Harnessing AI for Development: Can It Work?

1

KEY MESSAGES

- The future trajectories and impacts of artificial intelligence (AI) are highly uncertain, bringing new opportunities and new risks:
  - AI's future capabilities and impacts are highly uncertain. In the short term, modest benefits may come with uneven distributional effects. Long-term gains hinge on both technological advancements and societal factors, including infrastructure and skills, organizational and process changes, development of complementary innovations, evolving regulations, and social norms.
  - AI further lowers barriers to accessing information, opportunities, services, and expertise; unlocks productivity gains; expands market access; and spurs trade. Over time, it may enable the creation of new products, jobs, and industries, offering a pathway to higher living standards and economic transformation.
  - AI can potentially intensify competition among workers and firms, devalue human capital, and exacerbate inequality between and within countries. Developing economies that lag in AI adoption risk premature de-professionalization, with shrinking space to create high-skilled, high-income jobs.
- Although AI innovation is dominated by high-income countries, AI adaptation and localization are key to relevant and effective adoption.
  - The United States has led innovation with 62 percent of notable AI models ever created and more than 73 percent of funding in AI start-ups globally.
  - China now leads the world in AI research paper publication and patent applications. Argentina (0.1 percent), China (13 percent), and India (0.2 percent) are the only three middle-income countries (MICs) creating notable AI models. Low-income countries (LICs) and most other MICs play a marginal role in AI innovation.
  - Many AI applications require deep localization to be relevant and effective. Most developing countries may not need to develop foundational AI models, and they can benefit from localized innovation and practical adaptation of existing open-source platforms.
- Generative artificial intelligence (GenAI) is spreading at record speed, but a stark usage divide exists between LICs and the rest of the world.

- GenAI tools have been used by more than half a billion people (13 percent of the global workforce) in more than 200 countries within 2 years of ChatGPT's launch.
- Over 40 percent of ChatGPT's global traffic came from MICs in April 2025, with Brazil, India, Indonesia, and Viet Nam among top users. Yet LICs lag—making up less than 1 percent of global traffic—highlighting a stark gap.
- Intentional AI adoption by firms remains limited (8 percent) even in advanced economies, and it is lower in developing countries.

- Promising use cases are emerging across developing countries and across sectors. Realizing the potential of AI hinges on the foundational 4Cs: connectivity, compute, context (training data, models, and applications), and competency (digital skills).
  - In Brazil, the Arab Republic of Egypt, India, Nigeria, the Philippines, and Uzbekistan, teachers use AI for lesson planning, nurses for patient tracking, contact center agents for coaching, and farmers for agronomic advice. In these roles, AI acts more as a co-worker or coach—enhancing human capacity—than a replacement.
  - However, the scale, equity, and sustainability of benefits hinge on the 4Cs.

AI definitions, ecosystem, and significant trends

Artificial intelligence (AI) is a branch of computer science dedicated to creating systems capable of performing tasks that typically require human intelligence. Although the concept dates to the 1950s, in recent years, AI capabilities have improved significantly due to the availability of massive data, better training data, and more powerful computer hardware.

AI is an umbrella term used for a set of loosely related technologies. No consensus exists on what is and is not AI (Narayanan and Kapoor 2024). As technological capabilities grow, tasks once exclusive to humans are now within the purview of machines. As a result, experts in the field often joke that AI is everything that computers cannot currently do. AI also has become a marketing buzzword, with companies rushing to label products as being AI based to ride the hype. Hence, it is challenging to provide a precise definition of AI or even of its types or subfields.

This publication defines AI broadly, encompassing subfields such as predictive AI based on machine learning (ML), natural language processing (NLP) and speech recognition, computer vision (CV), transformer model-based generative AI (GenAI), autonomous driving, and robotics. Each type has distinct characteristics and applications. Although the report's primary focus is GenAI—the latest frontier in AI development—it also explores other AI types, reflecting that the ongoing advancements in AI are a continuum and are reshaping the economy and society in diverse ways.

GenAI

GenAI marks a leap forward in AI innovation. Traditional AI approaches are typically task focused, primarily used for data analysis and predictive modeling. GenAI takes this a step further by not only performing tasks but also creating new content. Leveraging advancements in large language models (LLMs), GenAI can interpret natural language prompts and generate content in a wide variety of outputs, including text, images, audio, video, code, 3D models, and even synthetic data. This has directly translated to GenAI models outperforming the average human benchmark in common tasks, such as reading comprehension, language understanding, and image recognition. Such capabilities have enabled new applications across industries and domains, notably in office productivity, creative industries, design, and marketing.

GenAI also complements and enriches traditional AI approaches across various subfields. Together, GenAI and traditional AI approaches expand the tool kit of AI practitioners and researchers, working in tandem to push the boundaries of AI capabilities.

AI ecosystem

The modern AI ecosystem is a vast, intricate, and rapidly evolving landscape (refer to figure 1.1). The ecosystem relies on a foundational energy and telecommunications infrastructure. Capitalizing on this basic infrastructure is the compute layer, where specialized AI chips, servers, storage solutions, and data centers form the physical backbone to store and process data for AI workloads. Cloud computing platforms offer scalable access to compute and integrate AI services.

FIGURE 1.1 Key layers and segments of the AI ecosystem

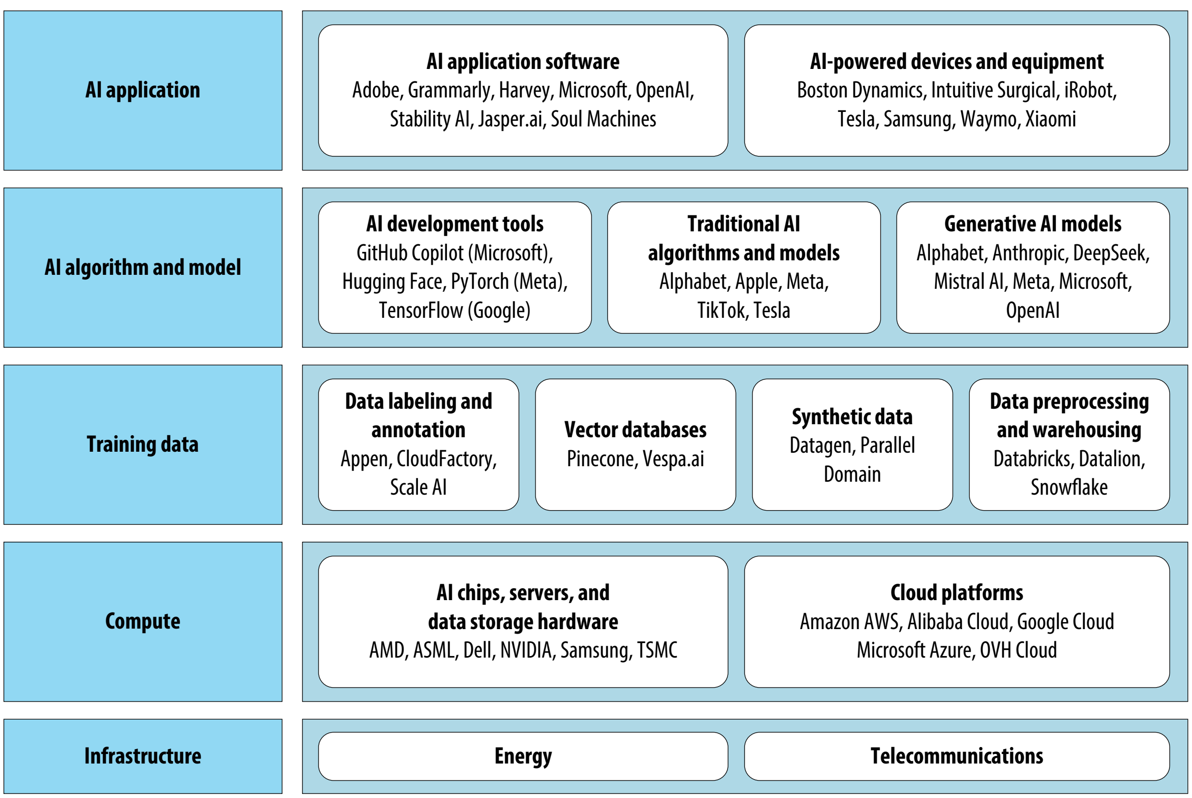

Source: Original figure for this publication.
Note: AI = artificial intelligence.

Training data are yet another critical part of this ecosystem, which can be gathered from various sources such as digital devices, sensors, digital platforms, business transactions, and public data sets. These data are then aggregated into central repositories, labeled, and annotated. Data preprocessing is key to ensuring accuracy and relevance, and it precedes data use for algorithmic training. Data feed into the training data and model layer, where AI development tools are used to build and train AI models.

Traditional AI training data such as ML are already deeply integrated into everyday technologies, powering digital devices, search engines, social media, and e-commerce platforms. Foundational models are large, pretrained models that serve as a general-purpose starting point for a variety of tasks across domains. At the application layer, AI applications leverage and fine-tune foundational models for specific tasks and use cases, and AI-powered devices and equipment embed AI capabilities for localized operation. The boundaries between these layers and segments are often fluid, with many companies operating across multiple segments within the ecosystem.

Significant trends

Although GenAI is still in the early stages of development and adoption, researchers are already exploring new frontiers. Two significant near-term trends are the move toward multimodality and the pursuit of advanced reasoning.

