returned to the end user on verification of all licensing arrangements. End users are free to upload their own data or their own model architecture to the platform.

applied to the clearinghouse. However, our proposal is oriented toward realizing the economic benefits of AI for U.S. workers and businesses, rather than solely supporting research efforts. Ultimately, our goal is to make AI as easy to use as Microsoft Excel.

The potential impact of our proposal can be illustrated by a hypothetical example. A small ice cream store seeking to expand its sales could use the clearinghouse to accurately forecast product demand. Via the clearinghouses, it could use a no-code AI platform to train an AI model on a dataset combining its daily ice cream sales with historical government weather data. By accessing pretrained AI models through the clearinghouse, the store would avoid the costs associated with designing its own model from scratch. More accurate sales forecasts could lead to increased efficiency, business growth, and expanded employment opportunities.

Platforms that host fake content could be required to not only establish a procedure for receiving complaints about deepfakes—as some have already done voluntarily—but to also provide a concise overview of the principles behind such standards. The Federal Trade Commission could then hold platforms accountable using its unfair trade practices authority. Platforms could also label content known or suspected to be machine generated, and the educators who train aspiring engineers could elevate policy and ethical literacy as important facets of technical education.

While none of these interventions will likely provide a quick fix to eroding trust in the information ecosystem, they offer a starting point for valuable discussions and provide a critical opportunity to affirm the values we hold most dear. Some considerations will undoubtedly lead to tradeoffs (both foreseen and unforeseen), but user research will be useful in finding best practices on implementation. How we distinguish reality from the synthetic in our evolving world of thinking machines presents one of the most pressing questions of our time.

Policymakers and the technical community are urged to embrace and address these challenges as readily as they're exploring the fascinating and exciting new uses of artificially intelligent systems.

Aniket Baksy and Avi Gupta are the first place winners of the 2022 Stanford Emerging Technology Writing Competition, funded by the Markle Foundation and co-hosted by the Stanford Digital Economy Lab, Stanford Institute for Human-Centered Artificial Intelligence, and the Stanford Institute for Economic Policy Research.

Aniket Baksy is a research scholar at the University of Sussex's Digital Futures at Work Research Centre and a Ph.D. graduate of Economics at Stanford University.

Avi Gupta is a J.D. candidate at Yale Law School and an M.S. graduate of Computer Science at Stanford University.

Stanford University's Institute on Human-Centered Artificial Intelligence (HAI), applies rigorous analysis and research to pressing policy questions on artificial intelligence. A pillar of HAI is to inform policymakers, industry leaders, and civil society by disseminating scholarship to a wide audience. HAI is a nonpartisan research institute, representing a range of voices. The views expressed in this policy brief reflect the views of the authors. For further information, please contact **HAI-Policy@stanford.edu.**

The Stanford Digital Economy Lab brings together an interdisciplinary group of passionate researchers to study how digital technologies are transforming work, organizations, and the economy. Our insights are helping companies, policymakers, students, and professionals rise to the challenges and opportunities created by AI and digitization. We believe in the power of human work and that augmenting human capabilities with machine capabilities will lead to the most profitable gains for everyone.

Endnotes

[1] Many private sector surveys on the diffusion of AI come from market research companies, consulting companies, or AI producers, and rely on self-reported estimates from a highly selective group of typically large firms (thus missing small and medium-sized businesses). Unsurprisingly, these surveys find high adoption rates. IBM's Global AI Adoption Index 2022 found that about 35% of firms worldwide have adopted AI technologies, based on a survey of 7,502 firms (https://www.ibm.com/downloads/cas/GVAGA3JP). Our estimates come from Acemoglu et al., "Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey," United States Census Bureau, April 2022, https://www.census.gov/library/working-papers/2022/adrm/CES-WP-22-12 .html. We note that survey responses may not account explicitly for AI embedded in tools that businesses operate, especially SaaS technologies like Salesforce: This implies that these surveys may underestimate overall adoption of AI technologies. § See https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/.

[2] Jae Song et al., "Firming Up Inequality," The Quarterly Journal of Economics 134, no. 1 (February 2019): 1–50, https://doi.org/10.1093/qje/qjy025.

[3] https://www.weforum.org/agenda/2022/03/diversity-inclusion-equity-business/

[4] On the one hand, AI technologies can automate tasks performed by some classes of workers, eliminating their jobs. On the other hand, the efficiency gains adopters see from AI adoption allow them to expand, creating jobs. The generalizability of these findings is further limited by current low levels of AI adoption. See Daron Acemoglu et al., "Artificial Intelligence and Jobs: Evidence from Online Vacancies," Journal of Labor Economics 40, no. 1 (April 2022): 293-340, https://www.journals.uchicago.edu/doi/full/10.1086/718327. See also Daron Acemoglu, "Harms of AI," National Bureau of Economic Research, September 2021, https://www.nber.org/papers/w29247.

[5] An example of such novel applications comes from the diffusion of the spreadsheet in the U.S. economy. In an interview with Quartz, Dan Bricklin, the inventor of VisiCalc, the first commercial spreadsheet program, describes what people did with them: "Early on, consultants told us they used it to help lay out slot machines on a casino floor. And doctors did calculations too. *I didn't know about those applications, or that those people would have even thought to use it that way*...[Spreadsheet users] were people who were able to figure out how to use a tool for a specific problem, even if it wasn't advertised to do it. That made them innovators." View interview at https://qz.com/578661/dan-bricklin-invented-the-spreadsheet-but-dont-hold-that-against-him

[6] David Autor et al., "New Frontiers: The Origins and Content of New Work, 1940–2018," National Bureau of Economic Research, August 2022, https://www.nber.org/papers/w30389.

[7] This could include private cloud computing providers, such as AWS, as well as government resources, such as supercomputers in use at national laboratories.

[8] Pretrained AI models (e.g., DALL·E 2, GPT-3) can be fine-tuned using transfer learning to exhibit strong performance on new tasks using relatively lower amounts of data.

[9] Metadata could include feature descriptions, feature summary statistics (including feature coverage), and the quantity of data input.

[10] Large technology companies with massive datasets internally impose similar controls on their engineers. The interface provides enough information to work with the data, but not enough to copy or de-anonymize it.

Stanford University Human-Centered Artificial Intelligence

STANFORD DIGITAL ECONOMY LAB

Stanford HAI: 353 Jane Stanford Way, Stanford, CA 94305-1234

T 650.725.4537 **F** 650.123.4567 **E** HAI-Policy@stanford.edu **hai.stanford.edu**