EXPERT INSIGHTS
PERSPECTIVE ON A
TIMELY POLICY ISSUE

Leading with Artificial Intelligence

Insights for U.S. Civilian and Military Leaders on Strengthening the AI Workforce

October 2024

Editors: Rachel Slama, Nelson Lim, and Douglas Yeung

Elie Alhajjar, Rushil Bakhshi, Julia Bandini, Dwayne M. Butler, Avery Calkins, Angela K. Clague, Arianne Collopy, Brandon Crosby, Brandon De Bruhl, Tuyen Dinh, Fernando Esteves, Susan M. Gates, Charles A. Goldman, Timothy R. Gulden, Wenjing Huang, Kelly Hyde, Rita T. Karam, Tracy C. Krueger, Jonah Kushner, Mary Lee, Nelson Lim, Maria C. Lytell, Laurie T. Martin, Nikolay Maslov, Michael G. Mattock, Skye A. Miner, Alvin Moon, Ojashwi Pathak, Neeti Pokhriyal, Carter C. Price, Sean Robson, Srikant Kumar Sahoo, Morgan Sandler, Anton Shenk, Rachel Slama, Éder M. Sousa, Tobias Sytsma, Ivy Todd, John Vahedi, Jessie Wang, Madison Williams, Jody Chin Sing Wong, Douglas Yeung

For more information on this publication, visit www.rand.org/t/PEA3414-1.

Cover designer: Santos W. Vega III

Cover photo: kieferpix/Adobe Stock

About This Paper

Artificial intelligence (AI) is poised to significantly affect the American workforce—both civilian and military personnel—through job displacement, augmentation, and the need for widespread upskilling. This publication is intended to inform the policymakers and leaders who are tasked with preparing civilian and military workers to create, use, and deploy AI in their jobs. The essays in this publication provide overviews of technical and organizational issues, challenges, and actionable insights to help organizations effectively integrate AI and equip personnel with AI-related skills.

This work was conducted by RAND Education and Labor in conjunction with the RAND Homeland Security Research Division and RAND Project AIR FORCE.

RAND Education and Labor

RAND Education and Labor is a division of RAND that conducts research on early childhood through postsecondary education programs, workforce development, and programs and policies affecting workers, entrepreneurship, and financial literacy and decisionmaking.

More information about RAND can be found at www.rand.org. Questions about this report should be directed to rslama@rand.org, and questions about RAND Education and Labor should be directed to educationandlabor@rand.org.

RAND Homeland Security Research Division

The RAND Homeland Security Research Division (HSRD) operates the Homeland Security Operational Analysis Center (HSOAC), a federally funded research and development center (FFRDC) sponsored by the U.S. Department of Homeland Security (DHS). HSRD also conducts research and analysis for other federal, state, local, tribal, territorial, and public- and private-sector organizations that make up the homeland security enterprise, within and outside the HSOAC contract. In addition, HSRD conducts research and analysis on homeland security matters for U.S. allies and private foundations.

For more information on HSRD, see www.rand.org/hsrd.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of RAND, is the Department of the Air Force's (DAF's) federally funded research and development center for studies and analyses, supporting both the United States Air Force and the United States Space Force. PAF provides the DAF with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future air, space, and cyber forces. Research is conducted in four programs: Strategy and Doctrine; Force Modernization and Employment; Resource Management; and Workforce, Development, and Health.

Additional information about PAF is available on our website:
www.rand.org/paf/

Funding

Funding for this research was made possible by gifts from RAND supporters and income from operations, as well as by the independent research and development provisions of RAND's contracts for the operation of its U.S. Department of Defense federally funded research and development centers.

Acknowledgments

We are grateful to RAND internal sponsors Darleen Opfer, Andrew Hoehn, James Chow, and Douglas Ligor. We are also grateful for valuable comments from our peers, John Pane, Daniel Egel, and Liz Nguyen, and for communications guidance from Melissa Bauman. We thank the RAND quality assurance management team, including Stephanie Rennane and Benjamin Master, and the publications and editing team, including Monette Velasco and Jaron Feldman. We are also appreciative for administrative and quality assurance support from Lea Ann Gerkin and Heather Gillispie.

Contents

About This Paper iii
Summary vii

Chapter 1
Introduction 1

SECTION A: Integrating AI into the Workforce 7

Chapter 2
Taxonomy of AI Adoption Risks and Effect on Adoption in the Workforce 8

Chapter 3
The AI Ecosystem 11

SECTION B: Use Cases—Applying AI in the Workforce 15

Chapter 4
Rapid-Development Chatbots for Workforce Training and Support 16

Chapter 5
To Chat or Not to Chat: Using AI to Communicate with Patients and Relieve the Burden on the Health Care Workforce 20

Chapter 6
Applying AI Tools to Complete Common Human Resource Management Tasks 24

Chapter 7
Navigating the AI Landscape: Choosing the Right Tool for the Job 28

SECTION C: Educating and Training the Workforce to Use AI 33

Chapter 8
Upskilling and Retraining the Federal Workforce for AI Adoption 34

Chapter 9
The Promise of AI to Transform Teaching Will Fail If School Systems Do Not Transform Too 38

Chapter 10
Helping Postsecondary Education and Training Institutions Overcome Barriers to Preparing the New AI Workforce 41

Chapter 11
Leveraging AI in the Military Leader Development Framework 45

SECTION D: Building a More Resilient and Diverse AI Workforce 49

Chapter 12
Retaining Workers with AI Skills in the Federal Workforce 50

Chapter 13
Is the Domestic AI Talent Pool Sufficient to Meet Public Sector Demand? 54

Chapter 14
Building a Diverse Pool of AI Talent for DHS Recruitment 57

Chapter 15
Identifying Resilient Skills in an AI-Enhanced Economy 62

Chapter 16
Engaging the Federal Workforce in AI Implementation 66

Chapter 17
Understanding and Addressing Resistance to the Adoption of New Technologies in DoD 70

Chapter 18
Principles on AI Implementation for Federal Leaders 75

Summary

Issue

Artificial intelligence (AI) is poised to significantly affect the American workforce—both civilian and military personnel—through job displacement, augmentation, and the need for widespread upskilling. President Biden's October 2023 executive order on AI emphasizes the commitment to upskilling the federal workforce in understanding, adopting, deploying, and using AI (Executive Order 14110, 2023). Many federal agencies and U.S. Department of Defense (DoD) entities have published AI guidance documents, several of which are referenced throughout this publication. Congress is also exploring the implications of advancements in AI in both the general and federal U.S. workforces (U.S. House of Representatives, 2024; U.S. Senate, 2023).

This publication is intended to inform the policymakers and leaders who are tasked with preparing civilian and military workers to create, use, and deploy AI in their jobs. The essays in this publication provide overviews of technical and organizational issues, challenges, and actionable insights to help agencies effectively integrate AI and equip personnel with AI-related skills.

Approach

To create this publication, we harnessed expertise across RAND to explore critical policy questions related to AI adoption. RAND researchers with diverse backgrounds contributed concise, evidence-based essays on key issues related to preparing the workforce to adopt AI technologies. These essays were selected by the editors, using their collective experience, to cover a breadth of considerations in adopting AI technology into the workplace. The essays were then refined through a process of feedback and peer review to ensure the highest quality and relevance to policymakers and leaders.

Key Findings and Policy Implications

Section A: Integrating AI into the Workforce

This section provides a high-level primer on key considerations for integrating AI into the workforce, including the types of risks to manage, the types of job functions and associated competencies in the broader AI ecosystem, and how to address worker resistance to AI adoption. Key findings include the following:

- Leaders should use a comprehensive taxonomy of AI adoption risks to assess and manage these risks when integrating AI into processes and operations. The taxonomy identifies technical, ethical, legal, economic, social, and existential threats, such as model inaccuracies, job displacement, and fracturing societal institutions (Chapter 2).
- The AI ecosystem requires a diverse workforce with tailored skills for different roles. From developers to end users, each role needs a basic understanding of AI's functionality and potential biases. Policymakers should ensure that worker training addresses these varied needs and that workers are made aware when they are interacting with AI (Chapter 3).

Section B: Use Cases—Applying AI in the Workforce

This section provides examples of three common use cases of AI-powered tools: processing reams of complex technical documents, improving user experience through chat bots, and serving as a job aid in human resource management tasks, such as processing high-volume job candidate materials. Key findings include the following:

- AI-powered chatbots can streamline workforce training by providing accurate, real-time support; reducing reliance on experienced personnel; and minimizing errors. To maximize a tool's effectiveness, organizations should prioritize data security, transparency, human oversight, and flexibility in adapting to policy changes (Chapter 4).
- AI chatbots can streamline administrative tasks in health care by reducing provider workload and enhancing patient access, but their integration must be carefully managed to avoid exacerbating biases and compromising patient trust. Prioritizing chatbot use for routine tasks, ensuring provider oversight, and implementing robust regulatory frameworks are essential to maintaining ethical and effective patient communication (Chapter 5).
- AI tools can enhance human resources (HR) processes by automating such tasks as screening resumes and creating job descriptions, freeing HR professionals to do more strategic work. Organizations should integrate these tools carefully, with regular reviews, to ensure accuracy, compliance, and alignment with organizational goals (Chapter 6).
- Organizations should choose AI tools based on their needs, considering such factors as task complexity, available resources, and the value added to operations. Decisionmakers must conduct a thorough cost-benefit analysis to ensure AI solutions align with organizational goals and are worth the associated risks (Chapter 7).

Section C: Educating and Training the Workforce to Use AI

This section highlights considerations for investing in AI tools and infrastructure, including upskilling the current civilian and military workforce to use AI. There are ways in which existing training systems must be adapted to prepare the workforce to use AI. Key findings include the following:

- The federal government should prioritize upskilling its existing workforce for AI adoption by implementing tailored training programs and investing in necessary tools and infrastructure. To sustain progress, fostering a culture of continuous learning is essential for updating employees on AI advancements and ensuring effective integration across operations (Chapter 8).
- For AI to effectively enhance education, school systems must realign their policies to prioritize the mastery of foundational skills over strict adherence to grade-level standards. Without this shift, AI tools will fail to reach their full potential; teachers will struggle to balance personalized learning and curriculum coverage (Chapter 9).
- Postsecondary institutions—including training and credentialing organizations—should form strategic partnerships with the AI industry, redesign curricula, and invest in infrastructure to align education with AI workforce needs. To ensure success, these institutions must also attract diverse students, retain skilled faculty, and leverage government support to overcome barriers and remain competitive in the evolving AI landscape (Chapter 10).
- The U.S. military should integrate AI into its professional military education by aligning AI education with existing training goals and using AI to enhance how training is delivered. A phased approach, including the continuous assessment of AI's costs and risks, will ensure responsible and effective AI integration in military leadership development (Chapter 11).

Section D: Building a More Resilient and Diverse AI Workforce

This section provides an overview of key considerations for recruiting and retaining diverse domestic AI talent, as well as building resilience among the current workforce through a focus on the skills in which humans still maintain a relative advantage over AI. Key findings include the following:

- To retain AI-skilled workers, the federal government should offer competitive compensation, including retention bonuses tied to service commitments, and create rewarding work environments. Additionally, providing ongoing AI training with associated service obligations and potential pay increases will help align federal opportunities with private sector demands and ensure skill relevancy (Chapter 12).
- The U.S. government struggles to recruit AI talent because of competition with the private sector, a limited domestic talent pool, and restrictions on hiring foreign-born workers. To overcome these impediments, the public sector should emphasize non-salary benefits, partner with academic institu-

tions, and consider expanding eligibility for foreign-born AI professionals (Chapter 13).

- The U.S. Department of Homeland Security (DHS) should prioritize recruiting AI talent from minority-serving institutions and underserved communities by leveraging partnerships, career pathways, and targeted investments. To ensure workforce diversity and retention, DHS must enhance its diversity metrics and address barriers that hinder the recruitment and advancement of women and minorities in AI roles (Chapter 14).
- Policymakers and educators should prioritize training in soft skills—such as interpersonal communication and complex problem-solving—because these are less likely to be automated. Investing in education programs that develop these skills will help create a more resilient workforce that is prepared for an AI-enhanced economy (Chapter 15).
- Federal agencies should adopt a worker-centered approach to AI. They should involve employees in developing and implementing AI tools to build trust and ensure that AI tools meet practical needs. To overcome such barriers as siloization, agencies should promote collaboration between AI developers and staff and should continue to track AI use cases to refine strategies and practices over time (Chapter 16).
- DoD must address workforce resistance to AI adoption, driven by fears of job displacement, lack of trust, and skill gaps. DoD can successfully integrate AI by focusing on transparent communication, comprehensive training, and ethical governance while keeping employees engaged (Chapter 17).

References

Executive Order 14110, “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence,” Executive Office of the President, October 30, 2023.

U.S. House of Representatives, *Toward an AI-Ready Workforce: Hearing Before the Subcommittee on Cybersecurity, Information Technology, and Government Innovation of the Committee on Oversight and Accountability*, U.S. Government Printing Office, January 17, 2024.

U.S. Senate, *AI and the Future of Work: Moving Forward Together: Hearing Before the Subcommittee on Employment and Workplace Safety of the Committee on Health, Education, Labor, and Pensions*, U.S. Government Printing Office, October 31, 2023.

Abbreviations

AI	artificial intelligence
COVID-19	coronavirus disease 2019
DHS	U.S. Department of Homeland Security
DL	deep learning
DoD	U.S. Department of Defense
FEMA	Federal Emergency Management Agency
HR	human resources
HRM	human resource management
K–12	kindergarten through grade 12
KSAOs	knowledge, skills, abilities, and other characteristics
LLM	large language model
ML	machine learning
MSI	minority-serving institution
OPM	U.S. Office of Personnel Management
PME	professional military education
RAG	retrieval augmented generation
STEM	science, technology, engineering, and math
STS	sociotechnical systems

CHAPTER 1

Introduction

Rachel Slama, Nelson Lim, Douglas Yeung, eds.

Artificial intelligence (AI)—in particular, recent technological advances in large language models (LLMs)—has already transformed many business sectors (U.S. General Services Administration, undated-a), changing the work landscape across industries, reshaping job roles, automating tasks, and augmenting human capabilities. This technological revolution presents a mix of opportunities and challenges for the federal civilian and military workforce.

Although there is no standard definition of AI, it commonly includes artificial (i.e., not human) systems or techniques that can perform a variety of tasks without significant human oversight. These systems demonstrate human-like qualities, such as learning from experience, by analyzing incoming data to improve how they perform a task. AI tools and systems can mimic—with varying degrees of success—other human attributes, such as reasoning, communicating, and decisionmaking (U.S. General Services Administration, undated-b). AI systems and techniques power such tools as some chatbots and copilots.

To date, advancements in AI are task-specific or narrow. Scientists have not yet achieved a *general intelligence* that can learn an array of tasks, equivalent to human capabilities, although some suggest that artificial general intelligence may be available soon. Given the rapid pace with which AI has been integrated into society and the likely technological advances, many scholars assert that AI is here to stay. Workplaces, workers, and employers must adapt.

Leveraging the Promise of AI in the Workforce and Mitigating Its Harms

AI's ability to augment human capabilities and enhance efficiency has the potential to significantly benefit the workforce, enabling employees to focus on higher-value, creative, and strategic work. AI can automate repetitive tasks, analyze large datasets to identify patterns and insights, personalize user experiences, and support data-driven decisionmaking. These advantages can lead to increased efficiency, better outcomes, and enhanced customer satisfaction. AI-driven tools have the potential to empower workers to make informed choices and improve their problem-solving skills.

At the same time, AI presents challenges, such as job displacement and a widening of the gap between the demand for advanced technical skills and the supply of workers who have them. The impact of these changes can vary across demographic groups, exacerbating income inequality (Manning, 2024). Workers may harbor concerns about AI that affect worker morale and well-being. For example, Americans are wary of specific workplace uses of AI, such as for hiring decisions (Rainie, 2023). Workers need reskilling and upskilling to remain employable in an AI-driven economy.

One step toward leveraging the promise of AI in the workplace is to develop better evidence for AI's benefits and risks and for implementation steps that mitigate concerns. The White House, by executive order, has tasked the U.S. government with identifying and addressing many implementation barriers to effectively and responsibly integrating AI into current federal workstreams (Executive Order 14110, 2023). Ethical concerns, such as the potential for biased decisionmaking and *AI hallucination* (generating incorrect or nonsensical information), underscore the need for robust governance, continuous monitoring, and transparent AI systems to mitigate negative impacts.

Federal Efforts on AI in the Workforce

Efforts on integrating AI into the federal workforce are well underway. For example, at the time of writing, almost one year has passed since the October 2023 release of the President's executive order on the "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence" (Executive Order 14110, 2023). The executive order requires the federal government to report on how agencies are using AI through an inventory of AI use cases in a searchable spreadsheet or through short vignettes by agency (AI.gov, 2023). In November 2023, the U.S. Department of Defense (DoD) released an AI Adoption Strategy with the aim of responsibly and rapidly leveraging decades of progress in data, analytics, and AI for business operations and warfighting (Clark, 2023).

Many of the deliverables mandated in the executive order have already been generated, while others are longer-term initiatives.[1] The initiatives touch on challenges across the employment life cycle. In the education realm, there are initiatives supporting AI education and workforce development (such as those led by the National Science Foundation) and pilot programs to bolster existing training programs for AI talent (such as those led by the U.S. Department of Energy and National Science Foundation). In the areas of recruiting and retention, there are initiatives to attract and keep more foreign-born AI talent (such as those led by the U.S. Department of State and the U.S. Department of Homeland Security [DHS]) and to identify and mitigate critical labor shortages in the science, technology, engineering, and math (STEM) workforce (such as those led by the U.S. Department of Labor). In addition, all agencies are directed to designate a chief AI officer to coordinate their agency's use of AI, promote AI innovation, and manage related risks. Since the executive order established a centralized front door to federal AI jobs, participating tech talent programs report a surge in applications and hiring, according to the government's AI and Tech Talent Task Force. However, the task force identified gaps in existing infrastructure, tools, and resources (AI and Tech Talent Task Force, 2024).

DoD's AI Adoption Strategy also emphasizes the need for "an educated, empowered workforce" that can use cutting-edge tools, conduct advanced research, and integrate with allies and partners (DoD, 2023, p. 4). At the core of this strategy is the need to expand digital talent management in DoD through increased hiring, training, upskilling, and retention of workers in data, analytic, and AI-related jobs. The DoD strategy is folded into a larger U.S. government strategy and must be coordinated within the broader AI ecosystems of academic and industry partners.

This collection of essays complements existing RAND research and stakeholder frameworks on federal action on AI with a deep dive into upskilling and retraining the existing federal workforce, leveraging the domestic AI talent pipeline, and evolving existing education and workforce training systems.

Purpose of This Publication

This essay collection provides objective and independent insights from diverse RAND experts to help policymakers, educators, and industry leaders navigate the challenges and opportunities of AI in the U.S. workforce. These essays were selected by the editors, using their collective experience, to cover a breadth of considerations in adopting technology into the workplace.

The goal of this publication is to foster an informed approach to integrating AI into the workforce, enabling leaders to create a future in which AI enhances human capabilities, drives innovation, and benefits society. This collection serves as a starting point for the conversations and collaborations needed to build a resilient and adaptable workforce in the age of AI.

Contributors and Their Perspectives

The contributors have expertise across every disciplinary area at RAND (RAND Corporation, undated), emblematic of the types of cross-disciplinary analysis needed to address the far-reaching implications of AI in the civilian and military workforce. The authors are trained in engineering and the applied sciences (e.g., information science, applied mathematics, computer science, engineering, biomedical engineering), economics, sociology, and statistics; behavioral and policy sciences (psychology, education, public policy, management); and defense and political sciences. Their backgrounds range from one year of work experience to nearly four decades at RAND. Author biographies are included at the end of this publication.

Structure of This Publication

This collection of essays is organized into the following four sections:

- **Section A ("Integrating AI into the Workforce")** examines the benefits of AI, such as increased efficiency, and the challenges it presents, including job displacement and ethical concerns.
- **Section B ("Use Cases—Applying AI in the Workforce")** provides real-world examples that demonstrate how AI enhances productivity and creates new roles across sectors.
- **Section C ("Educating and Training the Workforce to Use AI")** discusses the need to update education and training systems to equip workers with the skills necessary for an AI-driven job market.
- **Section D ("Building a More Resilient and Diverse AI Workforce")** explores strategies for developing a workforce that is resilient to technological changes and diverse in both skills and perspectives.

Together, these essays provide an expansive overview of the impact of AI on the workforce, the practical applications of AI across sectors, the necessary adaptations to education and training infrastructure, and strategies for building a strong and adaptable workforce in the face of technological advancements.

Notes

[1] A summary of provisions, action type, and department(s) tasked with executing the requirements is available from Freedman Consulting, undated.

References

AI and Tech Talent Task Force, *Increasing AI Capacity Across the Federal Government: AI Talent Surge Progress and Recommendations*, Executive Office of the President, April 2024.