Multimodal AI refers to systems that can process and synthesize information from multiple data types simultaneously. Instead of only "understanding" text or images, a multimodal model can watch a video, listen to the audio, read the subtitles, and generate a comprehensive summary of what occurred. This allows for a much richer, more contextual understanding of the world and enables more fluid and intuitive interactions with technology, moving it closer to how humans perceive reality.

Beyond better perception, the field is intensely focused on *advanced AI reasoning*. Although current models are excellent at pattern recognition, they can struggle with multistep problems that require logical deduction, strategic planning, or creative problem-solving. The next frontier is to develop AI that can not only retrieve information but can also reason with it to analyze a complex problem, break it into smaller steps, form a coherent plan, and execute it. This capability is essential for tackling grand challenges in science, medicine, and engineering.

Agentic AI allows computers to go beyond chat and content generation and carry out all sorts of tasks by delegating them to other software and even people. Traditional AI requires human input to navigate complex scenarios or shifting priorities. Agentic AI autonomously sets subgoals, resolves conflicts, and adjusts its strategies using real-time feedback—often anticipating needs before they arise. Agentic AI can handle complex, multistep reasoning, orchestrate entire workflows, negotiate trade-offs, and integrate cross-domain knowledge to achieve higher-order objectives. Instead of static retraining cycles, agentic AI models evolve continuously, learning from new environments and outcomes, leading to sustained adaptability in fluid and uncertain contexts. Despite agentic AI's great potential, meaningful deployment and adoption remain rare. Agentic AI also creates several critical challenges, including reliability, accountability, trust, and bias amplification risks.

Beyond virtual agents, another promising trajectory involves *physical AI*—notably, autonomous vehicles (AVs) and robotics. Physical AI captures data directly from its environment, for example, through sensors or Internet of Things devices. These capabilities present immense potential for AI to create economic and societal value, with applications ranging from industrial process automation to AVs and personal assistive robots in health care.

The global adoption of AVs is rapidly accelerating, particularly in the ride-hailing sector, with key players demonstrating significant growth. Waymo, Alphabet's self-driving division, recorded more than 700,000 monthly paid trips in California as of March 2025, a remarkable 55-fold increase since August 2023. In April 2025, Waymo was reportedly serving 250,000 paid trips per week in the United States, a fivefold increase year over year (Kolodny and Elias 2025). Similarly, Baidu's Apollo Go robotaxi service has seen substantial deployment in China, with Wuhan alone boasting a fleet of more than 500 vehicles and becoming a major operational area for autonomous driving services. Baidu reported more than 7 million cumulative robotaxi rides across 11 Chinese cities as of July 2024, with 899,000 rides in second-quarter 2024, representing a 26 percent year-over-year increase ("In Leaked Reports" 2024). Although the broader market for fully autonomous private vehicles is still in its nascent stages, the growth in robotaxi services underscores a significant shift in urban mobility.

Scope of this report

This report focuses on recent trends in AI development and the foundational building blocks required to support AI adoption and innovation in developing countries. To provide context and motivation, this chapter briefly discusses the potential economic and social implications of AI, emphasizing the urgency of global coordination and investment to strengthen the foundational 4Cs of AI readiness: connectivity, compute, context (training data, model, and applications), and competency (digital skills). The report acknowledges the wide range of expert views and the uncertainty surrounding AI's long-term development impacts. A more in-depth exploration of these implications will be provided in the *World Development Report 2026*.

This report draws on a range of data sources to track trends in AI innovation and adoption. It benchmarks country performance on the 4Cs, identifies key market failures and externalities that underpin public intervention, and reviews emerging policy approaches—highlighting their benefits, risks, and trade-offs. More detailed and context-specific policy recommendations will be developed in forthcoming reports.

Potential impacts of AI on the global economy and society

Expert perspectives on the pace of future advances in AI and their impacts vary significantly. Some AI researchers and technology entrepreneurs anticipate the arrival of *artificial general intelligence* (AGI), when AI surpasses human intelligence and can solve any intellectual task that a human being can, within 5–20 years (Hinton 2023; Kurzweil 2024). Although AGI could boost output growth substantially, it will necessitate an inexorable decline in the role of human labor in the economy. This shift will require a fundamental reevaluation of economic structures, income distribution mechanisms, social systems, the economic value of education, and the meaning of work (Korinek 2024).

On the one hand, figures such as Sam Altman, Geoffrey Hinton, Elon Musk, and Mustafa Suleyman have talked up AI's superhuman potential and warned about its existential threats to humanity. On the other hand, skeptical market analysts such as Jim Covello from Goldman Sachs have argued that the technology is not designed to solve the kind of complex problems that would justify its enormous cost (Mickle 2024). Scientist and writer Gary Marcus has criticized that the LLMs powering GenAI have too many flaws to be transformative (Zinn 2025). AI experts such as Yann LeCun have argued that today's AI models are far from rivaling the intelligence of people's pets (Mims 2024). Sayash Kapoor and Arvind Narayanan have contended that current AI benchmarks do not measure real-world utility—that much of the world's knowledge is tacit and uncodified and, thus, can hardly be passively learned by machines. They believe AI's economic impacts are likely to be gradual, unfolding over decades as complementary innovations emerge and organizational practices and institutional frameworks adapt.

AI differs from previous general-purpose technologies such as electricity, internal combustion engines, and computers in two key aspects:

1. *Learning and adaptability, which allow it to complement or substitute for human decision-making and judgment.* Unlike traditional technologies, which are rule based and static, AI systems can learn from data, adapt to new information, and improve over time. This makes AI more dynamic and capable of handling complex problems. Hence, AI systems can operate with a degree of autonomy, making decisions based on the data they have learned from and tackling more complex cognitive problems. This autonomy in optimization and decision-making allows AI to either substitute for or complement human decision-making and judgment.
2. *Data dependency.* AI relies heavily on large data sets for training and operation. This contrasts with technologies such as the steam engine, electricity, internal combustion engine, telecommunications, computers, and the internet, which require infrastructure but not vast data inputs. Computers, the internet, social media, and digital platforms all use and generate data, but they can operate without much data. They also primarily process or transmit data according to predefined rules. They do not learn and adapt in the same way AI does.

GenAI further differs from previous AI in the following aspects:

- Prediction with unstructured data, enabling content generation;
- Simplified human-machine interaction through natural language, enhancing user accessibility;
- Enhanced capabilities in cognitive and creative tasks that primarily affect high-skilled, nonroutine white-collar jobs;
- Flexibility and versatility, allowing more general applications than previous AI approaches;
- Stronger scaling laws, requiring even bigger data sets and more compute power; GenAI applications are also more computer intensive than predictive AI; and
- *Hallucination,* generating plausible but incorrect or entirely fictitious content, making reliability a challenge.

Economic implications and the potential magnitude of impacts

Given the significant uncertainties surrounding AI's future capabilities, adoption patterns, and potential regulatory and societal responses, predicting its precise longer-term impact is almost impossible.

Furthermore, the diverse nature of AI technologies, their varied applications, and the complex channels through which they exert influence add to this complexity. This section does not attempt a comprehensive analysis of various possible scenarios. Instead, it selectively focuses on four interconnected key aspects—(1) labor markets and human capital; (2) industrial organizations; (3) economic geography and international trade; and (4) aggregate productivity, living standards, and income distribution—discussing some of the likely scenarios and their potential consequences to guide policy actions.

Labor markets and human capital

AI can help democratize access to education and health care, particularly in underserved countries and regions. In education, the integration of GenAI into schools, colleges, and universities can enable skill-adaptive and personalized learning, instantaneous feedback, and on-demand student guidance and support. In health care, AI can support clinicians and doctors with diagnosis, screening, prognosis, and triaging, alleviating the burden on health care practitioners. These uses could be meaningful in developing countries and regions where access to quality education and health care remains limited. Yet challenges like unequal access, data fragmentation, privacy laws, and distrust can delay widespread adoption.

GenAI could potentially reduce the incentives for people to invest in higher education and skills acquisition. Traditionally, high-skilled, well-paid jobs have demanded specialized expertise, with the promise of higher income and more fulfilling careers motivating individuals to invest in college degrees and advanced training. GenAI tools could commodify expertise by allowing average workers to emulate their top-performing peers. This will reduce the returns on skills (Bloom et al. 2024), making individuals less incentivized to acquire advanced expertise because of the diminished prospects of securing rewarding employment, leading to further downsides to productivity and wages (Capraro et al. 2024). However, the future impact of AI on the demand for expertise is highly uncertain, and how AI will affect people's incentives to acquire various expertise remains to be seen.

The impact of AI on labor demand is complex and uneven. AI can automate certain tasks, reducing the demand for workers primarily engaged in such tasks and occupations. AI can also augment some tasks by making workers more productive. Augmentation has an ambiguous effect on labor demand, hinging on the balance between AI-driven productivity growth and growth in final product demand. Product demand is influenced by multiple factors and general equilibrium effects, including own-price effects, cross-price dynamics with substitutes and complements, the emergence of new products, and changes in consumer income and preferences.

AI may also shift demand toward less-experienced or lower-cost workers by transferring tacit knowledge from experts, enhancing productivity, and narrowing performance gaps (Brynjolfsson, Li, and Raymond 2025; Dell'Acqua et al. 2023; Noy and Zhang 2023). This could spur services outsourcing and offshoring, benefiting workers from developing countries. However, the opposite can occur: Experienced workers may use AI to automate entry-level tasks, displacing novices and reducing jobs outsourced to low-cost countries (Hui, Reshef, and Zhou 2024; Yilmaz, Naumovska, and Aggarwal 2023). Studies show that AI can also widen performance disparities in some tasks, amplifying the output of top performers while reducing the productivity of others (Otis et al. 2023)—potentially increasing demand for elite talent while squeezing out the rest.

