

WORLD DEVELOPMENT REPORT 2026

ARTIFICIAL INTELLIGENCE FOR DEVELOPMENT

Concept Note

Table of Contents

Introduction.....	4
Part I: Decoding AI for Development.....	5
1. Decoding AI.....	5
AI can enhance capabilities.....	5
AI is not new, but it is increasingly more powerful and pervasive	6
AI has heralded a wave of general-purpose technological change	7
Some characteristics of AI distinguish it from previous GPTs	8
AI’s promise for development.....	11
AI may deliver slowly on its promise for development.....	13
2. Leveraging AI for Development: Conceptual Framework.....	15
3. Empirical Contributions.....	18
Part II: Assessing AI’s Implications for Development.....	19
1. Economic Growth and Jobs	20
AI will displace workers in certain tasks, expand jobs in other tasks	20
AI can improve firm capabilities and worker productivity	21
AI can lower the barriers to entry for workers to complete cognitive tasks	22
AI can spur innovation and entrepreneurship	22
AI may reduce the jobs dividend of certain sectors	23
AI might benefit owners of capital more than workers.....	24
AI might benefit skilled workers more	24
The lack of complementary factors is likely to blunt AI’s impact on productivity.....	24
2. Government Services.....	25
AI can detect patterns and generate forecasts to guide the allocation of scarce resources	26
AI can improve transparency and accountability.....	26
AI can customize digital solutions for service delivery at a lower cost.....	26
AI use by governments can have economy-wide gains	27
The lack of complementary factors can blunt AI’s impact on improving government services	27
AI adoption might exacerbate biases, dependency, and accountability risks.....	28
3. Sociopolitical Change	28
AI can accelerate the creation of new ideas	28
AI is expanding consumer benefits and reshaping social interaction	29
AI can make it easier to deepen repression, spread propaganda and enable polarization.....	29
AI can reinforce social biases or lead to discriminatory outcomes.....	29

AI could erode the accumulation of human capital over time	30
AI might have adverse consequences for the environment.....	30
AI can intensify geopolitical contestation.....	30
Part III: Identifying Policy Priorities to Leverage AI for Development	31
1. Empowering Entrepreneurs to Leverage the Value of AI	32
Supporting competition and a local innovation ecosystem that spurs AI adaptation.....	33
Supporting AI adaptation through government procurement and other fiscal tools	34
2. Educating People to Enhance Workers’ Capabilities	34
Emphasizing broad-based skills development that will enable AI adoption and help adjust to AI- related disruption	35
Investing in AI-related technical skills that matter for AI adaptation	35
Leveraging diaspora and foreign AI talent.....	35
3. Ensuring Trust by Building Effective Governance and Institutional Frameworks to Manage AI’s Broader Societal Impacts	36
Experimenting with different AI governance approaches.....	36
Adopt good practices from around the world	37
Investing in the evaluation of AI solutions	37
Investing in the development of institutional agility.....	38
Consultations, Timeline, and WDR 2026 Team.....	38
Consultations and Timeline.....	38
The WDR 2026 Team	38
External Consultations	39
Regional Initiatives	40
Academic Advisory Committee and Other Academic Consultations	40
High-Level Advisory Panel.....	41
Notes	42
References.....	44

Introduction

Artificial intelligence (AI) is no longer the stuff of science fiction; it is unleashing a new wave of technological change. In a few short years, the application of AI models has evolved from machine-learning algorithms that generate forecasts and predict optimization strategies to AI chatbots that generate content based on text, audio, and video inputs by humans. The potential economic impacts are manifold.

There is optimism about AI enabling developing countries to leapfrog development challenges by reducing human error, optimizing complex production and distribution processes, and facilitating decision-making. Examples abound. AI can help address long-standing market failures—for example, in credit markets, where banks can provide loans to more people by using AI algorithms to trace people’s digital presence and assess their creditworthiness. AI also has the potential to fill in skills gaps in areas such as education and health services, which lack skilled teachers and doctors. The use of AI-generated business advice, powered just through smartphones and high-speed internet, can make small enterprises more profitable and productive.

Such optimism should not be unbridled. AI could widen the gap between high- and lower-income countries because of its onerous requirements for computing power, data, and skills. Advances in AI could automate some tasks—resulting in job losses for some and lower wages for others. There are also concerns about competition and dependency—namely, that a few large technology companies headquartered in high-income countries could have an advantage in creating and deploying AI. Moreover, without appropriate safeguards, AI can reinforce biases, misdiagnose needs, or result in flawed decisions that could make the solution worse than the problem.

World Development Report 2026 (WDR 2026) will investigate the development implications of AI as a general-purpose technology (GPT)¹ and assess what might be the best policy choices to leverage the benefits of AI while offsetting potential risks. In doing so, it will focus on the institutional and governance arrangements needed to ensure inclusive and responsible deployment of AI. The aim of the Report is to provide a developing-country perspective on a topic for which the academic and policy discussion has focused primarily on high-income countries.

The first part of the Report will decode AI as a technology to identify its salience for development. In doing so, it will define the scope of the technology and identify its core economic functions. It will also assess how AI is similar or dissimilar to previous waves of GPTs and in what ways AI is salient for development. Further, it will present a conceptual framework to assess how AI can improve, skew, or lower societal welfare through its economic, social, and political impacts and what policies would enable low- and middle-income countries to leverage AI to achieve better development outcomes.

The second part of the Report will assess the development implications of AI. In doing so, it will focus on how AI can affect people’s welfare in low- and middle-income countries through its impact on economic growth and jobs (economic), the delivery of government services (social), and broader sociopolitical change (political). These impacts of AI on economic, and social development will fundamentally draw on the role of AI in enhancing capabilities of people, including many low-skilled workers, owners of small enterprises, and governments that are often capacity constrained.

This also includes the role of AI in bringing about wider sociopolitical transformation, both positive and negative. AI is driving the production of new ideas, democratizing access to personalized information, and reshaping the nature of social interactions, but also raising cybersecurity risks, eroding fundamental cognitive learning, and obfuscating political processes by enabling surveillance and the spread of propaganda. The market dynamics of AI as a product are also relevant to the discussion about potential

societal harm. For example, AI could result in significant environmental harm through its demand for energy resources. Similarly, AI could amplify geopolitical contestation with its “winner-takes-most” market dynamic and the possibility that the benefits of AI are concentrated among a few advanced economies.

The third part of the Report will identify policy priorities for low- and middle-income countries to enhance the value of AI for development while mitigating adverse effects. Policy choices made today will shape countries’ preparedness to leverage AI for development tomorrow. Governments around the world are committed to investing considerable resources on developing AI. Yet numerous challenges, especially in developing countries, can limit the technology’s potential benefits. First, there is the difficulty of leapfrogging previous GPTs: the use of even basic information and communication technologies (ICT) is far from universal because of the lack of finance, skills, institutional capabilities, and other complementary factors in many developing economies. Second, there is the challenge of extracting greater value from the use of AI for businesses in developing economies, and the perennial challenge of equity versus efficiency: specifically, of balancing AI-generated efficiency gains with ensuring that large numbers of people do not get left behind. Finally, there is the possibility of governments in low- and middle-income countries having to always play catch-up: it is difficult to implement regulations, especially with weaker institutional capabilities, that need to keep pace with a rapidly evolving technology.

Part I: Decoding AI for Development

1. Decoding AI

AI can enhance capabilities

Artificial intelligence (AI) means different things to different people. Descriptions range from AI being a prediction technology and a tool that increases the value of human expertise to AI referring to intelligent machines that are capable of reasoning. All these perspectives suggest that AI’s key function is implicit in the name of the technology. AI has unleashed a new wave of technological change that enables a machine to impart expertise associated with intelligent human behavior. Unsurprisingly, the *Oxford English Dictionary* defines AI as “the theory and development of computer systems able to perform tasks normally requiring human intelligence.”

In principle, AI can raise capabilities and lower barriers to entry in terms of what kind of work people can do (Autor and Thompson 2025). There is evidence of such “leveling up” across activities such as customer support (Brynjolfsson et al. 2025), professional writing (Noy and Zhang 2023), and software development (Cui et al. 2024). When AI surpasses human capabilities to scale up activity and do things faster, AI can also enable people to do new tasks. The work of two recently crowned Nobel laureates in Economics, Philippe Aghion and Joel Mokyr, also emphasizes the fundamental role that AI will play in generating new ideas and knowledge.

In practice, “AI is an umbrella term for a set of loosely related technologies” (Narayanan and Kapoor 2024). ChatGPT and other large language models can write text, generate software code, summarize documents, and draw images while machine learning algorithms can optimize decision-making through, for example, better predicting weather forecasts or evaluating loan applicants.” AI has been described as prediction machines that depend on data, statistical models, and judgement (Agrawal et al. 2024). The use of AI in real-world applications can be organized into three categories based on the underlying cognitive function that it performs. This categorization mirrors the technology’s evolution from predictive AI to generative AI and agentic AI.

- *Predictive AI* uses statistical and machine learning (ML) models for prediction or classification tasks. Machine learning models can analyze historical as well as real-time data to detect patterns, generate forecasts, make recommendations, or guide optimization decisions. For example, ML models extract insights from unstructured data, such as images, audio, and text, to analyze radiology scans and use drone imagery to detect crop disease. Similarly, ML predictions based on past data have optimized responses to food crises, detected fraud in financial markets, and guided resource allocation in government programs.
- *Generative AI (GenAI)* creates new content, such as text, images, or audio. Most generative AI in use today consists of large language models (LLMs), such as ChatGPT, Gemini, Claude, and DeepSeek-R1. These LLMs are used as general-purpose AI chatbots.² LLMs can be trained or finetuned to build AI chatbots (such as for medical diagnosis or personalized tutoring) that can provide specific advice or assistance. These AI tools can, for example, summarize documents, draft legal contracts, generate educational materials, or engage in customer service interaction.
- *Agentic AI* combines predictive and generative AI with capabilities of reasoning and interacting with the broader operating environment to execute a series of tasks independently based on goals and context and learning gained during the process. Early-stage examples of AI agents include self-driving cars and software programming assistants such as Claude Code and OpenAI’s Codex.

