AI UPSKILLING: THE FUTURE OF WORK IN THE AGE OF AI

After-Action Summary

AI Adoption Initiative – Carnegie Mellon University Workshop

Pittsburgh, June 24–25, 2026

Carnegie Mellon University

AIAI
THE AI ADOPTION INITIATIVE

EXECUTIVE SUMMARY

This report synthesizes findings from a two-day workshop co-hosted by the AI Adoption Initiative and Carnegie Mellon University in Pittsburgh on June 24–25, 2026. More than seventy experts from government, industry, academia, and civil society participated, drawing on perspectives from North America, Europe, Asia, South America, Australia, and multilateral institutions.

The workshop took the AIAI pre-read paper, "The AI Labor Stack: Rethinking the Workforce from the Ground Up," as its analytical foundation. The paper argued that effective AI workforce policy requires attention to three distinct types of workers: innovators who develop AI infrastructure and models; facilitators who deploy, integrate, and adapt AI for specific organizational contexts; and users whose work is transformed by applying the technology. The paper's central finding was that facilitators represent the "missing middle" of national AI strategies, receiving less policy attention than either frontier innovation or broad public AI literacy.

Workshop participants agreed that the AI Labor Stack framework was a useful lens for workforce skills policy, but also extended it in a number of ways. For example, participants noted that the facilitator layer is broader than the pre-read suggested. In addition to more technically-focused roles, the facilitator layer also includes organizational leaders and middle managers who can gain buy-in for adoption, and integrate expert domain knowledge and organizational insight into efforts to enable domain-specific AI use.

Participants also flagged concerns that AI might make it harder for less experienced workers to cultivate deeper expertise and judgement that have typically been built through performing less glamorous work that may increasingly be offloaded to AI. Employers and workforce skills policies may need to find other ways to help junior workers cultivate these skills, while improving communication between students, educators, and employers about the workforce's evolving skills needs.

UNPACKING THE FACILITATOR LAYER

Workshop participants supported the pre-read paper's central contention, that the facilitator layer is the "missing middle" of national AI workforce strategies. Discussions focused on extending the concept, including developing a more nuanced understanding of who facilitators actually are.

In addition to the systems architects, data engineers, cloud engineers, and AI site reliability engineers identified in the pre-read, workshop discussions noted the importance of a second facilitator category: middle managers and other organizational champions. These workers often lack deep technical skills, but they possess valuable and hard-to-replicate knowledge about internal workflows and business processes. This means they are well situated to bridge the gap between a capable AI tool and the messy reality of how a specific firm operates. Middle managers may also have greater organizational standing to gain leadership buy-in for AI adoption than end-users, while also facilitating a culture of trust and experimentation at the workforce level. Several participants argued that these types of manager-facilitator roles, which rely on softer skills that may not currently be well-incentivized or captured in more technology-focused AI skills credentialing systems, could be an important component of workforce strategies.

SCIENTIFIC DISCOVERY AS A CASE STUDY FOR FACILITATORS

An in-depth discussion of the scientific workforce offered a parallel argument. Participants argued that AI-enabled science depends on the interaction of several roles: innovators who build models and tools; domain experts who understand the scientific questions, data and methods; and facilitators who help connect tools, data infrastructure, workflows, governance and research practice. The closer these roles are brought together, the more likely AI is to support genuine scientific discovery rather than simply accelerate existing processes.

Participants also noted that humanities and social-science disciplines can contribute important capabilities to this facilitator function, including critical thinking, contextual judgment, ethics, interpretive skill and reflection on how knowledge is produced. The broader point was not that scientific facilitators must come from any one discipline, but that AI-enabled science will require interdisciplinary teams that preserve and value translational capabilities alongside technical and domain expertise.

The discussion also raised a related concern about institutional incentives. If funding, publication metrics or organizational structures reward visible outputs while undervaluing data stewardship, validation, translation, ethics and other enabling work, the facilitator layer in science may be weakened. Participants suggested that science policy should pay closer attention to how these roles are aligned within teams, funding programs and research institutions.

The workshop also surfaced a potential tension for skills policy. The pre-read emphasized modular, stackable and updateable training pathways as one mechanism for building facilitator capacity. Discussions were broadly consistent with this direction, with participants considering examples such as short courses, industry certifications, skills savings accounts and other mechanisms that would allow workers to update their skills without leaving the labor market for extended periods.