The ultimate impact also depends on AI's ability to spur new products, industries, and occupations. If AI catalyzes such breakthroughs, it could generate demand for entirely new jobs and raise incomes.

AI has generated some new occupations and jobs, although these jobs tend to be highly polarized, and the high-income roles are limited in scope and number. High-skilled, high-income roles such as AI scientists, AI developers, robotics specialists, and data scientists are concentrated in a small number of AI firms, cities, and countries. The number of workers needed to meet global demand is negligible to overall employment. The new occupations created specifically by GenAI, such as prompt engineers, have so far been underwhelming in scale, as have peripheral roles such as AI governance and ethics specialists. Conversely, AI has created hundreds of thousands of low-skilled, low-income gig jobs, such as data labeling, further contributing to labor market polarization.

AI is likely to shift bargaining power away from workers to firms, potentially exerting downward pressure on wages and labor income share. As GenAI lowers the barrier to entry in previously elite occupations and potentially devalues some forms of cognitive and analytical expertise, it could suppress wages for a wider range of workers, excluding those at the very top.

AI could further constrain the potential for creating quality jobs in high-skilled services, particularly in developing countries. If the productivity growth in high-skilled services jobs outpaces growth in labor demand, the employment share of high-skilled services will likely stagnate or even decline because of AI. Similar to how computer-driven automation has led to premature deindustrialization, AI may lead to premature de-professionalization, where employment share in high-skilled services peaks earlier, lower, and at lower income levels (Liu 2024).

Industrial organizations

Data have become a critical resource and production factor in the AI age, creating new markets for data collection, storage, and processing. However, the unique characteristics of data complicate data exchange and AI usage. AI and recent GenAI have ignited a surge in investments in AI chips and data centers. In addition, organizations and governments are increasingly recognizing data as a critical asset, prompting a significant ramp-up in training data investment.

Because data are nonrival and can be reused limitlessly, their economic value varies significantly depending on size, quality, timeliness, specific application, context, and user. The challenge to determine the economic value of data, coupled with privacy and security concerns, has impeded the development of data exchange markets. Ambiguities surrounding data property rights and intellectual property (IP) of AI-generated content could also slow AI adoption.

Data exhibit strong economies of scale and scope, leading to a concentration of economic, social, and political power in the hands of those with privileged access to vast data sets. Data-driven businesses can leverage these data sets to solidify their market power and potentially stifle innovation. Furthermore, the increasing reliance on data in decision-making can introduce bias, discrimination, and systemic risks if the data are not representative or the algorithms are flawed.

AI solution providers with large data sets, abundant computational resources, top talent, and superior AI capabilities can benefit from significant economies of scale and scope, entrenching their competitive advantages and potentially leading to higher market concentration. However, GenAI tools tend to require much more computing power than traditional software, particularly during the inference phase. The nonnegligible marginal cost per additional user or service, coupled with the fact that LLMs do not necessarily exhibit network effects, may weaken economies of scale and scope for GenAI providers. This could limit the monopoly-like power often enjoyed by digital platforms. The increased demand for computing power—called *compute*—strengthens the position of compute suppliers such as NVIDIA and benefits cloud service providers who distribute LLMs. This dynamic could squeeze LLM providers such as OpenAI from both sides ("OpenAI's Latest Model" 2025).

The returns on digital and AI adoption tend to increase disproportionately with firm size and data intensity, potentially leading to higher market concentration across various sectors. Conversely, the cost of GenAI adoption tends to be lower than that of predictive AI because of GenAI's versatility and flexibility. This lower adoption barrier could mitigate the increasing returns to scale associated with traditional information technology (IT) investments, potentially acting as a counterforce to rising market concentration.

Economic geography and international trade

If AVs and robots eventually match or exceed human flexibility and dexterity, their widespread adoption could expand suburban areas and create more efficient urban environments. AVs would enable longer commutes, allowing people to live farther from city centers while continuing to work during travel, potentially leading to the expansion of suburban areas. By reducing the need for proximity to workplaces, AVs and remote work technologies could encourage the growth of larger, more dispersed cities. In addition, robots could take over many physical and manual tasks, which would allow workers to control or monitor these tasks remotely. This shift could reduce the need for workers to be physically present at traditional job sites, further decentralizing cities. Moreover, AVs and robots could optimize road usage, reduce traffic congestion, and decrease the need for parking, thereby creating more efficient urban environments (Legêne et al. 2020).

AI could shift global value chains, reshaping global trade patterns and labor demand in the outsourcing and offshoring of digitally deliverable services. The effects, however, remain uncertain. On the one hand, if AI significantly boosts the productivity of workers in advanced economies, there may be less demand to outsource jobs to workers in developing countries. Tasks that were once cost-effectively outsourced might be automated or performed more efficiently by domestic workers in rich countries using AI tools. On the other hand, if AI enables workers in developing countries to emulate the performance of their counterparts in more advanced economies, the demand for outsourcing and offshoring of white-collar jobs could increase. In this scenario, AI could help equalize global high-skilled labor markets, creating new opportunities for developing countries to attract higher-value jobs in sectors such as professional services and IT. GenAI could also bridge language barriers and facilitate international trade.

Aggregate productivity, living standards, and income distribution

AI has the potential to boost efficiency and accelerate innovation across various industries. Its ability to create and scale solutions without explicit programming, combined with the shift toward general-purpose AI applications, promises broader productivity gains. AI can optimize resource allocation, reduce inefficiencies, increase personalization, and accelerate scientific discovery and innovation. Automation of routine tasks frees human workers to focus on higher-value activities. Improved accessibility of GenAI tools can further enhance productivity, especially for smaller firms and a wider range of employees. GenAI could even take on some managerial functions, enabling more efficient organizational structures. A few experimental studies have documented substantial productivity gains for white-collar workers using GenAI, especially for less productive or lower-performing workers (Brynjolfsson, Li, and Raymond 2025; Dell'Acqua et al. 2023; Noy and Zhang 2023).

However, current AI's inherent flaws may limit its broader economic significance. First, the hallucinations that GenAI tools produce are rooted in the mathematical and logical structure of LLMs (Banerjee, Agarwal, and Singla 2024), making them unreliable in business environments where mistakes can be costly. Second, current AI tools remain pattern recognition engines without true understanding, common sense, or logical reasoning (Mirzadeh et al. 2024). Third, much economically valuable knowledge is tacit, context specific, and difficult to codify in a form that AI can learn from (Narayanan and Kapoor 2025).

AI will become more useful when it can "understand" and interact with the physical world. At present, AI models are used mainly to optimize software and generate information, text, and images, affecting a narrow range of activities that are mostly confined to virtual spaces. For AI to have a broader impact, it needs to be able to perceive, comprehend, and respond to physical environments reliably even in novel and unique situations. By bridging the gap between digital intelligence and physical action, AI could revolutionize more industries and solve practical challenges in ways that go far beyond current applications.

However, it is unlikely that robots and AVs will soon match the dexterity, emotional intelligence, and judgment that human workers bring to physical and manual jobs such as driving, child care, and nursing. Beyond the physical realm, emotional intelligence is crucial in jobs such as child care and health care, where reading nonverbal cues, offering empathy, and adjusting on the basis of emotional states are essential for building trust and providing support.

In addition, human judgment, shaped by experience and a deep understanding of context, ethics, and split-second decision-making, plays a vital role in navigating unpredictable situations. Although AI excels at data-driven tasks, replicating the intricate combination of physical dexterity, emotional intelligence, and ethical judgment required for these roles remains a monumental technical and ethical challenge.

Furthermore, low risk tolerance in the physical world, as well as legal and regulatory hurdles, complicates the widespread adoption of robots and AVs. Despite years of development and testing, no clear timeline exists on when AVs will be adopted widely and replace human drivers. However, within months of the release of ChatGPT, companies and people were using it widely in producing and editing communication materials and generating digital content. A crucial difference between the two is that people are much more risk averse when algorithms are brought into the physical world and, in particular, public spaces (Frey and Osborne 2023).

Therefore, AI's overall effect on aggregate productivity and living standards is likely to be modest in the short term, and even more limited in lower-income countries. AI primarily boosts labor productivity in high-skilled, white-collar services. Acemoglu (2024) estimates that fewer than 5 percent of all tasks in the United States can be profitably automated in the next decade. Therefore, AI will raise gross domestic product (GDP) by only 1 percent over the next decade, with even more limited positive effect on wages. The lack of infrastructure and digital skills further diminishes AI's impact in low-income countries (LICs) (Gmyrek, Winkler, and Garganta 2024).

AI is also likely to exacerbate inequality between capital owners and labor (Acemoglu 2025). Its distributional effect among workers is complex, potentially exerting downward pressure on wages for many workers. Autor (2024) argued that, unlike previous digital technologies, GenAI offers an opportunity to expand the middle class by enabling middle-skilled workers to perform higher-stakes decision-making tasks currently reserved for elite experts. In contrast, Frey and Osborne (2023) suggested that, by lowering barriers to entry in cognitive occupations, GenAI will increase competition, eventually driving down wages across the income ladder.

In the medium to long term (>10 years), AI's potential to drive more substantial positive impacts is greater but depends on several critical factors: advancements in AI capabilities, co-inventions and the creation of transformative new products, breakthroughs in energy supply, the development of complementary infrastructure and skills, organizational and process changes, and concomitant shifts in institutions and regulations.

