MAPPING THE MISSING MIDDLE IN THE AI WORKFORCE: FACILITATORS & THEIR ROLE IN AI ADOPTION

Consultation Paper

September 2026

By Lindsey Lacey, Dustin Ferrone, Christophe Combemale

With foreword by Ed Teather

Carnegie Mellon University

AIAI

THE AI ADOPTION INITIATIVE

CONTENTS

Acknowledgments:

We are grateful to many colleagues including Claude Dinsmoor (FANUC American, retired), Lisa Masciantonio (ARM Institute), Peta Skujins (Future Skills Organization Australia), Ajit Manocha (SEMI), Shari Liss (SEMI) and Jessica Montgomery (Cambridge University) for their guidance and support through interview participation or correspondence. We also extend our gratitude to all participants in the June AI Adoption Initiative at Carnegie Mellon University, especially those who participated in our breakout sessions and contributed written feedback.

FOREWORD

Artificial intelligence represents a once-in-a-generation opportunity to raise productivity across our economies, businesses, public services and workforce at a time when countries urgently need it. Since the Great Financial Crash, the G7 has been in a near-universal productivity slowdown at 1% annual productivity growth, compared to 2% in the decade preceding 2008. Ageing societies, combined with growing budgetary pressures for national defence, healthcare and social care costs, means that the pre-crash growth rate needs to be urgently restored.

However, an AI adoption policy agenda focused solely on the goal of productivity feels insufficient in today's political economy. It needs to be married to a pro-worker agenda that also encompasses the goals of participation - so that the opportunities for workers to acquire new skills and secure higher wage premiums are shared - and resilience, so that workers are supported when some jobs may change or disappear. Done right, these goals can be mutually reinforcing.

That was the argument of our last paper, *The AI Labor Stack: Rethinking the Workforce from the Ground Up*. The AI Labor Stack introduced three layers of the AI workforce: the innovators who build frontier AI models; facilitators who deploy, integrate and govern AI in real-world settings; and users whose work changes as AI is adopted. Each layer presents opportunities for workers to participate in the AI economy, and each requires a different policy response.

Too often however, national AI strategies have focused on frontier innovators or broad user literacy, while overlooking the facilitators - the "missing middle" between AI capability and adoption.

The history of general-purpose technologies shows why facilitators matter. In the second industrial revolution for example, greater access to facilitators - as opposed to leading on the scientific frontier - was how the U.S overtook the UK and Germany as the world's pre-eminent economic power in its application of machine tools. For the UK's part, it had trained too few in mechanical engineering, while Germany did deliver the numbers, but offered a curriculum that was too theoretical and too divorced from practical applications. U.S training institutions by comparison were much more geared towards practical applications and real-world exercises, and creating this large cohort of mechanical engineers ready to apply that GPT across the economy.

This paper develops the facilitator layer for AI in more detail across three overlapping domains: technical, social and infrastructural. Technical facilitators integrate models, data and software into usable systems. Social facilitators - the managers, domain experts and governance professionals - build trust, reshape processes and connect tools to organisational needs. Finally, infrastructural facilitators provide the compute, power and connectivity needed to serve inference workloads.

What may be surprising is that facilitator roles are already a significant driver of the U.S AI economy, going well beyond the oft-discussed capex for training frontier AI models. While the *Economist* recently estimated that AI has created five times more jobs than it has displaced since mid-2023, our analysis finds that 60 to 80% of AI-related job postings by U.S S&P 500 companies are already facilitator-level roles.

Postings signal demand rather than hires, but the direction is striking: U.S firms are investing not only at the frontier, but in the workforce needed to translate capability into use.

Meeting that demand will require both programmatic and institutional responses. Programs build specific capabilities through training, apprenticeships and practical experience. Institutions provide the continuing capacity to identify changing demand, bring employers and training providers together, update provision and connect skills development to adoption. Governments can support both.

The case studies in this paper illustrate how different interventions can build capabilities across the three facilitator domains. Singapore's AI Apprenticeship Programme for example combines intensive training with industry projects to develop technical facilitators, supported by a stipend that reduces financial barriers to participation. Ireland's Telecommunications and

Data Network Technician Apprenticeship, developed with TU Dublin and employers, combines employment-based learning with classroom and laboratory training to build connectivity skills relevant to infrastructural facilitation, although the programme is not AI-specific. Germany's Mittelstand-Digital illustrates how an adoption intermediary can develop social and technical facilitation, helping smaller businesses understand new tools, learn from peers and work through the organisational changes needed to adopt them.

An institutional approach gives these efforts the capacity to evolve. As we argued in *The AI Labor Stack*, policy should begin with an empirical baseline; provision should be organised around local institutions and sectoral needs; and learning should be modular, stackable and employer-recognised, with feedback loops that keep it aligned with changing demand.

This is intentionally an interim paper, published for consultation rather than as a final framework. Over the coming months, we will refine the definitions and profiles of social, technical and infrastructural facilitation, deepen the sectoral evidence and develop the measurement agenda for facilitator demand and supply.

One implication of the AI Labor Stack is that not every place seeking to benefit from AI need, or should want, to look like Palo Alto or King's Cross. For many cities and regions, the opportunity lies in building facilitator capacity around the institutions and sectors that already shape the local economy. Done well, investing in facilitators can make AI adoption a source not only of national productivity growth, but also of wider opportunity and renewed pride in local communities.

Ed Teather

Global Lead, AI Policy Research and Partnerships,
Amazon Web Services

Affiliated Researcher, AI Adoption Initiative

INTRODUCTION

Artificial intelligence has shown strong evidence of being a general-purpose technology (GPT), akin to technologies such as the steam engine and the internet in its reach and ability to reshape productivity, global comparative advantages, and the structure of labor.

In *The AI Labor Stack: Rethinking the Workforce from the Ground Up*, we introduced the concept of an AI labor pyramid: the adoption and diffusion of AI relies upon three distinct groups of workers: innovators who drive the technological frontier, users who deploy the technology broadly, and facilitators who deploy, integrate, and enable the use of artificial intelligence (or other GPTs), and bridge the gaps between innovators and users.