These categories of AI are not mutually exclusive in terms of the cognitive functions they perform. In fact, each category expands upon the capabilities of the previous category. For example, the outputs of generative AI, such as images, text, and audio, also involve prediction but are accessible to people on their smartphones without any specialist skills or training in a particular task. They can serve as a “front end” to specialist predictive AI systems whose outputs may have been previously difficult to access, understand, or interpret by many potential users—such as machine learning models to assist diagnosis of malaria and typhoid fever (Attai et al. 2024). Similarly, the ability of agentic AI to execute a series of tasks rests on both prediction and content generation, but there is little evidence so far on its potential impacts. As a result, much of the discussion in this Report will focus on predictive and generative AI tools.

AI is not new, but it is increasingly more powerful and pervasive

The conceptual foundations of AI can be traced back to the “Turing Test,” which evaluates whether a machine’s responses are indistinguishable from those of a human (Turing 1950). The formative era during the 1960s was based on symbolic systems that used hand-coded rules. ELIZA, a chatbot that mimicked conversation through simple pattern matching, is a much-cited example (Weizenbaum 1966). ELIZA’s language capabilities largely came from individual “scripts.” The most famous script, *Doctor*, engaged users with open-ended questions and responses characteristic of an empathic psychologist. Such chatbots illustrated technical feasibility but could not be scaled up without exhaustive manual input.

Subsequently, data-driven approaches in which AI models learned patterns directly from data overcame the reliance on humans to define rules. In the 1990s, ML techniques—such as decision tree algorithms that predicted the value of a target variable based on several input variables—enabled applications like spam filtering and credit scoring (Quinlan 1986). However, the reliance on human experts to engineer inputs specific to certain tasks, especially for unstructured data like images or text, limited their scalability. The 2010s marked a turning point with advances in computing hardware, large-scale labeled data sets like ImageNet (Deng et al. 2009; Krizhevsky et al. 2012), and improved data training techniques (LeCun et al. 2015). These developments enabled deep learning: brain-inspired neural network models with tens of millions of parameters that could automatically learn abstract features directly from data. This paradigm shift led to major breakthroughs in image recognition, speech recognition, and natural language processing

and enabled a range of task-specific applications, such as customer service chatbots or AI assistants, such as Apple’s Siri, Amazon’s Alexa, Google’s Assistant, and Microsoft’s Cortana.

The transition to the modern era of AI has been marked by a shift from narrow, task-specific models to flexible, large language models that learn patterns from massive, unlabeled, and often multimodal data sets such as text, images, and code (Bommasani et al. 2021; Vaswani et al. 2017). These models underpin the success of chatbots such as ChatGPT, Gemini, and Claude, which can generate text, images, code, and other outputs with simple prompts (Radford et al. 2018). Importantly, they can generalize across tasks they were not explicitly trained for (Brown et al. 2020). These chatbots are versatile and can be used across a variety of applications, including information retrieval, customer support, e-commerce, and the provision of education and health services.

Looking ahead, there is a great deal of uncertainty about the future trajectory of AI. Most discernibly, the technological frontier is shifting toward advanced AI agents that not only generate content but also plan actions, interact with tools, and execute tasks (Yao et al. 2023). While still unreliable over long sequences of tasks, the performance of AI agents is improving rapidly. Existing benchmarks suggest that the length of tasks they can complete is doubling every seven months (Kwa et al. 2025; Wang et al. 2024).

AI has heralded a wave of general-purpose technological change

Previous general-purpose technologies (GPTs), including the steam engine, electricity, computers, and the internet, were characterized by widespread use, the potential for continuous innovation, and complementary innovation in application industries (Bresnahan and Trajtenberg 1995). Based on these metrics, a cluster of AI-related technologies—including big data, data mining, data science, machine learning, and natural language processing—consistently ranks at or near the top among a large set of innovative technologies (Goldfarb et al. 2023).³ As Erik Brynjolfsson (2022) notes, “a good case can be made that AI is the most general of all general-purpose technologies: after all, if we can solve the puzzle of intelligence, it would help solve many of the other problems in the world.”

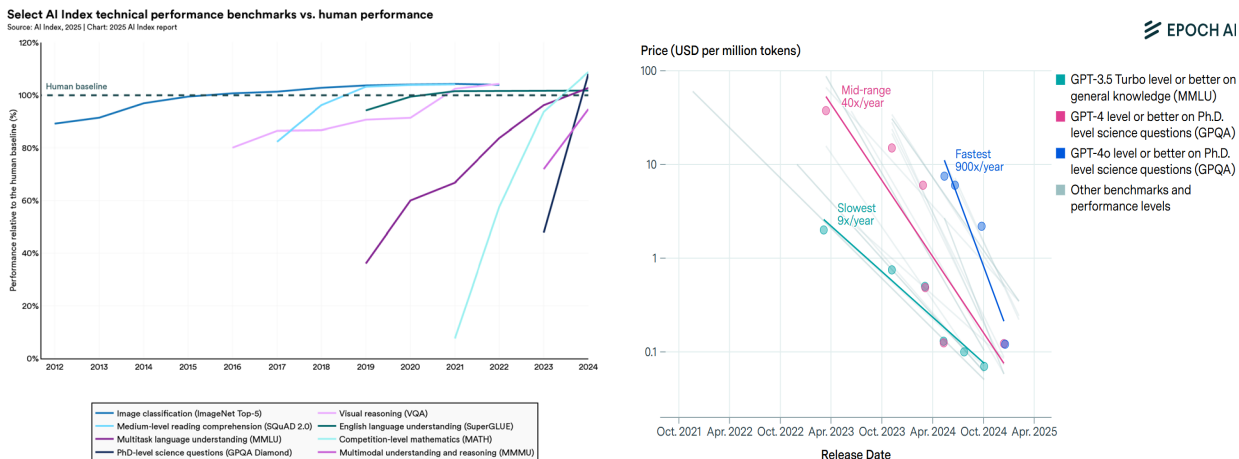
First and foremost, the potential applications of AI are widespread, at least in advanced economies where the technology is likely to diffuse faster than in developing countries. Evidence from the United States shows that about 80 percent of the workforce has at least 10 percent of their work tasks exposed to the capabilities of LLMs, and approximately 19 percent of workers have at least 50 percent of their tasks exposed (Eloundou et al. 2024). Evidence from the United States also shows that generative AI is used across occupations and sectors, even though it is used more by people employed in computer-related, management, and business-related occupations and among ICT, professional, finance, and business services (Bick et al. 2024).

Second, the potential for continuous innovation is manifested in the rising capabilities and falling costs of AI models. The capabilities of AI systems increasingly surpass average human performance on an increasing set of tasks, such as image classification, language and visual understanding, advanced mathematics, and reasoning (refer to figure 1, panel a). The pace of improvement has accelerated sharply in recent years. In competition-level mathematics, AI leaped from 10 percent of the best human performers to parity in less than two years (2021–23), and AI performance on PhD-level science questions exceeded human parity within less than one year (2023–24) (Maslej et al. 2025). AI usage costs have also been declining rapidly. For example, processing one million words with GPT-3.5, the model that powered the first released version of ChatGPT, fell from about US\$20 in late 2022 to just US\$0.07 by late 2024 (Maslej et al. 2025). The sharpest reductions in price are occurring in more advanced generative AI models (refer to figure 1, panel b).⁴

Figure 1 AI models are becoming more capable and cheaper to use

a. Benchmarking AI models against human performance

b. Usage costs of generative AI models



Sources: Panel a: Maslej et al. 2024. Panel b: Epoch AI 2025.

Note: AI = artificial intelligence; GPQA = Graduate-Level Google-Proof Q&A Benchmark; GPT = general-purpose technology; MMLU = Measuring Massive Multitask Language Understanding benchmark; Ph.D. = Doctor of Philosophy (advanced degree); VQA = Visual Question Answering dataset.

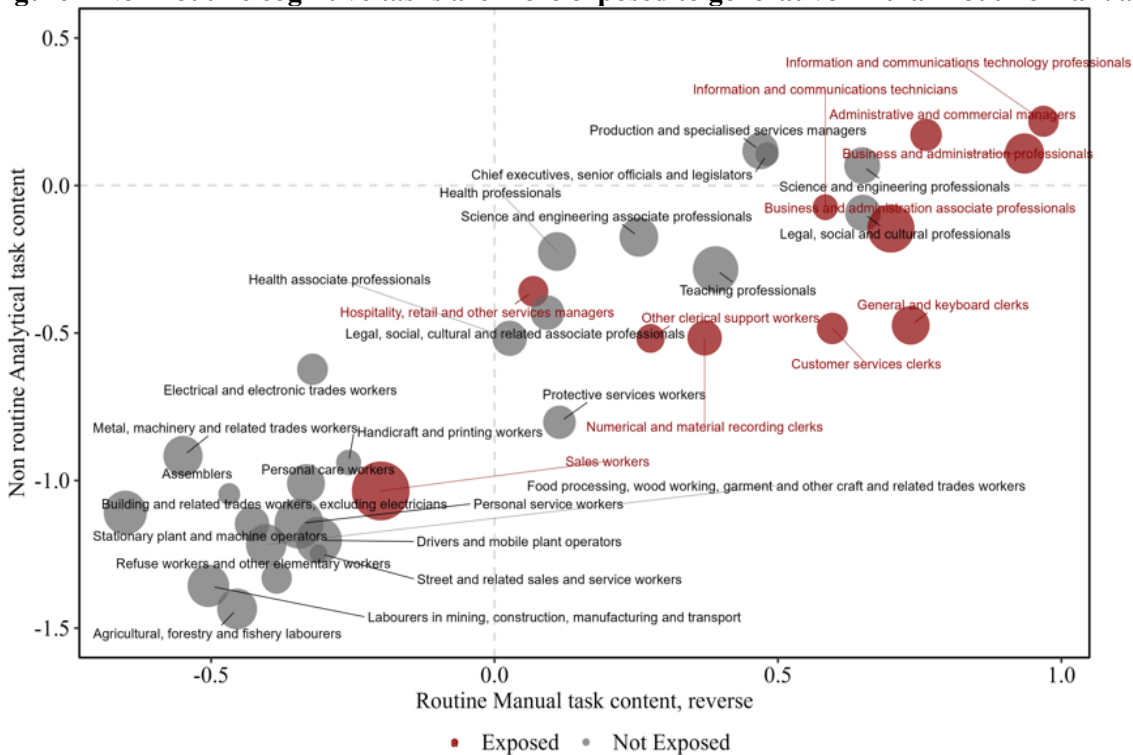
Third, the deployment of AI has been accompanied by complementary innovations across a range of application industries, including in developing countries. The feedback loop of innovation in producing and using industries has the potential to generate sustained productivity growth. Such innovation in using industries is carried out by entrepreneurs and software developers who generate and customize AI applications for specific contexts. The number of AI projects on the software development platform GitHub⁵ has ballooned from zero to nearly 150,000 in the past decade and more than doubled between just 2023 and 2024 (GitHub 2024). The number of AI start-ups globally exceeded 10,000 by mid-2024, with more than 70,000 AI-focused companies in total by 2025, highlighting massive growth in the AI application ecosystem (HubSpot 2025; Maslej et al. 2025).⁶