AI presents both opportunities and threats for LICs and middle-income countries (MICs). AI has the potential to enhance human capital and generate new employment, particularly in digitally deliverable services. However, uneven access and adoption could exacerbate existing inequalities in education, skills, and health outcomes. Furthermore, AI may lead to premature de-professionalization, shifting employment toward low-skilled service jobs and potentially worsening unemployment and underemployment. The prospect of declining returns to higher education could disincentivize investments in human capital development.

AI can greatly benefit small businesses in developing countries by making knowledge and expertise more accessible, improving decision-making, expanding market access, and boosting productivity. However, increased economies of scale in AI-driven industries may also make it more challenging for LICs and MICs to develop competitive domestic firms and participate effectively in global markets. Although AI could stimulate a new wave of services outsourcing and boost services exports, it could also have the opposite effect, making export-led growth more difficult to achieve. Finally, although AI may elevate productivity and living standards, the income gap between rich and poor economies may widen further.

Uneven progress in AI innovation, adaptation, and adoption

This report defines AI innovation, adaptation, and adoption as follows:

- *AI innovation: The creation of new AI technologies, models, training data, or methodologies that advance the state of the art.* This includes publishing fundamental research (for example, in scientific publications), developing and launching notable AI models, patenting novel AI techniques, commercializing original AI products, and forming AI-focused start-ups.
- *AI adaptation: The modification or customization of existing AI technologies for specific contexts, sectors, or applications.* Examples include building industry- or domain-specific

AI tools on top of general-purpose models (for example, services built using the ChatGPT application programming interface), fine-tuning open-source AI models for local languages or government service delivery, or integrating AI components into specialized products.

- *AI adoption: The use of AI technologies without in-house modification by public organizations, private enterprises, or individuals.* This can include deploying AI-powered software, using off-the-shelf AI tools such as Gemini, or integrating ready-made AI features into workflows.

Although analytically distinct, these categories are interconnected. Adaptation often builds on innovation, and adoption may involve minor adjustments that edge into adaptation. In practice, many AI activities sit along a continuum rather than in discrete boxes. Investment is a measure for AI innovation and adaptation.

AI innovation and adaptation

Advancements in AI research have been driven primarily by high-income countries (HICs) notably the United States, and increasingly by industry over academia. Researchers in all LICs and MICs combined produced only 46 percent of the global AI scientific publications during 2000–24 and received only 31 percent of global citations (refer to figure 1.2). Articles from upper-middle-income countries (UMICs) and lower-middle-income countries (LMICs) are driven by China and India, respectively.

Although the United States accounts for only 26 percent of HICs' academic articles and 35 percent of the overall citations, China and India account for 70 percent and 68 percent of UMIC and LMIC publications, respectively. Among the more than 900 notable AI models[1] published since the 1950s, 62 percent of lead contributors come from the United States; 25 percent from other HICs; 13 percent from China; 0.2 percent, or only 2 models, from India and 0.1 percent, or only 1 model, from Argentina.

FIGURE 1.2 AI innovation and adaptation activities, by country income group

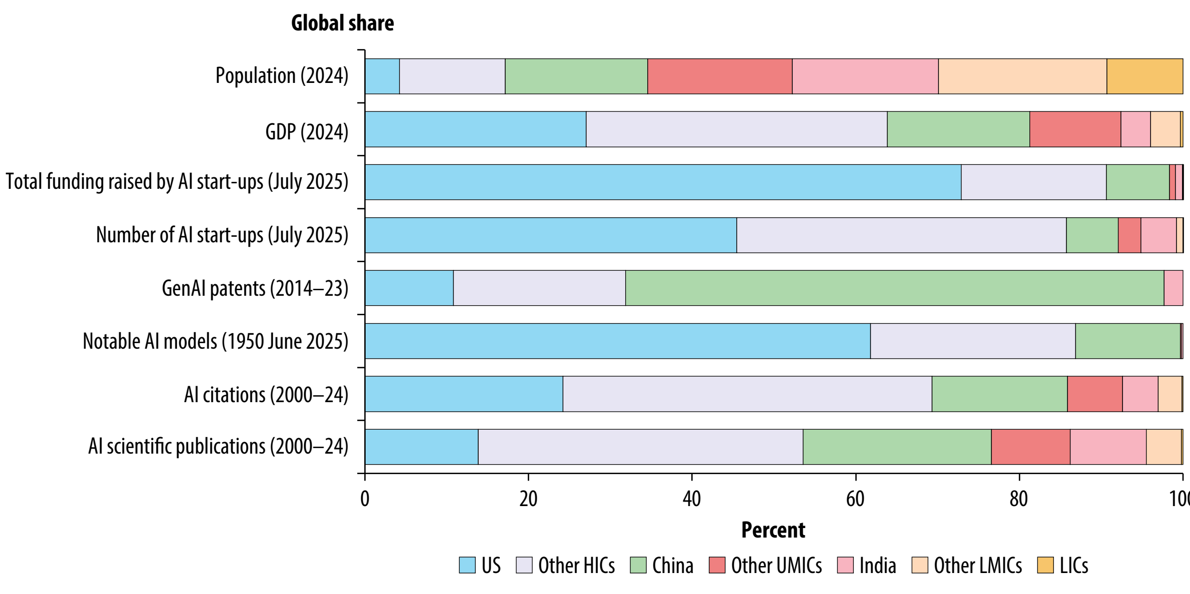

Sources: Original figure for this publication based on analysis of various indicators from Epoch (https://epoch.ai/), OECD.AI (https://oecd.ai/en/), and WIPO (https://www.wipo.int/en/web/ip-statistics).
Note: For models with multiple contributors, only the nationality of the first or leading contributor is counted. AI = artificial intelligence; GDP = gross domestic product; GenAI = generative artificial intelligence; HICs = high-income countries; LICs = low-income countries; LMICs = lower-middle-income countries; UMICs = upper-middle-income countries; US = United States.

The creation of notable AI models is dominated increasingly by industry over academia. Before 2022, 51 percent of the models' lead contributors were from academia, compared with 47 percent from industry; however, since 2022, nearly 80 percent of notable AI models' lead contributors are from industry, whereas only 20 percent are from academia.

Patent filings in AI, particularly GenAI, have experienced ninefold growth since 2017, primarily driven by China. The geographic landscape of GenAI patenting is highly concentrated, with five countries emerging as dominant players: China, India, Japan, the Republic of Korea, and the United States. Over the past decade, Chinese companies and research institutions have led the world in patenting AI applications, with a meteoric rise in filings since 2019. China accounts for 66 percent of global total GenAI patent filings during 2014–23 (WIPO 2024). Although initial concerns were raised about the quality of these Chinese patents, recent research suggests that patent quality has improved, and the approval rate is now aligned with other major patent systems (Thomas and Murdick 2020).

China dominates the overall GenAI patenting landscape, but India has the highest growth rate at 56 percent per year. In GenAI specifically, European countries, notably Germany and the United Kingdom, lag significantly behind, together accounting for less than 3 percent of GenAI patent families (Shoaib 2024). These findings suggests a potential innovation gap for these countries, with potential long-term implications for their technological competitiveness.

No LICs filed patent applications pertaining to AI in the past decade. Importantly, data on patenting trends across LICs and MICs remain sparse. This paucity of data underscores the need for improved monitoring and reporting of AI innovation activities in the developing world. The underrepresentation of diverse groups in AI research and patenting can impede the creation of more inclusive AI models.

The world's leading tech companies—all headquartered in the United States—are leading investments in AI infrastructure and development, reflecting the tech industry's strategic pivot toward AI capabilities across their product portfolios. In 2024, Microsoft allocated US$19 billion in capital expenditures, with half dedicated to data center expansion and the remainder to processing units. Meta has reportedly spent US$37–US$40 billion in 2024 on capital expenditures, including the acquisition of 350,000 NVIDIA graphics processing units. Google has invested US$3 billion in data centers and reportedly spent US$60 million on training AI models using Reddit data. Amazon has committed to a US$230 million investment in GenAI start-ups and plans to invest US$150 billion in data centers over 15 years, while also developing its own AI chips to reduce dependency on NVIDIA.

Although most big technology investments are focused on the United States, some are diversifying to developing countries, too. For example, between January and July 2024, Amazon, Google, and Microsoft announced 21 investments across 15 countries for a total of US$130 billion (Schoeberl and Corrigan 2024), including in Indonesia and Malaysia in Southeast Asia. This US$130 billion in announced investments represents approximately 23.2 percent of their combined first-half 2024 revenues.

Although big technology firms are leading the charge in large-scale AI investment, venture capital (VC) and private equity firms are pouring funds into AI start-ups. However, just as with academic and technical innovations, LICs and MICs represent only a fraction of venture-backed AI start-ups and an even smaller share of funding. Start-ups and the flow of VC investments measure the commercialization of innovation in AI across countries. Start-ups often serve as agile innovators, quickly adapting to new technologies and market needs, and VC investments can provide a quantitative measure of the confidence in a country's start-up ecosystem.

However, data on AI start-ups in LICs and MICs remain limited. One collection of AI global start-ups compiled by CB Insights, consisting of around 21,000 start-ups as of July 2025, counts 45 percent of them as from the United States. Other HICs house another 40 percent of AI start-ups, with the United Kingdom ranking second with 6.8 percent (1,456 start-ups). China ranks third with 6 percent (1,352 start-ups), and India ranks fourth with 4 percent (27 start-ups). Most LICs and MICs have few AI start-ups.