This report unpacks the facilitator layer, focusing on the workers who translate cutting-edge technologies into resources for creating broad economic gains. We propose three different domains of facilitation: social facilitation, technical facilitation, and infrastructural facilitation. All three domains serve important functions in the broader ecosystem of AI adoption, and we offer suggestions for how governments can measure the growth in demand for facilitator roles across their economies. Using the frameworks of the three domains of facilitation and the cycle of facilitation, we offer a preliminary empirical approach to measuring facilitator capacity-building. We use a novel identification method for AI-related job postings (Ferrone, Combemale, and Fleming 2026) to evaluate how firms are building human capital around AI adoption and deployment. We apply our facilitator taxonomy to 94 thousand AI-related job postings made by firms in the S&P 500, and identify emerging AI job types. We find that facilitator hiring varies from sector to sector, and show how AI-related jobs map to clear social, technical, and infrastructural niches.

Productivity gains from AI are not automatic; instead, they depend on the strength of the facilitator layers and on complementarities between the facilitator-user and facilitator-innovator layers of the AI labor pyramid. Complementarities between these layers occur across what we call a *cycle of facilitation*: an iterative process in which facilitators, users, innovators work together to apply raw underlying technologies for particular uses, using feedback between the different layers to drive successful adoption.

We combine preliminary firm-level evidence on the composition of social, infrastructural, and technical facilitators with our industry- and national-level examples of facilitation to identify potential strategic models for building and using facilitator capacity. Lastly, we discuss how measurement must be improved for evidence-driven policy to be effective.

TOWARD A FACILITATOR TAXONOMY

Facilitation is the process of translating frontier technological capabilities into usable, safe, and value-generating applications for users. Facilitation occurs, potential individually or jointly, on social, technological, and infrastructural dimensions. Socially, facilitators translate complex AI systems to use cases, build user trust, and capture user feedback that informs refinements and innovations in models and modeling processes. Technologically, facilitators integrate and govern AI tool usage. Infrastructure-wise, they build and maintain the middleware and ecosystems required to deploy AI at scale. Facilitators are gap-bridgers; they connect technological advances at the innovator layer with end-user tool usage. In this section, we propose a framework for classifying facilitators and facilitating activities that may be vital complements to AI-driven productivity.

Classifying Facilitators

Not all forms of facilitation require AI-specific knowledge, and facilitators across social, technical, and infrastructural domains use a variety of skills and talents that shape facilitation across multiple dimensions. For instance, occupations such as middle managers and AI Ethics and Compliance Analysts are critical to social facilitation. They foster trust, normalize AI adoption, and address ethical concerns surrounding AI deployment and adoption, and thus require specialized knowledge of model limitations and capabilities. Conversely, while occupations such as HR Directors and Labor Union Liaisons may play a critical role in normalizing AI adoption in the workplace, they do not require specialized AI knowledge.

In Figure 1, we suggest occupations from the three domains (social, technical, infrastructural) that represent facilitators requiring varying levels of specialized AI knowledge. While some facilitation occurs through AI-centric occupations (i.e., GPU Cluster Architects and AI Ethics and Compliance Analysts), other facilitation is embedded at the task level within traditional occupations. For instance, electricians building power lines for data centers or database engineers structuring and cleaning data perform facilitator functions, but are not strictly dedicated to performing functions on AI systems. However, some roles that require less specialized AI knowledge to date are evolving to require upskilling in AI technologies. Facilitation within electrician work is likely to remain embedded at the task level; electricians are unlikely to need extensive, specialized AI knowledge. In contrast, a growing share of database engineers' tasks (such as tuning LLM data pipelines) demands in-depth AI skills.

Figure 1: Facilitator Types

Type of Facilitation	More Specialized AI Knowledge Required	Less Specialized AI Knowledge Required
Social *(Primarily In-House)*	Middle Managers, AI Ethics and Compliance Analysts	HR Directors, Labor Union Liaisons
Technical *(Fluid/Hybrid)*	MLOps Engineers, RAG & Vector Database Architects	Database engineers, IT Support Services
Infrastructural *(Primarily Third-Party)*	GPU Cluster Architects, AI System Hardware Engineers	Electricians, Fiber Installers

Facilitators vary in whether they sit within or outside the firm. The types of occupations we map to social facilitation (such as HR directors and labor union liaisons) require relationship building, hands-on demonstrations of AI tools, and organizing firm-level processes and structures around AI deployment. Consequently, social facilitators are primarily in-house employees. However, some social facilitators, such as AI ethics and compliance analysts, may be employed by third parties. Evidence from a global survey of ~1300 firms across 105 countries (McKinsey & Company, 2024) shows that social facilitator functions (such as change management or process reform) are much less likely to be fully centralized or outsourced (23%) than technical facilitator functions such as data governance and risk compliance (where 57% of firms report full centralization).

Social facilitators	Social facilitators build user trust, mitigate anxiety, and help firms navigate regulations. For instance, a middle manager who sets clear rules for AI usage in the workplace, and provides examples of when to use AI, serves as a social facilitator.
Technical facilitators	Technical facilitators integrate, streamline, and troubleshoot AI systems. For example. a database engineer who automates data extraction and cleans raw data for machine learning models serves as a technical facilitator.
Infrastructural facilitators	Infrastructural facilitators build, maintain, and scale the hardware and compute necessary for AI adoption. For example, an AI Systems Hardware Engineer who designs custom accelerators serves as an infrastructural facilitator.

In our framework, technical facilitators may operate along a more fluid boundary than social facilitators. A firm may have both an internal IT support services team and MLOps engineering consultants. We hypothesize that the extent to which technical facilitators sit inside or outside a firm's boundary depends on the level of AI integration; employers that use sophisticated AI systems, such as proprietary fine-tuned models, are more likely to build internal technical facilitator capabilities. Conversely, infrastructural facilitators might reside outside firm boundaries; third-party cloud hyperscalers (GPU Cluster Architects) and specialized physical trade contractors (electricians) typically provide this facilitation.