AI.gov, "The Government Is Using AI to Better Serve the Public," webpage, last updated September 1, 2023. As of August 19, 2024: https://ai.gov/ai-use-cases

Clark, Joseph, "DOD Releases AI Adoption Strategy," U.S. Department of Defense, November 2, 2023.

DoD—*See* U.S. Department of Defense.

Executive Order 14110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," Executive Office of the President, October 30, 2023.

Freedman Consulting, "AI Executive Order Provision Tracker," spreadsheet, undated. As of August 19, 2024: https://airtable.com/appgXuAye4DqzHzo2/shrb56uYZzjQwX69H/tbltimdyx8sCktxK0/viwoDI3Udvsgt4HLr?blocks=hide

Manning, Sam, "AI's Impact on Equality in the US," Brookings Institution, July 3, 2024.

Rainie, Lee, Monica Anderson, Colleen McClain, Emily A. Vogels, and Risa Gelles-Watnick, "AI in Hiring and Evaluating Workers: What Americans Think," Pew Research Center, April 20, 2023.

RAND Corporation, "Research and Analysis Jobs at RAND," webpage, undated. As of August 19, 2024: https://www.rand.org/jobs/job-categories/research-and-analysis-jobs.html

U.S. Department of Defense, *Data, Analytics, and Artificial Intelligence Adoption Strategy: Accelerating Decision Advantage*, 2023.

U.S. General Services Administration, "Introduction to the AI Guide for Government," webpage, undated-a. As of August 19, 2024: https://coe.gsa.gov/coe/ai-guide-for-government/introduction/index.html

U.S. General Services Administration, "Key AI Terminology," webpage, undated-b. As of August 19, 2024: https://coe.gsa.gov/coe/ai-guide-for-government/what-is-ai-key-terminology/index.html

Section A

INTEGRATING AI INTO THE WORKFORCE

CHAPTER 2

Taxonomy of AI Adoption Risks and Effect on Adoption in the Workforce

Jessie Wang, Jody Chin Sing Wong

Artificial intelligence (AI) is the elephant in the breakroom at many workplaces. As employees incorporate AI into their daily tasks, it can be a daily reminder of the risks that AI poses to their livelihoods. However, the threat to employment is just one of many risks that AI poses to the workforce. We propose a risk taxonomy to help stakeholders understand the multiple threats that AI poses to the workforce.

A risk taxonomy can serve as a guide to identify, assess, and address various risk factors more effectively, so that workforce leaders can take a comprehensive approach to risk management while integrating AI into the workforce. In Table 2.1, we provide an example AI adoption risk taxonomy derived from the literature. We highlight a broad spectrum of risks that goes beyond simple technical failures, both in scope and severity. Because AI adoption risks can vary across sectors and worker groups, leaders should start with the sample taxonomy shown in Table 2.1, which showcases numerous risk categories that should be examined. The risk categories documented in the table are not exhaustive; leaders should refine the categories appropriately to fit their respective sectors.

TABLE 2.1
Taxonomy of AI Adoption Risks with Examples

Category	Types	Examples
Technical risks	• Model inaccuracy • Data privacy concerns • System vulnerabilities	• Biases or inconsistencies in training data • Bugs or errors in algorithms • Incorrect threat assessments, misidentification of individuals, non-optimal decisions • Increased susceptibility to cyberattack and exploitation
Ethical risks	• Lack of fairness, transparency, or accountability • Deliberate deception	• Deep fakes and other deceptive content • Widened inequity in access to AI • Manipulation of public opinion and erosion of trust in information sources
Legal and regulatory risks	• Failed compliance • Unauthorized use of intellectual property	• Lack of clarity in whether the use of AI or related data is authorized • Discriminatory hiring practices or unfair monitoring • Increased socioeconomic inequality • Unfair distribution of financial gains from AI use
Economic risks	• Job displacement • Skills mismatch • Market manipulation	• Workers replaced with automation • Workers cannot meet evolving demand for skills • Market monopolization • Increased socioeconomic inequality • Unfair distribution of financial gains from AI use
Social and psychological risks	• Lack of public trust in AI technologies • Negative emotional responses to working with AI	• Fear of being replaced by AI • Widened social stratification • Increased job dissatisfaction • Decreased human interaction • Undesirable workplace dynamics • Widened social unease from labor market restructuring
Existential risks	• Emergence of superintelligence • Loss of human control • Misalignment of AI goals and human values	• Creation of harmful biological agents, either intentionally or accidentally • Superintelligent AI pursues goals that are misaligned with human values or detrimental to human survival and well-being • Global instability • Doubt in human scientific and technological progress • Ethical considerations

SOURCE: Authors' analysis of sources listed in the "References" section to this chapter.

References

Abbott, Ryan, "The Reasonable Computer: Disrupting the Paradigm of Tort Liability," *George Washington Law Review*, Vol. 86, No. 1, January 2018.

Acemoglu, Daron, and Pascual Restrepo, "Artificial Intelligence, Automation, and Work," National Bureau of Economic Research, Working Paper 24196, January 2018.

Amodei, Dario, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané, "Concrete Problems in AI Safety," arXiv, arXiv:1606.06565, July 25, 2016.

Binns, Reuben, "Fairness in Machine Learning: Lessons from Political Philosophy," *Proceedings of Machine Learning Research*, Vol. 81, 2018.

Bostrom, Nick, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press, 2014.

Brynjolfsson, Erik, and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W. W. Norton & Company, 2014.

Calo, Ryan, "Artificial Intelligence Policy: A Primer and Roadmap," *UC Davis Law Review*, Vol. 51, No. 2, December 2017.

Doshi-Velez, Finale, and Been Kim, "Towards a Rigorous Science of Interpretable Machine Learning," arXiv, arXiv:1702.08608, March 2, 2017.

Gasser, Urs, and Virgilio A. F. Almeida, "A Layered Model for AI Governance," *IEEE Internet Computing*, Vol. 21, No. 6, November–December 2017.

Goodfellow, Ian J., Jonathon Shlens, and Christian Szegedy, "Explaining and Harnessing Adversarial Examples," arXiv, preprint arXiv:1412.6572, March 20, 2015.

Hancock, Peter A., Deborah R. Billings, Kristin E. Schaefer, Jessie Y. C. Chen, Ewart J. de Visser, and Raja Parasuraman, "A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction," *Human Factors*, Vol. 53, No. 5, October 2011.

Manyika, James, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott, and Martin Dewhurst, "A Future That Works: Automation, Employment, and Productivity," McKinsey Global Institute, January 12, 2017.

Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi, "The Ethics of Algorithms: Mapping the Debate," *Big Data & Society*, Vol. 3, No. 2, July–December 2016.

Russell, Stuart, *Human Compatible: Artificial Intelligence and the Problem of Control*, Viking, 2019.

Slovic, Paul, "Perception of Risk," *Science*, Vol. 236, No. 4799, April 17, 1987.

Slovic, Paul, Melissa L. Finucane, Ellen Peters, and Donald G. MacGregor, "Risk as Analysis and Risk as Feelings: Some Thoughts About Affect, Reason, Risk, and Rationality," *Risk Analysis*, Vol. 24, No. 2, April 2004.

Voigt, Paul, and Axel von dem Bussche, *The EU General Data Protection Regulation (GDPR): A Practical Guide*, 1st ed., Springer International Publishing, 2017.

West, Sarah Myers, Meredith Whittaker, and Kate Crawford, *Discriminating Systems: Gender, Race, and Power in AI*, AI Now Institute, April 2019.

Yudkowsky, Eliezer, "Artificial Intelligence as a Positive and Negative Factor in Global Risk," in Nick Bostrom and Milan M. Cirkovic, eds., *Global Catastrophic Risks*, Oxford University Press, 2008.

CHAPTER 3

The AI Ecosystem

Carter C. Price, Madison Williams

The AI ecosystem requires a workforce with a diverse set of skills, and the required understanding of AI systems varies by an individual's position relative to the AI system. To upskill the workforce to develop and use AI, policymakers, educators, and others need to develop and deploy training programs. However, one size does not fit all; the requisite skills and competencies will differ for workers based on their positions in the AI ecosystem. That said, there are some common things that everyone should understand about AI systems. In this essay, we describe the AI ecosystem, explain how people relate to AI systems, and conclude with some considerations for policymakers.

Supply and Demand in the AI Ecosystem

Figure 3.1 is a simplified portrayal of the AI ecosystem that captures both the supply and demand sides. The supply side represents those entities that feed the algorithms, computing power, and data used to produce AI systems. The supply

FIGURE 3.1
The AI Ecosystem

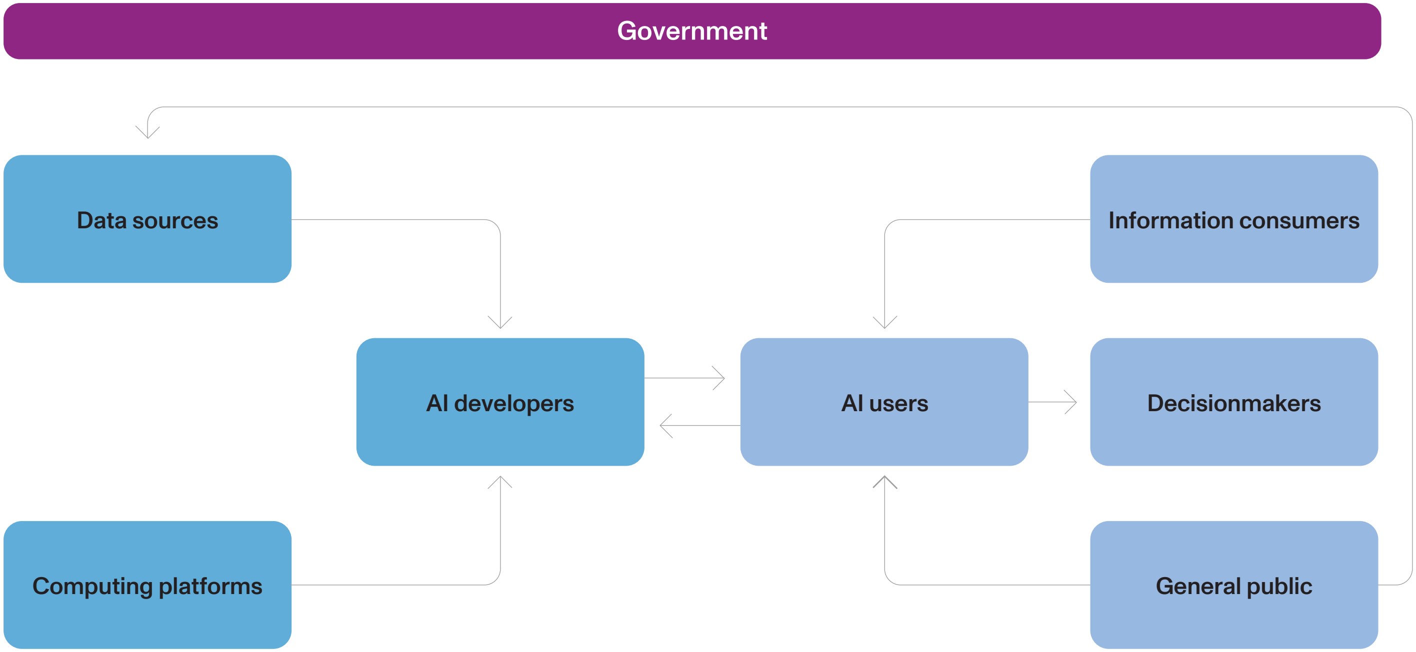

side includes individuals with such careers as data entry professionals, computer scientists, and data center workers. On the demand side are users of the AI system—typically people involved with knowledge work who use AI tools to analyze data, produce written text, or generate images, but also decisionmakers and the general public.[1]

Although AI developers are the highly visible tip of the iceberg, they are supported by a variety of other firms and organizations, including a wide array of data providers and computation suppliers and a workforce with specialized skills. Workers in the AI ecosystem are, thus, not restricted to the most visible segment of AI algorithm developers.

Likewise, workers who use AI as part of their jobs will not just include those who perform knowledge work tasks. These are only a narrow slice of workers involved in the AI ecosystem. The top half of Table 3.1 presents the key parts of the supply side of the AI ecosystem, an explanation of what they do, and the skills or knowledge they need to do their jobs. Because each component of the supply side fulfills a unique role, the skills required vary substantially.

The bottom half of the table shows the same information for the demand side. Individuals' roles vis-à-vis AI systems can change throughout the day; sometimes they may be users, and at other times they could be making

TABLE 3.1
Key Parts of the AI Ecosystem

Component	Description	Workforce Skills or Knowledge
Supply side		
AI developers	Companies and workers that make AI systems by developing and applying algorithms to datasets that train AI models	A strong background in computer science or a related field
Data owners	People and organizations that collect data from a variety of sources	Data literacy to ensure data standards are met and privacy is protected
Computing providers	Companies and workers that develop, build, and maintain computing platforms, ranging from cloud services to edge devices, such as phones, laptops, and desktops	Varies by role but could include information technology to maintain systems, electrical engineering for chip design, or vocational training relevant to building infrastructure
Demand side		
AI users	People who directly interface with AI systems	A minimum level of AI literacy that includes an understanding of how these systems work and how they can go wrong
Information consumers	People who receive information from AI systems (may also be AI users)	An ability to critically assess information and determine whether it is AI-generated
Decisionmakers	People who use AI outputs to make informed decisions (may also be AI users and information consumers)	A knowledge of potential biases or other limitations of information
General public	People who are subject to AI decisions and supply AI systems with data	An understanding of personal interactions with AI and the potential biases and other flaws in AI systems

decisions or consuming AI information. But data are essentially always being collected about individuals, and individuals are nearly always subject to AI algorithms. One key thing to acknowledge is that not all AI interactions are consensual or known to the person interacting with the AI system. For example, AI is widely used in hiring decisions (Myers, 2023) and has been found to exhibit bias (e.g., see Dastin, 2018). Similarly, not all data are collected consensually or in a way that is transparent to the public about whom data are collected (Thorbecke, 2024).

Lastly, government policymakers could target any of the entities described above directly, or the relationships and interactions between organizations could be targeted (e.g., as of this writing, there are export restrictions on some segments of computing platforms, primarily the highest-performing processors). Government can also operate in most of the roles described: supplying data, developing models, housing computing infrastructure, using models, and other roles.

Conclusion

When making decisions about AI and the workforce, policymakers should consider the full breadth of workers in the AI ecosystem—not just computer scientists. Workers with vocational training are necessary as well. The skills that these workers need vary substantially, but there are things that everyone who interacts with AI systems need to understand, such as how the systems operate (at least at a high level), the types of problems that AI is suitable to address, and how the systems fail.

Workers also need to be aware that many of their interactions with AI systems are hidden. They may be supplying data to AI models unknowingly, and they are subject to decisions made or informed by these systems, which have biases and other flaws.

By having a baseline level of understanding, AI users, information consumers, and decisionmakers can be safe, savvy consumers of AI products and informed makers of AI policy.

Notes

[1] *Knowledge work* can be thought of as tasks that require knowledge rather than physical labor to accomplish the desired outcome.

References

Dastin, Jeffery, "Insight—Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women," Reuters, October 10, 2018.

Myers, Sydney, "2023 Applicant Tracking System (ATS) Usage Report: Key Shifts and Strategies for Job Seekers," *Jobscan* blog, October 2, 2023.

Thorbecke, Catherine, "Google to Delete Billions of Browser Records to Settle 'Incognito' Lawsuit," CNN, April 1, 2024.

Section B

USE CASES—APPLYING AI IN THE WORKFORCE

CHAPTER 4

Rapid-Development Chatbots for Workforce Training and Support

Timothy R. Gulden, Dwayne M. Butler, Wenjing Huang, Nikolay Maslov, Anton Shenk

Workforce training and support can be time-consuming, labor-intensive, and costly for large organizations, but AI-powered chatbots could help provide some answers. Document-based chatbots can provide instant, knowledgeable assistance to employees navigating complex procedures and regulations, and they can be quick and easy to put in place. If designed properly, such chatbots have the potential to revolutionize workforce training and support, but poor design can waste resources while providing inaccurate guidance.

Drawing from a prototype tool under development at RAND, we provide principles for success and an architecture for implementing document-based chatbots into workforce training.

The Prototype Tool

The RAND Federal Emergency Management Agency (FEMA) Public Assistance program for the coronavirus disease 2019 (COVID-19) pandemic illustrates the potential benefits of AI-powered chatbots. Modern AI systems can identify the concepts involved with answering a question and respond accurately and comprehensively to questions, even when those questions do not use the same wording as the documents that contain the answers. The AI tool can also present answers in a way that may be more tailored and comprehensible to the workers who are asking the questions.

In the FEMA project, dozens of RAND analysts are tasked with navigating complex guidance documents and training materials to help FEMA avoid granting duplicative benefits. This effort involves a combination of accounting expertise and legal reasoning to ensure that decisions involving hundreds of millions of dollars are consistent, fair, and based on evolving policies that are documented in hundreds of pages of written material.

Currently, experienced analysts are responsible for training newcomers and providing expert support, but their availability is limited because of demanding caseloads. AI-powered chatbots can offer continuous support by accurately applying complex FEMA policies, cross-referencing claims against past decisions for consistency, and flagging potential issues for human review.

The tool that we are developing shows promise for enhancing and standardizing human judgment, streamlining processes, reducing errors, and ensuring consistent outcomes, thereby better serving health care providers and the public in the aftermath of the pandemic. At the time of this writing, the tool consists of (1) several hundred pages of documentation and training materials that cover the vast majority of questions that analysts might have and (2) a retrieval augmented generation (RAG) system that can find, summarize, and provide links back to the authoritative documentation in which the question is addressed. Although this system is still under development, prototype experiments and pilot testing have brought several principles into sharp focus that are likely to be essential to the success of projects of this sort.

Principles for Success

Security First

Protecting sensitive internal government data is essential. Many LLM tools use uploaded data for future model training. This means that proprietary information may become part of the next public-facing model's "common sense" understanding of the world. Workforce training chatbot projects should establish comprehensive data-sharing rules, consider in-house models for sensitive data, and regularly update data-protection policies. Our project uses a secure version of OpenAI/Azure GPT-4 as its basis.

Transparency Matters

The chatbot should always show its work, linking back to source documents to build trust and enable verification. These links should be convenient to use, and workers should be trained to use them routinely.

Maintain Human Responsibility

AI is a powerful tool but, in this context, not a replacement for human judgment. The goal is to augment and train the human workforce to make better, faster, and more-consistent decisions, not to replace it. Our system is not designed to produce final work products but rather to support the analysts who are producing those products.

Train on Policy Documents

Documents should be comprehensive, internally consistent, and authoritative (although they need not be fully organized or indexed as they would need to be to support direct human use). Modern LLM-based tools excel at finding relevant information in human-readable documents and surpass human performance in navigating complex and poorly organized material—as long as the correct answers are in the documents somewhere. Our current prototype does not reason beyond what is contained in the training corpus, although other tools—such as OpenAI's ChatGPT and Google's NotebookLM—can combine specific training documents with general knowledge. This means that the training documents for these tools can be less comprehensive, but this has trade-offs with respect to both reliability and data security.

Flexibility Is Key

The chatbot should be as adaptable as the policies it interprets and easily updatable as guidelines change. Current LLMs can often handle retraining on hundreds of pages of material in seconds, allowing the models to be retrained each time guiding documents are updated. This allows humans to maintain the human-generated guiding documents and the chatbot tool to remain up to date. An update strategy should be explicitly addressed, for instance, by adding a list of frequently asked questions to the corpus to address specific shortcomings in chatbot responses. Our prototype requires a collection of PDFs (including PDF versions of webpages) to be kept up to date. A next-generation design would enable the use of a mix of PDFs and webpages, and the model would update automatically when these webpages change.

Implementation Architecture

Several general architectures are available for implementing rapid-development AI chatbots. As of this writing, major approaches include RAG, fine-tuning, and large context windows. Each of these has strengths, and the best fit will depend on the nature of the specific task.

RAG involves slicing each document into smaller (often one-page) parts, indexing them with mathematical embedding vectors, and storing them in a vector database. This technique is extremely powerful for finding and linking back to specific facts and concepts in the text, but it currently has real limitations in terms of its ability to draw inferences from different bits of retrieved information or to synthesize information that is indirectly relevant to the query. We chose this approach because it can be done with off-the-shelf tools and does not require specialized programming for the LLM to ingest the training documents and provide useful answers.

Fine-tuning involves creating large numbers of query-response pairs based on material in the documents and training the model to respond accurately to a wide array of queries. This can produce excellent responses but makes it difficult to link back to the specific places in the text that support each part of an answer. It may also result in meaningfully slower and more expensive training and updating compared with other approaches.

Large context windows are increasingly available and allow LLMs to actively consider hundreds of pages at once when generating a response. As of this writing, models with extremely large context windows can have issues with attention and positional encoding that can lead them, for example, to pay more attention to the beginning and end of the training corpus than to the middle of it. This can cause these models to miss important information. They also tend to have less reasoning power than models that have more-focused context windows. The ability of this technology to both reason coherently about the whole contents of a corpus and to link back to key information is evolving rapidly and may eventually become a preferred approach.