The cumulative funding received by these AI start-ups is much more uneven. AI start-ups in the United States received more than $500 billion in VC funding as of July 2025, 73 percent of the global total. Other HICs received another 18 percent. China accounts for 8 percent of global VC funding of AI start-ups. Other MICs and LICs combined make up less than 2 percent of global VC funding of AI start-ups, indicating limited local AI innovation, adaptation, and commercialization.

The analysis of AI innovation and adaptation globally suggests that few developing countries are engaged and dominant in AI innovation and significant adaptation. China's and India's exceptional positions in AI development leverages three critical foundations: a vast science, technology, engineering, and mathematics talent pool; government support for AI research and development; and mature technology industry ecosystems. Both countries have large-scale technical workforces and established digital infrastructure to provide the necessary foundation for AI innovation.

However, both countries have leveraged these advantages differently. China's AI innovation has benefited from coordinated state support and integrated digital ecosystems, whereas India capitalizes on its global IT services leadership and entrepreneurial culture. Both countries' large domestic markets and strong digital adoption rates create natural testing grounds for AI applications. These structural advantages create self-reinforcing cycles: Success attracts more talent and investment, which further strengthens AI capabilities. This cycle helps explain why they have pulled ahead of other developing nations in AI development and commercialization. Other low- and middle-income economies will also benefit from investments in these fundamentals to develop a domestic entrepreneurial and research ecosystem that leverages the power of AI for their own economic ends.

AI adoption

The adoption and global diffusion of GenAI by individuals has been much faster than that of previous technologies.[2] It took 75 years for the telephone to reach 100 million users, 33 years for cars, 7 years for the internet, 3.3 years for WhatsApp, and only 2 months for ChatGPT (Ebert and Louridas 2023). By June 2025, at least hundreds of AI tools were available, with the top 60 most-visited tools attracting nearly 6 billion monthly visits. Chatbots, because of their versatility and wide range of applications, dominate this space, accounting for 95 percent of traffic among these top tools.

Notably, ChatGPT alone commands nearly 80 percent of the total traffic, with 500 million monthly users—equivalent to 13 percent of the global workforce. By March 2024, ChatGPT had reached 209 of 218 economies worldwide (refer to map 1.1, panel a). The top five economies for ChatGPT traffic in March 2024 were Brazil, India, Indonesia, the Philippines, and the United States. Colombia, Mexico, and Viet Nam have also seen impressive total traffic.

However, LICs and MICs lag far behind HICs in GenAI tool penetration. By April 2025, 24 percent of internet users in HICs used ChatGPT, compared with 5.8 percent in UMICs, 4.7 percent in LMICs, and 0.7 percent in LICs (Liu, Huang, and Wang, 2025). When adjusted for internet users, a clear divide exists between advanced economies and the rest of the world (refer to map 1.1, panel b). Although each internet user in Singapore visits ChatGPT 2.7 times on average each month, most developing countries record less than 0.4 visit per internet user.

MAP 1.1 ChatGPT traffic, by country, 2024

a. Total traffic, March 2024

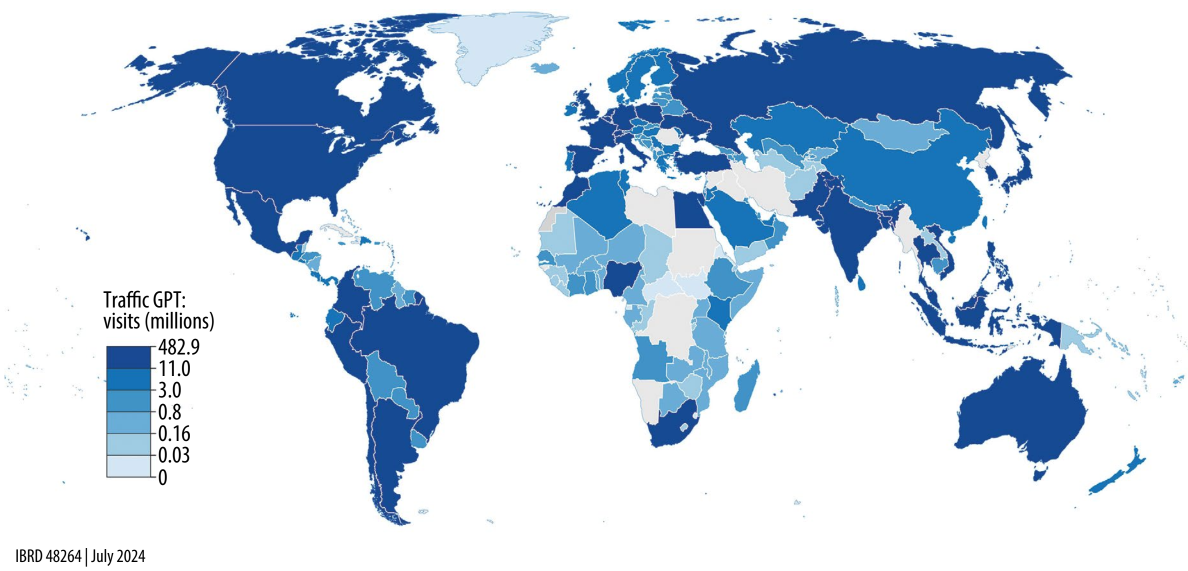

b. Traffic per internet user, March 2024

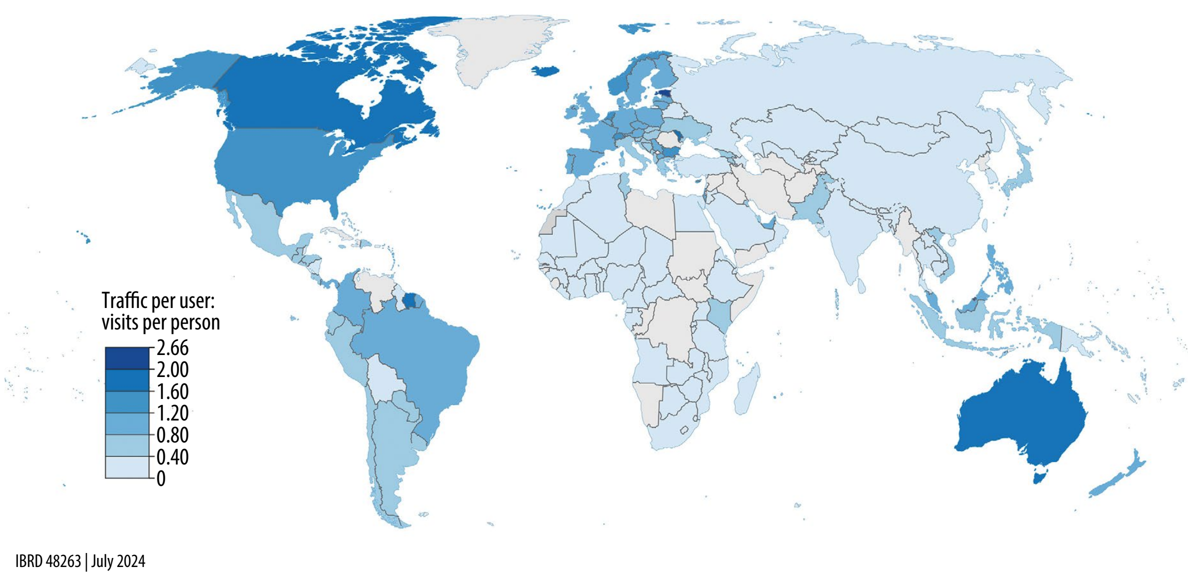

Source: Liu and Wang 2024.
Note: GPT = general-purpose technology.

Still, MICs show disproportionately high GenAI adoption relative to their economic size. The share of US traffic to ChatGPT plummeted from 70 percent to 25 percent within 1 month after its debut, but middle-income economies contributed more than 50 percent of global traffic only 6 months later (refer to figure 1.3). MICs have a disproportionately high adoption rate of GenAI relative to their GDP, electricity consumption, and search engine traffic, although low-income economies remain underrepresented, accounting for less than 1 percent of global ChatGPT traffic.

Robust digital infrastructure, specialization in digital services, higher human capital, and a large youth population are key factors driving higher GenAI adoption at the country level. Specifically, higher fixed broadband penetration, faster internet speed, a large share of white-collar workers, comparative advantage in digitally deliverable services, widespread English fluency and digital skills, and a higher share of youth population predict higher GenAI adoption when GDP per capita, population, and internet penetration are controlled for.

The demographics of GenAI users as captured by Semrush[3] reveal a clear skew toward young, highly educated men. Women account for only one-third of ChatGPT users, a stark contrast to the more balanced gender splits seen on platforms such as Google (48 percent female) and Wikipedia (52 percent female). The most-active GenAI user group is those ages 18–24, particularly for video-based GenAI tools. Nearly half of chatbot users hold a college degree, exceeding the educational attainment of Google's user base.