Facilitation in past waves of technological advances followed similar domain patterns. The steam engine parallels the structure of Figure 1: technical facilitation relied on mechanical engineers and millwrights (with specialized knowledge of the steam engine) and blacksmiths (who helped with tasks such as producing intermediate inputs to mechanical frames but didn't need specialized knowledge of steam technology). Professional managers who didn't have a deep understanding of steam technology but understood its criticality played a vital role in socializing and normalizing its use on factory floors (Chandler, 1977). Similarly, internet adoption relied heavily on infrastructural facilitators such as DNS engineers (more specialized knowledge) and fiber cable installers (less specialized knowledge). Dedicated training and education programs were one important feature of social facilitation that often lay outside the boundary of the firm. In each of these technological revolutions, social and infrastructural facilitators needed domain-specific understanding of the technology, whereas technical facilitators needed detailed knowledge of how the technology functioned.

Social, technological, and infrastructural facilitation was historically a prerequisite for adopting a broad range of innovations. AI enables worker-led adoption, but productivity gains may hinge on complementarities between facilitators, users, and innovators. Bottom-up technology adoption cannot reliably drive changes to core business processes without top-down orchestration. Google's AI ATLAS study, which analyzed conversations with 15 million Gemini interactions, found that AI usage within the workplace is broad, but shallow. Importantly, the ATLAS study does not include enterprise data, and so while it differentiates between household and work usage, it is likely that the data tend to capture user-driven behavior (i.e. likely users as self-facilitators, rather than enterprise-level facilitation). While AI usage spanned 68% of occupations that represent 90% of employment in the United States, only 21% of core worker tasks were augmented using AI technologies (Metir AI Team, 2026). This evidence suggests broad user-led adoption.

The Cycle of Facilitation

We propose a cycle of facilitation, in which facilitators expand the frontier of capturable AI gains. Figure 2 illustrates the iterative cycle of facilitation in four stages: (1) *preconditions*, such as base foundation models and fine-grained data, interact with (2) *structure*, such as API pipelines and safety guardrails, in order to foster (3) AI applications such as adoption, user training, and user enablement. Crucially, this cycle is iterative; as AI users interact with outputs, they discover hallucinations, errors, and unaddressed needs from AI tools. Feedback (stage 4) directly refines the inputs (such as fine-tuning data or selecting different models) that facilitators gatekeep. Evaluation and validation occur as part of a feedback stage that feeds into and refines *both* the preconditions and structures. As inputs are updated, the structures that govern, integrate, and maintain them must be scalable and mutable, and facilitators must fine-tune these structures iteratively.

Figure 2: The Cycle of Facilitation

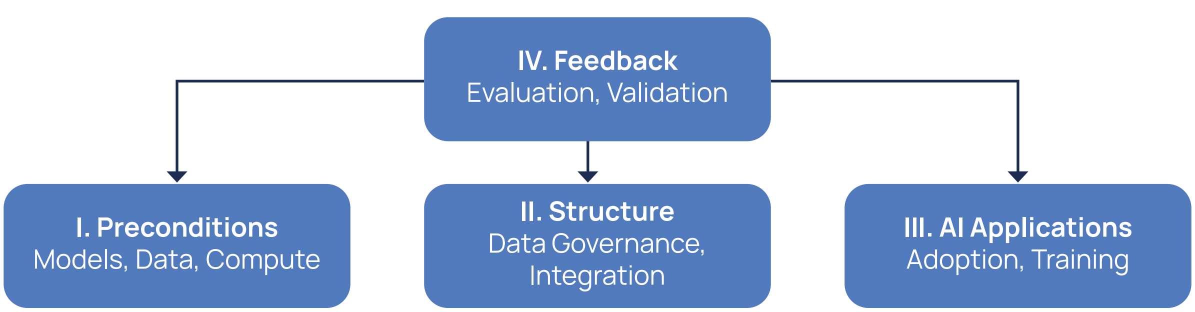

Complementarities between the user, facilitator, and innovator layers of the AI labor pyramid occur at the precondition, AI applications, and feedback stages of the cycle of facilitation. At the preconditions stage, facilitation centers on curating and context-matching frontier capabilities with system-level problems. There is a gradient along which facilitation and innovation occur; innovators directly invent new model architectures, while facilitators fine-tune existing models, refine data inputs or collect data to fit domain-specific contexts, and work on the hardware infrastructure required to scale AI systems. While innovators expand the frontier of available models and inputs, the marginal productivity of upstream model investments depends *directly* on accumulated facilitator capital. Any increase in facilitator capabilities or productivity raises the value of innovators' breakthroughs.

For example, despite a breakthrough in reasoning and acceptance of image inputs produced by innovators, GPT-4 had minimal value to healthcare professionals due to hallucination risks and (U.S.) HIPAA restrictions on uploading data directly to LLMs. Recognizing this barrier, facilitators created Epic's In Basket AI, an augmented-response system that dynamically maps patient medical histories and clinical records directly to model inputs behind strict data security and safety guardrails. The integration of GPT-4 with Epic allowed physicians to automate routine documentation and inbox tasks, freeing up time for direct patient care. Facilitation unlocked GPT-4's value in healthcare settings by integrating an innovative breakthrough into software providers already used daily.

Cambridge's AI for Cultural Heritage Hub (ArCH) exhibits another important example of user-facilitator complementarities. Data from archives, libraries, galleries, and museum artifacts at the university were largely inaccessible to many and underexplored because of lack of transcription and fragmentary or dispersed artifacts. ArCH created a secure "sandpit" where non-technical users could easily analyze cultural heritage data from artifacts, and embedded expert cultural knowledge into algorithms for data processing and integration. For end-users, this project unlocked new data and research potential; however, without expert user input, facilitators would have lacked the deep domain context needed to train algorithms capable of accurately processing, integrating, and interpreting complex historical artifacts.

At the AI applications level of the cycle of facilitation, user-facilitator complementarities are also at play. Users' marginal productivity increases when combined with facilitator capital; without facilitators, user costs and frictions are very high, possibly leading to productivity losses rather than gains. Conversely, without a broad user base, facilitator investments yield lower returns because users are active drivers of the iterative facilitation cycle, and are responsible for many of the productivity gains incurred from AI.

A major point of friction in early deployments of Epic's InBasket AI was overly formulaic notes or messages that did not match physicians' stylistic preferences. This led physicians to reject or heavily edit content, ultimately driving refinement of the input. Continuous iteration between facilitators and clinical users led to features like "Apply My Style" in Epic, which allows physicians to define granular structural preferences, such as conciseness, bulleting, and narrative tone, that automatically reformat AI drafts. Feedback loops, or the lack thereof, are critical to the cycle of facilitation. In the case of Epic, user critiques and suggestions were collected through qualitative surveys and focus groups with physicians (Hong et al., 2025). User-driven, facilitator-implemented solutions that emerged from these groups reduced physician edit times.