Some characteristics of AI distinguish it from previous GPTs

AI affects a different set of tasks

The spread of computers largely affected routine cognitive tasks and was associated with falling numbers of middle-skilled, middle-wage jobs and increasing numbers of high- and low-skilled jobs across the United States and Europe (Autor and Dorn 2013; Autor et al. 2006; Goos et al. 2014). The spread of physical robots mostly affected routine manual tasks. For example, evidence from advanced economies shows that the share of hours worked by low-skilled workers fell with increased use of industrial robots, while the share of hours for middle- and high-skilled workers rose (Graetz and Michaels 2018). AI has increased capabilities in many areas, such as prediction, speech recognition, image processing, and content generation. It is therefore not surprising that non-routine cognitive tasks are more exposed to generative AI than routine manual tasks (refer to figure 2). This impact spans white-collar occupations that require advanced degrees, such as software developers, financial analysts, designers, and marketing professionals, and is concentrated among high-skilled services industries such as finance, accounting, insurance, and legal services (Bick et al. 2024; Felten et al. 2023). The exposure of a task to AI could result in either automation, such as AI replacing a human in that task, or augmentation, such as AI complementing a human in that task.

Figure 2 Non-routine cognitive tasks are more exposed to generative AI than routine manual tasks



Sources: WDR 2026 team based on Gmyrek et al. 2025.


Note: The bubbles capture the share of employment at the 2-digit ISCO-08 level for the 46 countries with available data. The larger the bubble, the greater the share of employment in that occupation category. The 2-digit ISCO-08 level refers to occupations listed at the 2-digit level of the International Standard Classification of Occupations published in 2008. The x- and y-axes show two task content scores. The scores are harmonized following Caunedo et al. (2023) and computed at the 2-digit ISCO-08 level across four data sources: STEP 2012, STEP 2013, STEP 2015/16, PIAAC first wave, and PIAAC second wave. All scores across all countries are standardized to have a mean of zero (highlighted by the two gray dashed lines) and a standard deviation equal to that of the United States. Here, the routine manual task content score is reversed (that is, when routine manual task content = 0.5, it is 0.5 standard deviations below the US level). PIAAC = Programme for the International Assessment of Adult Competencies; STEP = Standard for the Exchange of Product Model Data.

Access barriers are likely high for frontier AI innovation but low for AI adoption

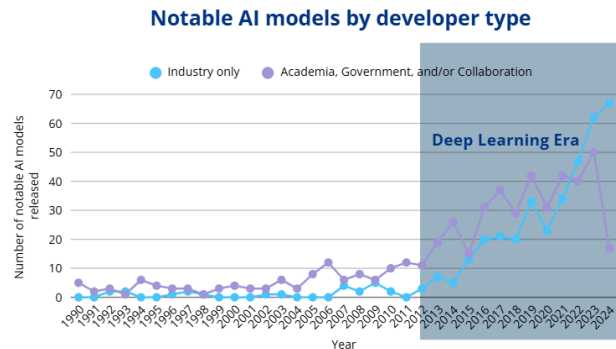
Frontier AI model development is almost entirely based in high-income countries and in China (refer to figure 3, panel a). This pattern is not dissimilar to innovation in previous waves of technological change. Within high-income countries, frontier AI model development is increasingly concentrated among firms (refer to figure 3, panel b). These are largely the AI research labs in big tech companies (including OpenAI, Anthropic, Google, Microsoft, Meta, DeepSeek), largely headquartered in the United States and China. This concentration among a few firms is explained by the growing financial cost of developing frontier AI systems that requires hardware inputs such as semiconductors, massive computing infrastructure, large amounts of big data, and engineering talent to gain an advantage (Çorba et al. 2024; Kaplan et al. 2020; Maslej et al. 2025).⁷ Many of these big technology companies are vertically integrating or investing in other companies across the entire AI value chain, which includes activities such as customizing hardware components, constructing data centers, developing frontier LLMs, and deploying end-user applications. There is early evidence that the returns to scaling in driving AI advances might diminish and reduce access barriers to frontier AI innovation in the future.⁸

Figure 3 AI innovation is concentrated in high-income countries and in the private sector

a. Number of notable AI models (cumulative, lead country method)



b. Prominence of private companies in developing most modern AI models



Sources: Panel a: Maslej et al. 2025. Panel b: Maslej et al. 2024.

Note: Panel a: The figure shows the number of total notable AI models released since 1960 (using the lead country method). AI models are identified as notable if they advanced the state of the art, achieved many citations in an academic publication, had more than one million monthly users, were highly significant historically, or were developed at a cost of more than one million US dollars. Panel b: The figure shows the number of notable AI models developed since 1990, broken down by the type of model developer—split between solely industry (private-sector) and all other types of developers, including academia, government, and/or a collaboration between multiple developer types, such as the private sector.

At the same time, access barriers to adopting AI are low, at least for generative AI. Generative AI advisory tools can be accessed through plain (local) language prompts in chatbots on smartphones, enabling use by people without formal technical skills or expensive hardware. In contrast, the use of personal computers typically needed specialized training, such as programming and spreadsheet literacy, as well as multiple software subscriptions and more expensive hardware (Brachman et al. 2025). Generative AI has also lowered entry barriers for building AI applications. Coding done by AI, such as through no-code software development platforms, enables the development of AI apps without sophisticated coding skills and computing power. People can code in plain English rather than a computer programming language and can do so accessing AI models through cloud platforms.

The local context is central for AI applications

The effectiveness of AI solutions, by construction, depends on the quality and relevance of the data they are trained on. This distinguishes AI from previous digital technologies. Much of the training data used to train the largest foundation models is sourced from the internet. And a handful of languages dominate most internet content. The imbalance is most pronounced in text training data. As of April 2024, 56 percent of open-source data sets on Hugging Face, a leading AI model repository, were in English (World Bank 2025a). Similarly, nearly 98 percent of scientific papers are estimated to be written in English, even by researchers from countries where English is not the native language (Gordin 2015). This imbalance in data availability has contributed to AI models being trained largely on data sets originating from high-income countries and China (Longpre et al. 2024). In contrast, more than 2,000 languages spoken in Africa are neglected in these LLMs. For example, ChatGPT recognizes only 10 percent to 20 percent of sentences written in Hausa, a language spoken by 94 million people in Nigeria (Wild 2025).⁹

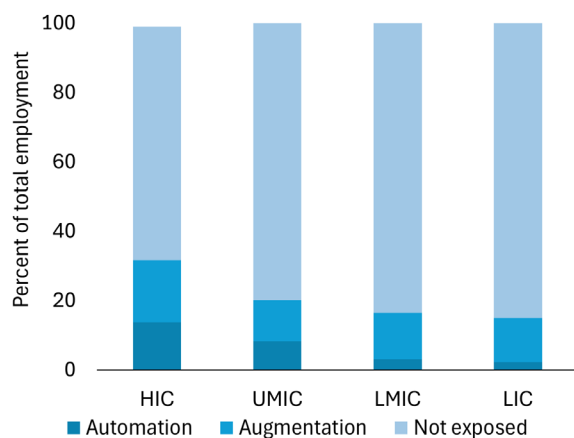
AI's promise for development

AI is more likely to complement workers than displace them in developing countries today


AI is more applicable to cognitive tasks than manual tasks, and cognitive tasks account for a smaller share of tasks in developing countries. It is thus not surprising that a smaller share of tasks in developing countries is exposed to AI (Cazzaniga et al. 2024; Demombynes et al. 2025). In fact, a lower share of jobs is exposed to automation by AI in developing countries than the share that is exposed to augmentation by AI (refer to figure 4, panel a). Further, a lower share of jobs is exposed to automation by AI in developing countries than in high-income countries (refer to figure 4, panel b).¹⁰ This implies that the augmentation potential of AI may not materialize in developing countries with low AI adoption, which has three implications.

Figure 4 The share of jobs exposed to AI automation is lower in developing countries than in advanced economies, and is lower than the share exposed to AI augmentation

a. Jobs exposed to GenAI (automation versus augmentation)



b. Jobs exposed to automation



Sources: WDR 2026 team based on Gmyrek et al. 2025 and World Development Indicators.

Note: Panel a: According to Gmyrek et al. (2025), automation = the share of employment with GenAI Exposure Gradients 1–2; augmentation = the share of employment with GenAI Exposure Gradients 3–4; not exposed = the share of employment with low/no GenAI Exposure. GenAI = generative artificial intelligence; HIC = high-income countries; LIC = low-income countries; LMIC = lower-middle-income countries; UMIC = upper-middle-income countries. In panel b, data labels use International Organization for Standardization (ISO) country codes.

However, the process of structural transformation, whereby people typically move from less to more productive activities and from less skilled to more skilled tasks, might encounter greater hurdles, such as with call centers and business process outsourcing. Second, AI can augment the capabilities of workers in developing countries, thereby raising efficiency. This could be manifested in better-paying jobs and higher productivity among many small businesses, improved public service delivery (such as tax collection, social protection, and grievance redressal) that especially matters for the well-being of the poor, as well as higher-quality education and health care, which improve income-earning opportunities in the future. In this way, AI can accelerate progress toward overcoming development challenges. Third, realizing the potential of AI to augment workers' capabilities in developing countries depends on the actual adoption of AI.


AI use might diffuse faster than previous GPTs

The adoption lag between high- and lower-income countries has been shrinking with each wave of GPTs. Evidence suggests that this technology adoption lag was cut from about 80 years for the steam engine to 40 years for electricity, and then to about 20 years for computers and the internet (Comin and Mestieri 2018). The low access barriers to AI adoption means that the use of AI tools has the potential to diffuse rapidly from high- to lower-income countries.¹¹


When ChatGPT was launched on November 20, 2022, more than 70 percent of user traffic originated from the United States, where the model was developed, but within six months, middle-income countries collectively accounted for half of ChatGPT traffic following the release of GPT-3.5 (Liu and Wang 2024). India, Brazil, the Philippines, and Indonesia comprised four of the top five countries in terms of ChatGPT traffic by the end of March 2024 (refer to figure 5, panel a). Middle-income countries also exhibit disproportionately high AI adoption compared to other economic indicators (refer to figure 5, panel b). For instance, India’s GDP is 14 percent of US GDP, India’s electricity consumption is 35 percent of US electricity consumption, and India’s internet search engine traffic is 30 percent of US search engine traffic; yet India’s generative AI website traffic is more than 50 percent of US generative AI website traffic.