Conclusion

Integrating AI-powered chatbots into workforce training and support has the potential to enhance efficiency, knowledge-sharing, and decisionmaking within complex regulatory environments. These chatbots can be stood up quickly and inexpensively and can be trained using the same documents that are used to train human analysts. By providing timely, accurate assistance, these chatbots can reduce the burden on experienced personnel and minimize errors, ultimately empowering the federal workforce to serve the public with greater consistency and accuracy.

CHAPTER 5

To Chat or Not to Chat: Using AI to Communicate with Patients and Relieve the Burden on the Health Care Workforce

Skye A. Miner, Rushil Bakhshi, Julia Bandini, Laurie T. Martin

As the number of patient messages sent through patient portals dramatically increases, health care systems and providers are looking to AI chatbots as a way to increase efficiencies in patient-provider communication while maintaining provider responsiveness and reducing burnout (Holmgren et al., 2022; Tai-Seale et al., 2019). We weigh the promise of AI chatbots in improving workforce efficiencies and for reducing burnout in the health care workforce alongside ethical, social, and quality concerns.

AI Chatbots May Assist with Administrative Tasks to Decrease Burdensome Touchpoints

Given the limited amount of time that health care providers can spend with patients and the increasing number of messages that providers receive, chatbots could assume responsibility for replying to some of the messages to allow the health care workforce to spend more time providing direct patient care. AI chatbots are already being used to assist patients and staff with appointment scheduling, billing, and insurance inquiries. In the future, AI chatbots may also be able to effectively distinguish among types of messages, removing messages from a physicians' queue that are better suited to billing, scheduling, or other departments. This could help streamline health care providers' workflows so that providers would no longer be responsible for reviewing all messages.

AI Chatbots May Be Used to Better Streamline Patient Care Interactions

Providers may routinely use chatbots to directly answer patient questions, formulate example responses, and create and refine educational materials. Chatbots have been shown to be effective for answering patients' dermatology and diabetic foot ulcer management questions (Reynolds et al., 2024; Shiraishi et al., 2024), creating patient educational material for appendicitis and orthopedic support (Ghanem et al., 2024; Morya et al., 2024), writing patient discharge summaries (Patel and Lam, 2023), and supporting patients' mental health by providing education and self-support methods (Vaidyam, Linggonegoro, and Torous, 2021). In the near future, chatbots may help providers communicate with their patients in language that is simpler, is more empathetic, and uses less medical jargon (Ayers et al., 2023).

Chatbots, if trained appropriately, may also be able to collate large amounts of patient data to provide more-nuanced responses based on a patient's medical history and other characteristics, which may alleviate provider burden. AI-driven chatbots could enable predictive analysis and proactive health care to further augment remote patient monitoring and telehealth capabilities. Early detection and intervention may help patients manage their conditions at home, decrease hospital admissions, and reduce time-consuming interventions—all of which ease the workloads of health care providers. Moreover, initial triaging by AI-driven chatbots may help assess whether self-care is sufficient, a telehealth consultation is needed, or an in-person visit is required. Chatbots can also make telehealth or in-person appointments more efficient by gathering preliminary information from the patient beforehand. This approach may help improve access for patients who live in rural areas, have mobility issues, or may not have the time or resources for in-person medical appointments.

AI Chatbots May Create Problems or Exacerbate Existing Problems in Health Care Communication

Although AI chatbots may reduce provider workload and improve communication, the use of AI in patient communication could exacerbate existing racial and other biases within the health care system, reduce patient trust in their provider, create privacy and transparency concerns, and increase the time that providers spend reviewing electronic documentation (a source of burnout) rather than focusing on patient care. For example, because of training biases, chatbots may favor particular linguistic patterns or cultural norms, so their responses could be less effective and relevant for communicating with marginalized patients. As a result, the American Medical Association has suggested that governments—as well as health care institutions, practices, and professional societies—should take a risk-based approach to AI, meaning that the level of scrutiny, validation, and oversight should be proportionate to the potential harms and consequences of the technology (American Medical Association, 2023). Thus, the integration of AI chatbots to address administrative health care tasks may be able to immediately decrease provider workload, while more-complex communication tasks may require more regulatory and provider oversight and resources in the short term.

Although this framework may help to evaluate the risks of chatbots, it does not address the concern that patients will not readily accept the widespread use of chatbots for communication (further creating a disconnect between patients and providers). Patients may be concerned with sharing sensitive health information, and providers may have reservations about the accuracy and reliability of AI chatbots because the algorithms may be biased and can overgeneralize or hallucinate. These biases in AI algorithms could be caused by such issues as a lack of diversity in training data to ensure adequate representation of racial or ethnic groups, biased feature selection, and model overfitting.

In the near term, so that providers can spend more quality time with patients, health system integration of chatbots in health care should prioritize reducing administrative burden, including repetitive tasks, and providing educational materials. As chatbots become more widely used for more-complex communication, health systems will need to provide training and resources for providers so that they can recognize and document biases, inefficiencies, and other areas in which chatbots may struggle with health care situations (e.g., situations that require more personalized care or appropriate responses to the nuanced emotional or psychological needs of patients).

Considerations for the Future

AI chatbots may provide greater access to key health information that patients need and may lessen the burden on health care providers to engage heavily in communications. However, chatbots cannot provide health care solutions on their own and will require provider oversight to ensure that information is correct and appropriate. Moreover, mitigating data privacy and transparency risks with a risk-based approach while also developing more-robust regulatory and trust frameworks will be crucial to ensure the ethical, equitable, and reliable application of AI in health care communication.

References

American Medical Association, “Principles for Augmented Intelligence, Deployment, and Use,” November 14, 2023.

Ayers, John W., Adam Poliak, Mark Dredze, Eric C. Leas, Zechariah Zhu, Jessica B. Kelley, Dennis J. Faix, Aaron M. Goodman, Christopher A. Longhurst, Michael Hogarth, and Davey M. Smith, “Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum,” *JAMA Internal Medicine*, Vol. 183, No. 6, June 2023.

Ghanem, Yazid K., Armaun D. Rouhi, Ammr Al-Houssan, Zena Saleh, Matthew C. Moccia, Hansa Joshi, Kristoffel R. Dumon, Young Hong, Francis Spitz, Amit R. Joshi, and Michael Kwiatt, “Dr. Google to Dr. ChatGPT: Assessing the Content and Quality of Artificial Intelligence-Generated Medical Information on Appendicitis,” *Surgical Endoscopy*, Vol. 38, No. 5, May 2024.

Holmgren, A. Jay, N. Lance Downing, Mitchell Tang, Christopher Sharp, Christopher Longhurst, and Robert S. Huckman, “Assessing the Impact of the COVID-19 Pandemic on Clinician Ambulatory Electronic Health Record Use,” *Journal of the American Medical Informatics Association*, Vol. 29, No. 3, March 2022.

Morya, Vivek K., Ho-Won Lee, Hamzah Shahid, Anuja Gajanan Magar, Ju-Hyung Lee, Jae-Hyung Kim, Lang Jun, and Kyu-Cheol Noh, “Application of ChatGPT for Orthopedic Surgeries and Patient Care,” *Clinics in Orthopedic Surgery*, Vol. 16, No. 3, June 2024.

Patel, Sajan B., and Kyle Lam, “ChatGPT: the Future of Discharge Summaries?” *Lancet Digital Health*, Vol. 5, No. 3, March 2023.

Reynolds, Kelly, Daniel Nadelman, Joseph Durgin, Stephen Ansah-Addo, Daniel Cole, Rachel Fayne, Jane Harrell, Madison Ratycz, Mason Runge, Amanda Shepard-Hayes, Daniel Wenzel, and Trilokraj Tejasvi, “Comparing the Quality of ChatGPT- and Physician-Generated Responses to Patients’ Dermatology Questions in the Electronic Medical Record,” *Clinical and Experimental Dermatology*, Vol. 49, No. 7, July 2024.

Shiraishi, Makoto, Haesu Lee, Koji Kanayama, Yuta Moriwaki, and Mutsumi Okazaki, “Appropriateness of Artificial Intelligence Chatbots in Diabetic Foot Ulcer Management,” *International Journal of Lower Extremity Wounds*, February 28, 2024.

Tai-Seale, Ming, Ellis C. Dillon, Yan Yang, Robert Nordgren, Ruth L. Steinberg, Teresa Nauenberg, Tim C. Lee, Amy Meehan, Jinnan Li, Albert Solomon Chan, and Dominick L. Frosch, “Physicians’ Well-Being Linked to In-Basket Messages Generated by Algorithms in Electronic Health Records,” *Health Affairs*, Vol. 38, No. 7, July 2019.

Vaidyam, Aditya Nrusimha, Danny Linggonegoro, and John Torous, “Changes to the Psychiatric Chatbot Landscape: A Systematic Review of Conversational Agents in Serious Mental Illness: Changements du paysage psychiatrique des chatbots: une revue systématique des agents conversationnels dans la maladie mentale sérieuse,” *Canadian Journal of Psychiatry*, Vol. 66, No. 4, April 2021.

CHAPTER 6

Applying AI Tools to Complete Common Human Resource Management Tasks

Sean Robson, Maria C. Lytell, Tracy C. Krueger

Human resource management (HRM) professionals perform a variety of essential yet repetitive and time-consuming tasks to support employee productivity, performance, and satisfaction. With advancements in technology, organizations everywhere are exploring how AI could expedite many traditional human resources (HR) tasks and free up that workforce to perform more meaningful tasks. In this essay, we examine the potential uses for two types of AI in HRM and some practical tips for using such tools.

HR Tasks That Lend Themselves to AI

AI is most useful for HR tasks that require

- processing and analyzing high volumes of data (e.g., reviewing resumes from thousands of applicants)

- standardizing inconsistencies across different cases, such as producing position descriptions for different roles in the organization
- repeating steps, such as parsing knowledge, skills, abilities, and other characteristics (KSAOs) from position descriptions
- analyzing information in real time, such as summarizing employee survey results for organizational leadership.

Two main types of AI technologies could be useful for HR tasks and have minimal start-up costs: embeddings and generative AI (e.g., ChatGPT).

Overview of Embeddings and Generative AI for HRM

Embeddings and generative AI use models that are pretrained, often with publicly available data (e.g., content on the internet), which reduces start-up costs. In almost all cases, these technologies lack the precision and consistency to replace HR expertise. Instead, HR professionals can use them to augment current practices and workflows.

Embeddings

Embeddings are numerical representations of text, such as words, sentences, and paragraphs. Each piece of text is represented in a multidimensional space. The closer that two pieces of text are in this space, the more similar they are in context. For example, terms that are similar but that use different words (e.g., *machine learning* and *statistical modeling*) would be closer in proximity than two unrelated terms (e.g., *machine learning* and *autobiography*). Embeddings are therefore useful to find similarities between different sets of text. Embeddings can be useful for the following types of HR tasks:

- matching a job candidate's skills with the required skills for a position
- finding employees within an organization who have specific skill combinations
- creating an internal skills database by grouping similar skills into categories
- updating an existing internal skills database by evaluating the similarity of new or emerging skills
- directing employees to relevant policies based on their questions.

Figure 6.1 illustrates the steps that can be used to create and use embeddings for HR tasks.

Generative AI

Generative AI, such as ChatGPT, is useful for extracting and organizing data, as well as creating new content based on guidelines. ChatGPT can be used for the following types of HR tasks:

- summarizing required qualifications and key responsibilities from a job description
- suggesting labels for KSAOs based on job descriptions
- organizing KSAOs into a hierarchical structure
- writing job announcements.

If an HR professional wants to use ChatGPT to create a job announcement, the steps in Figure 6.2 can be used as a general guide.

FIGURE 6.1

Steps for Creating and Using Embeddings for HR Tasks

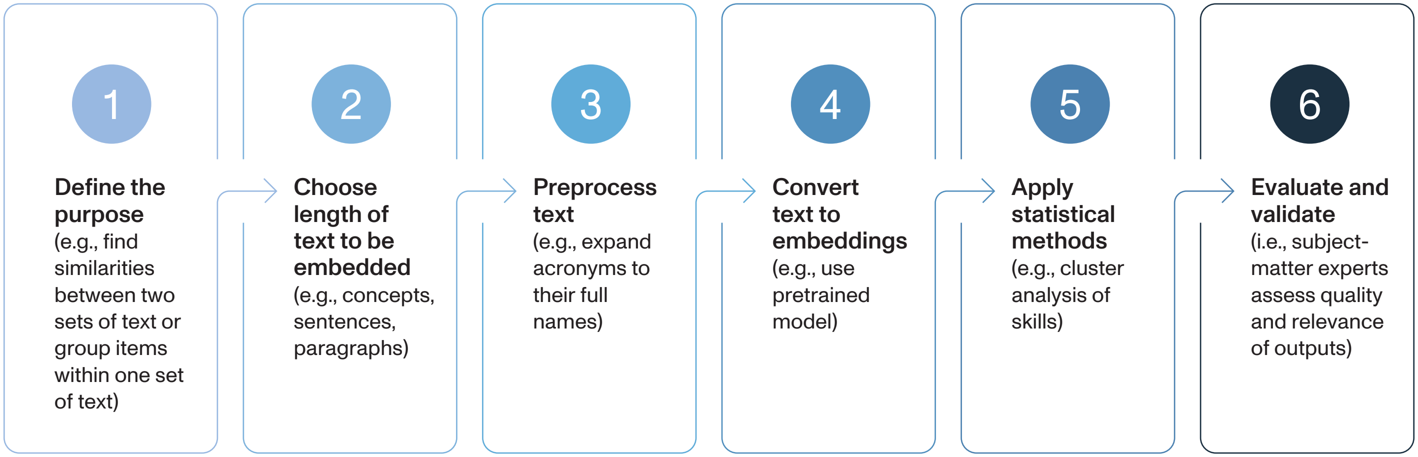

FIGURE 6.2

Steps for Using ChatGPT to Create a Job Announcement

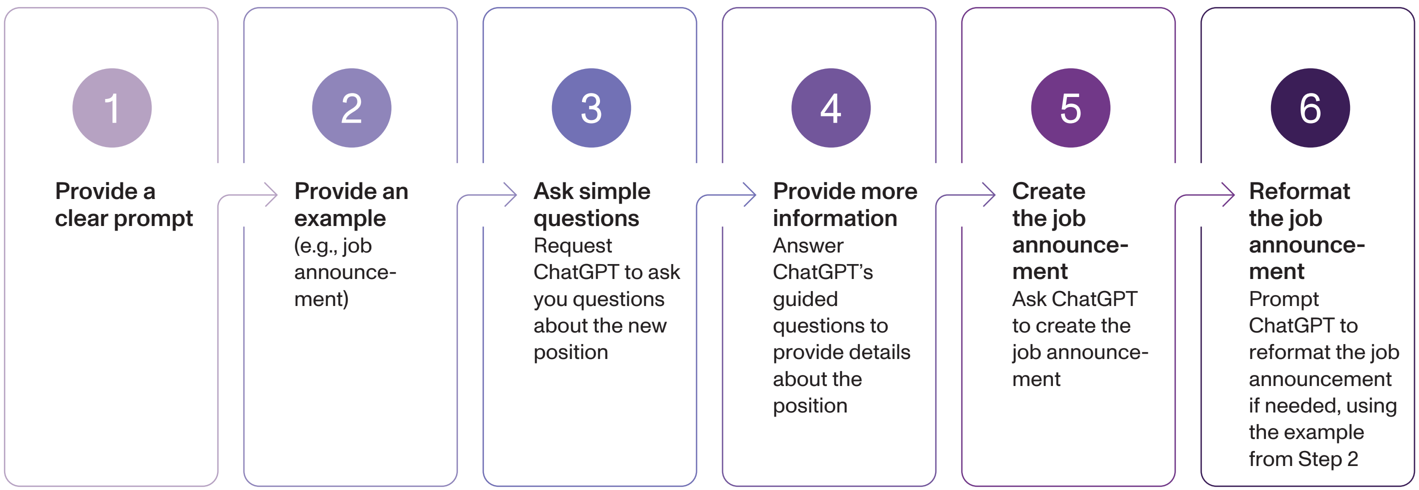

Conclusion

AI technologies offer opportunities to expedite a variety of HR tasks. To get the best results, some experimentation will be needed to determine which HR tasks AI technologies can be applied to and what prompts work best. Most importantly, implementation of any AI into HRM should be coupled with regular, systematic reviews of outputs to ensure that the results are in line with expectations and meet organizational and legal requirements.

CHAPTER 7

Navigating the AI Landscape: Choosing the Right Tool for the Job

Ojashwi Pathak, Morgan Sandler, John Vahedi, Kelly Hyde

Many organizations think that AI could be the answer to their problems,[1] but are they asking the right questions? We provide an overview of factors to consider and questions to ask to aid organizations in selecting an appropriate AI for their needs.

The three types of AI that we discuss in this essay are machine learning (ML), deep learning (DL), and generative AI. ML methods focus on systems that learn to perform predictive and classification tasks by analyzing input (i.e., training) data (IBM, undated-b). Some ML methods involve assumptions about the structure of the relationship being modeled (similar to such traditional methods as linear regression), while other methods flexibly learn these relationships from the training data. DL is a subset of ML that focuses primarily on the use of neural networks designed to recognize patterns in large input data (Holdsworth and Scapicchio, 2024). Because of their flexibility, DL methods are particularly suitable in cases in which the underlying relationships between data and outputs are highly complex. Finally, generative AI is a type of DL that is designed to generate new content (Holdsworth and Scapicchio, 2024). Generative AI methods respond to user prompts by imitating the human-generated content in training data, making these methods capable of creating text, images, video, audio, and other outputs that appear as though they could have been created by a person.

Each of these AI methods has limitations and risks. For instance, in some cases, ML methods that appear accurate in training may be much less accurate when applied to new data. DL methods used for image classification sometimes may mistake some objects for others even when the difference is obvious to humans, such as mistaking a Chihuahua for a blueberry muffin (Cloudera, 2022). And generative AI can hallucinate, making seemingly authoritative statements that are factually incorrect (IBM, undated-a). The risks of these errors can be amplified by how and where AI is used. For example, using generative AI to help write a routine work email is much less risky than generating patient diagnoses in a hospital's intensive care unit. In summary, there are many factors for decisionmakers to consider in identifying the right AI method for their organizations and use cases.

How to Choose the Right Type of AI

Table 7.1 provides some illustrative considerations to help determine which types of AI might be most suitable for a particular use case. In the table, we divide AI into three broad categories: generative AI methods, such as LLMs and diffusion models; types of DL that are not considered generative AI, such as convolutional neural networks and recurrent neural networks; and types of ML that are not considered DL, such as Random Forest, Naive Bayes, or K-Nearest Neighbor Classifiers. The considerations in the table are not exhaustive but represent some fundamental factors that are easy to ascertain and can help distinguish which AI methods are most suitable early in the development of a use case. The table reflects recommended tools for particular use cases based on the likelihood of superior performance or efficiency in that case with currently available methods. However, as the field evolves, new methods may emerge that prove to be effective, highlighting the importance of staying flexible and adapting to these advancements as they come. In cases in which multiple types of methods are recommended, it is common practice to test several different options and compare their performance to determine the best tool.

Apart from choosing the right AI method, decisionmakers should consider whether implementing AI tools aligns with business goals. Many organizations have adopted some AI tools, but the field of AI is under rapid and intense development, and generative AI is new and largely unexplored. The cost of using AI can be large or small, depending on whether an organization chooses an already existing model and applies it to their business or develops a model from scratch. For example, according to a study published by Meta, LLaMa 3 (an LLM) was trained on millions of dollars' worth of high-power computing equipment that also required significant (and likely costly) labor from technical experts (Llama Team, 2024). Therefore, along with the considerations in Table 7.1, we suggest three guiding questions to identify use cases for AI technology. If the answer to these three questions is yes, it may be worth considering AI use cases if the cost permits:

- Does it add value in your industry or domain?
- Does it fit with the operations of the organization?
- Is it worth the risk in your industry or domain?

TABLE 7.1
Considerations for Employing Different Types of AI

		Recommendation		
Consideration	Case	Generative AI	Other DL (not generative AI)	Other ML (not generative AI or DL)
Task[a]	Statistical analysis (e.g., predictive modeling)		X	X
	Pattern recognition in text or numerical data		X	X
	Image recognition		X	
	Speech/audio recognition		X	
	Tracking objects in an image		X	
	Creating new content (image, text, audio, video) based on a text, audio, video, or image prompt	X		
Resource availability[b]	High computing power (computer chips that can handle multiple large tasks at one time, distributed cloud computing)	X	X	X
	Low computing power (main processor in a personal computer or laptop)			X
Explainability[c]	Prioritize simple explainability of tool outputs over raw performance			X
	Prioritize raw performance over simple explainability of tool outputs	X	X	X

[a] Duda, Hart, and Stork, 2006; Goodfellow, Bengio, and Courville, 2016.
[b] Lawton, 2023.
[c] Onose, 2023.