Without feedback loops (such as focus groups, surveys, in-person check-ins), the productivity of end-users and facilitators is diminished; users and facilitators operate as strong economic complements in that their marginal productivities are intertwined. If facilitators are unable to access user feedback, refinements of preconditions or structure may not match user needs, leading to lower adoption or productivity for users. Additionally, innovators incur productivity losses when they invest in features or products that solve theoretical problems, instead of realized bottlenecks for users. Firms also experience additional operational costs and frictions when user feedback is delayed or absent.

One example of such low-latency feedback loops in action is the UK's AI for Operations program, which funded an AI initiative at the University of Cambridge to automate degree exam paper verification. To keep the tool aligned as regulations evolve, the team developed a step-by-step "Haynes manual." This structured feedback mechanism enables domain experts (the end-users) to directly convey exam regulation process updates to system facilitators, ensuring continuity even in an evolving environment.

Firms as Facilitators

Facilitation work is performed by individual workers, and facilitation may be represented by entire occupations (such as MLOps Engineers) or by specific roles within occupations (such as via manager-led training in AI tools). However, an entire firm can also serve as a facilitator within its respective industry. Third-party specialist firms that experiment with new AI tools, pilot industry-specific use cases, and package tools for firm end-use can act as facilitators for larger firms. Whether a firm chooses to develop internal facilitation capabilities or hire third-party facilitator firms is influenced by production interdependencies, coordination costs, and validation timeframes.

Interviews with experts from leading advanced manufacturing technology verticals, including robotics and semiconductor manufacturing, highlighted two contrasting examples. First, within the semiconductor industry, where costs associated with errors propagate quickly between production steps and validation is difficult until late in the production process, facilitation often emerges through small, third-party firms. Chipmakers often cannot experiment with unproven AI models, as small errors can ruin large wafer batches.

Due to this structure, many facilitator firms have emerged to discover use cases for AI in semiconductors. For example, Synopsys absorbed R&D risk by developing DSO.ai, an AI platform that executes hundreds of parallel cloud compute trials to learn which parameter adjustments optimize Power, Performance, and Area (PPA) without violating design rules. Third-party facilitators like Synopsys iterate on discrete steps at lower costs, discovering AI use cases that eventually propagate across the broader market.

Conversely, robotics is an industry where validation timelines are faster, and there are fewer interdependencies between steps, reducing the risks and costs associated with in-house facilitation. For example, companies such as Amazon (and its Sparrow Robot) have built internal MLOps teams to deploy vision-guided reinforcement to robotic arms. If a robotic arm moves a package to a wrong location or mislabels one, this error is able to be quickly detected and rarely disrupts downstream process steps, allowing internalization of facilitation at lower risk.

On the whole, large firms can often create internal facilitation if costs of failure are low, but may turn to small, facilitator firms in cases where validation is difficult or time-intensive, or strong interdependencies between process steps exist. All else being equal, larger firms often realize greater productivity gains and have more opportunities to discover high-value use cases of AI. However, structural complexity moderates this dynamic and, in extreme cases, flips the relationship in favor of smaller, specialized third-party firms.

MEASURING FACILITATORS

To effectively design policies to produce facilitators, strengthen facilitator-user and facilitator-innovator complementarities, and improve cycles of facilitation, a performant, low-latency taxonomy and measurement approach is imperative. Understanding who facilitators are, where they work, and their national endowments can help inform national policy agendas. In this section, building on a methodology developed by Ferrone, Fleming and Combemale (2026) to identify AI jobs in job postings data, we offer a preliminary approach to distilling innovator, facilitator and user roles, further differentiating facilitator roles into combinations of the social, technical and infrastructural domains, and finally evaluating clusters of titles or functions within those domains. These classification methods allow us to track AI job demand signals over time and by firm and explore possible complementarities or evidence of hiring strategies.

When developing a taxonomy of facilitators, we consider strategies firms might employ to adopt AI, where they draw the line between a facilitator and a user, and how these strategies influence the firm's labor needs. To this end, we deploy natural language processing techniques to 3.6 million job postings made by firms in the S&P 500 from 2017 to the present, and ask:

- Which firms are adopting AI?
- Which jobs require a worker to have knowledge of AI systems?
- How is this knowledge being deployed?

While job descriptions capture most skills a firm demands, postings alone make it hard to isolate those skills from their broader context, so some skills appear together as an artifact of company policy rather than complementarity. To address this, we look to online course data. Online course designers are incentivized to teach skills that are marketable to employers. While course development might be expected to lag behind firms that are early adopters of a technology, we find that related courses often roll out well before the modal firm implements a large-scale adoption strategy. For example, many generative AI courses were available before 2024, when we see the first significant uptick in AI skill terms in S&P 500 job postings. Building on this, we used online Coursera course descriptions to develop a list of skills commonly associated with AI. We identified terms ranging from "Large Language Model" to "Feature Engineering" to "Semantic Web" by:

1. Quantifying how frequently skills appeared alongside known AI skills.
2. Quantifying how much more likely a skill was to appear in AI literature when compared to its likelihood to appear in any academic literature.
3. Validating those results with human reviewers.

After removing duplicate job postings and boilerplate text (e.g., company descriptions, salary information, benefits packages), we found that 99 thousand of our original sample of 3.6 million job postings referenced AI skills. We employ a multi-agent voting strategy to remove false positives, leaving 94 thousand AI-related job postings. As expected, the presence of these skills in job postings increased moderately after the technology became broadly commercially available at the end of 2022, and then substantially each following year as more powerful models and more accessible tooling were released (see Figure 3).