Figure 5 Adoption of generative AI by middle-income countries is disproportionately high

a. GenAI traffic, by country, top 20



b. GenAI traffic, search engine traffic, electricity consumption, and GDP of leading countries relative to the United States



Source: WDR 2026 team based on Liu and Wang 2024.

Note: In panel a, data labels use International Organization for Standardization (ISO) country codes. GenAI = generative artificial intelligence.

AI models can be adapted and customized to context

Adapting and customizing AI models to incorporate local data are likely needed for AI to deliver value and address bias in diverse contexts (Christian 2020; Liu and Wang 2024). The availability of open-source foundation AI models¹² developed in advanced economies is a prerequisite to facilitate such adaptation. Cloud-based software development platforms, such as Hugging Face and GitHub, have expanded access to open-source AI models by allowing developers to use them over the internet (Bommasani et al. 2021). Further, these “low-code” or “no-code” software development platforms allow domain experts and non-technical users to interact with AI systems using plain language rather than software code (Brachman et al. 2025). This has enabled smaller organizations, start-ups, and research institutions to build AI applications without frontier-scale infrastructure and skills.

Recent developments show that smaller open source AI models that rely on the judgments of the original AI model but are trained efficiently on local data can match or even exceed the performance of much larger proprietary systems at a fraction of the cost (Abdin et al. 2024; Sevilla and Roldán 2024).¹³ These small AI models can also be run on laptops or smartphones, enabling offline use in low-connectivity settings (Hinton et al. 2015). This shift opens new opportunities for fine-tuned and mobile-ready models that are increasingly viable for real-world use, particularly in low- and middle-income countries.

Most developing countries are poised to benefit from AI adoption and adaptation rather than AI advancement

AI's value for transformational change can be leveraged along its value chain from developing core AI models (advancement), to developing and customizing AI solutions (adaptation), and using off-the-shelf AI tools (adoption). There is a hierarchy of costs around adopting, adapting, and advancing AI technologies. Investments in terms of internet connectivity, computing infrastructure, training data, frontier skills, and model testing are the least for adopting off-the-shelf AI solutions, the most for advancing core AI technologies, and somewhere in the middle for adapting AI solutions to local contexts.

Most low- and middle-income countries are less likely to be involved in AI advancement—that is, the development of frontier AI models—given the high requirements for computing infrastructure, training data, and talent. The choice will largely be between AI adoption and AI adaptation. On the one hand, the adoption of ready-made AI solutions is likely to be the most cost-effective way to expand the productive use of AI and enhance the capabilities of firms, workers, and governments. On the other hand, the adaptation of AI solutions can enable developing economies to extract more value from the AI value chain while also better aligning AI models to the local context. The role of small AI models that are trained on local data for specific use cases and can be used in low-connectivity settings will play a central role in the adaptation of AI solutions. Adaptation may also spur adoption, for example, if the AI solutions align better with local norms and regulations.


AI may deliver slowly on its promise for development

It is difficult to leapfrog previous GPTs

While AI may enable countries to leapfrog development challenges, leapfrogging waves of technological change is practically impossible. While leapfrogging over special-purpose technology is possible (such as skipping landline telephones in favor of mobile phones), it is difficult for countries to leapfrog general-purpose technologies. Technological advancement is a cumulative process, and countries must first catch up. The steam engine was necessary for firms to take advantage of electric power, and electricity was a precondition for information technology. Similarly, the deployment of AI models relies, in large part, on having access to electricity, internet connectivity, and affordable telecommunication services. It is hard to see countries leapfrogging into the use of AI without sufficient diffusion of electric power, internet connectivity, and affordable telecommunication services.¹⁴

However, the lack of digital infrastructure might constrain the wider adoption of AI in low- and middle-income countries where even the use of smartphones is far from universal (refer to figure 6). The same is true for the digital connectivity of businesses. For example, less than 50 percent of firms in developing countries use a website, up from 26 percent in the late 2000s. In contrast, more than 82 percent of firms in advanced economies use a website, up from 55 percent in the late 2000s (Nayyar et al. 2024). Similarly, more than 80 percent of advanced economies have an online government e-services portal, compared to only 38 percent of developing countries (Nayyar et al. 2024).

Figure 6 Access to smartphones is far from universal in lower-income countries



Source: World Bank Findex.

Note: Data labels use International Organization for Standardization (ISO) country codes.

Lessons from previous GPTs show that economic impacts of technology adoption are gradual

Evidence from previous GPTs shows that the wider diffusion of AI will likely be limited by the slow pace of change in organizational processes. For example, the productivity impacts of electric power in the United States were experienced long after the technology first came to the market. Notably, the pace of factory electrification was slowed by the cost of changing factory layouts and instituting the organizational processes necessary to engineer the technology switch (David 1990). Similarly, in nineteenth-century France, the productivity impacts of mechanized cotton spinning took time to materialize, given a lengthy period of “trial and error” in reorganizing production (Juhasz et al. 2024). More recently, the slowness in the rise in productivity that has accompanied the spread of computerization has been attributed to the need for complementary investments in organizational processes and other intangible capital (Bresnahan et al. 2002; Brynjolfsson et al. 2021).

The market dynamics of AI as a product might slow the adaptation of AI applications

Many developing countries currently lack the foundational building blocks that enable AI applications to be customized to the local context. Developing countries account for only 5 percent of high-performance computing capacity and only 12 percent of secure internet servers. In 2024, more than 85 percent of AI training data start-ups were located in high-income countries and less than 10 percent of venture capital funding in training data start-ups was being channeled to firms outside the United States and China (World Bank 2025a). Last, but not least, large flows of AI talent originate in China and India and converge in the US innovation ecosystem.

Further, the vertical integration of big technology companies across developing frontier LLMs and deployment of end-user applications might limit opportunities for developing AI applications in low- and middle-income countries. In particular, there is a risk of technological lock-in with early AI applications developed by the big technology companies. The high sunk costs and network effects associated with the choice of hardware, software, and algorithms can make these decisions hard to reverse. For example, when

a small number of AI solution vendors dominate, governments can become dependent on them for procurement and the dominant players can shape standards and regulatory frameworks (World Bank 2025c).¹⁵

The lack of trust due to the societal harms of AI may be a barrier to wider diffusion

Wider AI adoption might be limited if users mistrust AI. First, AI models can produce confident but incorrect outputs, known as hallucinations, especially in technical domains or when asked about recent events. In fact, AI solutions might perform well on recognized benchmarks but may fail when applied to specific tasks within firms. Second, similar to how the internet increased crime by making harmful acts easier to commit and harder to detect (Bhuller et al. 2013; Moore et al. 2009), AI can increase criminal activity through deep fakes and cybersecurity risks. AI can also amplify social media to spread propaganda and enable societal polarization. For example, predictive AI algorithms in social media platforms often prioritize content consistent with users' existing beliefs to maximize user engagement—which, in extreme cases, could even fuel civil unrest.¹⁶ Third, AI models can encode historical and structural biases, thereby perpetuating them (Bender et al. 2021). Further, AI models trained on English language data sets might be misaligned in certain contexts and cultures. Such risks are exacerbated in low-capacity environments, especially if AI is integrated into critical infrastructure or the delivery of essential public services. Fourth, the adverse environmental consequences of AI's development and deployment may slow down the pace of AI-related investments.

2. Leveraging AI for Development: Conceptual Framework

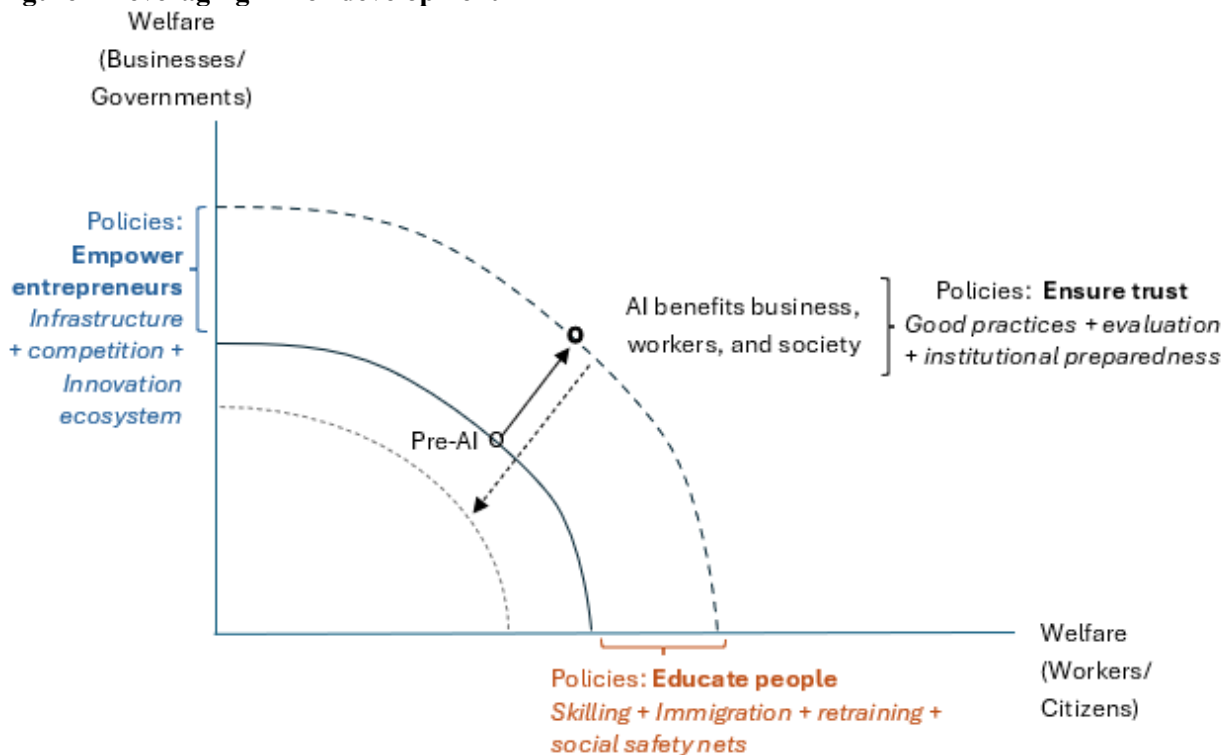
AI enhances capabilities across a range of cognitive tasks. As a result, the *productive use of AI* by firms, workers, and governments can, on average, improve welfare (refer to figure 7). However, *unequal returns to AI* can raise the welfare only of a certain group of individuals, such as owners of capital rather than people as workers and consumers. Further, the potential *misuse of AI* that results in social harm can lower societal welfare. These opportunities and risks are present along the economic (growth and jobs), social (government services), and sociopolitical (political) pillars of the development process.