This gender gap in GenAI usage is significantly more pronounced than the global gender divide in internet use overall. According to International Telecommunication Union data, 70 percent of men are internet users worldwide, compared with 65 percent of women. In HICs and UMICs, this gap has nearly closed, with 94 percent of men and 93 percent of women online in HICs and 81 percent of men and 80 percent of women online in UMICs. Around 32 percent of US men use

FIGURE 1.3 ChatGPT and GenAI traffic over time, by income level and country

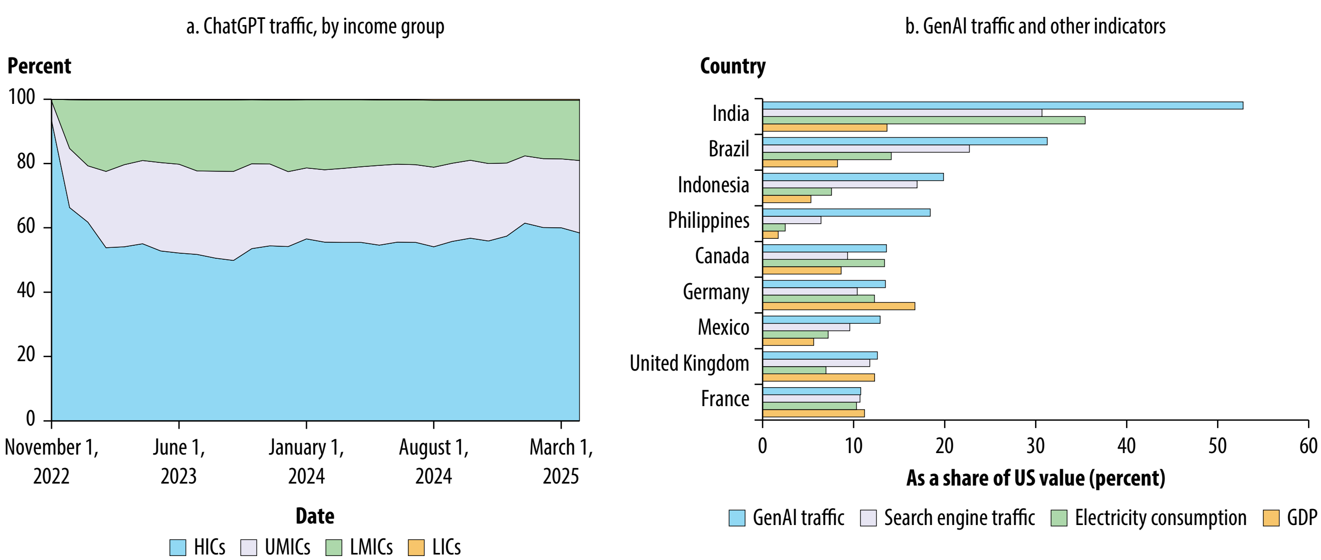

Sources: Liu, Huang, and Wang 2025; Liu and Wang 2024.
Note: GDP = gross domestic product; GenAI = generative artificial intelligence; GPT = general-purpose technology; HIC = high-income countries; LICs = low-income countries; LMICs = lower-middle-income countries; US = United States.

GenAI at work, compared with just 23 percent of women (Bick, Blandin, and Deming 2024). In Denmark, women are 18 percentage points less likely to use ChatGPT than comparable men in the same occupation, of the same age and experience, and at the same earnings level (Humlum and Vestergaard 2024). The gender usage gap is likely much greater in developing countries.

GenAI tools are primarily accessed via desktop computers during weekdays, reflecting their role in professional and academic environments. Although Netflix sees binge-watching spikes over the weekend, ChatGPT usage peaks during the work week. Weekend traffic drops to about two-thirds to one-half of its weekday highs, indicating that people primarily turn to GenAI for work-related tasks.

People primarily use GenAI for writing communications both at work and for personal use. In the United States, nearly 40 percent of people list writing as the top one or two tasks for which GenAI is most useful, and 27 percent use GenAI for writing outside of work (Bick, Blandin, and Deming 2024). Other common tasks for which people most often use GenAI include administrative tasks, interpretation, translation, and information summarization.

Employees are moving faster than companies to integrate GenAI into the workplace. A global survey of 31,000 workers across 31 countries by LinkedIn and Microsoft in 2024 (Microsoft Source 2024) indicated that 75 percent of global knowledge workers are already using GenAI. Without guidance or clearance from the top, 78 percent of employees are bringing their own AI to work. This is even more common at small- and medium-size companies, where 4 of 5 workers bring their own AI. GenAI adoption is high in administrative support, real estate, and retail industries, with countries such as China, India, Indonesia, and Thailand reporting usage rates above 90 percent.

GenAI adoption at work is highest for IT, management, business, and finance occupations and industries. In the United States, nearly half of workers in these occupations and industries use GenAI frequently (Bick, Blandin, and Deming 2024). Software developers report the highest GenAI adoption. A survey conducted in May 2024 of 65,000 developers on Stack Overflow suggested that 76 percent are using or plan to use GenAI tools in their development process in 2024—a 6 percentage point increase from 2023 (Stack Overflow 2024). In the next year, surveyed developers believed that AI tools will be more integrated into documenting code (81 percent), testing code (80 percent), and writing code (76 percent). Of all the available AI tools, ChatGPT is used most—and 74 percent want to keep using it next year.

However, GenAI usage is not limited to white-collar workers. Interestingly, 22 percent of workers in blue-collar jobs—construction and extraction, installation and repair, skilled production, and transportation and moving occupations—use GenAI at work in the United States.

In Denmark, IT-prone and high-skilled occupations similarly report the highest GenAI adoption at both the extensive and the intensive margins. Around 80 percent of software developers, marketing professionals, and journalists used ChatGPT in 2024, and 20 percent of them have a ChatGPT Plus subscription (Humlum and Vestergaard 2024).

In contrast, AI adoption by firms—for use in firms' business processes, for instance—remains limited even in advanced economies, with large firms leading the way. Firm surveys in OECD countries show that, on average, only 8 percent of firms used AI in 2023 (refer to figure 1.4). In contrast, allied technologies such as cloud computing were used by nearly half of the firms in OECD countries (OECD 2024). Korea leads OECD economies in firm AI adoption, with 28 percent of all firms reporting AI adoption. At the other end of the spectrum, only 2 percent of firms in Romania use AI.

FIGURE 1.4 Firm AI adoption rate across OECD countries, 2024

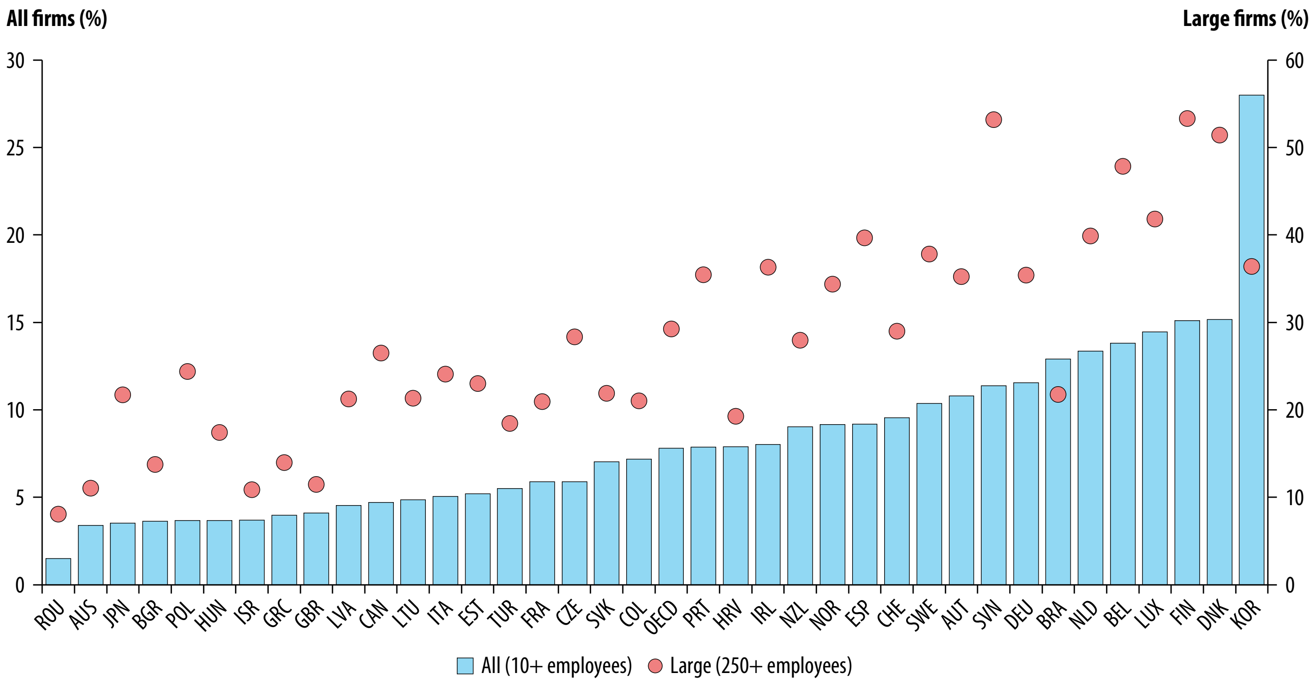

Source: Original figure for this publication using data from OECD 2024 (https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/05/oecd-digital-economy-outlook-2024-volume-1_d30a04c9/a1689dc5-en.pdf).
Note: For a list of country codes, refer to https://www.iso.org/obp/ui/#search. AI = artificial intelligence; OECD = Organisation for Economic Co-operation and Development.

Large firms with at least 250 employees tend to have much higher AI adoption rates. In Belgium, Denmark, Finland, and Slovenia, around half of large firms are using AI. In the United States, fewer than 6 percent of firms are currently using AI-related technologies.