Figure 3: Firm-Level Share of Postings with AI Indicators, by Year

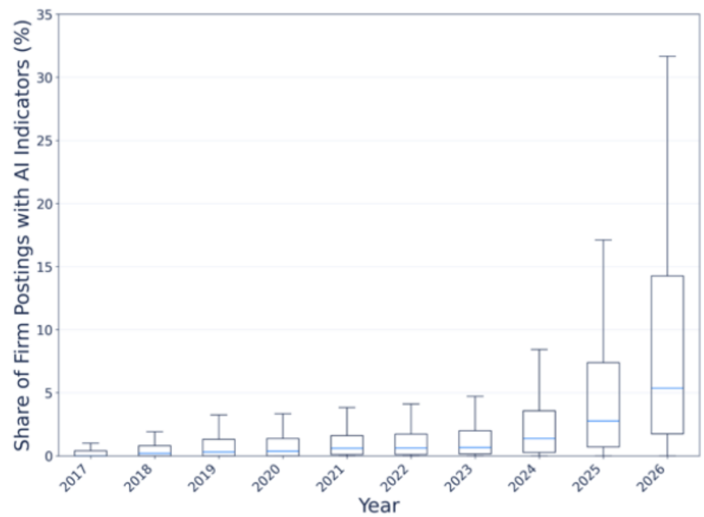

This trend mirrors advances in LLM capabilities. However, postings should not be confused with hires; multiple analysts might be hired under the same job posting, or a posting might be made without any subsequent hires. As such, Figure 3 should not be interpreted as a precise measure of firm AI adoption. However, changes over time are still informative of the direction of hiring activity. Furthermore, the ratio of hiring by job type within firms can control for some variability in postings strategy, allowing us to then compare behavior across firms for evidence of potential complementarity across roles (from co-variance of hiring ratios) or evidence of sectoral-level differences.

AI adoption may take different forms, resulting in variations in workforce composition strategies. Employing the same multi-agent voting strategy as was used to identify false positives, we classify postings as describing innovators, facilitators, and users, as well as the degree to which their relation to AI is social, technical, or infrastructural. After aggregating postings at the firm and sector levels (see Figure 4), we find that the Information Technology sector has the highest adoption rate and the highest representation of AI Innovators. The Communications Services sector also has a considerable number of innovators, but almost all of these are employed by Meta, a firm that (among other things) develops commercially available large language models. Innovators are rare outside these cases, with other sectoral hires split between users and facilitators; Real Estate focuses the most on AI users, and the industrial sector focuses the most on facilitators (followed closely by the Information Technology and Financials sectors).

Figure 4: AI adoption intensity and role tier by sector (2024-2026) –

Mean Share of Firm Postings that are AI is calculated by averaging the adoption rate of each firm in a sector. Firms are only included in their sector's Share of AI Postings that are User, Facilitator, or Innovator if at least 10 AI postings were identified for that firm.

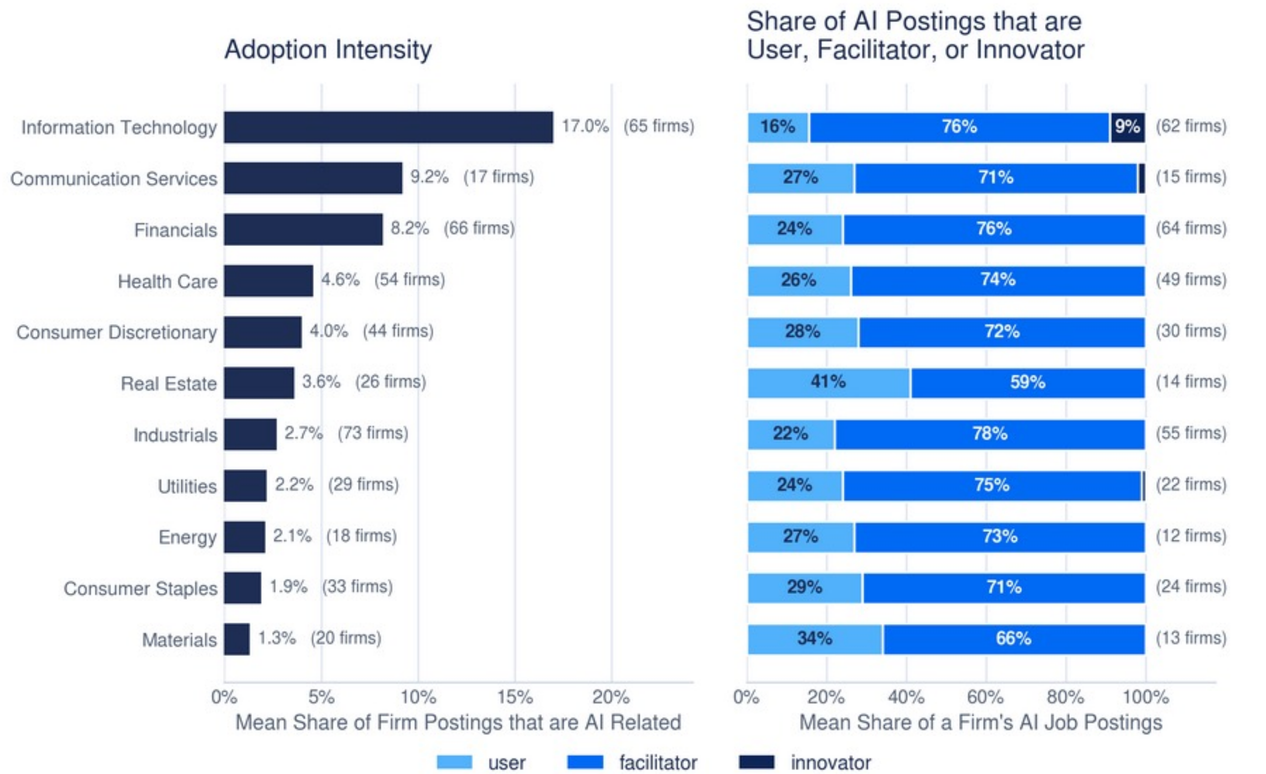

Figure 5: Combinations of Technical, Social, and Infrastructural Facilitator Roles