In principle, the *productive use of AI* can improve the capabilities of firms, workers, and governments, which are typically lower in low- and middle-income countries than in high-income countries (this appears as an upward shift of the welfare possibilities frontier from the black curve to the dotted black curve in figure 7). These enhancements of capabilities create opportunities for economic, social, and political development. In the economic pillar, AI advice can, for example, raise the productivity of farmers and small businesses with few skills and limited expertise, while AI coding assistants can enable non-specialists to engage in software development. In the social pillar, AI can optimize government responses to food crises and guide the allocation of scarce resources. In the sociopolitical pillar, AI can democratize access to personalized advice that can raise people's welfare as consumers and citizens. In such ways, AI can help address long-standing market failures, fill in skill gaps, and bridge information gaps at lower costs than before. The resulting gains in capabilities can raise the welfare of both businesses, workers, and consumers or governments and citizens in low- and middle-income countries. At the macro level, stronger capabilities among firms, people, and governments can help countries leapfrog development challenges by boosting economic growth, creating jobs, improving the delivery of public goods and services, and bringing about wider societal transformation.

In practice, the *productive use of AI* in low- and middle-income countries can be facilitated by infrastructure requirements and market structures that govern the diffusion of AI. For starters, barriers to adopting off-the-shelf AI solutions are low; these solutions can be accessed through smartphones connected to the internet and with basic digital literacy. Further, there is a business opportunity to add value to goods,

services, and processes in low- and middle-income countries by developing AI applications: entrepreneurs can adapt open-source foundation AI models to customize AI solutions for the local context.

Figure 7 Leveraging AI for development



Source: WDR 2026 team.

Note: AI = artificial intelligence.

However, there may be *unequal returns to AI* that result in a concentration of benefits. In the economic pillar, AI might only benefit owners of capital if it makes firms more profitable but lower the welfare of workers, either because it results in job losses or lowers real wages. In principle, the diffusion of AI solutions that automate tasks in certain industries, such as call centers and business process outsourcing, in high-income countries could, in principle, lower the welfare of both businesses and workers in low- and middle-income countries if, for example, it reduces exports that derive on cheap labor as a source of advantage in global markets. In the social pillar, AI tools might disproportionately benefit government schools and health facilities with better infrastructure, skilled staff, and reliable connectivity. In the political pillar, the control of a few countries and corporations over high-value hardware inputs (including semiconductors), data centers, frontier AI models, as well as AI products and services can concentrate value among a few and create dependencies.

Government policies can play a crucial role in enabling countries to leverage the productive use of AI to raise welfare and ensure that as many people as possible have the necessary tools to benefit. The first set of policies focus on empowering businesses and entrepreneurs. First, the adoption of AI solutions depends on electricity and internet infrastructure but will need to be accompanied by organizational restructuring within firms to realize the full potential of AI in raising productivity. Second, enabling local AI innovation ecosystems, including through establishing access to key inputs such as compute infrastructure and data, can level the playing field for entrepreneurs to adapt AI solutions and increase the profits associated with AI use in low- and middle-income countries. Third, government procurement of AI solutions can also be a relevant instrument of industrial policy to jumpstart the development of local AI applications.

The second set of policies focus on skills upgrading. First, the large-scale adoption of AI solutions by workers will require the expansion of foundational literacy, numeracy, digital, and soft skills—whether acquired through formal schooling, in communities, or learning on the job. Second, the adaptation of AI solutions to local contexts will require investments in science, technology, engineering, and mathematics (STEM) education and AI-specific competencies (such as data science, ethics), as well as industry-academic partnerships to develop sector-specific AI skills. The immigration of high-skilled workers and engagement with their diaspora could also be a relevant tool. Third, governments may consider combining social safety nets with active labor market policies (ALMPs) such as reskilling, training programs and job search assistance, to support workers who are displaced by AI.

Implementing both sets of policies to raise the welfare of businesses and workers is necessary but it is not sufficient for AI to be able to raise the welfare of society. This is due to possible societal harms of AI use. In the economic pillar, a breakdown in trust over using digital payments, could lower welfare possibilities for both firms and workers in low- and middle-income countries (even relative to the pre-AI scenario) (*this appears as a downward shift of the welfare possibilities frontier from the larger dotted black curve to the smaller dotted black curve in figure 7*). In the social pillar, there is a risk of perpetuating biases that are encoded in past data when applied to sensitive citizen-oriented services, such as education, health, and judicial services. In the sociopolitical pillar, the substantial depletion of energy resources from the operation of data centers can raise environmental costs. Therefore, policies that establish societal trust around AI ethics will also be required to ensure that AI improves societal welfare. First, governments can adopt good practices from around the world and experiment with different regulatory approaches. Second, governments can prepare for the uncertain impacts of AI by investing in the evaluation of AI solutions, engaging multiple stakeholders, raising capabilities of bureaucrats, and fostering global cooperation to build trust in AI systems that operate across borders.

Each of these sets of government policies will differentiate between fostering AI adoption, adaptation, and advancement. In the first set of policies that focus on empowering businesses and entrepreneurs, for example, investments in electricity, internet, and telecommunication infrastructure can spur AI adoption while investments in data centers and open data can enable local AI adaptation. Similarly, in the second set of policies that focus on skills upgrading, broad-based skills development will spur AI adoption, while investments in AI-related technical skills will support local AI adaptation. In the third set of policies that focus on trust in AI, baseline testing of AI solutions to ensure compliance with local sectoral regulations can spur AI adoption, while frameworks for AI model testing and evaluation can support local AI applications.

The emphasis on AI adoption, adaptation, and advancement varies across the economic, social, and political pillars of the development process (refer to table 1). For the economic pillar, AI adaptation matters for entrepreneurs and startups that drive the development of local AI applications. However, the adoption of off-the-shelf AI solutions by large numbers of small businesses will matter most to bring widespread productivity gains. For the social pillar, AI adoption and AI adaptation go together to improve the delivery of government services because AI solutions are applied to sensitive government functions and for use cases where local context is crucial. For the political pillar, AI advancement also becomes relevant to address concerns around over-dependence and sovereignty.

Table 1 AI adoption, adaptation, and advancement

		DEVELOPMENT PILLARS		
		Economic (growth and jobs)	Social (government services)	Political (political economy)
TYPE OF AI SOLUTION	Adopt (using off-the-shelf AI software)	Most	Most	Yes
	Adapt (building AI applications)	More	Most	Yes
	Advance (establishing building blocks to develop frontier AI models)	Less	Less	Yes

Source: WDR 2026 team.

3. Empirical Contributions

Data on the use and adoption of AI technologies in developing economies remains scarce. To address these critical gaps, the empirical contributions of this Report will be based on new sources of data that track the use of AI—a considerable value added to the thin evidence base in this field. There are several initiatives underway in this regard.

First, to better understand the use of AI within governments, the WDR 2026 team is undertaking a survey across 20 countries, spanning all World Bank regions.¹⁷ The survey includes questions on the use of AI within government, the types of AI adopted or applied to core government data systems, the enablers of AI (from strategies to data-readiness) and potential barriers, the modalities of implementation, and the perceptions and expectations about the productivity potential of the technology. The survey will follow the World Bank’s *Data for Better Governance* methodology and will be implemented by a consortium of World Bank staff, University College London, and local public administration associations. The team will also leverage ongoing country-level surveys being conducted by World Bank staff on governance (such as in Brazil and Malaysia, and possibly other countries) and health facilities (in several countries in Africa) to better understand how public servants use AI in their daily work. In addition, the team will compile a global database of AI use cases in the public sector by merging existing datasets with newly collected information.

Second, to better understand the adoption of AI technologies within firms, the team is undertaking a follow-up survey to the World Bank Enterprise Survey (WBES) in eight countries, also spanning all World Bank regions.¹⁸ This survey includes questions on the use of AI within businesses, the types of AI adopted, the tasks performed by AI technologies, the modalities of implementation, and the implications of AI on jobs, as well as on the expectations of the productivity potential of the technology.¹⁹ Following the WBES methodology, the survey covers formal establishments with five or more employees in the manufacturing sector and selected services sectors and is based on representative samples obtained from statistical agencies, business registers, and other business listings (depending on data availability). Stratification will be applied based on firm size as well as sector, separating sectors with high degrees of AI intensity, such as knowledge-intensive services and high-tech manufacturing, from other sectors.

Third, to measure the creation of new jobs related to AI, the team will analyze online job openings in more than 80 countries. These data will be used to identify jobs requiring AI-related skills, estimate the wage

premium associated with AI skills across occupations, and measure the relationship between AI exposure and labor demand. Finally, the team will match these job postings data at the firm level to data on AI and digital technology investments to evaluate the drivers of AI adoption and job creation at the firm level. The team will also expand the work on exposure to AI based on task and occupational level data, drawing on detailed microdata from 135 countries, covering 69 percent of the global population, providing detailed demographics of potentially affected groups. In estimating AI exposure using a sample of 46 countries, this work will also use the PIAAC (Programme for the International Assessment of Adult Competencies) and STEP (Skills Toward Employment and Productivity) skills surveys to account for systematic differences in task content within the same occupation across countries.

Fourth, to better understand the impacts of AI on human capital accumulation, the team is carrying out systematic research syntheses of the evidence in the literature that is nascent yet rapidly expanding.

Fifth, in partnership with the Development Data Partnership, the team is currently exploring the development of a Global Generative AI Adoption Index—a standardized, country-level indicator that will measure the adoption of generative AI by consumers, developers, and potentially enterprises.