Value added in an industry or domain represents the potential benefits of adopting AI. For example, AI adoption can lead to task automation, increased productivity, and improved efficiency in business processes. Overall, implementing AI can add value by reducing costs over time and streamlining operations. To determine whether the chosen AI solution aligns with the organization's operations, decisionmakers should carefully consider whether their organization is prepared to handle the intended and unintended consequences of AI adoption, taking into account budgetary constraints and available resources (e.g., acquiring the potentially expensive high computing power indicated in Table 7.1 as a requirement for generative AI and DL methods), as well as talent gaps (e.g., ensuring the organization has the expertise needed to interpret the outputs of the tool, both in the simpler case of non-DL or non-ML methods and in the more complex, performance-oriented cases of generative AI and other DL as reflected in Table 7.1). The adopted AI should support the organization's mission and overall goals without hindering growth. Finally, the decisionmakers should analyze the potential risks associated with AI adoption. Such analysis includes considering the risks of AI methods re-creating human biases in their results (Holdsworth, 2023), leaking the sensitive data (such as personally identifiable information) on which they are trained, or being exploited by adversaries to harm the organization (Tummalapenta, 2023). Therefore, it is imperative for decisionmakers to ask themselves the three critical questions shown on page 29 once they have selected a potential AI solution for their organization.

Conclusion

There are two things that a decisionmaker should consider while assessing the appropriate AI for their organization. First, they need to appraise their business need and determine the right type of AI, depending on the availability of resources, complexity of the tasks, and interpretability of the output provided by AI. Second, the decisionmaker should carefully perform a cost-benefit analysis for adopting AI for their business need, also ensuring that the implementation of AI aligns with the organization's mission and goals. Because the cost of adopting different types of AI tools can vary significantly, it is important to simultaneously consider both of these factors to determine the optimal balance of cost and performance. By choosing the type of AI that suits one or more of an organization's many needs (e.g., HR, operations, production, sales), decisionmakers can effectively leverage AI to improve productivity, increase efficiency, and promote creativity and innovation within their organizations.

Notes

[1] AI is a broad field that enables computers to imitate human intelligence tasks, such as reasoning, learning, and adapting (Stryker and Kavlakoglu, 2024).

References

Cloudera, “Of Muffins and Machine Learning Models,” *Cloudera Blog*, February 16, 2022.

Duda, Richard O., Peter E. Hart, and David G. Stork, *Pattern Classification*, Wiley India, 2006.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, 2016.

Holdsworth, James, “What Is AI Bias?” webpage, IBM, last updated December 22, 2023. As of August 30, 2024: https://www.ibm.com/topics/ai-bias

Holdsworth, Jim, and Mark Scapicchio, “What Is Deep Learning?” webpage, IBM, last updated June 17, 2024. As of August 30, 2024: https://www.ibm.com/topics/deep-learning

IBM, “What Are AI Hallucinations?” webpage, undated-a. As of August 30, 2024: https://www.ibm.com/topics/ai-hallucinations

IBM, “What Is Machine Learning (ML)?” webpage, undated-b. As of August 30, 2024: https://www.ibm.com/topics/machine-learning

Lawton, George, “Compare GPUs vs. CPUs for AI Workloads,” TechTarget, December 28, 2023.

Llama Team, *The Llama 3 Herd of Models*, arXiv, arXiv:2407.21783, August 15, 2024.

Onose, Ejiro, “Explainability and Auditability in ML: Definitions, Techniques, and Tools,” *Neptune Blog*, August 17, 2023.

Stryker, Cole, and Eda Kavlakoglu, “What Is Artificial Intelligence (AI)?” webpage, IBM, last updated August 16, 2024. As of August 30, 2024: https://www.ibm.com/topics/artificial-intelligence

Tummalapenta, Srinivas, “How to Establish Secure AI+ Business Models,” IBM, September 18, 2023.

Section C

EDUCATING AND TRAINING THE WORKFORCE TO USE AI

CHAPTER 8

Upskilling and Retraining the Federal Workforce for AI Adoption

Elie Alhajjar, Alvin Moon, Mary Lee

Artificial intelligence (AI) is rapidly transforming the landscape of the modern workforce, offering unprecedented opportunities for efficiency, innovation, and strategic decisionmaking (Frost and Sullivan, 2024). However, fully realizing these benefits requires a strategic approach to workforce development—one that focuses on developing existing workers to smooth the transition to an AI-enabled federal workforce.

Investing in upskilling the existing federal workforce rather than relying mainly on hiring new talent offers numerous advantages (discussed in Chapter 12). For one, current employees have a deep understanding of their agencies' operational nuances, cultures, and missions. This institutional knowledge is invaluable and cannot be easily replicated by new hires. Another advantage is that retraining existing employees fosters a culture of continuous learning and adaptability, which is crucial in the rapidly evolving field of AI.

Focusing on the existing workforce also represents a more cost-effective and time-efficient strategy (Walsh, Yuen, and Raman, 2024). Recruiting and onboarding new talent, especially in such highly competitive fields as AI (Chapter 13), can be a lengthy and expensive process. In contrast, upskilling and retraining current employees can be done incrementally and more flexibly, allowing agencies to adapt training programs to specific needs and existing workflows (Tamayo et al., 2023). We provide several recommendations to help military and civilian leaders develop an AI-capable workforce with their existing employees.

Implement a Broad Collection of Training Programs

The cornerstone of preparing the federal workforce for AI adoption is education and training (Wilson and Daugherty, 2019). Employees at all levels must gain a foundational understanding of AI technologies, including their capabilities, limitations, practical applications, and ethical considerations. This foundational knowledge can be imparted through introductory courses and workshops that cover the basics of AI, ML, and data analytics.

For employees in technical roles, advanced training is essential. This training should delve into the specific AI methodologies, programming languages, and tools used in AI development and deployment. Courses on data science, neural networks, natural language processing, and computer vision will equip these employees with the skills needed to build and maintain AI systems. Partnerships with academic institutions and online education platforms can provide valuable resources for these advanced training programs. For employees in nontechnical roles, AI literacy programs should focus on how AI tools affect their specific roles and how they can collaborate effectively with AI systems. These programs can include interactive workshops, online courses, and hands-on projects to foster a deeper understanding of AI applications.

Role-specific training tailored to the unique needs of different departments will ensure that employees can apply AI technologies effectively within their specific contexts. For instance, procurement officers might need training on AI tools for supply chain optimization, while health care administrators could benefit from courses on AI applications in patient care and medical research (Chapter 5). Similarly, law enforcement officers might gain insights from AI applications in predictive policing, crime pattern analysis, and facial recognition technologies, while federal financial employees can put AI training to good use in fraud detection and financial analysis.

Invest in the Right AI Tools and Infrastructure

Providing employees with access to the latest AI tools and platforms is crucial for fostering innovation and enhancing productivity (Zirar, Ali, and Islam, 2023). Modern AI tools enable employees to experiment with various applications, develop new solutions, and seamlessly integrate AI into their daily workflows (Chapter 3). Access to these tools allows for practical, hands-on experience, which is essential for understanding and effectively leveraging AI capabilities. Creating a one-stop shop for AI tools—such as servers, code libraries, and development environments—will provide a central resource for employees' needs (U.S. General Services Administration, undated).

Adopting AI requires significant upgrades to the existing infrastructure (Chapter 16). Hence, the federal government must invest in robust computing resources, cloud computing services, and secure data storage solutions to meet the intensive processing demands of AI algorithms. For the workforce, such investments ensure that AI systems have the necessary computational power and data access to function effectively. Organizations should also establish AI centers of excellence or innovation labs to support the deployment and management of AI technologies. These centers can serve as hubs for best practices, providing guidance and support to departments as they integrate AI into their workflows.

Foster a Culture of Continuous Learning

The rapid pace of AI development means that continuous learning and adaptation are essential for the federal workforce. Establishing a culture of continuous learning involves creating opportunities for employees to update their skills and stay informed about the latest advancements in the AI field. This can be achieved through regular training sessions, workshops, and seminars on emerging AI trends and technologies. Encouraging employees to participate in AI-related conferences and industry events can also help them stay abreast of new developments and best practices. Furthermore, providing access to online learning platforms and resources allows employees to learn at their own pace and convenience.

Mentorship programs and knowledge-sharing initiatives can facilitate the exchange of expertise and experience within the federal workforce. By pairing less experienced employees with AI experts, agencies can accelerate the learning curve and foster a collaborative environment in which knowledge is freely shared. In the same vein, creating interdisciplinary teams that combine expertise from various fields can foster a more comprehensive approach to AI adoption. These teams can work collaboratively to identify AI opportunities, address challenges, and ensure that AI solutions are both technically sound and socially responsible. Encouraging cross-departmental collaboration can also help break down silos and promote the integration of AI across different functions of the government.

Conclusion

Upskilling and retraining the federal workforce for AI adoption is a multifaceted endeavor that requires comprehensive training, infrastructure investment, and a strong focus on continuous learning. By equipping employees with the necessary skills and resources, the federal government can effectively leverage AI to enhance its operations, improve public services, and maintain a competitive edge in an increasingly digital world. Through careful planning and execution, the potential of AI can be fully realized, transforming the federal workforce into a model of innovation and efficiency.

References

Frost and Sullivan, *Global AI Maturity Survey 2024: Analysis of AI Technology, Adoption Priorities, and Progress of Enterprises Across Industries*, 2024.

Tamayo, Jorge, Leila Doumi, Sagar Goel, Orsolya Kovács-Ondrejkovic, and Raffaella Sadun, "Reskilling in the Age of AI," *Harvard Business Review*, September–October 2023.

U.S. General Services Administration, "AI Guide for Government, Chapter 2: How to Structure an Organization to Embrace AI," webpage, undated. As of August 30, 2024: https://coe.gsa.gov/coe/ai-guide-for-government/organizing-teams-and-resources/index.html

Walsh, James, Goldie Yuen, and Shankar Raman, "AI Roles Become Most Sought-After Digital Job Discipline," Willis Towers Watson, June 20, 2024.

Wilson, H. James, and Paul R. Daugherty, "Creating the Symbiotic AI Workforce of the Future," *MIT Sloan Management Review*, October 21, 2019.

Zirar, Araz, Syed Imran Ali, and Nazrul Islam, "Worker and Workplace Artificial Intelligence (AI) Coexistence: Emerging Themes and Research Agenda," *Technovation*, Vol. 124, June 2023.

CHAPTER 9

The Promise of AI to Transform Teaching Will Fail If School Systems Do Not Transform Too

Ivy Todd

If local and state education agency leaders want to use AI to support teachers, they must also consider the policy changes needed to give AI tools a chance of working. In this essay, I address the question of how administrators can best support innovative technology in schools.

The one-size-fits-all approach to K–12 education fits worse than ever: Before the COVID-19 pandemic, teachers were charged with the challenging task of providing grade-level instruction to classes whose background knowledge and skills varied widely. Teachers now contend with teaching those same grade-level standards, plus the foundational skills that students missed because of pandemic disruptions, all while supporting students' social and emotional development (Dougherty and Fleming, 2012; Lewis and Kuhfeld, 2024).

The challenge of catching up students from learning loss should not be underestimated. Addressing students' foundational skill gaps can ultimately help students develop strong mastery of grade-level material, but doing so requires teachers to diagnose the deficit(s), provide targeted instruction, and continually assess learning. Every student brings different strengths and academic needs that teachers must plan for. Many teachers lack the resources, time, and expertise required for this type of ongoing differentiation (Pauketat et al., 2023). Moreover, teachers across the United States report feeling burned out, overworked, and underpaid (Doan, Steiner, and Pandey, 2024).

Education and technology leaders point to AI as a potential solution to help students make up lost learning time without imposing additional burdens on teachers (e.g., Chen, 2023). Systems that leverage AI to blend teacher- and computer-led instruction to personalize student learning have existed for decades, been well-studied, and have often been found to be effective (Escueta et al., 2020). These systems vary in their approaches: They can serve as curriculum supplements, fully replace the main curriculum, or even facilitate new classroom models that reimagine teachers' role in instruction. But they share the promise to identify, address, and continually monitor students' academic progress and needs, potentially facilitating deeper learning and saving teachers time.

Prior education technology adoptions offer useful guidance for education leaders who are interested in AI systems. One of the primary challenges teachers face in implementing personalized learning systems, including those facilitated by AI, is known as the *mastery versus coverage dilemma*: the incompatibility between personalized approaches that strive to develop foundational curriculum mastery and coverage of the grade-level standards to which students and teachers are held accountable (Slavin, 1987; Rose, 2023).

Consider ALEKS, a tutoring software that uses machine learning to map students' existing subject-area knowledge and offer personalized instruction that is tailored to students' readiness. RAND researchers rigorously evaluated ALEKS's effectiveness as an algebra supplement but did not find significant benefits to the ALEKS group's learning outcomes (Phillips et al., 2020).

Implementation challenges associated with the mastery versus coverage dilemma can help explain the disappointing result. Teachers in the study struggled to implement ALEKS as the developers intended; only one class out of 40 met the usage and personalization expectations. Teachers worried that allowing students to build foundational skills with ALEKS was "giving up" time for the traditional algebra curriculum. The vast majority of teachers prioritized preparing students for the end-of-year algebra test over using ALEKS.

Nitkin, Ready, and Bowers (2022) documented similar challenges in their study of a personalized blended learning program for middle schoolers. The program provides students with daily instructional assignments that are customized to build on their existing knowledge with an appropriate level of challenge. Daily assignments could be technology-based, teacher-led, or small-group collaborative learning experiences, and teachers' roles shift from planning and leading instruction to focusing more on facilitation.

Even in this radically different learning setting, teachers *still* grappled with the mastery versus coverage dilemma. Leading up to the standardized testing season, teachers overrode the software recommendations for personalized instruction in favor of grade-level material that would be tested.

As they contemplate whether and how to embrace AI, education leaders must learn from prior studies that demonstrate the difficulty of implementing innovations without corresponding policy changes. Promising innovations cannot deliver impacts to students if the students do not actually receive the innovations. Teachers cannot implement innovative programs if the programs do not align with the priorities established by school-, district-, and state-level policies (Bingham et al., 2018).

Education leaders must also ensure that assessment strategies are sensitive enough to measure the effectiveness of new programs. In personalized systems, students with many learning gaps might spend substantial time addressing those gaps before progressing to grade-level topics. Even if student skills grow tremendously over the course of a year, traditional year-end grade-level assessments might not capture this growth if it occurred mostly in covering topics from prior grade levels. As school systems work to catch up students after the COVID-19 pandemic, they must consider whether their existing assessments capture student growth and progress and restrategize if not (Pane, 2018).

If school systems want to embrace the personalized approaches that AI is well-positioned to deliver, they must ask teachers to prioritize mastery of learning gaps over traditional curriculum coverage, and they must ensure that incentives and structures are in place to allow teachers, students, and AI tools to succeed. School systems that do not carefully plan for and support strong program implementation will fail the teachers and students whom they aim to help.

References

Bingham, Andrea J., John F. Pane, Elizabeth D. Steiner, and Laura S. Hamilton, "Ahead of the Curve: Implementation Challenges in Personalized Learning School Models," *Educational Policy*, Vol. 32, No. 3, May 2018.

Chen, Claire, "AI Will Transform Teaching and Learning. Let's Get It Right," Stanford Accelerator for Learning, March 10, 2023.

Doan, Sy, Elizabeth D. Steiner, and Rakesh Pandey, *Teacher Well-Being and Intentions to Leave in 2024: Findings from the 2024 State of the American Teacher Survey*, RAND Corporation, RR-A1108-12, 2024. As of August 30, 2024: https://www.rand.org/pubs/research_reports/RRA1108-12.html

Dougherty, Chrys, and Steve Fleming, *Getting Students on Track to College and Career Readiness: How Many Catch Up from Far Behind?* ACT, November 2012.

Escueta, Maya, Andre Joshua Nickow, Philip Oreopoulos, and Vincent Quan, "Upgrading Education with Technology: Insights from Experimental Research," *Journal of Economic Literature,* Vol. 58, No. 4, December 2020.

Lewis, Karyn, and Megan Kuhfeld, *Recovery Still Elusive: 2023–24 Student Achievement Highlights Persistent Achievement Gaps and a Long Road Ahead*, NWEA, July 2024.

Nitkin, David, Douglas D. Ready, and Alex J. Bowers, "Using Technology to Personalize Middle School Math Instruction: Evidence from a Blended Learning Program in Five Public Schools," *Frontiers in Education*, Vol. 7, April 13, 2022.

Pane, John F., *Strategies for Implementing Personalized Learning While Evidence and Resources Are Underdeveloped*, RAND Corporation, PE-314-RC, October 2018. As of August 30, 2024: https://www.rand.org/pubs/perspectives/PE314.html

Pauketat, Regena, Elaine Lin Wang, Julia H. Kaufman, Allyson D. Gittens, and Ashley Woo, *Teachers' Perceptions of Coherence in K–12 English Language Arts and Mathematics Instructional Systems*, RAND Corporation, RR-A279-3, 2023. As of August 30, 2024: https://www.rand.org/pubs/research_reports/RRA279-3.html

Phillips, Andrea, John F. Pane, Rebecca Reumann-Moore, and Oluwatosin Shenbanjo, "Implementing an Adaptive Intelligent Tutoring System as an Instructional Supplement," *Educational Technology Research and Development,* Vol. 68, No. 3, June 2020.

Rose, Joel, "Artificial Intelligence Will Not Transform K-12 Education Without Changes to 'the Grammar of School,'" The 74, May 30, 2023.

Slavin, Robert E., "Mastery Learning Reconsidered," *Review of Educational Research,* Vol. 57, No. 2, Summer 1987.

CHAPTER 10

Helping Postsecondary Education and Training Institutions Overcome Barriers to Preparing the New AI Workforce

Charles A. Goldman, Rita T. Karam, Brandon De Bruhl

Postsecondary institutions—including colleges, universities, and standalone career training providers—have a crucial role to play in preparing the next generation to develop and use AI tools in the workforce. Yet, without an integrated strategy that prepares the education and training *system* to respond to workforce changes, traditional education and training institutions will sacrifice impact and relevance and will risk being outcompeted by alternative certification providers. Those alternative providers have already contributed to the general public's questioning of the value of traditional college pathways. In this essay, we recommend coordinated strategies, depicted in Figure 10.1, that postsecondary institutions can use to meet the new demands of the AI workforce.

FIGURE 10.1

Coordinated Strategies Needed to Adapt Postsecondary Education and Training to Meet New AI Workforce Needs

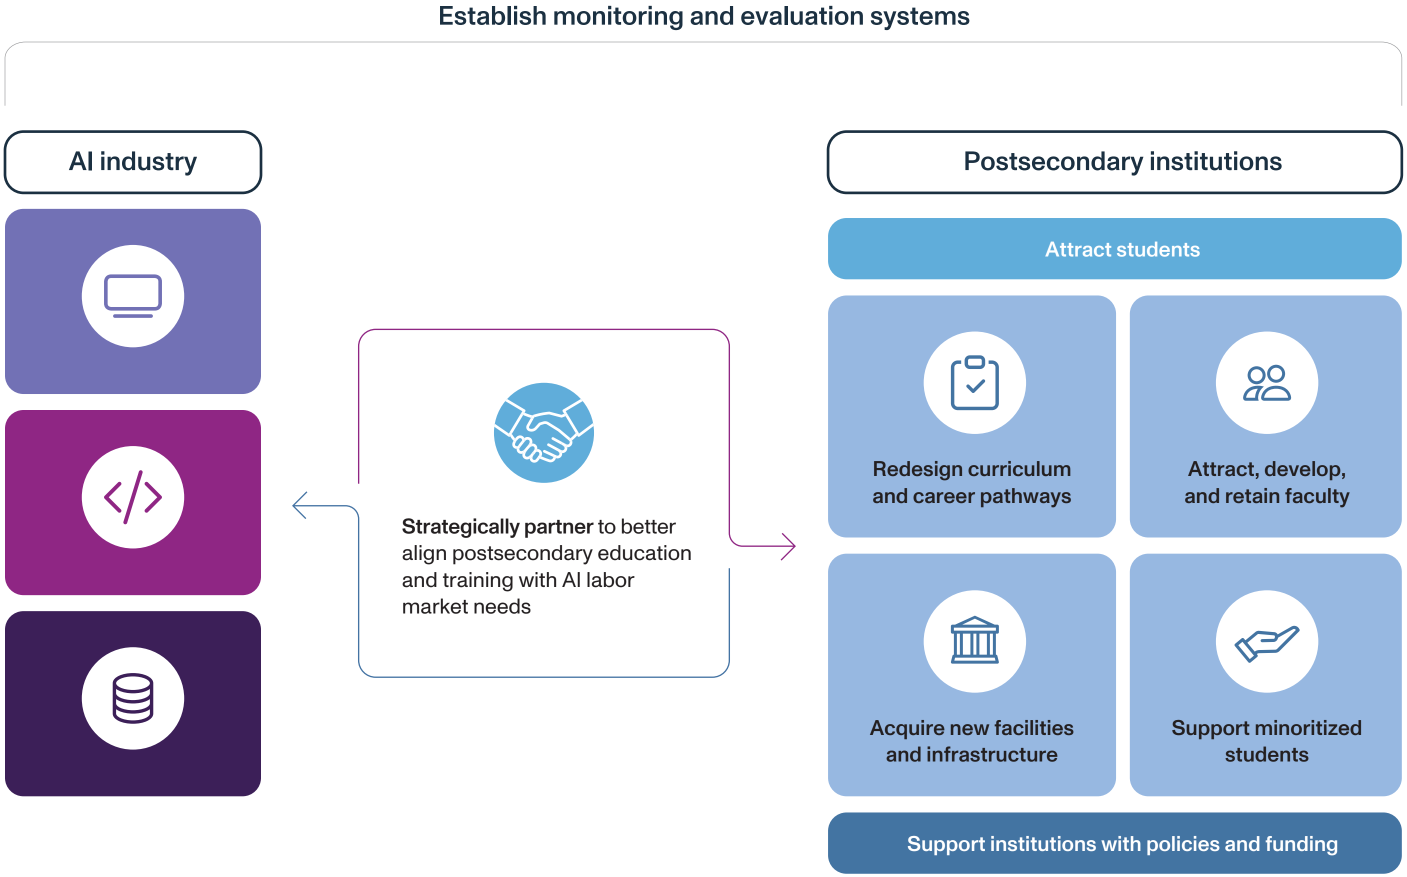

Strategically Partner with AI Industry

As of this writing, AI industries and postsecondary institutions generally have limited engagement, such as participating on program advisory boards or having a faculty member join industry associations. Many postsecondary institutions also lack the resources and know-how for developing partnerships. It is critical for institutions and industry to strategically partner and develop a strategic vision, common goals, and coordinated efforts that identify essential AI skills and jointly develop cutting-edge career pathways and curricula.