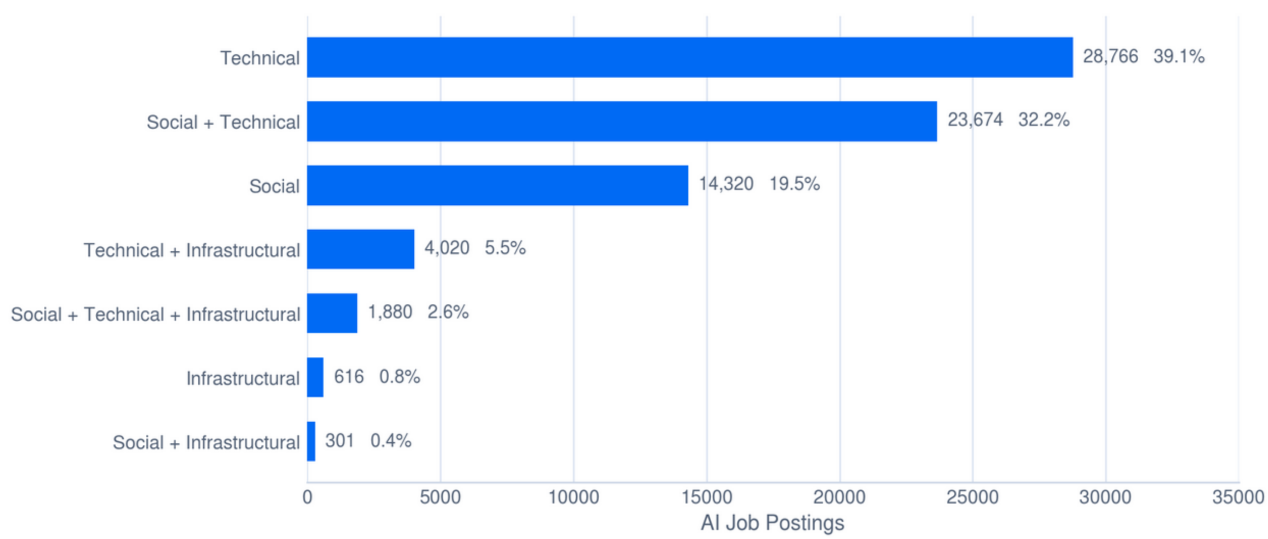

The job postings that our model classifies as facilitators are most often mapped to a technical role. However, we also find a significant contingent of social facilitators, with nearly 20% of postings having their *only* AI role as social facilitation. As expected given their likely concentration outside of the boundary of AI adoption firms, we find fewer infrastructural facilitators – we believe our sampling method undercounts infrastructural facilitators (e.g., working in telecommunications or energy) whose work is not explicitly tied to AI, but is necessary for an economy to adopt it at scale.

While S&P 500 firms are not representative of the broader U.S. or global economy, we might expect their hiring practices to encompass most of the domain of AI-related job types. To identify these job types, we conduct a topic modeling analysis to flag emergent clusters of postings with semantically similar descriptions. We identified 34 emergent AI Job categories. Of these, *Software Engineering and AI Platforms* are the most common, but almost all innovators were categorized as belonging to *Machine Learning Engineering and Research or Semiconductor and Accelerated Computing Engineering,* the workers pushing the frontier in model and hardware design (see Figure 6). *Semiconductor and Accelerated Computing Engineering* are also among the most likely to take on an infrastructure-related role.

Importantly, our method finds that facilitators dominate "AI job" postings relative to innovators and users in almost all topic areas. Across the sectors covered by our analysis, facilitators represent approximately 60–80% of identified AI-related job postings. Technical facilitation represents the largest component of facilitator demand, while more than one-third of facilitator postings span multiple domains (technical, social and infrastructural). These findings suggest that facilitation is a substantial component of advertised AI-related demand among the large firms studied, and that the boundaries between the three domains are often porous in practice.

Further work is needed to characterize the extent to which these facilitators are also users of AI, but the evidence suggests that large firms in the current phase of AI adoption are investing heavily in facilitation in-house, possibly more so (overall) than the innovation layer, while pure users may not have the same salience of AI skills (especially if AI capabilities are embedded tacitly in other enterprise technologies).

Figure 6: AI Job Types and Facilitator Categories

Job Cluster	Share of all AI Jobs	AI Job Types			Facilitators by Domain		
		User	Fac.	Inn.	Soc.	Tech.	Inf.
Software Engineering and AI Platforms	12.2%	9%	88%	3%	21%	90%	9%
Product Management and Marketing	11.0%	20%	78%	2%	69%	54%	7%
Data Science and Business Analytics	8.7%	8%	92%	0%	44%	88%	0%
Machine Learning Engineering and Research	8.3%	0%	75%	25%	15%	99%	11%
Semiconductor and Accelerated Computing Engineering	7.0%	4%	43%	53%	11%	95%	83%
Enterprise Technology Sales	4.8%	20%	80%	0%	78%	18%	7%
Enterprise and Solution Architecture	4.5%	3%	97%	1%	56%	94%	17%
Healthcare and Life Sciences Analytics	4.0%	24%	75%	1%	48%	64%	0%
Data Engineering and Analytics Platforms	3.5%	3%	96%	0%	17%	97%	3%
Cybersecurity Engineering and Operations	3.5%	15%	83%	2%	32%	80%	4%
Finance and Accounting	2.7%	47%	53%	0%	42%	35%	1%
IT Leadership and Digital Strategy	2.6%	6%	93%	1%	76%	60%	11%
Human Resources, Talent and Learning	2.5%	37%	63%	0%	58%	25%	0%
Risk Management and Regulatory Compliance	2.4%	26%	74%	0%	61%	45%	0%
Technical Program and Project Management	2.0%	11%	85%	3%	79%	49%	17%
Product and User Experience Design	1.7%	24%	75%	0%	50%	52%	0%
Legal and Privacy Counsel	1.6%	21%	79%	0%	79%	13%	0%
Business Strategy and Process Improvement	1.5%	22%	78%	0%	76%	15%	3%