At the same time, participants noted that short, applied training pathways work best when learners already have sufficient baseline judgment, domain knowledge and practical experience to apply what they learn. Several discussions raised concerns that AI could make it harder for students and early-career workers to acquire that foundation if it absorbs the difficult, repetitive or lower-status tasks through which judgment has historically been developed. Participants described this as a version of the “struggle” problem: the very tasks organizations may be tempted to automate can also be part of how workers learn.

This reinforces the case for a more detailed facilitator taxonomy. Mapping facilitator roles by sector, skill profile, career stage and delivery model, including whether roles are typically performed in-house, by domestic third parties or through global services, was the research priority that participants returned to most often. This research would help policymakers understand not only where facilitator capacity is needed, but how the pipeline into those roles may be changing as AI adoption advances.

AI ADOPTION AS AN ORGANIZATIONAL AND CULTURAL CHALLENGE

Workshop participants noted an additional social, cultural, and organizational dimension to the AI adoption challenge. Although user figures floated by leading AI labs suggest wide uptake of AI, this has yet to translate into noticeable increases in total-factor productivity. A key insight from AIAI's 2024 paper, "AI, Everywhere, All at Once" was that broad-based productivity gains will require embedding AI across a wide range of economic sectors and inside both large companies and small and medium enterprises (SMEs). This will depend not only on advancements underlying AI capabilities, but on businesses re-designing workflows and business processes for AI.

Building on the framework presented in the pre-read paper, the workshop discussion noted that this type of cultural and organizational change is likely to require sustained engagement between facilitators and users, and that workforce skills policies should account for this.

Participants noted that worker anxiety may limit willingness to use AI even where skills exist. Fear of being blamed for AI-driven mistakes, uncertainty about what AI tools can and cannot do reliably, and fears about job displacement may lead workers to underuse tools they nominally have access to. One industry participant noted that internal sandboxes where employees could experiment without the pressure of observation could help reduce this kind of friction.

Participants noted a structural gap in adoption between large firms and small and medium enterprises (SMEs). Large technology companies invest heavily in internal upskilling and have the organizational capacity to manage complex technology transitions. SMEs, which employ the majority of workers in most economies, lack comparable resources, internal expertise, and risk tolerance. Policy design frequently defaults to large-firm models that do not transfer at SME scale. Participants from multiple countries noted that peer learning through trusted intermediaries, such as chambers of commerce, industry associations, and community organizations consistently outperforms formal institutional programs, because it operates at a speed and scale SMEs can absorb. Government support for these types of intermediaries may help accelerate skills acquisition and adoption among SMEs and other groups that may not traditionally have access to the same resources as larger firms.

BUILDING FACILITATOR CAPACITY

Two ideas emerged as particularly promising for building facilitator capacity at scale. India's program to open hundreds physical AI and data labs and to establish sectoral centers of excellence that link industry, academia, and government around domain-specific curricula and use-case inventories attracted significant interest for its potential to help build a facilitator pipeline that includes specific sector expertise.

Participants also discussed the potential for AI prize competitions to be adapted to help address workforce skills and adoption challenges. Redirecting some of the prize money that is currently being aimed at technology breakthroughs could help facilitate growth of the facilitator or user layers, which may be more important for adoption.

GENERATIONAL CONSIDERATIONS

The pre-read paper highlighted a potential challenge in ensuring that the next generation of workers can acquire the judgement and expertise that employers value in higher-level roles, if AI automates the entry-level tasks through which workers historically developed these capabilities. **The workshop discussion extended this idea in three directions:**

1. The first is what participants called the struggle paradox. While businesses have incentives to use AI to replace difficult, unglamorous, or repetitive tasks, such tasks have often been a part of the "struggle" that helps students and early-career workers acquire deeper expertise. Participants noted the need for improved communication between employers and educators about longer-term skills needs, and how developmental experiences that are eventually replaced by AI could be preserve or replace in other context. There was lively debate about whether this was an issue for policymakers, or whether market forces would be sufficient to address this challenge.

2. Participants also discussed teacher education as an ongoing priority. Several noted that the education environment remains uneven, with some institutions experimenting actively and others lacking clear guidance. While many teachers have now had some exposure to AI training, participants stressed that one-off training will not be enough. Teachers will need continuing professional learning to understand where AI can help, where it remains unreliable, and how classroom practice, assessment and student judgment should evolve.