With strong partnerships, industry representatives can provide postsecondary institutions with information about the knowledge, skills, and competencies needed for careers in AI or that use AI technologies; offer real-world examples of problems and challenges for students to resolve; donate and provide AI technologies; and supplement instruction by offering AI-related internships and apprenticeships to students. Government can support these partnerships by providing postsecondary institutions with resources and guidance.

Better Align Education and Training with AI Labor Market Needs

Postsecondary institutions need to consider ways to attract students to new AI career pathways; redesign their curriculum and career pathways to account for the prevalence of AI; attract, develop, and retain faculty; acquire new facilities and infrastructure; and support minoritized and underrepresented students so that by the time these students graduate they will have developed the multifaceted AI skills needed for the new workforce.

Attract Students

Postsecondary institutions will need to market their new programs to students and workers by connecting curricula to career pathways and demonstrating positive returns on education investment. To promote equitable opportunities to gain AI expertise and diversify their enrollment, institutions should use marketing strategies that target students from racial and ethnic minority groups and should partner with K–12 schools that serve these students to expose them to AI-related programs and certificates as early as middle school. Institutions can also partner to implement dual enrollment courses in high schools to facilitate college transfer into their AI programs. Postsecondary institutions and training centers could also market their programs in community and employment centers to reach adult learners and unemployed workers.

Redesign Curriculum and Career Pathways

Postsecondary institutions will have to redesign their curriculum and career pathways to incorporate AI skills. Doing so will require more-nimble procedures for approving new courses and programs to keep pace with the evolution of AI.

New curricula must integrate experiential learning so that students are exposed to the changes AI brings to the workplace in real time. The curricula for technical programs should include increased emphasis on coding and programming skills, as well as ML and data analysis across many areas of study. Because AI will affect many job roles, curricula for all programs should address how AI tools can be used in different fields and how to do so in ways that mitigate bias and promote equity. To be competitive in the labor market, workers will also need foundational skills, such as data literacy and problem solving, as well as nontechnical skills, such as leadership, entrepreneurship, creativity, and understanding of diverse contexts.

Attract, Develop, and Retain Faculty

Postsecondary institutions will face challenges with attracting and retaining AI faculty unless they can offer strong compensation and benefits to compete with industry. Faculty across every content area will need ongoing professional development (which could be supported by industry partnerships) to understand AI tools and their uses and misuses as they apply to each content area and change over time.

Acquire New Facilities and Infrastructure

Postsecondary institutions will need to purchase computer hardware, software, and networks—or subscribe to increasingly popular cloud services—to support the use of AI tools across content areas. Technical programs would require additional investments in hardware and software or cloud services that are capable of the very high throughput necessary for developing AI models. Institutions may struggle with the costs of dedicated data centers or subscriptions to cloud services, although the increasing use of open-source tools can help with some of these cost pressures.

Support Minoritized Students

Even if they enter programs, students of racial and ethnic minority groups may not gain equitable access to opportunities in the AI workforce because they face a higher risk of dropping out of postsecondary programs. Therefore, postsecondary institutions should target academic and non-academic supports to meet these students' unique needs, which will help them access these opportunities.

Establish Monitoring and Evaluation Systems

With so much innovation occurring within and outside education and training providers, it is essential to define what constitutes effective AI programs and build systems to track the progress and success of programs at both the institutional level and in the larger education-workforce system.

Support Institutions with Policies and Funding to Overcome Barriers

Federal and state governments can support colleges, universities, and training providers in training the AI workforce by way of supportive policies (e.g., guidelines from government bodies regarding AI education, research ethics, and data use) and funding streams. For example, the U.S. Department of Labor could target technical assistance to help institutions develop their curriculum and career pathways and provide models of collaboration between postsecondary institutions and industry. Postsecondary education and training hold great promise in developing the future AI workforce, but it needs these sorts of supports to realize this promise.

CHAPTER 11

Leveraging AI in the Military Leader Development Framework

Dwayne M. Butler, Tuyen Dinh, Arianne Collopy

The U.S. military has expansive pillars of professional military education (PME) to augment what service members and civilians can do on their own or through the civilian education system. A holistic approach is needed to optimize these systems to ensure that the military can effectively and efficiently leverage the uses of AI and meet the internal and external threats associated with AI. In this essay, we provide DoD and military leaders with considerations for strategies to integrate AI into professional military education and training pathways.

DoD has a congressional mandate to "develop a strategy for educating service members in relevant occupational fields on artificial intelligence" (DoD, 2020, p. 1) so that service members have the skills to innovate and adapt to novel tools and technologies. Although some military services are already using AI technology in their combat training (Stilwell, 2020), a holistic approach is needed to meet the upskilled demands of an AI-centric work environment.[1]

Military officers are required to attend PME throughout their careers as part of their professional development and to prepare for the complex challenges of modern warfare—a natural first step in integrating AI education. To be clear, AI can augment and inform—but certainly not replace—the critical and strategic thinking required of leaders.

We offer strategic insights to help DoD integrate AI into the PME system. Although our focus is on the military, these principles could benefit any organization that has institutionalized training systems for its workforce development.

Setting Conditions for Integrating AI into PME

To insert an AI curriculum into the tiered levels of PME leader training, it is important to understand the training and education learning environment, which will inform how leaders at various levels might leverage AI and use it to deliver PME.

We consider PME and the potential applications for AI through four lenses, as shown in Table 11.1: (1) what enables PME delivery, (2) how PME is delivered, (3) what PME delivers, and (4) what PME is applied to in practice. Considering the actor and audience through each lens is helpful when enacting changes to PME systems. The variety of examples illustrates the range of doctrine or policy, organizational, and materiel considerations that underpin PME systems and delivery, which may need to be adapted to fully leverage AI within PME.

Future Considerations for AI in PME

PME provides a baseline of the what, who, and when that certain competencies must be developed across a service member's career to ensure mission readiness at each stage.

TABLE 11.1
Potential AI Applications Within PME

Lens	Examples	Examples of AI Enablers of PME	Examples of AI Topics in Curriculum
What enables PME delivery?	Pillars of PME and leader development, training structures and systems, policy and doctrine	AI-enabling training system (e.g., allocation of instructors and resources)	Alignment of AI tools and resources to pillars of PME and leader development
How is PME delivered?	Basic training, specialty or advanced training, unit training, promotion training	Custom delivery of training content or custom training content	AI as a tool to enable training in multiple areas
What does PME deliver?	Knowledge, skills, abilities, and competencies	LLMs or other algorithms to summarize and extract key information	AI knowledge, skills, and competencies
What is PME applied to?	Operational mission execution, mission support	Mission-enabling AI of allies	Mission-enabling AI of adversaries, how to identify disruption to DoD's or an ally's AI capability

Furthermore, DoD's 2020 *Artificial Intelligence Education Strategy* provides ample guidance for designing the training and education curriculum of which competencies to develop, to be taught at what depth, and for which AI-based role (DoD, 2020). We suggest that the current PME system integrate AI in ways that match the level of AI instruction with the level of leadership development.

Matching AI Instruction with the Level of Leadership Development

Matching the level of AI instruction to the level and tenets of leadership development in a mission-oriented ecosystem is an efficient approach to integrate AI instruction to existing curriculum. We introduce examples of AI topics in curriculum that draw on the examples in Table 11.1 for each of the five levels of leadership development in PME:

- **Precommissioning education** (e.g., Reserve Officers' Training Corps [ROTC]) currently focuses on foundational knowledge of U.S. defense. A foundational introduction to AI concepts relevant to the military could include as topics the general understanding and application of AI.
- **Primary education** (e.g., pilot training) focuses on specialized training within a service branch, at the tactical level of war. At this point, service members should have an intermediate to advanced level of understanding of the foundational concepts and begin developing a basic level of understanding of AI competencies. Examples of such AI applications are identifying trends and risks, as well as other technical concepts that may be relevant to a specialty (e.g., infrastructure, coding, and software development; performing analysis).
- **Intermediate education** focuses on the operational level of war. AI education at this level should aim for advanced knowledge in foundational concepts and AI applications but may retain a basic to intermediate level of knowledge in other aspects.
- **Senior education** focuses on the strategic level of war. This education should provide at least a basic level of understanding across all AI competencies and advanced knowledge in concepts relating to military applications of AI, AI doctrine, AI predictability, and responsible AI.
- **General officer and flag officer education** needs to encompass all the above levels of education and continue instruction in specialty roles so that topics of emerging AI applications and the risks and impacts on ally and adversary missions are appropriate.

Conclusion

Although there are ample opportunities for applying AI within the PME system, there are costs to implement and operate the tools, along with risks that reliance on AI tools may change how the military makes decisions and operates. There is also a risk of overemphasizing the promise of technological innovation by integrating AI without implementing the educational components of how to effectively, ethically, and responsibly use it. [2] We offer the following recommendations:

- AI is both a topic that should be included in PME and a potential way to deliver PME that should be explored at multiple scales (see Table 11.1).

- The desired attributes of military leaders that guide PME should also guide the integration of AI into PME.
- A four-step process should be used to phase the integration of AI into the PME system:
  1. Align AI to existing curricular goals and implement AI curriculum.
  2. Optimize the PME delivery system via AI to respond to today's mission environment considering both who is delivering PME and for whom.
  3. Identify future opportunities for AI to understand future mission environments and adapt the AI curriculum (such as using PME courses as a test and evaluation simulation for proposed AI solutions to minimize risks or unnecessary losses).
  4. Continually assess the costs and risks of AI use within PME and operational missions to inform and adapt the implementation of AI within the PME system and the curriculum to introduce AI topics and tools.

Notes

[1] In this essay, *AI* refers to LLMs embedded in chatbots and knowledge ecosystems, algorithms to process and synthesize information at size and scale, and autonomous warfighting capability.

[2] Military innovation requires evidence that the innovation is impactful (i.e., increases military power) and has widespread adoption (or uptake) among the military community, suggesting access and familiarity with the innovation (Horowitz and Pindyck, 2023).

References

DoD—*See* U.S. Department of Defense.

Horowitz, Michael C., and Shira Pindyck, "What Is a Military Innovation and Why it Matters," *Journal of Strategic Studies*, Vol. 46, No. 1, 2023.

Stilwell, Blake, "How the Army and Air Force Integrate AI Learning into Combat Training," Military.com, February 14, 2020.

U.S. Department of Defense, *2020 Department of Defense Artificial Intelligence Education Strategy*, September 2020.

Section D

BUILDING A MORE RESILIENT AND DIVERSE AI WORKFORCE

CHAPTER 12

Retaining Workers with AI Skills in the Federal Workforce

Michael G. Mattock, Avery Calkins

Federal agencies are planning to hire hundreds of AI experts by 2025 (Heckman, 2024). The type and level of AI skills likely to be used by federal workers will be diverse. This diversity is recognized in the new U.S. Office of Personnel Management (OPM) skill-based hiring guidance for federal AI positions (Ahuja, 2024), which emphasizes skills over formal educational attainment.

High demand for AI skills has led to a significant pay premium for these jobs in the private sector. In early 2024, jobs requiring AI skills were associated with a wage premium of approximately 25 percent relative to jobs that do not require AI skills (Kupelian et al., 2024), and on the job site Indeed, technology job postings that require skills in generative AI listed earnings that were 47 percent higher than other technology job postings (Gafner, 2024). The pay differential and budget constraints make it difficult for the federal government to compete for this talent, and, on hiring these workers, the government may struggle to retain them long-term unless it takes action.

The federal government has several tools to maintain or improve the retention of civilian workers who have or might acquire AI skills. These tools are in the areas of compensation, work environment, and training opportunities.

Raise Compensation

One option is to permanently raise pay for workers with AI skills—for instance by raising the grade of general schedule AI workers, moving AI workers to an alternative pay plan,[1] or introducing a special pay (i.e., a permanent wage differential) for AI workers. Another option is to offer retention bonuses tied to service commitments. This would give the government the flexibility to adapt compensation to current circumstances rather than permanently raise pay to a level that in the long run could potentially exceed private sector earnings. In addition, offering bonuses that are conditional on multiyear contracts is more efficient than offering unconditional special pays (Hosek, Mattock, and Asch, 2019). Furthermore, portions of compensation other than pay could be adjusted for AI workers to improve retention,[2] as could ensuring a more rewarding work environment.

Improve Future Opportunities

Another approach is to provide training in the AI skills needed by the federal government. Training can increase the opportunity wage available to workers outside their current organization. Figure 12.1, drawn from previously unpublished analysis in support of a RAND report on civilian cyber workers in DoD (Mattock et al., 2022), shows the cumulative retention rate at each year of service when cyber workers are (purple line) and are not (blue line) offered training with a payback period,[3] assuming that training raises external wages. Fewer workers leave the organization prior to training because they anticipate the expanded opportunities that training can provide, and retention remains higher through the end of the payback period. However, after the payback period, retention declines at a faster rate than under the baseline as these trained employees take advantage of new opportunities and the pay premium. The government could improve retention following the training payback period by increasing pay for workers who complete the training. Training will also be of benefit to the government if work roles are available to take advantage of the increased capabilities of the trained workers.

The Way Forward

The way forward depends on whether the wage premiums for AI skills in the private sector labor market are likely to be permanent. Wage premiums could shrink or disappear entirely as more workers gain AI skills, if demand shrinks, or if AI-related human capital depreciates more quickly after the current AI boom.[4] If so, the government should opt for short-term bonuses tied to service obligations because they are more flexible. If wage premiums are likely to be permanent, raising pay might be a more appropriate approach. Furthermore, regardless of how persistent wage premiums are, we suggest that the government pursue options for training workers in the AI skills most needed for federal work—with associated service obligations and potentially with higher pay and a more rewarding work

FIGURE 12.1

Training with a Payback Period but No Internal Wage Increase Might Result in a Temporary Increase in Retention

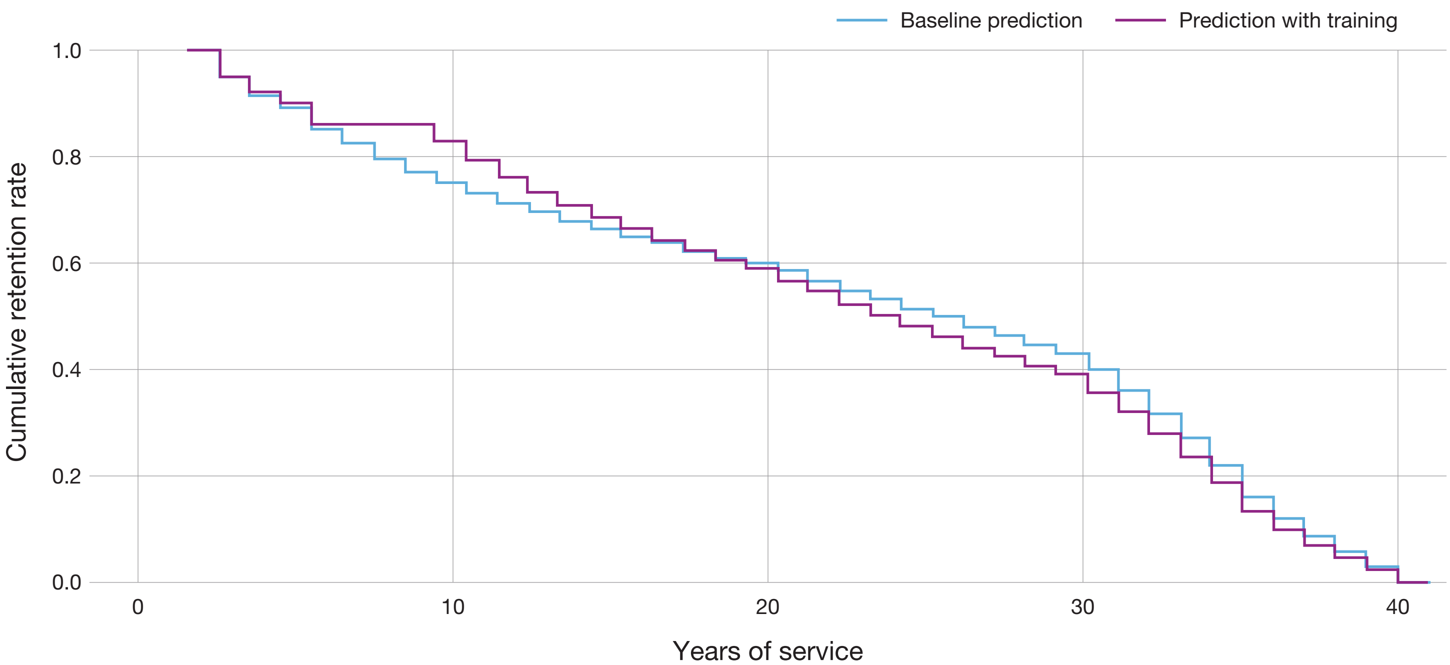

SOURCE: Authors' previously unpublished analysis in support of Mattock et al. (2022).

NOTE: The graph represents a dynamic retention model simulation of the retention behavior of cyber workers with a bachelor's degree only in occupation series 1550 (computer science) with one year of training starting at the beginning of the fifth year of service that incurs a three-year payback obligation and raises the external opportunity wage by 3 percent. The internal wage is assumed to remain unchanged in this simulation.

environment to ensure that the government can retain such workers after training is complete. AI will continue to evolve and serve increasingly many uses. Ongoing training will enable the federal government to reap the benefits of AI and provide federal civilian workers with the assurance that their job skills will not decay any faster than they would in the private sector.

Notes

[1] For instance, the Cyber Excepted Service and parts of the intelligence community are on a different grade-step plan than the general schedule (DoD Workforce Innovation Directorate, undated; Seacord, 2023). The government could also shift to a pay band approach, such as those used for DoD acquisitions workers and in some scientific research laboratories (Office of the Deputy Assistant Secretary of Defense [Civilian Personnel Policy], 2017; Under Secretary of Defense for Research and Engineering, 2020).

[2] For instance, the government already provides telework, alternative or flexible work schedules, student loan forgiveness, and other benefits to many workers (OPM, undated), and it could ensure that, when possible, these benefits are offered for positions that require AI skills.

[3] Extensive training in the federal government generally comes with a *payback period*, in which individuals are obligated to stay in their position for a certain number of years after the training is complete.

[4] For instance, there seems to have been greater human capital depreciation among information technology workers following the dot-com boom of the late 1990s; see Hombert and Matray, 2019.

References

Ahuja, Kiran, "Skills-Based Hiring Guidance and Competency Model for Artificial Intelligence Work," memorandum for heads of executive departments and agencies, U.S. Office of Personnel Management, April 29, 2024.

DoD Workforce Innovation Directorate, "DoD Cyber Excepted Service—TLMS Pay Rates," webpage, DoD Cyber Exchange, undated. As of July 9, 2024:
https://public.cyber.mil/wid/dod-cyber-excepted-service-ces/pay

Gafner, Jocelyne, "Highest Paid Skills in Tech," Indeed, February 21, 2024. As of June 27, 2024:
https://www.indeed.com/career-advice/news/highest-paid-skills-tech

Heckman, Jory, "White House Looks to Add 500 AI Experts to Federal Workforce by 2025," Federal News Network, April 29, 2024.

Hombert, Johan, and Adrien Matray, "Technology Boom, Labor Reallocation, and Human Capital Depreciation," Griswold Center for Economic Policy Studies, Working Paper No. 260, September 2019.