Job Cluster	Share of all AI Jobs	AI Job Types			Facilitators by Domain		
		User	Fac.	Inn.	Soc.	Tech.	Inf.
Manufacturing Automation and Process Control	1.5%	22%	76%	1%	39%	70%	9%
Enterprise Software Solution Consulting	1.3%	7%	93%	0%	89%	73%	3%
Autonomous Vehicle Software and Safety	1.1%	4%	75%	22%	19%	96%	15%
Cleared Defense and Intelligence Systems	1.1%	15%	84%	1%	22%	80%	6%
Mechanical, Thermal and Process Engineering	1.1%	35%	47%	17%	6%	50%	44%
Supply Chain Planning and Logistics	1.1%	24%	76%	0%	59%	54%	4%
Financial Data Sales and Service	1.0%	59%	41%	0%	38%	13%	1%
Customer Success and Service Operations	1.0%	17%	83%	0%	80%	28%	1%
Business Systems Analysis	1.0%	16%	84%	0%	55%	71%	1%
Quality Assurance and Test Automation	0.9%	24%	72%	4%	14%	77%	21%
Computer Vision for Semiconductor Inspection	0.8%	3%	78%	19%	7%	97%	26%
Robotic and Intelligent Process Automation	0.8%	1%	99%	0%	61%	96%	3%
Site Reliability and Infrastructure Engineering	0.7%	14%	85%	1%	15%	85%	28%
Internal and IT Audit	0.6%	26%	74%	0%	69%	39%	0%
Procurement and Strategic Sourcing	0.5%	47%	53%	0%	48%	16%	3%
CRM Platform Engineering and Administration	0.5%	5%	95%	0%	48%	89%	1%

FACILITATING THE FACILITATORS

National endowments across facilitation domains help characterize the policy landscape a nation can use to increase adoption and expand productivity gains from AI. Technical, social, and infrastructural facilitators are all imperative for national strategy, and AI diffusion depends on all three. The cycle of facilitation depends on each of these domains to create the preconditions and structure for adoption and to interact with users at the application and feedback stages. National agendas must treat facilitator capacity as a critical input, investing broadly in the three domains and in successful cycles of facilitation.

Building that capacity requires both effective programs and the institutions that sustain them. Programs develop particular capabilities; local and sectoral institutions identify changing demand, coordinate employers and training providers, update provision, and connect skills development to practical adoption. This extends the institutional emphasis of the AI Labor Stack (Ferrone et al., 2026a).

In our first paper, *The AI Labor Stack: Rethinking the Workforce from the Ground Up*, we discussed institution-level approaches to expanding workforce participation and resilience in an AI economy - in particular that these efforts should be tuned to local institutions and sectors, stackable and modular with employer feedback loops built in, and guided by an empirical baseline. In this section, having established the demand for facilitators in a diffusion-focused economy, we highlight program-level case studies of how to boost the supply and availability of facilitators to meet that demand across the technical, social and infrastructural dimensions.

Upskilling and Training Facilitators

Establishing a taxonomy of facilitators provides a framework for understanding how to systematically create a facilitator workforce. To date, most government funded upskilling and reskilling training programs focus on AI adoption and the creation of a broad user-base. While these programs integrate foundational AI knowledge and literacy, they often omit specialized skills required by facilitators. Because of the positive externalities that facilitators produce, cultivating a facilitator workforce requires government investments. Recognizing the criticality of this layer for national competitiveness, multiple strategies have been deployed across different countries to invest in the technical, infrastructural, and social layers of facilitation.

Singapore's AI Apprenticeship Program (AIAP) serves as one example of a program built to produce technical facilitators. The program is structured over 9 months and in two phases, the first of which consists of "deep-skilling": structured training in AI engineering that covers MLOps, AI governance, and a variety of other skills. The second phase of the program pairs students alongside industry experts and current AI engineers to apply the skills learned in the first phase to real-world problems. Workforce training programs like AIAP often suffer from undersubscription due to opportunity costs associated with enrollment, such as training costs or foregone wages during training. Crucially, AIAP offers a stipend to its trainees, creating high demand for the program.

Similarly structured workforce training initiatives exist for infrastructural facilitators. In response to rapidly increasing demand for compute, and thus the expansion of data centers, Ireland launched the Telecoms and Data Apprenticeship Program in partnership with TU Dublin. This 2-year program was designed with a focus on training technicians in high-density fiber-optic installation and electrical infrastructure maintenance. This apprenticeship pairs 8 weeks of off-the-job training at TU Dublin with extensive on-the-job training; importantly, apprentices are fully employed and paid during training, and the cost to participate in the program is low.

Programs building technical and infrastructural capacity combine classroom learning, hands-on practice, and industry apprenticeships. Their scope across the cycle of facilitation ranges widely, moving from targeted, single-stage interventions to comprehensive, end-to-end ecosystems. For instance, Ireland's Telecoms and Data Apprenticeship Program upskills infrastructural facilitators to create preconditions within the cycle of facilitation, establishing the physical foundations for AI systems. Skills in the structure, AI applications, and feedback stages of the cycle are not a focus of the program.

Singapore's AIAP, on the other hand, provides training on classical ML models and LLMs (preconditions), with some focus on AI governance (structure). Skills such as model selection, model tuning, object-oriented programming, and operationalizing AI risk management frameworks are core to the first three months (stage 1) of the program. Notably, AIAP also focuses on the abilities required for enablement and training of users, and systems to elicit and incorporate user feedback. In the second stage of the program, AIAP students work alongside AI engineers, project managers, principal investigators, and MLOps engineers, allowing them to upskill across the full cycle of facilitation.

Optimal training program designs for facilitators (narrowly focused or holistic) depend partially on whether or not facilitators are likely to be employed within or outside the boundaries of an AI-adopting firm. Because infrastructural facilitation (such as electricians or fiber cable installers) often occurs through third-parties, these facilitators are less likely to need skills across multiple stages of facilitation. External facilitators typically operate in discrete, transactional phases, requiring narrow, highly specialized skill sets. For technical and social facilitators, who are more often employed in-house, developing skills across multiple stages of facilitation is more critical. In-house facilitators are more likely to guide users through multiple stages, or incorporate feedback across stages, and thus necessitate a more diverse skill set.

Social facilitators, who are more likely to perform facilitation as an embedded task within an existing occupation rather than as a standalone role, frequently receive less targeted investment from both firms and governments. Despite this, research leveraging the UK Management and Expectations Survey (Bennett Institute for Public Policy, 2026) reveals that social facilitators (specifically, managers) play an important role in AI adoption. Firms with more structured performance monitoring and higher levels of goal-oriented management are quicker adopters. These practices impact adoption even when controlling for firm technical capabilities and data intensity, suggesting that increasing organisational capital and social facilitation is a policy lever for increased adoption.