Lastly, to better understand existing policy initiatives around AI, the team will analyze existing repositories of national AI strategies to document their distribution around the world, as well as the focus of their content. For a selection of low- and middle-income countries, the team will carry out an in-depth, qualitative analysis of these national AI strategies along different themes. These themes include objectives for AI adoption in the public sector, ethical concerns and responsible AI development/adoption, regulatory orientation, international cooperation approaches, public investment and funding mechanisms, business environment strategies, and skills development measures.

Part II: Assessing AI's Implications for Development

AI can affect people's welfare through its impact on economic growth and jobs (markets), the delivery of government services (state), and broader sociopolitical change (society). These impacts fundamentally draw on the role of AI in enhancing capabilities of people and on the market dynamics of AI as a product. This section of the Report will assess the evidence around the impacts of AI on markets, state, and society to bring about economic, institutional, and societal transformation.

First, the productive use of AI can enhance the capabilities of firms and workers to raise labor productivity and income-earning opportunities. This potential is especially salient in low- and middle-income countries, where capabilities are relatively low, and AI can enable large numbers of lower-skilled people to complete higher-skilled tasks. However, if constraints to the use of AI are higher in lower-income settings with weaker infrastructure, many small (often informal) firms, and large numbers of workers without basic literacy and digital skills, the potential of AI to boost growth and job creation may remain unrealized. Further, if AI use displaces workers from certain jobs or increases the returns to capital relative to wages and the returns to skills, this can result in a concentration of benefits.

Second, the productive use of AI can similarly enhance the capabilities of government officials to provide public goods and services. This potential is again especially salient in developing countries, where institutional capacity is relatively low. For example, AI can improve the efficiency of government services by guiding the allocation of scarce resources, improving transparency, and supporting citizen engagement. However, the benefits of AI use will be limited without accompanying organizational redesign that is particularly challenging in public administration. Similarly, AI tools will disproportionately benefit government schools and health facilities with more dependable infrastructure, skilled staff, and reliable internet connectivity.

Third, AI can impact the broader process of sociopolitical transformation. For one, AI can accelerate the production of ideas where each new discovery requires more researchers and resources. For another, AI can democratize access to customized information. AI tools can also intrinsically raise consumer welfare and reshape the nature of social interactions. On the flipside, AI's diffusion poses many risks to society. For instance, AI can erode the accumulation of fundamental cognitive learning and reinforce social biases in decision-making. AI can also significantly raise environmental costs, obfuscate political processes by enabling surveillance and the spread of propaganda, as well as amplify geopolitical contestation with its "winner-takes-most" market dynamic. The possibility of AI-related social harm is higher in developing countries where institutions and state capacity are weaker.

1. Economic Growth and Jobs

Improving the productivity of businesses and enlarging the pool of well-paying jobs is the immediate priority to raising living standards in low- and middle-income countries. Like with previous technologies, AI will displace workers in certain tasks, expand jobs in other tasks, and create new tasks, with displacement less likely in low- and middle-income countries given the current sectoral composition of jobs. At the same time, AI could increase the capabilities of large numbers of low-skilled workers employed in many small, unproductive firms. AI also has considerable potential to spur innovation and entrepreneurship. On the flipside, there is a risk that AI could reduce growth opportunities in specific industries, such as knowledge-intensive services, or that it could have limited effects on aggregate productivity due to weak management practices and insufficient organizational restructuring within firms. There is also the possibility that AI might benefit capital owners more than workers, and skilled workers more than unskilled workers.

Opportunities for productive use

AI will displace workers in certain tasks, expand jobs in other tasks, and create new tasks

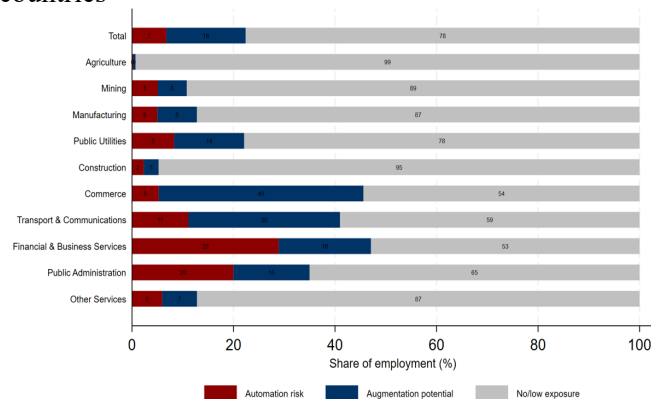
In previous waves of technological change, the negative displacement effect of automation has typically been offset by productivity gains that increase the demand for labor in complementary and new tasks (Acemoglu and Restrepo 2019).²⁰ There is some evidence of worker displacement with AI. For example, AI-related hiring in India has been linked to declines in job postings and wage offers for non-AI job postings within firms; these effects are most pronounced in managerial and professional occupations (Copestake et al. 2024). Similarly, evidence shows that AI may automate entry-level tasks (Brynjolfsson et al. 2025). At the same time, new tasks and occupations, including data scientists, machine learning specialists, and "prompt engineers" have emerged during the past three years (World Bank 2025a). These high-paying AI-related jobs remain scarce and concentrated in a small number of firms, cities, and countries. A larger number of low-paying AI-related gig jobs, such as data labelers, are mainly found in developing countries. Lastly, the share of tasks likely to be augmented by AI is higher in sectors, such as retail services and transportation services, that are large employers in developing countries. This suggests that AI-related productivity gains that increase the demand for labor in complementary tasks are likely to offset the negative displacement effect of automation in certain tasks.

AI can improve firm capabilities and worker productivity


In developing countries, many people are employed in small (and often informal) businesses that are unable to upgrade their capabilities and grow. Generative AI advisory tools can potentially be diffused widely and deliver benefits to small businesses with few skills and limited expertise but with access to smartphones and broadband internet connectivity. In the agriculture sector, for example, the Virtual Agronomist generative AI chatbot, implemented through WhatsApp, has increased profits and yields across seven African countries (World Bank 2025b). Retail services and transportation services are other large employers of low-skilled workers in developing countries and are highly exposed to AI. Much of this AI exposure carries the promise of augmentation (refer to figure 8). Generative AI tools have also improved productivity of larger businesses among knowledge-intensive services. In the Philippines, for example, access to an AI-based conversational assistant increased the productivity of customer service agents, as measured by issues resolved per hour, by 15 percent (Brynjolfsson et al. 2025).

Figure 8 Exposure to AI varies considerably by sector in developing countries

a. GenAI exposure by sector, upper-middle-income countries



b. GenAI exposure by sector, low-income and lower-middle-income countries



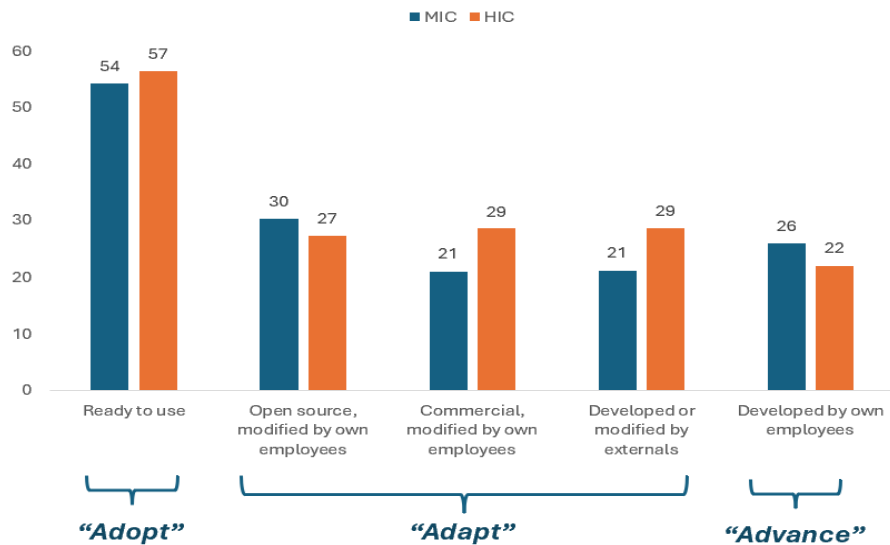
Source: WDR 2026 team based on Gmyrek et al. 2025.

Note: GenAI = generative artificial intelligence.

The need for relatively low investment by firms to access AI advisory tools is similar to other technologies that were able to spread quickly (such as mobile money) but is very different from previous technologies, such as computerization and robots, where diffusion was slower. Firms mostly adopt off-the-shelf AI solutions, with some adaptation through customization (refer to figure 9).

Figure 9 The diffusion of AI is led by ready-to-use AI solutions

Percentage, within firms using at least one AI technology



Source: Eurostat.

Note: Data are for various European countries as of 2024. Middle-income countries (MICs) include Bosnia and Herzegovina, Montenegro, Serbia, and Türkiye, as well as Bulgaria (which graduated to high-income status in 2025). High-income countries (HICs) include all European Union member states, Liechtenstein, and Norway. AI = artificial intelligence.

AI can lower the barriers to entry for workers to complete cognitive tasks

The expectation is that generative AI could automate certain highly skilled, decision-making tasks such as medical care, legal services, and software coding, potentially reinstating middle-skilled jobs that were previously diminished by ICT usage in advanced economies (Autor 2024). There is some evidence that suggests that AI may shift demand toward less-experienced workers by transferring tacit knowledge from experts, enhancing productivity, and narrowing performance gaps (Dell’Acqua et al. 2023; Noy and Zhang 2023).

For example, access to a GenAI-based conversational assistant was associated with an average productivity increase of 14 percent among customer support agents based around the world but mainly in the Philippines, with novice and lower-skilled employees experiencing gains as high as 34 percent (Brynjolfsson et al. 2025). Similarly, randomized controlled trials at large tech companies show that, on average, developers with access to an AI-based coding assistant increased their productivity by 26 percent, with the less experienced developers experiencing larger productivity gains (Cui et al. 2024). In an online experiment, college-educated professionals provided access to ChatGPT completed their writing tasks faster and produced higher-quality work than those without access (Noy and Zhang 2023). Advances like these could benefit workers in developing countries. At the same time, there is some evidence that AI could amplify the output of the top performers (Otis et al. 2024).

AI can spur innovation and entrepreneurship

Generative AI tools can support entrepreneurship by alleviating barriers to information, while predictive AI can improve access to credit. For example, evidence from India shows that the use of ML algorithms to develop an alternate credit scoring based on individuals’ digital presence resulted in lenders providing credit to more people (Agarwal et al. 2019; Björkegren and Grissen 2020). AI use can also spur innovation.