Hosek, James, Michael G. Mattock, and Beth J. Asch, *A Wage Differential Approach to Managing Special and Incentive Pay*, RAND Corporation, RR-2101-OSD, 2019. As of June 13, 2024:
https://www.rand.org/pubs/research_reports/RR2101.html

Kupelian, Barret, Sarah Brown, Mehdi Sahneh, Simon Oates, Justine Brown, Adam Deasy, Nabil Taleb, Harry Ingham, and Ilhan Guner, *AI Jobs Barometer: USA Findings*, PwC, 2024.

Mattock, Michael G., Beth J. Asch, Avery Calkins, and Daniel Schwam, *Civilian Cyber Workers in the U.S. Department of Defense: Demographics, Retention, and Responsiveness to Training Opportunities*, RAND Corporation, RR-A730-3, 2022. As of August 30, 2024:
https://www.rand.org/pubs/research_reports/RRA730-3.html

Office of the Deputy Assistant Secretary of Defense (Civilian Personnel Policy), "Civilian Acquisition Workforce Personnel Demonstration (AcqDemo) Project," *Federal Register*, Vol. 82, No. 216, November 9, 2017.

OPM—*See* U.S. Office of Personnel Management.

Seacord, James M., "2024 Defense Civilian Intelligence Personnel System Pay Rates," memorandum for senior pentagon leadership defense agency and DoD field activity directors, Office of the Undersecretary of Defense for Intelligence and Security, December 22, 2023.

Under Secretary of Defense for Research and Engineering, "Science and Technology Reinvention Laboratory (STRL) Personnel Management Demonstration Project in the Joint Warfare Analysis Center (JWAC) of the United States Strategic Command (USSTRATCOM)," *Federal Register*, Vol. 85, No. 95, May 15, 2020.

U.S. Office of Personnel Management, "Federal Employee Compensation Package," webpage, undated. As of July 10, 2024:
https://www.opm.gov/policy-data-oversight/pay-leave/pay-administration/fact-sheets/federal-employee-compensation-package

CHAPTER 13

Is the Domestic AI Talent Pool Sufficient to Meet Public Sector Demand?

Angela K. Clague, Srikant Kumar Sahoo

Demand for AI professionals—workers trained to build AI systems who typically work as data scientists, machine learning or deep learning engineers, software engineers, or software architects (Gafner, 2024)—is on the rise, and the U.S. government and private businesses are in fierce competition for their services. Is there enough talent to go around? Although degrees are not necessary for many AI positions, we estimate the supply of this talent using the number of recent graduates with AI-relevant degrees.[1] We find that the small number of Ph.D.'s, limits on foreign-born workers, and higher compensation in the private sector make that question hard to answer.

A Disparity in Degrees

There is a significant difference across graduate levels in the share of graduates who have AI-relevant degrees, especially at the highest level. As shown in Figure 13.1, the percentage of university graduates with an AI-relevant Ph.D. remained below 1 percent over a ten-year period. At the bachelor's-degree level, the percentage of graduates with an AI-relevant major doubled from 2 percent in 2011 to 4 percent in 2021. For graduates with master's degrees, this percentage increased fourfold, from 1 percent in 2011 to 4 percent in 2021, with a peak at 9 percent in 2016.

Limits on Foreign-Born Employment

A large share of graduates with AI-relevant degrees is foreign-born, which may limit these graduates' ability to work in the public sector. In 2021, 62 percent of AI-relevant Ph.D.'s were awarded to foreign-born graduates.

FIGURE 13.1
University Graduates with AI-Relevant Training by Year and Degree Level

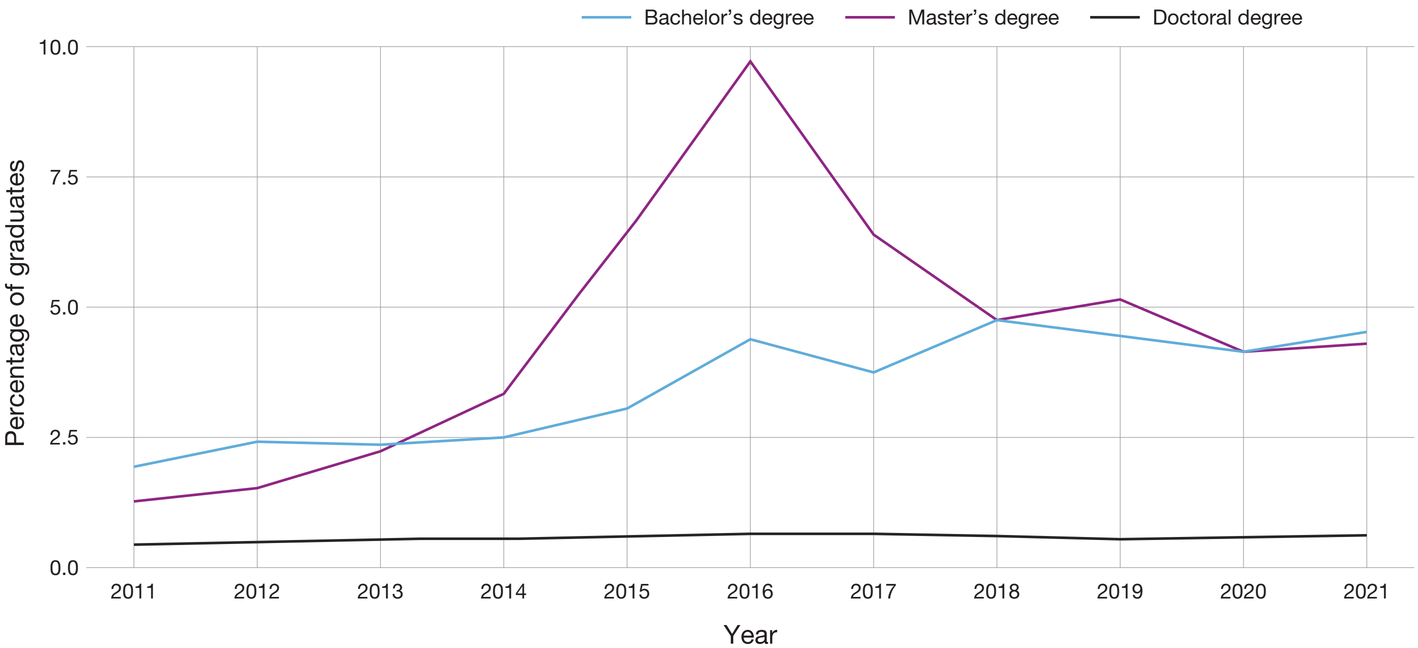

SOURCE: Features information from the 2012 to 2022 waves of the Integrated Postsecondary Education Data System (National Center for Education Statistics, undated).

At the bachelor's and master's degree levels in 2021, more than 80 percent of these graduates were born in the United States. In jobs for which a Ph.D. is not required, individuals with bachelor's and master's degrees might provide a sufficient talent pool to meet rising government demand.

Competition with Private Businesses

The government still faces steep competition from the private sector in recruiting AI professionals, and salary discrepancies with the private sector at all degree levels puts the government at a disadvantage.

In 2021, AI professionals in the private sector earned an average of $138,000 with a bachelor's degree; $184,000 with a master's degree; and $176,000 with a Ph.D. In government, earnings at those three levels averaged $71,000, $119,000, and $156,000, respectively.[2] Losing job candidates to the private sector potentially reduces the quantity and quality of the talent that may be available for and attracted to government service.

Ultimately, whether the talent pool is sufficient to fill the need may depend on whether the organization is public or private, the level of education the job requires, and how motivated individual candidates are by salary.

Recommendations

We provide some recommendations for recruiting AI talent in the public sector. Although the public sector may not be able to offer salaries that are competitive with those earned in the private sector, the public sector can **communicate the benefits of public sector careers** to job seekers that extend beyond salary, such as job stability, employee benefits (e.g., pensions), and community impacts. Public sector organizations can **partner with academic institutions** to educate students who are pursuing AI-relevant degrees about careers and internships in the public sector. Another way of increasing the public sector talent pool is to **expand eligibility requirements to some foreign-born AI professionals**, particularly for research and development AI positions that may require advanced degrees.

Notes

[1] We define *AI-relevant training* as degree fields (such as computer science and applied mathematics) that are typically listed on LinkedIn and Indeed applications for AI-building jobs.

[2] Data are from the 2002 to 2022 American Community Survey (U.S. Census Bureau, 2022).

References

Gafner, Jocelyne, "Highest Paid Skills in Tech," Indeed, February 21, 2024.

National Center for Education Statistics, Integrated Postsecondary Education Data System, database, undated. As of September 6, 2024: https://nces.ed.gov/ipeds/

U.S. Census Bureau, "Public Use Microdata Sample (PUMS)," webpage, last updated December 15, 2022. As of September 6, 2024: https://www.census.gov/programs-surveys/acs/microdata.html

CHAPTER 14

Building a Diverse Pool of AI Talent for DHS Recruitment

Neeti Pokhriyal, Angela K. Clague

Although DHS's current focus for building a diverse AI workforce for the United States seems to be on how to attract and maintain international talent (DHS, 2024a), we argue that DHS should also consider cultivating the domestic talent pipeline in AI and, more broadly, in STEM fields. Attracting, recruiting, and retaining a diverse STEM workforce is particularly challenging for DHS because the private sector can offer higher salaries than the public sector. These staffing challenges are exacerbated by limited diversity in the existing AI workforce and AI talent being concentrated in very few locations nationwide, such as San Jose, Seattle, and San Francisco (Fitzpatrick and Beheraj, 2024).

DHS has an imminent need to cultivate a domestic talent pipeline in AI so that it has a diverse pool of candidates to recruit from. In this essay, we delve into the specific challenges related to building a diverse AI workforce at DHS and provide some concrete recommendations for doing so. The *DHS Inclusive Diversity Strategic Plan for Fiscal Years 2021–2024* focuses on *inclusive diversity,* emphasizing belonging and individual experience (DHS, undated). Yet, the plan does not detail how to attract, recruit, and retain workers with diverse backgrounds (DHS, 2024a). The *Department of Homeland Security Artificial Intelligence Roadmap 2024* recognizes the need to develop a diverse AI-ready workforce that can sustain and advance its mission across a varied portfolio, including but not limited to anti-terrorism, border security, immigration, cybersecurity, disaster prevention, emergency management, law enforcement, and national infrastructure (DHS, 2024a, p. 5). However, what constitutes a diverse AI workforce and how to achieve it is less clear.

The DHS AI roadmap stresses that AI technologies not only deliver benefits to the public but also fuel strategic research and development efforts across all areas of homeland security. To fulfill these dual missions, the AI workforce that DHS seeks to build needs various levels of AI-related skills and competencies. Although some jobs may require AI experts with advanced degrees in computer science, others might be fulfilled by technical certifications and on-the-job training. The staffing demand could be considerable if DHS needs AI-ready staff across its 22 agencies and its offices that are geographically dispersed nationwide.

Thus, there is an imminent need to attract, cultivate, and match domestic AI talent unique to DHS's demands while ensuring that it is done equitably so that the economic and societal gains that AI technology has the potential to bring can be broadly dispersed throughout the country. A source of such talent is the students at minority-serving institutions (MSIs), historically Black colleges and universities (HBCUs), tribal colleges and universities, Hispanic-serving institutions, and other academic institutions that are classified as R2 (high research activity) institutions, as well as R1 (very high research activity) institutions (Carnegie Classification of Institutions of Higher Education, undated). Figure 14.1 shows the geographic proximity of a sample of DHS offices and MSIs, which can provide a source of AI talent for DHS workforce needs.

Some institutions have ongoing engagement with DHS in workforce development initiatives (DHS, 2024b; DHS, 2024c; AI and Tech Talent Task Force, 2024). Many of these academic institutions cater to students from rural areas and communities that are traditionally underserved in computing (Gershenfeld et al., 2021). Thus, building and attracting this talent will ensure that DHS prioritizes recruiting its AI workforce from a diverse pool of qualified candidates who are representative of the communities they serve and better understand the local contexts. Hiring qualified staff who live within geographic proximity of DHS offices and agencies might also help with retention (AI and Tech Talent Task Force, 2024), which is challenging in technical disciplines, where employees quickly change jobs to acquire rapidly evolving skills.

Focusing on the existing challenges that DHS faces in building a diverse workforce, we highlight the deficiencies in the metric it employs to evaluate diversity. Currently, DHS evaluates demographic diversity—the proportion of men, women, and various racial and ethnic groups—by rank across the agency (DHS, 2023a). DHS measures diversity by employees' responses to yes/no questions about whether diversity, equity, and inclusion (DEI) policies, procedures, or communications exist in each agency (DHS, 2023a). However, these questions do not ask whether DEI efforts are effective. There is no publicly available data on the demographic composition of the DHS AI workforce at present because these measures do not evaluate diversity by job or department.

FIGURE 14.1

Geographical Proximity of DHS Offices and Various MSIs

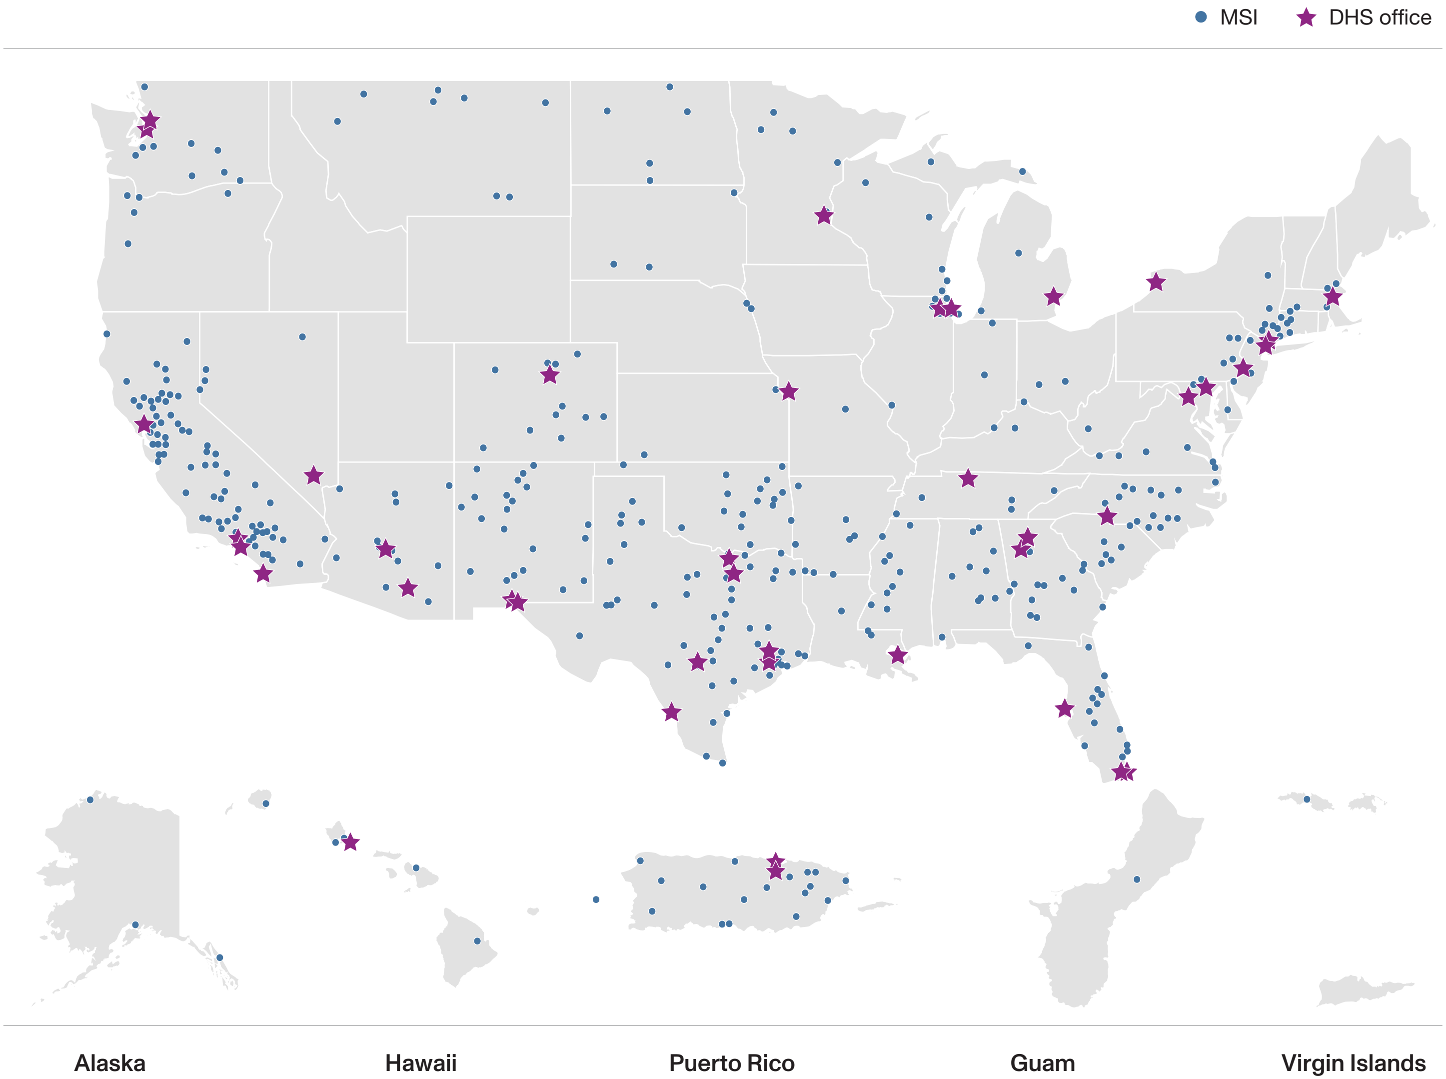

SOURCES: Authors' analysis of data from Homeland Security Investigations, 2024; Federal Emergency Management Agency, 1999; U.S. Customs and Border Protection, undated; and Minority Serving Institutions Exchange, 2024.

Additionally, DHS has difficulty retaining women and racial and ethnic minorities long enough for them to enter leadership positions (Curry Hall et al., 2019). Such factors as remote locations and discrimination have been cited as possible drivers of attrition (Lim et al., 2021; Sims et al., 2022). DHS needs to ensure that the existing difficulties in retaining a diverse workforce are not exacerbated in the context of AI. A nuanced understanding is needed on how job requirements and the current job climate at DHS would affect the hiring and retention of the AI workforce.

Recommendations

We provide the following concrete recommendations on how to build domestic AI talent focusing on students in rural and underserved communities so that DHS has a diverse talent pool of candidates to recruit from:

- **Redesign recruiting efforts**, recognizing that varied levels of AI-based skills and competencies are needed at different jobs across the DHS centers. Leverage DHS's existing partnerships with academic institutions (DHS, 2024b), especially MSIs, to further build pathways and onramps via career fairs, internships, etc., so that DHS can hire students trained in AI and STEM fields from academic institutions, especially institutions serving rural areas and underserved communities in geographically dispersed places.
- **Build robust funding mechanisms** to invest, engage, and train the AI workforce in these underserved communities and institutions and provide them with clearer and well-defined career pathways to hiring and retention. DHS can leverage existing programs, such as the Office of University Programs Minority Serving Institutions Program (DHS, 2023b) and the Centers of Excellence programs (DHS, 2024b). Understand the barriers faced by students from underserved communities in joining the DHS workforce and how to lower them.
- **Build robust partnerships with other federal agencies** (e.g., National Science Foundation, DoD) that fund academic centers to develop a workforce in emerging technologies and with industries that can provide certifications and training focused on DHS-relevant aspects (Office of Critical and Emerging Technologies, undated).
- **Enhance how DHS measures diversity**, including measures to evaluate whether its diversity policy is effective beyond reporting simple percentages. The DHS diversity strategic plan does not define what diversity means to the organization or propose how best to attract, recruit, and retain a diverse AI workforce. Identify the demographic composition of the AI workforce, challenges with recruiting, and drivers of attrition. Existing data on diversity in DHS does not capture this information, which is critical to understanding how to attract and retain a diverse AI workforce.

References

AI and Tech Talent Task Force, *Increasing AI Capacity Across the Federal Government: AI Talent Surge Progress and Recommendations*, Executive Office of the President, April 2024.

Carnegie Classification of Institutions of Higher Education, "Basic Classification," webpage, undated. As of August 30, 2024: https://carnegieclassifications.acenet.edu/carnegie-classification/classification-methodology/basic-classification

Curry Hall, Kimberly, Kirsten M. Keller, David Schulker, Sarah Weilant, Katherine L. Kidder, and Nelson Lim, *Improving Gender Diversity in the U.S. Coast Guard: Identifying Barriers to Female Retention*, Homeland Security Operational Analysis Center operated by the RAND Corporation, RR-2770-DHS, 2019. As of September 10, 2024: https://www.rand.org/pubs/research_reports/RR2770.html

DHS—*See* U.S. Department of Homeland Security.

Federal Emergency Management Agency, "Appendix C—FEMA Offices," in *Protecting Building Utilities from Flood Damage: Principles and Practices for the Design and Construction of Flood Resistant Building Utility Systems*, 1st ed., November 1999.