Germany's Mittelstand-Digital, a support initiative funded by the German Federal Ministry for Economic Affairs and Climate Action, presents one example of a policy implemented to produce social facilitator skills in a small business context. This program funds 25 "innovation hubs" across the country to assist small and medium sized enterprises (SMEs) in embracing digital transformation. The hubs host technical demonstrations from AI experts, networking events between SMEs trying to integrate AI tools, and interactive workshops for project managers and leaders. Direct training for middle managers that is embedded into Mittelstand-Digital cultivates social facilitation skills across a wide range of SMEs.

Replicating Successful Cycles of Facilitation

While upskilling and training facilitators themselves is critical for AI adoption and productivity, so is implementing successful cycles of facilitation. Cycles of facilitation differ greatly across different industries, firms, and contexts. How preconditions are derived or curated, how structures are maintained, how training and enablement occurs, and the processes implemented for feedback loops can vary, and thus so can their success.

The National Network for Microelectronics Education (NNME) exemplifies an important example of not only training and upskilling the workforce to become facilitators, but also of creating competency frameworks for cycles of facilitation that can be ported into new contexts. Created under the CHIPS and Science Act of 2022, the NNME connects federal strategy, private industry demand, and educational systems to address projected shortages of semiconductor workers. The NNME standardizes employer skill needs into competency maps, runs train-the-trainer programs for educators, manages earn-and-learn models, and more.

Importantly, the NNME takes successful programs (such as the partnership between Austin Community College and regional industry leaders like Samsung, Applied Materials, and the Texas Institute for Electronics) and scales them across its multi-state network. By sharing these pre-validated, competency-based curricula and train-the-trainer pathways through its regional nodes, the NNME enables other educational institutions to adopt proven models rapidly without building them from scratch. This continuous loop of industry feedback, structural scaffolding, and national distribution demonstrates how localized cycles of facilitation can be systematically packaged and ported into new academic, geographic, and industrial environments.

Recruiting and Attracting Facilitators from Global Talent Markets

Capacity-building also involves attracting, retaining and making effective use of skilled people. Recruitment and mobility policies can complement domestic training, while cross-domain learning can help technical specialists work with users, managers and governance teams. As an example, within the domains of facilitation, technical represents the largest share of firm-level hiring, followed by technical and social, and then social, but over one-third of facilitator postings are cross-domain, emphasizing the importance of policies and upskilling pathways that generate skills in multiple domains. Countries considering visa policies to attract top technical talent may therefore also want to emphasize social facilitation skills in their immigration criteria. Emergent job categories provide additional evidence on the industries and sectors that nations should prioritize for adoption and diffusion.

Economies and sectors begin with different combinations of social, technical and infrastructural capacity, but these are not fixed endowments. Policymakers can identify the constraints on adoption and address them through training, institutional adaptation, infrastructure investment and organisational capability. The objective is sufficient, accessible facilitator capacity to support users and sustain feedback, rather than a universally high facilitator-to-user ratio. Effective shared services and reusable tools (such as open-source tools) may allow each facilitator to support many users.

MEASUREMENT IMPLICATIONS FOR POLICY

Measurement of occupational AI-exposure, AI jobs, and AI adoption are key challenges for identifying occupations that are most likely to be transformed, augmented, or automated by AI. Displacement is a central focus of policymakers, some of whom have proposed policies such as AI-retraining funds or AI-linked unemployment insurance. While high-fidelity, low-latency measures of AI-exposure and adoption enable quicker identification of labor market disruptions, less policy attention has been paid to the potential for these measures to inform national workforce strategies aimed at strengthening the productivity gains from AI adoption.

In this report, we have presented a novel method for taxonomizing facilitators, innovators, and users within job postings, as well as their roles in social, technical, and infrastructural facilitation of AI adoption. This method allows us to characterize AI job demand signals over time, quantify national endowments of workers across domains, and explore national strategies for AI adoption. A precursor to any policies to capitalize on gains from AI, manage workforce disruptions, and increase AI adoption is a measurement system that is built on low-latency data, and that is able to characterize demands for AI skills and occupations.

At national or regional levels, our framework for quantifying facilitators, innovators, and users identifies where gaps exist, and thus informs strategies for hiring, retraining and upskilling, or attracting talent. Using our measurement framework, policymakers can transition to proactive, targeted workforce programs that build facilitator capacity and prioritize funding across domains where gaps exist. For instance, investing in AI ethics and usage training for managers, or dynamically adjusting immigration policies to focus on these skills may be important for contexts where social facilitators are a bottleneck.

Continued investment in measurement, such as in in-depth characterizations of domain-specific facilitators, is critical for policy. Facilitation, its domains, and the set of skills required for successful cycles of facilitation, will likely evolve as AI adoption matures. To make sure that policy keeps pace with this rapidly shifting landscape, continued investment in measurement is imperative not only for understanding exposure and displacement, but also for targeting investments and expanding opportunities.

REFERENCES

Bennett Institute for Public Policy. (2026, May). *More than plug and play*. University of Cambridge. https://www.bennettschool.cam.ac.uk/wp-content/uploads/2026/05/More-than-plug-and-play.pdf

Chandler, A. D. (1977). *The Visible Hand: The Managerial Revolution in American Business.* Harvard University Press.

Coursera Courses Metadata for Analytics 2025 [Data set]. (n.d.). Kaggle. Retrieved June 9, 2026, from https://www.kaggle.com/datasets/longnguyen3774/coursera-courses-metadata-for-analytics-2025

Hong, C., Chowdhury, A., Sorrentino, A. D., Wang, H., Agrawal, M., Bedoya, A., Bessias, S., Economou-Zavlanos, N. J., Wong, I., Pean, C., Li, F., Pollak, K. I., Poon, E. G., & Pencina, M. J. (2025). Application of unified health large language model evaluation framework to In-Basket message replies: bridging qualitative and quantitative assessments. *Journal of the American Medical Informatics Association, 32*(4), 626–637. https://doi.org/10.1093/jamia/ocaf023

Metir AI Team. “Google’s ATLAS Study: What 15 Million AI Conversations Reveal About How People Actually Use AI.” *Metir AI,* 31 Jul. 2026, www.metirai.com/blog/google-atlas-ai-economy-how-people-use-ai-2026

McKinsey & Company. (2024). *The state of AI: How organizations are rewiring to capture value*. QuantumBlack, AI by McKinsey. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value