Adoption is concentrated among very large firms; those in "superstar" cities; and businesses owned by younger, more educated individuals. Recent data from the Business Trends and Outlook Survey (BTOS) indicate a modest increase in AI adoption, from 3.7 percent in September 2023 to an expected 6.6 percent by fall 2024 (Bonney et al. 2024). Among developing countries, 13 percent of firms in Brazil use AI, nearly twice the 6 percent and 7 percent in Türkiye and Colombia, respectively. AI adoption is higher for firms with more than 1,000 employees. Nearly 60 percent of firms in India, 50 percent of those in China, and 47 percent of those in Latin America report actively deploying AI, based on the IBM Global AI Adoption Survey 2023 (IBM 2024).

Firms in IT, professional services, and financial services are most likely to adopt AI. The US BTOS shows that more than 20 percent of firms in IT services use AI, followed by about 17 percent in professional services and 9 percent in finance and insurance (Bonney et al. 2024). The IBM Global AI Adoption Survey 2023 and McKinsey's GenAI adoption survey show similar results (McKinsey 2025). These three industries tend to top all other industries in the share of workers with advanced degrees, average earnings, and dependence on digital inputs, and they show high trade intensity. Thus, they are at the forefront of AI adoption.

When surveyed, firms report inaccuracy as the top risk in their GenAI use, and nearly one-quarter have already experienced negative consequences. As businesses begin to see the benefits of GenAI, they are also recognizing associated diverse risks. These can range from data management risks such as data privacy, bias, or IP infringement to model management risks, which tend to focus on inaccurate output or lack of explainability.

A third big risk category is security and incorrect use. GenAI hallucinations and inaccurate output are the top concern for firms, with 63 percent of firms reporting it as a relevant risk in 2024, up from 56 percent in 2023. Inaccuracy—which can affect use cases across the GenAI value chain, ranging from customer journeys and summarization to coding and creative content—is the only risk that firms are significantly more likely to mitigate in 2024 than in 2023. Nearly one-quarter of firms also reported negative consequences from GenAI's inaccuracy. However, given the inherent nature of hallucinations in GenAI, this will likely limit its use cases and effectiveness for firms.

Firms most often see meaningful cost reductions from AI in human resources (HR), risk, legal, and compliance functions. The McKinsey Global Survey on AI in 2024 (McKinsey 2025) indicates that 31 percent of firms report cost reductions of at least 10 percent in HR because of GenAI adoption, by far the area in which firms see the largest cost reduction among all business functions. Of firms, 15 percent also see cost reductions of at least 20 percent in risk, legal, and compliance because of GenAI adoption. Similarly, analytical AI adoption reduces costs in HR and compliance by at least 20 percent in more than 8 percent of firms.

Developing countries face significant challenges in fostering AI innovation and adaptation because of a complex interplay of economic, educational, and infrastructural factors. The lack of substantial research funding from both government and private sectors hampers long-term AI projects—including basic and applied research relevant to developing use cases within countries. Limited access to VC and angel investors creates a fundamental funding gap. High interest rates and stringent collateral requirements from traditional banks make debt financing challenging for AI start-ups, which typically have few tangible assets.

The innovation ecosystem is further strained by limited access to advanced computing resources, with insufficient high-performance computing infrastructure and prohibitively expensive cloud computing services. Moreover, market constraints, such as smaller domestic markets and limited ability to monetize AI solutions, affect revenue potential. Many local firms may lack the digital maturity to adopt AI solutions, resulting in a narrower customer base. This creates a chicken-and-egg problem in which limited market demand inhibits investment in development of applied AI solutions.

In addition, the inadequacy of specialized AI education and training programs in universities creates a skills gap at the entry level, hindering the growth of a robust AI workforce. Weak IP protection in many developing countries can also discourage innovation and commercialization, because researchers and companies may fear that their AI innovations may be stolen or copied without legal recourse.

The exodus of AI talent from developing to developed nations represents a severe impediment to building local AI applications. The salary disparity driving this migration is substantial—for example, in Latin America, the average AI engineer earns approximately US$30,000 annually, compared with US$150,000 in the United States. Career growth limitations further accelerate this drain. These barriers collectively create a challenging environment for AI innovation, limiting the ability of developing countries to contribute to and benefit from advancements in AI technology.

The adoption of AI technologies in developing countries is similarly hindered by a range of structural, economic, and social challenges. Limited digital infrastructure, characterized by unreliable internet connectivity and insufficient access to smart devices, forms a fundamental barrier to widespread AI adoption. A significant skills gap in the workforce can make it difficult for organizations to effectively implement and manage AI systems. Cultural and language barriers also play a crucial role, because many AI systems are not adapted to local languages or cultural nuances, leading to reduced effectiveness and acceptance.

Financial constraints pose another significant hurdle, with the high initial costs of AI implementation often out of reach for local businesses and government agencies. Moreover, a general lack of AI awareness and trust among decision-makers and the public can result in skepticism about AI's potential benefits and concerns about its impact on employment. These multifaceted barriers create a complex landscape that significantly slows the adoption and integration of AI technologies in developing countries, potentially widening the global digital divide.

4C pillars of AI foundations

This report highlights the importance of strengthening the foundational 4Cs that underpin a country's capacity to adopt, adapt, and innovate with AI: connectivity, compute, context, and competency (refer to figure 1.5). These four pillars represent the critical enablers of an inclusive and effective AI ecosystem.

- *Connectivity*—including reliable energy and digital infrastructure and access to and ownership of digital devices—is essential to ensuring that individuals, firms, and institutions can access and use AI technologies.

FIGURE 1.5 The 4Cs of AI foundations

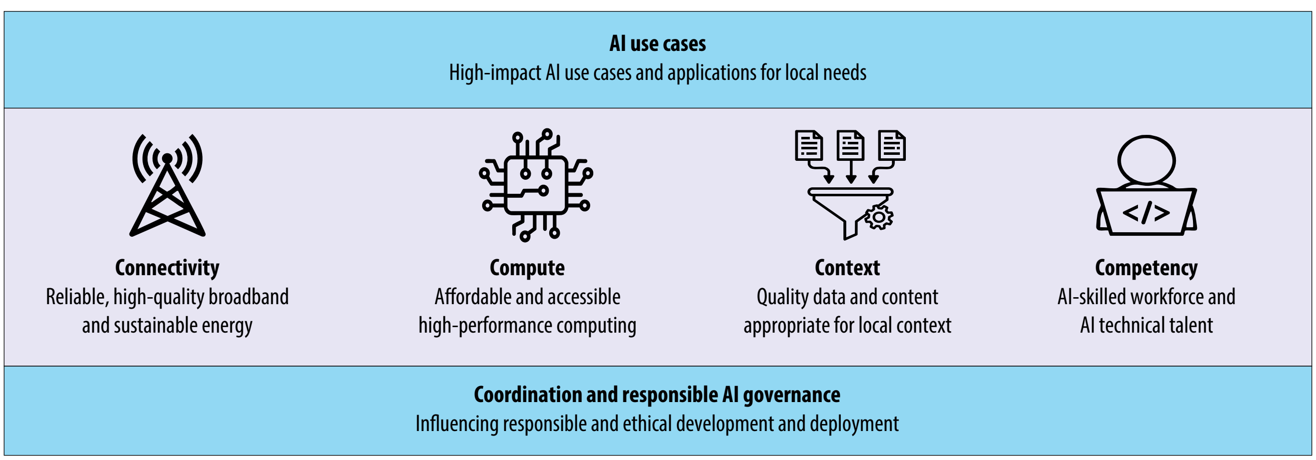

Source: Original illustration for this publication.
Note: The "coordination and responsible AI governance" topic is covered in a companion publication (World Bank 2024, https://hdl.handle.net/10986/42500). AI = artificial intelligence.

- *Compute* refers to the availability and affordability of processing power, such as AI chips, data centers, and cloud computing, which are necessary to train and deploy models at scale.
- *Context* encompasses the training data, models, and applications tailored to local languages, cultures, needs, and governance frameworks, without which AI tools may remain irrelevant or untrusted.
- *Competency* captures the digital skills and broader capabilities required to integrate AI into workflows and build new solutions.

Although all four components are necessary across the board, their relative importance and required sophistication differ depending on whether a country is focused on adoption, adaptation, or innovation. For basic adoption of off-the-shelf AI tools, countries must prioritize widespread connectivity and foundational digital skills. Adaptation—modifying AI tools to local contexts—demands more advanced data governance, locally relevant content, and specialized talent. Innovation, such as developing frontier models or novel applications, requires high-end computer infrastructure, robust venture and research ecosystems, and globally competitive expertise.

Thus, the 4Cs offer a practical policy framework for diagnosing readiness, identifying binding constraints, and tailoring interventions. They can help governments set priorities based on their AI ambitions and capabilities while also fostering coherence across infrastructure, innovation, education, labor, and data strategies. A balanced and coordinated investment in the 4Cs is essential not just for enabling AI but also for ensuring it delivers broad-based economic and social impact.

About this publication

After this introduction, this publication proceeds as follows: chapters 2–5 delve into the four foundational pillars crucial for AI success. For each of the 4Cs, these chapters examine supply and demand dynamics, benchmark country performance, pinpoint market failures and other obstacles, and offer policy insights. Chapter 6 synthesizes these findings and proposes high-level policy considerations for how developing countries can leverage the 4Cs to unlock the potential benefits of AI. Finally, Appendix A summarizes key findings from three thematic case studies; the full cases are available online at https://hdl.handle.net/10986/43822.