Fitzpatrick, Alex, and Kavya Beheraj, "America's AI Job Hotspots, Mapped," Axios, April 16, 2024.

Gershenfeld, Joel E. Cutcher, Alan Blatecky, Damian Clarke, Deborah Dent, Rebecca Hipp, Ana Hunsinger, Al Kuslikas, and Lauren Michael, *The Missing Millions: Democratizing Computation and Data to Bridge Digital Divides and Increase Access to Science for Underrepresented Communities*, National Science Foundation, October 3, 2021.

Homeland Security Investigations, "Contact Your Local SAC Office," webpage, last updated May 9, 2024. As of September 12, 2024: https://www.dhs.gov/hsi/contact/sac-offices

Lim, Nelson, Kimberly Curry Hall, Kirsten M. Keller, David Schulker, Louis T. Mariano, Miriam Matthews, Lisa Saum-Manning, Devon Hill, Brandon Crosby, Leslie Adrienne Payne, Linda Cottrell, and Clara A. Aranibar, *Improving the Representation of Women and Racial/Ethnic Minorities Among U.S. Coast Guard Active-Duty Members*, Homeland Security Operational Analysis Center operated by the RAND Corporation, RR-A362-2, 2021. As of August 30, 2024: https://www.rand.org/pubs/research_reports/RRA362-2.html

Minority Serving Institutions Exchange, "Institution Categories," webpage, National Aeronautics and Space Administration, last updated September 12, 2024. As of September 12, 2024: https://msiexchange.nasa.gov/categories

Office of Critical and Emerging Technologies, "Supercharging America's AI Workforce," webpage, U.S. Department of Energy, undated. As of August 30, 2024: https://www.energy.gov/cet/supercharging-americas-ai-workforce

Sims, Carra S., Coreen Farris, Terry L. Schell, Miriam Matthews, Barbara Bicksler, Owen Hall, and Lisa Wagner, *Harassment and Discrimination on the Basis of Gender and Race/Ethnicity in the FEMA Workforce: 2021 Survey Follow-Up*, Homeland Security Operational Analysis Center operated by the RAND Corporation, RR-A383-2, 2022. As of August 30, 2024: https://www.rand.org/pubs/research_reports/RRA383-2.html

U.S. Customs and Border Protection, "List of Field Offices," webpage, undated. As of September 12, 2024: https://www.cbp.gov/about/contact/ports/field-office

U.S. Department of Homeland Security, *DHS Inclusive Diversity Strategic Plan for Fiscal Years 2021–2024: Elevating the Human Experience*, undated.

U.S. Department of Homeland Security, *U.S. Department of Homeland Security EEOC Management Directive 715: Equal Employment Opportunity Program Status Report Fiscal Year 2022*, July 2023a.

U.S. Department of Homeland Security, "Office of University Programs (OUP) Minority Serving Institutions (MSI) Program," fact sheet, November 2023b.

U.S. Department of Homeland Security, *Department of Homeland Security Artificial Intelligence Roadmap 2024*, March 2024a.

U.S. Department of Homeland Security, "Centers of Excellence," webpage, last updated August 2, 2024b. As of August 30, 2024: https://www.dhs.gov/science-and-technology/centers-excellence

U.S. Department of Homeland Security, "Workforce Development Initiatives," webpage, last updated August 8, 2024c. As of August 30, 2024: https://www.dhs.gov/science-and-technology/workforce-development-initiatives

CHAPTER 15

Identifying Resilient Skills in an AI-Enhanced Economy

Jonah Kushner, Tobias Sytsma, Éder M. Sousa

As concerns about AI's impact on the workforce are becoming more prevalent across industries—from software development to entertainment—policymakers, educators, and workers are looking for ways to prepare for a future in which AI is increasingly integrated into the economy. The focus of these discussions has been on the tasks that AI might automate; the ones it struggles with get less attention. Preparing people to perform these more-difficult tasks may become more urgent as AI automates other aspects of people's jobs.

To mitigate the disruptive impact that AI may have on the workforce and prepare workers for the labor market of the next decade, educators and policymakers must identify the tasks and work activities in which humans hold a comparative advantage over AI. In this essay, we address the following policy problems:

- How can organizations identify the skills and work activities that are the least likely to be disrupted by AI?
- How can researchers and policymakers use this information to develop policies and strategies that foster resilience in the workforce?

Proactive investment in training with a focus on the skills required for activities that are less likely to be automated can help mitigate AI's negative impacts on the workforce. Building on prior RAND research (Sytsma and Sousa, 2023), we use natural language processing to compare written descriptions of tasks performed in different occupations with the capabilities of AI tools.[1] Tasks with more AI patents are considered more susceptible or "exposed" to AI automation.

Using occupational wage data and task importance ratings from federal surveys, we estimated the economic value of tasks and the broad work activities to which the tasks belong.[2] This approach identifies the types of work activities that are most and least exposed to AI automation and their economic value, and it can provide valuable information to policymakers, educators, employers, and workers to help build a more resilient workforce.

Key Findings

Figure 15.1 shows the average number of AI patents for each work activity, indicated by the length of the bars (longer is greater), and the average economic value of tasks for each work activity, indicated by the color of the bars (light blue is low value, dark blue is high value). We find that

- Work activities involving soft skills are generally less exposed to AI technologies. Activities that involve helping or influencing others show relatively low AI exposure and require interpersonal interactions and emotional intelligence.
- Higher-value activities that have lower AI exposure appear to combine soft skills with some level of specialized content knowledge or *nonroutine manual skills*, defined as the ability to manipulate physical objects in unstandardized and unpredictable environments.

What do these patterns imply for the future of work and the demand for workers to perform certain kinds of activities? As today's AI tools diffuse throughout the economy, they will raise productivity among certain occupations and drive increased needs for complementary skills. Demand may decrease for workers capable of monitoring, analyzing, and optimizing, with a small core of employees remaining to develop and implement AI and other tools to perform these activities. Simultaneously, demand may increase for workers capable of performing soft skill–intensive activities. Fortunately, many soft skill–intensive work activities appear to be of high value. Moreover, AI tools could be directed to prepare workers for high-value, soft skill–intensive activities and increase their productivity by helping educate and train workers in the specialized knowledge that they need for these activities and by providing guidance and feedback on nonroutine manual tasks.

FIGURE 15.1

The Least Exposed Work Activities Appear to Involve Soft Skills

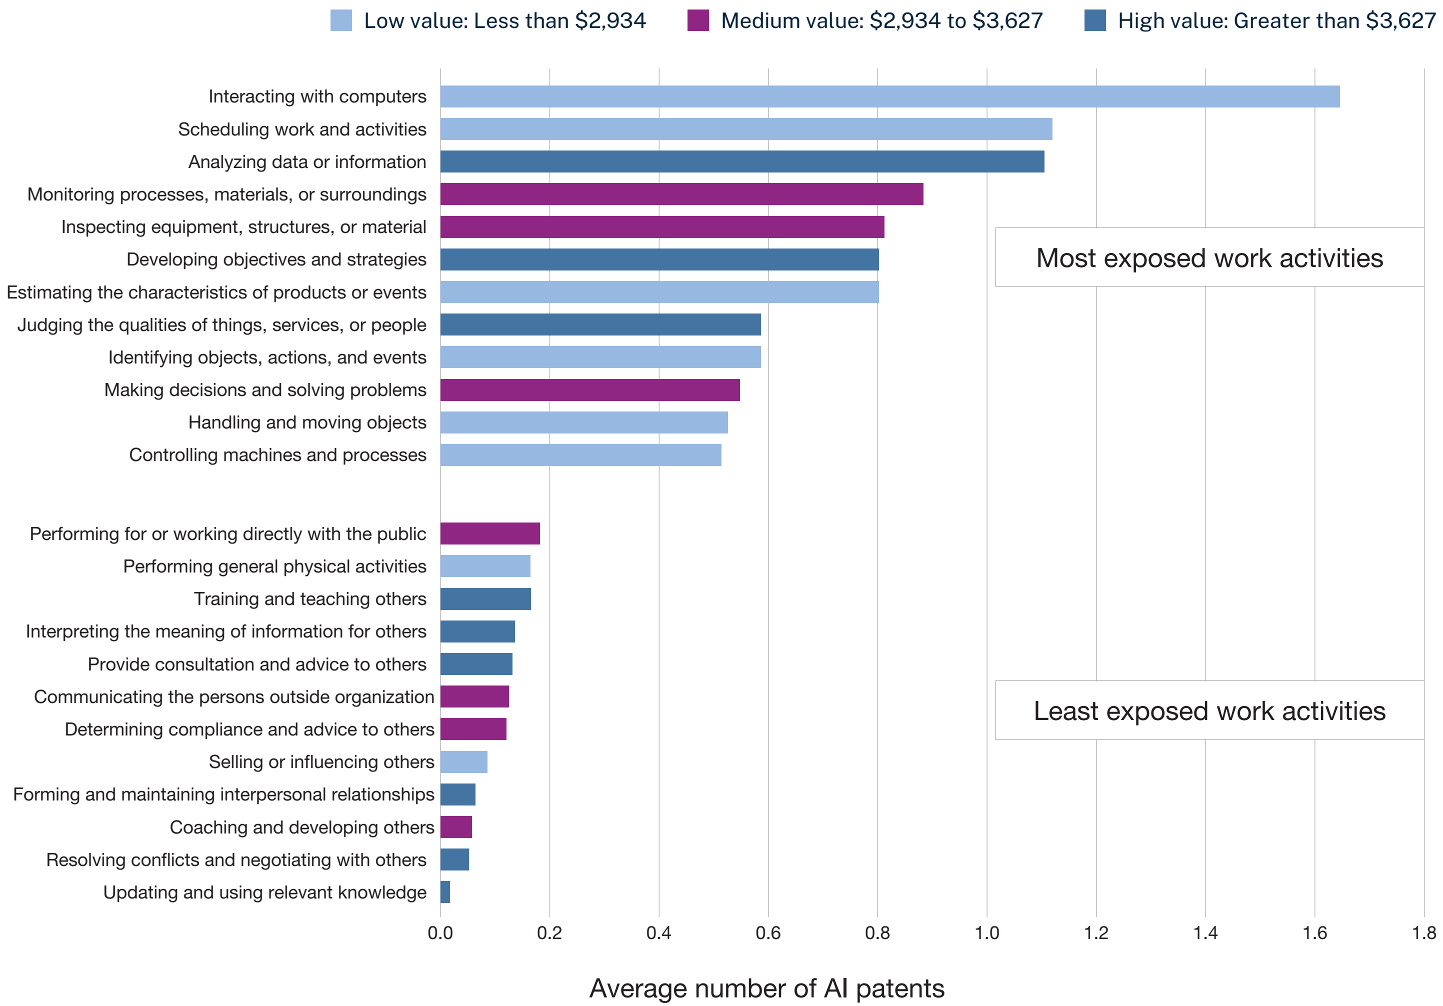

SOURCES: Authors' analysis of patent titles from U.S. Patent and Trademark Office (2020), task descriptions from U.S. Department of Labor (2024), and earnings data from U.S. Bureau of Labor Statistics (undated).

Recommendations

We make the following recommendations:

- Policymakers should consider investing in education and training programs that prioritize the development of interpersonal skills, complex problem-solving, and adaptability in the face of uncertainty and ambiguity.
- Education and training institutions should ensure that they offer training in soft skills. This may involve adding to existing courses and curricula or building in opportunities for learners to acquire and practice soft skills, for example, through collaborative project-based learning, presentations and public speaking assignments, and greater use of peer review and peer feedback.

Notes

[1] Task descriptions were drawn from the Department of Labor's O*NET database (U.S. Department of Labor, 2024), which contains descriptions of more than 19,000 tasks performed in 923 occupations; descriptions of AI capabilities were drawn from AI patents. We used only detailed occupations in the Standard Occupational Classification system because they can be mapped to wage data from the Occupational Employment and Wage Statistics (OEWS) (U.S. Bureau of Labor Statistics, undated). We identified AI patents using the U.S. Patent and Trademark Office's AI Patent Dataset, which classified all patents granted from 1976 to 2020 as AI or non-AI (U.S. Patent and Trademark Office, 2020). Our final sample consists of 16,817 tasks across 715 occupations from the May 2024 version of the O*NET database.

[2] O*NET contains ratings of importance for each task performed in an occupation (which are provided by a survey of workers in the occupation) and maps tasks to generalized work activities (GWAs) performed by workers across occupations and industries. For each task in an occupation, we constructed a task weight equal to the task's importance rating divided by the sum of importance ratings for all tasks in the occupation. To estimate the economic value of each task, we multiplied the task weight by the average earnings for the occupation, drawn from OEWS. To estimate the economic value of each GWA, we averaged the economic value of tasks that map to each GWA in each occupation and then across occupations.

References

Sytsma, Tobias, and Éder M. Sousa, *Artificial Intelligence and the Labor Force: A Data-Driven Approach to Identifying Exposed Occupations*, RAND Corporation, RR-A2655-1, 2023. As of August 30, 2024: https://www.rand.org/pubs/research_reports/RRA2655-1.html

U.S. Bureau of Labor Statistics, "Occupational Employment and Wage Statistics," webpage, undated. As of August 1, 2024: https://www.bls.gov/oes

U.S. Department of Labor, O*NET 28.3 Database, May 2024. As of August 1, 2024: https://www.onetcenter.org/database.html

U.S. Patent and Trademark Office, "Artificial Intelligence Patent Dataset," dataset, 2020. As of August 1, 2024: https://www.uspto.gov/ip-policy/economic-research/research-datasets/artificial-intelligence-patent-dataset

CHAPTER 16

Engaging the Federal Workforce in AI Implementation

Fernando Esteves, Susan M. Gates

Although the adoption of AI promises to enhance the productivity of workers and the organizations that employ them, there is growing recognition that effective adoption must be worker-centered. With more than 2.2 million civilian workers and a similar number of military personnel, the federal government is the largest U.S. employer. Operating across diverse missions and occupations, the federal government is an ideal testing ground for worker-centered AI adoption. In this essay, we describe what worker-centered AI adoption looks like, identify key implementation barriers in federal agencies, and suggest how federal agencies could build on existing guidance to support worker-centered AI adoption.

When AI adoption is worker-centered, workers (1) generate ideas about where work processes could be improved; (2) collaborate with AI tool developers to conceptualize, develop, test, and refine tools; and, once tools are implemented, (3) provide feedback on the tools and monitor the implications of their use (Leonardi, 2023; De Cremer, 2024; Bajwa et al., 2021). These steps help ensure that tool development resources are used efficiently, that tools meet the real-world needs of those who use them, and that tool use is ethical, accurate, and improves over time.

Unfortunately, barriers to worker-centered AI adoption exist at both the worker and organizational levels. At the individual level, key barriers include trust and gaps in expertise. Lack of trust largely stems from the fear that AI adoption will lead to job loss. Gaps in expertise include AI tool developers who do not understand the processes they are asked to support and process experts who lack expertise in AI tool development. At the organizational level, key barriers include resource constraints, siloization,[1] and a lack of appreciation for the importance of the worker perspective. In the federal context, resource constraints include not having enough of the resource (time, people, skills, money) and not having the flexibility to trade one type of resource for another—a barrier that is less salient for private sector employers.

These limitations stem from the fact that AI initiatives and AI tool developers are often funded and managed by organizational units that are separate from those doing the day-to-day work that the organization hopes to improve through AI tools. Siloization reinforces a tendency for organizations to prioritize hiring individuals with AI expertise while overlooking the essential contributions to tool development from those who have been performing the day-to-day work for years. Taken together, these organizational barriers make it difficult for federal agencies to create a sandbox in which those with technical skills and those with the needed process expertise can collaborate productively.

Creating such a sandbox requires planning and creativity. Reflections from AI experts, guidance issued by the U.S. Department of Labor (U.S Department of Labor, undated; U.S Department of Labor, 2024), the General Services Administration's Centers of Excellence initiative, and practical guidelines by the U.S. Department of Agriculture summarize some of the key workforce-related challenges to the worker-centered approach to AI adoption and recommend actions to address these barriers.

How to Clear the Hurdles

Organizations can address individual-level barriers through transparency and by supporting worker training. For instance, conducting workshops, feedback sessions, and pilot projects where workers test AI tools and provide insights during all phases before full deployment can be beneficial (De Cremer, 2020; Fountaine, McCarthy, and Saleh, 2019). Supporting collaboration between AI developers and nontechnical employees can help ensure that AI tools meet an organization's needs. This can be achieved by funding ground-up projects involving AI specialists and nonspecialist staff who grapple with AI's practical implications and limitations, or by pairing AI developers with nontechnical employees at specific stages of tool development to gain users' perspectives and ensure AI tools meet users' needs.

Robust communication mechanisms can help address both individual and organizational barriers. Input from those who understand and execute existing workflows is key for maintaining trust and for developing and improving AI solutions (U.S. General Services Administration, undated; U.S. Department of Agriculture, 2024). These mechanisms allow employees to propose projects, report issues, suggest improvements, and share experiences to help refine AI tools.

Given the complexity in defining strategies and coordinating actions, federal agencies are likely to face difficulties when implementing these suggestions, especially in engaging on-the-ground workers and fostering collaboration between AI tool developers and the workers whose efforts those tools are designed to support. For example, the U.S. Office of Management and Budget provides guidance on AI governance structures within federal agencies (Young, 2024), but it is extremely high level and does not compel organizations to consider worker-centered features. Additionally, OPM has developed a competency framework for AI positions and offers tools for skills-based hiring and the development of AI professionals. Although OPM acknowledges that "AI work is multidisciplinary and is not limited to one occupation" (OPM, 2024, p. 17), its framework primarily focuses on positions in which AI work constitutes at least 25 percent of a worker's responsibilities. These initiatives are steps in the right direction, but leveraging the expertise and experience of those who handle the day-to-day work is the crucial link to achieving effective, worker-centered AI adoption in federal agencies.

Conclusion

To inform the development of more-concrete supports that bridge the gap between AI efforts at the highest levels of an agency and the workers who perform the activities, we recommend that federal agencies begin tracking the success and failure of use cases, along with the administrative barriers they experience when implementing those use cases, tying those barriers to characteristics of the agency and the workforce involved in implementing the use case. Over time, this information can generate lessons about the relative advantages and disadvantages of certain management structures or features (e.g., the military personnel system or civilian personnel demonstration projects) when it comes to implementing certain AI use cases. With so much variation in the work environment and workforce management structures, there is much to learn from the experience of federal agencies. The Department of Labor and OPM could help bridge the knowledge gap by creating a clearinghouse of examples or case studies about the promising practices outlined above and strategies that agencies have used to address the administrative barriers to their use.

Notes

[1] *Siloization* refers to the organizational barriers to interactions between tool developers and with on-the-ground workers, as well as barriers to interaction between on-the-ground workers in different functional areas within the organization.

References

Bajwa, Junaid, Usman Munir, Aditya Nori, and Bryan Williams, "Artificial Intelligence in Healthcare: Transforming the Practice of Medicine," *Future Healthcare Journal*, Vol. 8, No. 2, July 2021.

De Cremer, David, *Leadership by Algorithm: Who Leads and Who Follows in the AI Era?* Harriman House, 2020.

De Cremer, David, "For Success with AI, Bring Everyone on Board," *Harvard Business Review*, May–June 2024.

Fountaine, Tim, Brian McCarthy, and Tamin Saleh, "Building the AI-Powered Organization," *Harvard Business Review*, July–August 2019.

Leonardi, Paul, "Helping Employees Succeed with Generative AI," *Harvard Business Review*, November–December 2023.

OPM—*See* U.S. Office of Personnel Management.

U.S. Department of Agriculture, *Framework for State, Local, Tribal, and Territorial Use of Artificial Intelligence for Public Benefit Administration*, April 29, 2024.

U.S. Department of Labor, "Artificial Intelligence and Worker Well-Being: Principles for Developers and Employers," webpage, undated. As of August 30, 2024:
https://www.dol.gov/general/AI-Principles

U.S. Department of Labor, "Artificial Intelligence and Equal Employment Opportunity for Federal Contractors," webpage, last updated April 29, 2024. As of August 30, 2024:
https://www.dol.gov/agencies/ofccp/ai/ai-eeo-guide

U.S. General Services Administration, "AI Guide for Government," webpage, undated. As of July 31, 2024:
https://coe.gsa.gov/coe/ai-guide-for-government/introduction/index.html

U.S. Office of Personnel Management, *Artificial Intelligence (AI) Competency Model*, April 2024.

Young, Shalanda D., "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," memorandum for the heads of executive departments and agencies, Office of Management and Budget, Executive Office of the President, March 28, 2024.

CHAPTER 17

Understanding and Addressing Resistance to the Adoption of New Technologies in DoD

Brandon Crosby, Tuyen Dinh

DoD is adopting AI to enhance capabilities across various domains (Hicks, 2024). It will likely face two significant obstacles that typically stand in the way of successful AI implementation. First, there is a general workforce resistance to AI because of concerns about job displacement, trust, and skill gaps (Xu et al., 2023). Second, DoD has historically experienced substantial resistance to new technologies, stemming from cultural inertia, risk aversion, and integration challenges (Birkler, Bracken, and Lee, 2021; Read et al., 2015). How can DoD address this likely resistance to AI?