Among other opportunities, the possibility of low-cost (no code/low code) software development enabled by AI assistants creates ways for start-ups and entrepreneurs to develop AI applications that are suited to local contexts.

AI can also benefit traditional sectors, including manufacturing and agriculture, by spurring innovation. For example, the ability of ML algorithms to process large amounts of data to predict the efficacy and toxicity of drug compounds with a high degree of accuracy has improved the efficiency of processes to discover new drugs and medicines (Blanco-Gonzalez et al. 2023). The same is true for the process to discover new crops in agriculture. ML algorithms enable the analysis of large amounts of data collected from satellites, drones, and ground-based cameras to accurately identify desirable traits such as plant height, leaf area, and disease symptoms for crop breeding programs.

Risks

AI may reduce the jobs dividend of certain sectors

ICT, professional services, and business services are among the sectors that are most suitable for machine learning (Nayyar et al. 2021). These knowledge-intensive services are also among the sectors most exposed to automation by generative AI. Predictive and generative AI, by performing a range of cognitive tasks, might dampen the prospects of developing countries exporting back-office ICT and professional services because they reduce the importance of labor costs in determining competitiveness. This may reduce well-paying jobs in these knowledge-intensive services, which has previously enabled a path of structural transformation.²¹

At the same time, AI can lower communication costs, enable movement up the knowledge-intensive services value chain (for example, from outsourcing business processes to outsourcing knowledge processes), and facilitate greater offshoring to countries where English is not the native language through machine translation. There is also a possibility that key knowledge-intensive parts of the AI value chain can be produced offshore in countries with abundant and lower-cost STEM workers, like India, raising demand for labor in these countries. Indeed, India is one of the countries with the highest share of new jobs demanding AI skills—a phenomenon that could be consistent with this hypothesis.

The evidence so far is mixed. Stapleton and O’Kane (2021) find that ML deployment has led to an increase in offshoring of services from the United Kingdom to lower-income countries, with the largest increase in ICT, business services, and professional services. This evidence is consistent with the fact that the productivity effect of AI that reduces trade costs has exceeded the labor substitution effect. Similarly, Bastos and Nayyar (2025) find no evidence of a negative effect of adoption of generative AI in the United States on services imports from developing countries. However, some emerging evidence shows that the use of generative AI tools has reduced services offshoring through online labor market platforms (Betai and Chen 2024; Hui et al. 2024).

AI might benefit owners of capital more than workers

AI can reduce the share of national income allocated to labor if it displaces workers from jobs. Even if workers do not lose their jobs, the productivity gains may not be shared equally by owners of capital and workers, either because wages do not rise as much as profits or because cost reductions are not passed on to workers. Further, AI can be used by large businesses to charge consumers the maximum they might be willing to pay for a service, such as through predictive AI that analyzes people's behavior on ridesharing apps. AI might also be used by large businesses to better target advertising to consumers, which does not necessarily increase productivity and lower costs.

In the past, the share of labor income was relatively stable even during an era of widespread automation (Drozd and Tavares 2024). Early evidence from the European Union shows that for every doubling of regional AI innovation, the labor share of income declines by 0.6 percent to 1.6 percent, potentially reducing it by 0.20 to 0.53 percentage points from an average of 52 percent. And AI has a particularly negative association with the labor income share of high- and medium-skilled workers, primarily through wage declines (Minniti et al. 2025). As AI becomes more widely adopted, it may augment workers and create new jobs, potentially reversing its initial impact on labor income shares at the aggregate level.

AI might benefit skilled workers more

Earlier forms of automation, such as with computerization and industrial robots, have been associated with skill-biased technological change. For example, a sizable body of evidence from high-income countries shows that computerization was associated with falling numbers of medium-skilled, middle-wage jobs and increasing numbers of high-skilled and low-skilled jobs (Autor and Dorn 2013; Autor et al. 2006; Goos et al. 2014). Cross-country evidence suggests that exposure to AI is higher in sectors, such as ICT, finance, and professional services, that employ a large share of skilled workers with higher average incomes (Liu 2024). This suggests that generative AI is currently less likely to negatively affect less-skilled workers. One study estimates that US workers' exposure to AI will reduce wage inequality—defined as the ratio of the 90th to the 10th percentile of wages—by 4 percent (Webb 2020).

However, other early evidence of generative AI suggests that skilled workers might benefit more from AI use. In the United States, high-wage occupations experienced employment and wage growth following the introduction of AI, while low-wage occupations experienced wage losses (Felten et al. 2021). Similarly, there is evidence of the negative effects of AI exposure on employment, especially for low-skilled workers in the services sector (Bonfiglioli et al. 2025). Across 16 European countries, there is a positive relationship between the use of generative AI and employment, particularly for highly skilled workers (Albanesi et al. 2023). Across countries, evidence shows that jobs demanding AI skills offer a wage premium of as much as 25 percent (World Bank 2025a).

The lack of complementary factors is likely to blunt AI's impact on productivity

The adoption of AI has led to large productivity gains when targeted to specific tasks like customer service interaction or language translation. However, evidence of impacts at the level of the firm or the economy is more limited so far.

The benefits of AI adoption may be gradual for many reasons. First, weak market competition inhibits the reallocation of resources from less to more productive firms, which, in turn, lowers technology adoption. Similarly, weak entrepreneurship ecosystems might constrain innovation and entrepreneurship with respect to the development of AI solutions.


Second, even with the lower costs of adopting AI advisory tools, the lack of information or the lack of digitalization might constrain small businesses from realizing the benefits of AI solutions. For example, an AI assistant via WhatsApp benefited only high-performing small-scale businesses in Kenya by more than 15 percent, based on how entrepreneurs selected and implemented the AI advice they received (Otis et al. 2024). Similarly, smaller firms might benefit less from AI solutions either because they do not use complementary digital technologies across a range of business functions or because they lack multiple devices that power AI use.

Third, productivity gains will only be realized with sufficient organizational restructuring and stronger firm capabilities that enable the more widespread diffusion of AI. Putting these enablers in place takes time. This situation is similar to that of previous GPTs (Agrawal et al. 2023). Many firms in developing countries are characterized by less structured managerial practices (Bloom et al. 2012), persistently small size (Bento and Restuccia 2017), and informality—which typically inhibits the degree of technological adoption.

2. Government Services

Improving the delivery of public goods and services is central to the welfare of people. Yet governments in developing countries face persistent challenges in completing core government functions such as taxation and regulation, and service delivery such as health, education, social protection, and transportation, given tighter fiscal space, critical staffing gaps, and lower administrative capacity compared to high-income countries. AI holds considerable potential to address these persistent bottlenecks in government service delivery. Take the example of education and health services that have been plagued by critical shortages of trained teachers and health care providers, especially in rural and remote regions.²² The public administration sector has a larger set of tasks exposed to both automation and augmentation by (generative) AI than other sectors of the economy (refer to figure 10). This is indicative of AI’s potential to increase productivity in the delivery of government services by raising the capacity of public officials and frontline workers. In doing so, AI can detect patterns to guide the allocation of scarce resources, enable greater transparency and accountability, and expand engagement with citizens through customized digital solutions. However, AI’s potential to improve the delivery of government services may remain unrealized without the necessary skills, infrastructure, and organizational restructuring. The use of AI might also exacerbate biases in government decision making as well as raise dependency and accountability risks.

Figure 10 Public administration has greater exposure to AI than other sectors



Sources: WDR 2026 team based on Gmyrek et al. 2025.

Note: AI = artificial intelligence.

Opportunities for productive use

AI can detect patterns and generate forecasts to guide the allocation of scarce resources

Predictive AI is well suited to solving allocation and coordination problems that governments routinely face, such as efficiently dispensing social benefits, prioritizing inspections, or managing workloads under resource constraints. For example, during the COVID-19 pandemic, machine learning algorithms were combined with satellite imagery and cell phone metadata to improve the targeting of beneficiaries (Lawson et al. 2023). Similarly, machine learning can improve demand forecasting and optimize investment decisions in transport and energy systems. Computer vision algorithms use property images to improve property tax valuations in Pakistan (Asad et al. 2019). Similarly, machine learning models incorporate real-time news streams to forecast food crises up to a year in advance (Balashankar et al. 2023).²³

AI can improve transparency and accountability

AI tools can alleviate institutional barriers in the delivery of public services, by automating oversight, increasing transparency and accessibility, and flagging anomalies in government procurement. For example, machine learning algorithms can detect irregular patterns in procurement transactions, verify supplier credentials based on historical performance, and track whether goods or services procured (such as school construction and road works) are delivered as specified. Brazil's procurement monitoring system, *Alice*, demonstrates this approach by collecting daily information about federal procurement processes, evaluating risks, and raising red flags for auditors and managers to investigate (OECD 2024).

AI can customize digital solutions for service delivery at a lower cost

AI can personalize earlier digital innovations for service delivery, which have demonstrated significant potential to improve, for example, education and health outcomes.²⁴ AI-powered tutoring, with targeted instruction and personalization, has proven effective in improving mathematics' scores in Ghana (Henkel et al. 2024) and English language outcomes for secondary school students in Nigeria (De Simone et al. 2025). In Kenya, *Jacaranda Health PROMPTS*, an AI-enabled SMS (text message) service that has reached

2.74 million mothers across more than 1,200 health centers, provides personalized pregnancy advice, and triages high-risk cases. Further, AI is improving early detection and intervention in low-resource settings, for instance by assisting community health workers where access to specialized labs is limited (Adapa et al. 2025; Khan et al. 2022). AI solutions can not only personalize service delivery, but also reduce barriers related to digital access. For example, evidence from Sierra Leone shows that generative AI chatbots offer a cost-effective alternative for teacher support—87 percent cheaper than a web search (because they can be operated on basic mobile devices and low-bandwidth networks—or even no-data networks)—and are perceived by teachers as providing more relevant, helpful, and accurate responses than web search results (Björkegren et al. 2025).

AI use by governments can have economy-wide gains

AI solutions that bring greater efficiency in the delivery of government services, such as streamlined tax administration, improved infrastructure planning, and responsive social services, will yield economy-wide productivity gains because they enable work in the private sector. For instance, predictive AI using real-time sensor data may anticipate infrastructure failures in grids, water systems, and transport networks, reducing maintenance costs and increasing the quality of public infrastructure. Further, government procurement of AI solutions can spur AI adaptation whereby local entrepreneurs adapt global AI models to the local context. Procurement choices also influence technology standards, guide private investment, and set norms that extend well beyond government. If governments lag or apply AI poorly, these actions become bottlenecks, undermining trust, slowing diffusion, and dampening private AI investment.