Notes

1. Creation of notable AI models measures both AI innovation and adaptation.
2. Due to the fast pace of innovation, a consistent definition and measurement of AI adoption is difficult. AI is often an invisible component of existing products and services, leading to "unintentional adoption" and significant underreporting in surveys. The challenges are amplified by the lack of consistent, cross-country, and longitudinal surveys, particularly in developing nations with limited resources. Given these issues, this section focuses on the intentional adoption of GenAI, which is more easily tracked.
3. Semrush (https://www.semrush.com/) calculates user demographics through data gathered via hundreds of partnerships with clickstream data providers, bringing in over 2 million events every minute. All the data collected is anonymized to protect individual identities but provides insights on socioeconomic characteristics of users (such as age, gender, and location) that are made available for analytics. With this anonymized clickstream data, Semrush uses a neural network algorithm to create realistic estimates of various subgroups and conducts error testing to validate these estimates.

References

Acemoglu, D. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40 (121): 13–58.

Autor, D. 2024. "Applying AI to Rebuild Middle Class Jobs." Working Paper w32140, National Bureau of Economic Research, Cambridge, MA.

Banerjee, S., A. Agarwal, and S. Singla. 2024. "LLMs Will Always Hallucinate, and We Need to Live With This." Preprint, submitted September 9, 2024. https://arxiv.org/abs/2409.05746.

Bick, A., A. Blandin, and D. J. Deming. 2024. "The Rapid Adoption of Generative AI." Working Paper 32966, National Bureau of Economic Research, Cambridge, MA.

Bloom, D. E., K. Prettner, J. Saadaoui, and M. Veruete. 2024. "Artificial Intelligence and the Skill Premium." Working Paper 32430, National Bureau of Economic Research, Cambridge, MA.

Bonney, K., C. Breaux, C. Buffington, E. Dinlersoz, L. S. Foster, N. Goldschlag, J. C. Haltiwanger, Z. Kroff, and K. Savage. 2024. "Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey." Working Paper 32319, National Bureau of Economic Research, Cambridge, MA.

Brynjolfsson, E., D. Li, and L. Raymond. 2025. "Generative AI at Work." *Quarterly Journal of Economics* 140 (2): 889–942.

Capraro, V., A. Lentsch, D. Acemoglu, S. Akgun, A. Akhmedova, E. Bilancini, J. F. Bonnefon, P. Brañas-Garza, L. Butera, K. M. Douglas, and J. A. Everett. 2024. "The Impact of Generative Artificial Intelligence on Socioeconomic Inequalities and Policy Making." *PNAS Nexus* 3 (6): 191.

Dell'Acqua, F., E. McFowland, III, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Unit Working Paper 24-013, Harvard Business School, Technology and Operations Management, Boston.

Ebert, C., and P. Louridas. 2023. "Fig. 1: Time to Reach 100 Mio Users." https://www.researchgate.net/figure/Time-to-reach-100-million-users-for-different-technologies-in-months-after-initial_fig1_372212988.

Frey, C. B., and M. Osborne. 2023. "Generative AI and the Future of Work: A Reappraisal." *Brown Journal of World Affairs* 30: 161.

Gmyrek, P., H. Winkler, and S. Garganta. 2024. "Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America." Working Paper 121, International Labour Organization, Geneva.

Hinton, G. 2023. "I Now Predict 5 to 20 Years but without Much Confidence." X, May 3, 2023, 5:08 a.m. https://x.com/geoffreyhinton/status/1653687894534504451.

Hui, X., O. Reshef, and L. Zhou. 2024. "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market." *Organization Science* 35 (6): 1977–89.

Humlum, A., and E. Vestergaard. 2024. "The Adoption of ChatGPT." Becker Friedman Institute for Economics Working Paper 2024–50, University of Chicago, Chicago.

IBM. 2024. "Data Suggest Growth in Enterprise Adoption of AI Is Due to Widespread Deployment by Early Adopters, but Barriers Keep 40% in the Exploration and Experimental Phases." News, January 10, 2024. https://newsroom.ibm.com/2024-01-10-Data-Suggests-Growth-in-Enterprise-Adoption-of-AI-is-Due-to-Widespread-Deployment-by-Early-Adopters.

"In Leaked Reports, Baidu Shows It's in the Global Robotaxi Race." 2024. *Bamboo Works,* October 14, 2024. https://thebambooworks.com/in-leaked-reports-baidu-shows-its-in-the-global-robotaxi-race/.

Kolodny, L., and J. Elias. 2025. "Waymo Reports 250,000 Paid Robotaxi Rides per Week in U.S." https://www.cnbc.com/2025/04/24/waymo-reports-250000-paid-robotaxi-rides-per-week-in-us.html.

Korinek, A. 2024. "Economic Policy Challenges for the Age of AI." Working Paper w32980, National Bureau of Economic Research, Cambridge, MA.

Kurzweil, R. 2024. *The Singularity Is Nearer: When We Merge with AI*. New York: Penguin.

Legêne, M. F., W. L. Auping, G. Homem de Almeida Correia, and B. van Arem. 2020. "Spatial Impact of Automated Driving in Urban Areas." *Journal of Simulation* 14 (4): 295–303.

Liu, Y. 2024. "Generative AI: Catalyst for Growth or Harbinger of Premature De-Professionalization?" Policy Research Working Paper 10915, World Bank, Washington, DC. http://documents.worldbank.org/curated/en/099520009172451039/.

Liu, Y., and H. Wang. 2024. *Who on Earth Is Using Generative AI?* Washington, DC: World Bank.

Liu, Y., J. Y. Huang, and H. Wang. 2025. *Who on Earth Is Using Generative AI? Global Trends and Shifts in 2025*. Washington, DC: World Bank.

McKinsey. 2025. "The State of AI: How Organizations Are Rewiring to Capture Value." *QuantumBlack: AI by McKinsey,* March 12, 2025. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai.

Mickle, T. 2024. "Will AI Be a Bust? A Wall Street Skeptic Rings the Alarm." *Seattle Times,* September 23, 2024. https://www.seattletimes.com/business/will-ai-go-bust-a-wall-street-skeptic-rings-the-alarm/.

Microsoft Source. 2024. "Microsoft and LinkedIn Release the 2024 Work Trend Index on the State of AI at Work." *Microsoft Company News,* May 8, 2024. https://news.microsoft.com/source/2024/05/08/microsoft-and-linkedin-release-the-2024-work-trend-index-on-the-state-of-ai-at-work/.

Mims, C. 2024. "This AI Pioneer Thinks AI Is Dumber Than a Cat." *Wall Street Journal,* October 11, 2024. https://www.wsj.com/tech/ai/yann-lecun-ai-meta-aa59e2f5.

Mirzadeh, I., K. Alizadeh, H. Shahrokhi, O. Tuzel, S. Bengio, and M. Farajtabar. 2024. "GSM-Symbolic: Understanding the Limitations of Mathematical Reasoning in Large Language Models." Preprint, last revised August 27, 2025. https://arxiv.org/abs/2410.05229.

Narayanan, A., and S. Kapoor. 2024. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton, NJ: Princeton University Press.

Narayanan, A., and S. Kapoor. 2025. *AI as Normal Technology*. New York: Knight First Amendment Institute, Columbia University, New York. https://knightcolumbia.org/content/ai-as-normal-technology.

Noy, S., and W. Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381 (6654): 187–92.

OECD (Organisation for Economic Co-operation and Development). 2024. *The Digital Economy Outlook (Volume 1)*. Paris: OECD. https://doi.org/10.1787/a1689dc5-en.

"OpenAI's Latest Model Will Change the Economics of Software." 2025. *The Economist,* January 20, 2025.

Otis, N., R. P. Clarke, S. Delecourt, D. Holtz, and R. Koning. 2023. "The Uneven Impact of Generative AI on Entrepreneurial Performance." SSRN. Unpublished manuscript, February 27, 2024. http://dx.doi.org/10.2139/ssrn.4671369.

Schoeberl, C., and J. Corrigan. 2024. "Funding the AI Cloud—Amazon, Alphabet, and Microsoft's Cloud Computing Investments, Part 1: If You Build Cloud, They Will Come." Center for Security and Emerging Technology, Georgetown University, Washington, DC. https://cset.georgetown.edu/publication/funding-the-ai-cloud-computing-investments-part-1/.

Shoaib, Z. 2024. "Big Tech Moves More AI Spending Abroad." *Wall Street Journal,* May 23, 2024. https://www.wsj.com/tech/ai/big-tech-moves-more-ai-spending-abroad-088988de.

Stack Overflow. 2024. "2024 Developer Survey." https://survey.stackoverflow.co/2024/.

Thomas, P., and D. Murdick. 2020. *Patents and Artificial Intelligence: A Primer.* CSET Data Brief. Washington, DC: Center for Security and Emerging Technology, Georgetown University. https://cset.georgetown.edu/wp-content/uploads/CSET-Patents-and-Artificial-Intelligence-1.pdf.

WIPO (World Intellectual Property Organization). 2024. *Generative Artificial Intelligence*. Patent Landscape Report. Geneva: WIPO. https://www.wipo.int/web-publications/patent-landscape-report-generative-artificial-intelligence-genai/assets/62504/Generative%20AI%20-%20PLR%20EN_WEB2.pdf.

Yilmaz, E. D., I. Naumovska, and V. A. Aggarwal. 2023. "AI-Driven Labor Substitution: Evidence from Google Translate and ChatGPT." Working Paper 2023/24/EFE, INSEAD, Fontainebleau, France. http://dx.doi.org/10.2139/ssrn.4400516.

Zinn, A. 2025. "Generative AI's Most Prominent Skeptic Doubles Down." Techxplore, May 29, 2025. https://techxplore.com/news/2025-05-generative-ai-prominent-skeptic.html.