In this essay, we aim to help DoD effectively implement AI by exploring resistance factors, proposing strategies, and considering DoD's unique context, ultimately providing guidance for successful AI adoption while prioritizing workforce well-being.

Factors Contributing to Resistance

As with any organizational change, employees' attitudes and beliefs play a critical role in shaping their acceptance and engagement with initiatives related to AI systems. To mitigate resistance and create a work environment that supports successful implementation, it is essential to examine both the psychological and human-computer interaction aspects of AI adoption. We used insights from *sociotechnical systems (STS) theory*, which views organizations as complex entities composed of interdependent social and technical subsystems, stressing that both subsystems must be optimized together for the best performance (Read et al., 2015; Walker et al., 2008).

First, we discuss three key factors contributing to resistance and then propose theoretically supported solutions to these challenges.

Fear of Job Displacement

A primary challenge in AI adoption is that employees fear their jobs might be replaced by automated systems (Frey and Osborne, 2013; Russo, 2020). Resistance is fueled by concerns about job security, reduced autonomy, and slower wage growth (Chowdhury, Link, and van Hasselt, 2022; Green et al., 2017).

To mitigate these concerns, AI systems could be introduced through a sociotechnical redesign process that makes job tasks both productive and human-centered, enhancing task meaningfulness, skill development, and autonomy. Integrating AI requires job design that influences meaningfulness, responsibility, and knowledge of results (Hackman and Oldham, 1976). Collaboratively designing work structures with employees and identifying tasks for automation could result in more-accurate task selection, enhanced understanding of new systems, and increased job satisfaction (Brachman et al., 2024; Waterson et al., 2015).

Lack of Trust in AI Systems

Trust in AI systems is crucial for their acceptance and continued use (McKnight et al., 2011). Complex and confusing technologies can lead to negative attitudes, affecting trust and willingness to use such technologies. High error rates and visibility of errors can also undermine confidence in AI's reliability, especially among employees who are familiar with computer science (Ryseff et al., 2022).

Research identifies two dimensions of trust in technology: human-like aspects (e.g., ability, benevolence) and system-like aspects (e.g., reliability, functionality) (Lankton, McKnight, and Tripp, 2015). Both types of trust are essential for positive attitudes and behavioral intentions toward AI use (Choung, David, and Ross, 2023). AI designers should ensure that the technology aligns with the desired level of human-like interaction to build trust. For example, conversational AI requires more human-like attributes to foster trust compared with such tools as Microsoft Word (Chandra, Shirish, and Srivastava, 2022). This alignment can enhance user engagement and adoption.

Skill Gaps

Employees are more likely to adopt new technologies when they believe that the tools are useful and easy to use (Choung, David, and Ross, 2023). If there is a mismatch between the workforce's existing skills and the skills required to operate new AI systems, employees may resist adopting the technology. Therefore, organizations should emphasize training and quickly adopt policies to prepare for AI usage. Rapid technological changes require employees to be adaptable and open to learning new skills.

Strategies to Address Resistance

Transparent Communication and Collaborative Implementation

Clear communication about AI goals, benefits, and impacts is vital to reducing resistance within DoD. Regular updates and opportunities for feedback build trust and engagement (Walker et al., 2008). Involving employees in AI design and implementation ensures that solutions align with their needs, fostering ownership and reducing resistance (Read et al., 2015).

Comprehensive Training and Support

Training programs could cover both technical skills and promote adaptability and continuous learning. Ongoing support—such as helpdesks, mentoring, and communities of practice—enhances confidence and proficiency in using AI (Hurley, 2018). Comprehensive training programs that empower employees to work with AI can boost acceptance and engagement (Marler, Liang, and Dulebohn, 2006). These programs would provide a clear rationale for the change, offer hands-on experience, encourage collaboration, and allow customization to meet individual needs.

Ethical and Safety Concerns

Developing governance frameworks and clear guidelines for AI use, along with rigorous testing and validation, is critical to ensuring ethical standards, enhancing system reliability, and safeguarding against potential misuse. Addressing ethical and safety concerns builds workforce trust and ensures that AI use aligns with DoD values (Coppersmith, 2021). Promoting transparency, minimizing algorithmic bias, and ensuring human oversight are essential for the responsible and effective implementation of AI systems.

Conclusion

AI adoption in DoD offers significant opportunities but also challenges, particularly workforce resistance because of concerns about job displacement, a lack of trust in AI, and skill gaps. By applying STS theory insights, DoD can understand and address these factors through a multifaceted approach that involves transparent communication, collaboration, training, and ethical considerations. Successful AI integration requires a sustained effort that recognizes the interdependence of social and technical factors while prioritizing the human dimensions of adoption. By leveraging STS principles and focusing on employee well-being, DoD can harness AI's potential while ensuring its workforce remains empowered and engaged.

References

Birkler, John L., Paul Bracken, and Gordon Lee, "Innovating Innovation in the U.S. Department of Defense Looking to Large Corporations for Inspiration," *Journal of Management Policy and Practice*, Vol. 22, No. 4, 2021.

Brachman, Michelle, Amina El-Ashry, Casey Dugan, and Werner Geyer, "How Knowledge Workers Use and Want to Use LLMs in an Enterprise Context," *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, 2024.

Chandra, Shalini, Anuragini Shirish, and Shirish C. Srivastava, "To Be or Not to Be . . . Human? Theorizing the Role of Human-Like Competencies in Conversational Artificial Intelligence Agents," *Journal of Management Information Systems*, Vol. 39, No. 4, 2022.

Choung, Hyesun, Prabu David, and Arun Ross, "Trust in AI and Its Role in the Acceptance of AI Technologies," *International Journal of Human–Computer Interaction*, Vol. 39, No. 9, 2023.

Chowdhury, Farhat, Albert N. Link, and Martijn van Hasselt, "Public Support for Research in Artificial Intelligence: A Descriptive Study of U.S. Department of Defense SBIR Projects," *Journal of Technology Transfer*, Vol. 47, No. 3, June 2022.

Coppersmith, Frank, "The Department of Defense AI Ethical Principles," *Scitech Lawyer*, Vol. 17, No. 3, March 2021.

Frey, Carl Benedikt, and Michael Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerization?" Oxford Martin School, University of Oxford, September 17, 2013.

Green, Francis, Alan Felstead, Duncan Gallie, and Golo Henseke, "Work Intensity in Britain: First Findings from the Skills and Employment Survey 2017," Cardiff University, 2017.

Hackman, J. Richard, and Greg R. Oldham, "Motivation Through the Design of Work: Test of a Theory," *Organizational Behavior and Human Performance*, Vol. 16, No. 2, August 1976.

Hicks, Kathleen, "Remarks by Deputy Secretary of Defense Kathleen Hicks on 'The Global AI Contest' at the Advantage DOD 2024: Defense Data & AI Symposium in Washington, D.C. (as Delivered)," speech, February 21, 2024.

Hurley, J. S., "Enabling Successful Artificial Intelligence Implementation in the Department of Defense," *Journal of Information Warfare*, Vol. 17, No. 2, Spring 2018.

Lankton, Nancy K., D. Harrison McKnight, and John Tripp, "Technology, Humanness, and Trust: Rethinking Trust in Technology," *Journal of the Association for Information Systems*, Vol. 16, No. 10, 2015.

Marler, Janet H., Xiaoya Liang, and James Hamilton Dulebohn, "Training and Effective Employee Information Technology Use," *Journal of Management*, Vol. 32, No. 5, October 2006.

McKnight, D. Harrison, Michelle Carter, Jason Bennett Thatcher, and Paul F. Clay, "Trust in a Specific Technology: An Investigation of Its Components and Measures," *ACM Transactions on Management Information Systems*, Vol. 2, No. 2, June 2011.

Read, Gemma J. M., Paul M. Salmon, Michael G. Lenné, and Neville A. Stanton, "Designing Sociotechnical Systems with Cognitive Work Analysis: Putting Theory Back into Practice," *Ergonomics*, Vol. 58, No. 5, 2015.

Russo, Amanda, "Recession and Automation Changes Our Future of Work, but There Are Jobs Coming, Report Says," press release, World Economic Forum, October 20, 2020.

Ryseff, James, Eric Landree, Noah Johnson, Bonnie Ghosh-Dastidar, Max Izenberg, Sydne Newberry, Christopher Ferris, and Melissa A. Bradley, *Exploring the Civil-Military Divide over Artificial Intelligence*, RAND Corporation, RR-A1498-1, 2022. As of September 10, 2024: https://www.rand.org/pubs/research_reports/RRA1498-1.html

Walker, Guy H., Neville A. Stanton, Paul M. Salmon, and Daniel P. Jenkins, "A Review of Sociotechnical Systems Theory: A Classic Concept for New Command and Control Paradigms," *Theoretical Issues in Ergonomics Science*, Vol. 9, No. 6, 2008.

Waterson, Patrick, Michelle M. Robertson, Nancy J. Cooke, Laura Militello, Emilie Roth, and Neville A. Stanton, "Defining the Methodological Challenges and Opportunities for an Effective Science of Sociotechnical Systems and Safety," *Ergonomics*, Vol. 58, No. 4, 2015.

Xu, Wei, Marvin J. Dainoff, Liezhong Ge, and Zaifeng Gao, "Transitioning to Human Interaction with AI Systems: New Challenges and Opportunities for HCI Professionals to Enable Human-Centered AI," *International Journal of Human–Computer Interaction*, Vol. 39, No. 3, 2023.

CHAPTER 18

Principles on AI Implementation for Federal Leaders

Rachel Slama, Nelson Lim, Douglas Yeung, eds.

There are several principles for the leaders charged with integrating AI into their workforce that emerge from this collection of essays.

AI integration into the civilian and military workforce will happen within systems. Leaders are tasked with integrating new tools into existing workplace cultures, workflows, and operating procedures. They will make key decisions about which new tools, if any, to acquire; how to integrate them into existing systems; and the risks and benefits to consider. They will benefit from using our taxonomy outlining the technical, ethical, legal, economic, social, and existential risks involved in AI adoption and our primer on the AI ecosystem organized by supply (e.g., AI developers) and demand (e.g., AI users) sides with a summary of the types of skills and competencies required of each group.

Much AI innovation and experimentation is already underway across federal agencies, presenting opportunities to leverage learning from these early adopters. Three common use cases of AI-powered tools across sectors are (1) the use of AI to process reams of complex technical documents (in this case, for FEMA); (2) the use of AI-powered chatbots to improve user experience (e.g., patient-provider communication in health care); and (3) the use of AI as a job aid, as in applying AI to HR tasks to process high-volume text data (such as candidate resumes), to standardize position descriptions, and to summarize employee survey results. Federal agencies can participate in the AI Community of Practice to leverage best practices, tools, and resources across agencies (U.S. General Services Administration, undated).

In tandem with AI talent recruitment, federal agencies will need to reskill the current workforce through the adaptation of existing education and training pipelines. Leaders will need to upskill and retrain the existing workforce in tandem with hiring new AI talent. For example, DoD contends with obstacles in hiring from the private sector, so it must focus significant efforts on upskilling and reskilling service members and civilians with basic domain knowledge and digital skills while also experimenting with innovations in hiring and retention. The essays in this collection provide a rationale and considerations for investing in AI tools and infrastructure, which is critically important for government agencies that have limited human and technology resources compared with the private sector.

The essays also describe ways in which K–12, postsecondary, and military training institutions will need to adapt to meet new education and workforce training needs. Agencies should also evaluate the appropriateness of partnering with external providers that have established credentialing systems.

Human workers will play a central role in the success of AI integration for the foreseeable future. Federal agencies must recruit and retain diverse domestic AI talent. These essays discuss mechanisms to recruit and retain workers, such as compensation, work environment, and training opportunities. The authors also offer insight into the types of skills that workers need to remain employed as jobs change and automate, and they discuss strategies for overcoming barriers to technology adoption, such as worker resistance, skill gaps, and integration challenges. Perhaps most importantly, the essays give concrete guidance on how to engage workers in AI implementation because AI integration will not succeed without their buy-in.

References

U.S. General Services Administration, "Community of Practice: Artificial Intelligence," webpage, undated. As of August 30, 2024: https://coe.gsa.gov/communities/ai.html

About the Authors

Elie Alhajjar is a senior scientist at RAND with a strong technical background in AI and ML. His current work revolves around STEM workforce development on the local and federal levels. He holds a Ph.D. in applied mathematics.

Rushil Bakhshi is a technical analyst at RAND who works mostly for the RAND Homeland Security Research Division. He has also done work in the biotechnology, health, and community-centered technology spaces. He holds an M.S. in biomedical engineering.

Julia Bandini is a behavioral/social scientist at RAND whose work focuses on palliative and end-of-life care, including the integration of technology in serious illness care. She holds a Ph.D. in sociology.

Dwayne M. Butler is a senior management scientist at RAND and professor of policy analysis at the Pardee RAND Graduate School. He holds a Ph.D. in organizational psychology.

Avery Calkins is an economist at RAND whose work focuses on issues related to recruiting, retention, and compensation in the U.S. military and DoD civilian workforce. She holds a Ph.D. in economics.

Angela K. Clague is an associate behavioral/social scientist at RAND who works mostly for RAND Project AIR FORCE, concentrating on talent management. She holds a Ph.D. in sociology.

Arianne Collopy is a senior technical analyst at RAND who works on capability development and systems design topics for DoD and DHS. She holds a Ph.D. in design science.

Brandon Crosby is a social scientist at RAND who studies the employee life cycle, diversity, recruitment, and retention issues throughout DoD and DHS. He holds a Ph.D. in organizational psychology.

Brandon De Bruhl is a senior technical analyst at RAND who works on issues related to AI, financial technology, and federal budget policy. He is completing a Ph.D. in public policy.

Tuyen Dinh is an adjunct policy analyst at RAND who works on organizational change leadership and workplace training and development topics across various areas, including health care, defense, intelligence, and military. She is a Ph.D. candidate in psychology.

Fernando Esteves is a Ph.D. candidate at the Pardee RAND Graduate School. His research focuses on labor, climate change, and developing-nation economics.

Susan M. Gates is a senior economist at RAND, specializing in personnel economics with a focus on talent and resource management issues for DoD. She is also a professor at the Pardee RAND Graduate School. She holds a Ph.D. in economics.

Charles A. Goldman is a senior economist at RAND who studies the organization and performance of postsecondary education institutions and their interactions with labor markets, primarily in the education and national security domains. He holds a Ph.D. in economic analysis and policy.

Timothy R. Gulden is a senior policy researcher at RAND who specializes in geographic analysis, complex systems, and social simulation. He holds a Ph.D. in public policy.

Wenjing Huang is a senior behavioral and social scientist at RAND who applies innovative modeling in her workforce, well-being and personnel management research. She holds a Ph.D. in quantitative psychology.

Kelly Hyde is an associate economist at RAND who specializes in quantitative data analysis with a broad tool kit of methods, including econometrics, machine learning, natural language processing, and geospatial techniques. He holds a Ph.D. in economics.

Rita T. Karam is a senior policy researcher at RAND who concentrates primarily on postsecondary education and workforce projects. She holds a Ph.D. in educational policy.

Tracy C. Krueger is a senior behavioral scientist at RAND working primarily on talent management topics within RAND Project AIR FORCE. She holds a Ph.D. in industrial-organizational psychology.

Jonah Kushner is an assistant policy researcher at RAND and a Ph.D. candidate at the Pardee RAND Graduate School. His research interests include postsecondary education, workforce development, and the future of work.

Mary Lee is a senior mathematician at RAND. Her work for RAND Project AIR FORCE includes topics on emerging technologies, force modernization and employment, and workforce development. She holds a Ph.D. in mathematics.

Nelson Lim is a senior social scientist at RAND who directs the Workforce, Development, and Health Program in RAND Project AIR FORCE. His recent work explores the application of emerging technologies to human resource management. He holds a Ph.D. in sociology.

Maria C. Lytell is a senior behavioral scientist at RAND. She works across RAND's military, defense, and homeland security research divisions. She holds a Ph.D. in industrial-organizational psychology.

Laurie T. Martin is a senior policy researcher at RAND. Her work examines topics related to access to health and social services, patient-provider communication, and the quality of care and health outcomes. She holds a Sc.D. in maternal and child health.

Nikolay Maslov is an assistant policy analyst at RAND and a student at the Pardee RAND Graduate School. His areas of interest and study include technology policy, security cooperation, and emergency response/operations. He holds an M.S. in national cybersecurity studies and an M.A. in mathematics.

Michael G. Mattock is a senior economist at RAND. Much of his work has focused on helping decisionmakers evaluate possible changes to compensation, incentive pays, and retirement for both military and civilian personnel. He holds an M.S. in social science.

Skye A. Miner is an associate behavioral and social scientist at RAND. Her work as a sociologist and bioethicist focuses on emerging issues in health care, including technologies' ethical, social, and legal implications. She holds a Ph.D. in sociology.

Alvin Moon is an associate mathematician at RAND. His work includes topics in emerging technology, workforce development, and modeling. He holds a Ph.D. in mathematics.

Ojashwi Pathak is a summer associate at RAND, focusing on machine learning and talent management in the U.S. Air Force. She is also a Ph.D. candidate.

Neeti Pokhriyal is an information scientist at RAND with expertise in research and policy aspects of developing and operationalizing responsible AI technologies. She holds a Ph.D. in computer science.

Carter C. Price is a senior mathematician at RAND and works on quantitative policy analysis on a variety of topics. He also teaches the use of advanced technologies for policy analysis at the Pardee RAND Graduate School. Price has a Ph.D. in applied mathematics.

Sean Robson is a senior behavioral scientist at RAND and associate director of the Workforce, Development, and Health Program in RAND Project AIR FORCE. He holds a Ph.D. in industrial-organizational psychology.

Srikant Kumar Sahoo is an assistant policy researcher at RAND whose work concentrates primarily on AI in education. He is pursuing his M.Phil in policy analysis from Pardee RAND Graduate School.

Morgan Sandler is a technical analyst at RAND, specializing in AI, particularly its applications in warfare and the workplace. He is currently pursuing a Ph.D. in computer science.

Anton Shenk is a quantitative research assistant at RAND whose focus is on the economics of emerging technologies. He has contributed to studies about AI-enhanced biological weapons and the future of economic warfare. He holds B.S. degrees in mathematics and economics.

Rachel Slama is the director of Labor and Workforce Development in RAND Education and Labor. Her work focuses on strengthening the STEM talent pipeline through the development of technology and workforce training infrastructure. She holds a Ph.D. in education policy analysis.

Éder M. Sousa is a computational scientist at RAND. His work focuses on the impact of advanced technologies in the workforce and defense domains, modeling and simulation, and space policy and acquisition. He holds a Ph.D. in aeronautics and astronautics.

Tobias Sytsma is an economist at RAND and a professor of policy analysis at the Pardee RAND Graduate School. He studies the economics of AI, globalization and international trade, and workforce development. He holds a Ph.D. in economics.

Ivy Todd is a policy analyst in RAND's Education and Labor division. Her research focuses on the coherence, implementation, and effectiveness of education initiatives. She holds an M.S. in public policy and management.

John Vahedi is a technical analyst at RAND who works on AI policy, tools, and research. He holds an M.S. in applied mathematics and an M.A. in STEM education.

Jessie Wang is an economist at RAND and a professor of policy analysis at the Pardee RAND Graduate School. She studies individual and aggregate responses to shocks and unintended consequences of policy interventions. She holds a Ph.D. in economics.

Madison Williams is a quantitative research assistant at RAND. She previously served as an AI researcher and technical writer at Virginia Tech as part of the Commonwealth Cyber Initiative. She holds a B.A. in political science.

Jody Chin Sing Wong is an associate policy researcher at RAND who studies decisionmaking under uncertainty, primarily in the areas of health, risk, and the environment. She holds a Ph.D. in communication.

Douglas Yeung is an associate director of the Management, Technology, and Capabilities Program in the RAND Homeland Security Research Division, a senior behavioral scientist, and a professor at Pardee RAND Graduate School. His recent work has explored how policymakers can use insights from emerging technologies. He holds a Ph.D. in psychology.

Artificial intelligence (AI) is poised to significantly affect the American workforce—both civilian and military personnel—through job displacement, augmentation, and the need for widespread upskilling. President Biden's October 2023 executive order on AI emphasizes the government's commitment to upskilling the federal workforce in understanding, adopting, deploying, and using AI. Many federal agencies and U.S. Department of Defense entities have published AI guidance documents. Congress is also exploring the implications of advancements in AI in both the general and federal U.S. workforces.

This publication is intended to inform the policymakers and leaders who are tasked with preparing civilian and military workers to create, use, and deploy AI in their jobs. The essays in this publication provide overviews of technical and organizational issues, challenges, and actionable insights to help organizations effectively integrate AI and equip personnel with AI-related skills.

$48.00

ISBN-10 1-9774-1406-0
ISBN-13 978-1-9774-1406-9

54800

9 781977 414069

PE-A3414-1