3. A more domain-specific idea was the apprenticeship problem in science. Agentic AI systems are decomposing scientific tasks into component steps and absorbing some forms of work that have historically helped junior researchers develop tacit scientific judgment. Agentic systems may produce results that make scientists look more productive on certain metrics (such as numbers of peer-reviewed articles published), without producing a proportional uplift in scientific understanding. Participants argued this issue deserved further attention and discussion amid ongoing shifts in science funding and institutional incentives.

The workshop produced concrete proposals for preserving career pipelines. Skills savings accounts, modeled on US-style health savings accounts, could create a tax-efficient opportunity for employers and employees to contribute resources to fund continuous upskilling; a bill on this was recently introduced in the US Congress.

Programs that require structured work experience, with government support for SME placements, would give younger workers access to practical experience that helps them build domain expertise and judgement, while also helping companies access younger workers' technical skills and more up-to-date knowledge of AI tools. Project-based learning programs, where companies that do not have the resources to host interns or work placement students in-house provide real problems for student teams to solve using AI tools, can help students gain experience and build project portfolios that can serve as labor market signals without requiring physical placements.

Competency frameworks that combine such portfolio evidence with testing, rather than relying on either alone, would give employers better signals about what applicants actually know, replacing the credentialing inflation that has made many existing credentials poor proxies for capability.

IMPLICATIONS FOR POLICYMAKERS

The pre-read identified labor market data and workforce skills measurement as critical gaps for effective policymaking. Without metrics for AI adoption at the firm and worker level, governments may lack the ability to discern how skill needs are changing as AI adoption accelerates and the underlying technology continues to evolve. Developing firm- and worker-level adoption metrics analogous to what existing technology-use surveys provide could help improve workforce policy design.

Participants were more skeptical of top-down mandates. Instead of trying to develop one-size-fits-all national strategies for workforce skills, governments may be better off focusing on reducing information asymmetries and creating feedback loops that help companies and employees understand evolving needs. By concentrating efforts on data, standards, and enabling infrastructure; governments can avoid overly-prescriptive solutions that could quickly become out of date.

The session comparing international approaches reinforced the pre-read's strategy matrix argument: countries occupy very different positions in terms of their access to facilitators and their relative labor costs. A high-income country with strong existing facilitator capacity is likely to need different policies than a lower-income economy focused on keeping its services exports globally competitive.

The discussion also reinforced that AI workforce policy should include transition resilience for workers whose roles change or disappear as firms reorganize work around AI. Some workers will use AI to become more productive in their current jobs; some may move into new facilitator roles; and others may face disruption. Participants pointed towards policy packages that combine labor-market dynamism with worker security, including retraining, transition services, income support and pathways back into work.

Participants also noted the importance of digital connectivity for workforce skills strategies. Programs focused on building AI literacy or strengthening the facilitator layer that do not coordinate with broadband deployment may fail in rural areas or in countries where populations struggle to access digital services.

PRIORITIES FOR FUTURE RESEARCH

The workshop produced several priorities for follow-on work.

There was general agreement that building a more detailed facilitator taxonomy should be a high priority for AIAI. Mapping facilitator roles by sector, required skills, job role, and whether they are likely to be provided in-house, by domestic third-parties, or through the global trade in services would help improve the AI Labor Stack framework's usefulness for policymakers.

Mapping how existing workforce skills policies map to the innovator, facilitator, and user layers of the AI Labor stack was another clear priority. Several participants argued this is a near-term deliverable requiring no new data collection, only a systematic review of existing program evaluations.

Strengthening the empirical baseline on AI adoption, tasks and labor-market change was another clear priority. Participants repeatedly noted that policymakers and economists are being asked to make decisions under conditions of high uncertainty, while current data sources remain too slow, incomplete or poorly structured to support confident forecasting. Follow-on work could explore how industry data, job postings, employer surveys, course content, hiring outcomes and public-sector administrative data might be combined to support better policy design.

Participants discussed hosting additional region- and sector-specific workshops to develop a more nuanced understanding of skills issues across different geographies and domains of economic activity.

They also suggested a deeper dive on scientific workforce issues.

This report was prepared from session notes taken during the workshop, which took place under the Chatham House Rule. The report reflects the content of discussions, but does not represent the official views of AIAI or any other participating organization.