Risks

The lack of complementary factors can blunt AI’s impact on improving government services

The adoption of AI solutions may not be sufficient to bring widespread gains in the delivery of government services without the necessary infrastructure and skills. For example, it is possible that AI tools will disproportionately benefit better-resourced schools and health facilities—that is, those with dependable infrastructure, skilled staff, and reliable connectivity (Cross et al. 2024). Similarly, it is likely that countries with robust digital public infrastructure (DPI), such as digital IDs, verifiable digital credentials, and digital payment wallets, will be better able to enable trusted transactions and scale up the use of AI in the delivery of government services.

The role of accompanying organizational reforms is just as important. This is illustrated by experimental evidence from the National Bank of Slovakia, where generative AI tools improved individual workers’ performance, but organization-wide productivity gains remained constrained because of the way workers were assigned to tasks (Marsal and Perkowski 2025). Such organizational redesign is particularly challenging in the public sector, where factors like employee tenure and institutional inertia often increase resistance to change (Finan et al. 2017; Kleiner et al. 2016). Legal frameworks are also often unprepared for the inclusion of AI in decision-making. Furthermore, organizational culture that often includes risk-aversion and hierarchical decision-making structures may dampen the adoption of AI in the public sector. These structural and cultural barriers mean AI tools are frequently integrated into existing, inefficient processes rather than leveraging its potential through institutional restructuring.

AI adoption might exacerbate biases, dependency, and accountability risks

Human bias embedded in different sources of data or poor data quality can increase the risk of perpetuating biases with algorithms which draw from these data. This is particularly problematic for predictive AI that is used in sensitive citizen-oriented services, such as education, health, and judicial services. For example, medical imaging AI has been shown to exhibit “hidden stratification,” performing well overall but missing critical subgroups—such as aggressive disease types—by more than 20 percent (Oakden-Rayner et al. 2019). There is also the risk of algorithmic bias based on using AI models trained on data from abroad with a different linguistic and cultural context. For example, an evaluation of the feasibility of deploying AI models developed in the United Kingdom within hospitals in Viet Nam found that model performance was low (Yang et al. 2024). Similarly, the integration of an LLM for clinical decision support in Nigeria tended to over-recommend laboratory testing, which is more in line with medical practice in high-income countries (McPeak et al. 2024).

Further, governments may lack domain-specific expertise or the training data to fine-tune and contextualize AI models. As a result, governments might become dependent on proprietary software controlled by a small set of firms, creating lock-in, and limiting bargaining power in sensitive government functions. In contexts with weak state capacity, AI use may also lead to overreliance, which could further erode institutional learning and accountability, shifting discretion away from public officials toward opaque algorithms or private vendors.

3. Sociopolitical Change

AI has broader implications for society beyond its impact on economic growth, jobs, and the delivery of government services. AI is rapidly expanding humanity’s capacity to generate new ideas. AI is democratizing access to personalized information. AI is also reshaping social interactions and redefining the way that people use their time beyond the workplace. At the same time, AI might impose significant costs on society through, for instance, obfuscating information flows, eroding cognitive abilities, and damaging the environment. This raises questions of trust around AI systems and use, especially in low- and middle-income countries where institutional capabilities are typically weaker. Further, AI is a source of power for those who can control or exploit it. Companies and countries with access to critical and rare minerals, data, computing resources, LLMs, and engineering talent, can wield extensive influence, with wide-ranging implications for geopolitics.

Opportunities for productive use

AI can accelerate the creation of new ideas

Economic growth arises from people creating ideas. However, there is evidence that the production of ideas has been slowing down because each new discovery requires more researchers and resources (Bloom et al. 2020). AI has the potential to ease this constraint. For example, recent AI advances have enabled the successful prediction of protein folding structures, which was a 50-year-old challenge in the field of molecular biology (Wang et al. 2023). Although frontier scientific research often takes place in high-income countries, there is the possibility of knowledge spillovers to low- and middle-income countries when new ideas in areas such as agriculture, medicine, and climate science, can be adapted to the local context.

AI is expanding consumer benefits and reshaping social interaction

AI products and AI-enabled goods and services are expanding consumer benefits by democratizing access to information and content that is tailored to individual preferences and needs. For instance, entertainment platforms are increasingly employing AI-driven recommendation systems to curate playlists and viewing suggestions that reflect user habits, thereby increasing consumers' satisfaction. AI chatbots are also reshaping social interaction. For example, in June 2025, among 700 million weekly active (non-enterprise) users of ChatGPT, approximately 70 percent used ChatGPT for non-work purposes, up from an equal split between work- and non-work usage in June 2024. Notably, "practical guidance" accounted for the biggest category of ChatGPT conversations, at roughly 29 percent of overall usage (Chatterji et al. 2025).

Risks

AI can make it easier to deepen repression, spread propaganda and enable polarization

While surveillance technologies, such as closed-circuit television (CCTV) cameras and biometric systems, have existed for decades, AI-enhanced surveillance (such as facial recognition systems, predictive policing, social media monitoring) offers unprecedented scale, precision, and anonymity for governments to deepen political repression and social control (Acemoglu 2024; Tirole 2021). Evidence suggests that local unrest increases government procurement of facial recognition AI and that governments are more likely to import surveillance AI technologies during periods of domestic political unrest (Beraja et al. 2023). Further, while disinformation techniques have long existed, advances in generative AI now make it easier to create propaganda (such as deepfakes, automated content generation), and target and spread it (such as through social media bots), potentially resulting in societal polarization (Acemoglu 2024; Harari 2024). A recent Freedom House report documents that generative AI was used in "at least 16 countries to sow doubt, smear opponents, or influence public debate" (Funk et al. 2023).

AI can reinforce social biases or lead to discriminatory outcomes

The proliferation of AI solutions may perpetuate existing biases or result in discrimination, especially in low- and middle-income countries given the lack of institutional capacity, lack of technical skills, and the dearth of high-quality data. For example, predictive AI tools can systematically exclude job applicants who do not align with biased historical hiring data and perpetuate biases through feedback loops (Chen 2023). Firms may also use AI to monitor workers' performance to improve productivity, which could come at the expense of workers' well-being if used excessively, especially in the informal sector with weak labor protection (Beraja and Yuchtman 2025). Further, the advancement of generative AI systems comprises tasks such as data labeling and content moderation that are essential for aligning algorithms with safety goals (Roberts 2019). Much of this work is outsourced to developing economies and exposes workers who already face poor working conditions to additional occupational hazards (Greenbaum and Gerstein 2025). Evidence suggests that some of these workers have experienced mental health effects from extreme exposure to explicit and graphic content, with limited access to appropriate health services (Perrigo 2023).

AI could erode the accumulation of human capital over time

Over time, AI could reshape how people learn and build skills, creating risks to the development of human capital. As AI becomes more capable in writing, translation, and problem-solving, students and workers may be tempted to rely on it instead of developing their own abilities. This may weaken investment in literacy, communication, cognitive tasks, and critical thinking—skills that are hard to automate and central to leadership, collaboration, and innovation (Acemoglu and Autor 2011; Brynjolfsson and McAfee 2014). Education systems may also find it harder to assess genuine learning if AI is doing much of the work behind assignments and exams, which could reduce the value of formal education to signal quality in the labor market and lower incentives to complete formal education (Goldin and Katz 2008).

AI might have adverse consequences for the environment

The development and deployment of AI, especially generative AI, might have adverse consequences for the environment and the use of natural resources. For one, the growing demand for rare minerals needed to produce AI infrastructure and increasing volumes of AI-related electronic waste generated in the process may harm the environment and the health of local communities. The bulk of rare earth minerals such as cobalt, copper, gold, and lithium are sourced from developing countries and most electronic waste ends up in developing countries (Wang et al. 2024). Further, data centers used in the development and deployment of generative AI consume large amounts of electricity (O’Donnell and Crownhart 2025) which is already scarce in many countries. Consumption of these vital resources of power and water could also have negative spillover effects by reducing citizens’ and firms’ access to these necessities that enable livelihoods. Recent developments in technology and infrastructure are contributing to a reduction in the environmental footprint associated with AI. For example, many data centers are adopting renewable energy sources, with some even constructing dedicated solar farms to supply clean power directly.

AI can intensify geopolitical contestation

The transformative potential of AI has led countries and corporations to increasingly compete over control of its inputs (e.g., rare minerals, data, computing) and outputs (e.g., patents, information). While low- and middle-income countries provide the critical and rare minerals that power AI, much of the value accrues to large companies in high-income countries that use these raw materials to produce high-value hardware (including semiconductors) and, in turn, frontier AI models, products, and services. Aligning with these geopolitics offers low- and middle-income countries access to AI infrastructure, knowledge, trade, and investment, but it also carries the risk of technological over-dependence.²⁵ One alternative—developing all the building blocks (large data sets, computing power, and top talent) required for frontier AI innovation domestically—is much too resource-intensive. The alternative of developing local AI applications might also not be a level-playing field; frontier AI companies may have an advantage in deploying AI applications due to scale economies or ownership of either related assets like proprietary data or distribution channels like mobile phones.²⁶ How low- and middle-income countries navigate these challenges will influence the extent to which they benefit from AI.

Part III: Identifying Policy Priorities to Leverage AI for Development

AI is reshaping how firms compete, how people work and how governments deliver services. But realizing its promise for development will depend less on the technology itself and more on whether governments in low- and middle-income countries can coordinate the right mix of policies to address a range of barriers and bottlenecks.

First, the diffusion of AI solutions will not automatically translate into productivity gains in the production process where foundational digital and analog capacities remain uneven. Governments will need to empower entrepreneurs and businesses that can implement solutions on a large scale to enhance the *productive use of AI*. Government actions may include investing in foundational infrastructure, such as electricity and internet connectivity, and enabling local innovation so firms and entrepreneurs can customize AI solutions that better align with the needs and realities of domestic industries in developing countries. This government role also calls for considering AI through an industrial policy lens: supporting firms to develop and deploy AI applications that serve local markets while enabling access to data and computing resources across borders where needed.

Second, the diffusion of AI across a variety of tasks will not necessarily result in better-paying jobs or mass unemployment in low- and middle-income countries. For most workers, the question is whether they have the skills and enabling environment to use AI productively and safely. Governments will need to prioritize educating workers to address the potential *unequal returns to AI* and ensure that as many people as possible benefit from the technology.

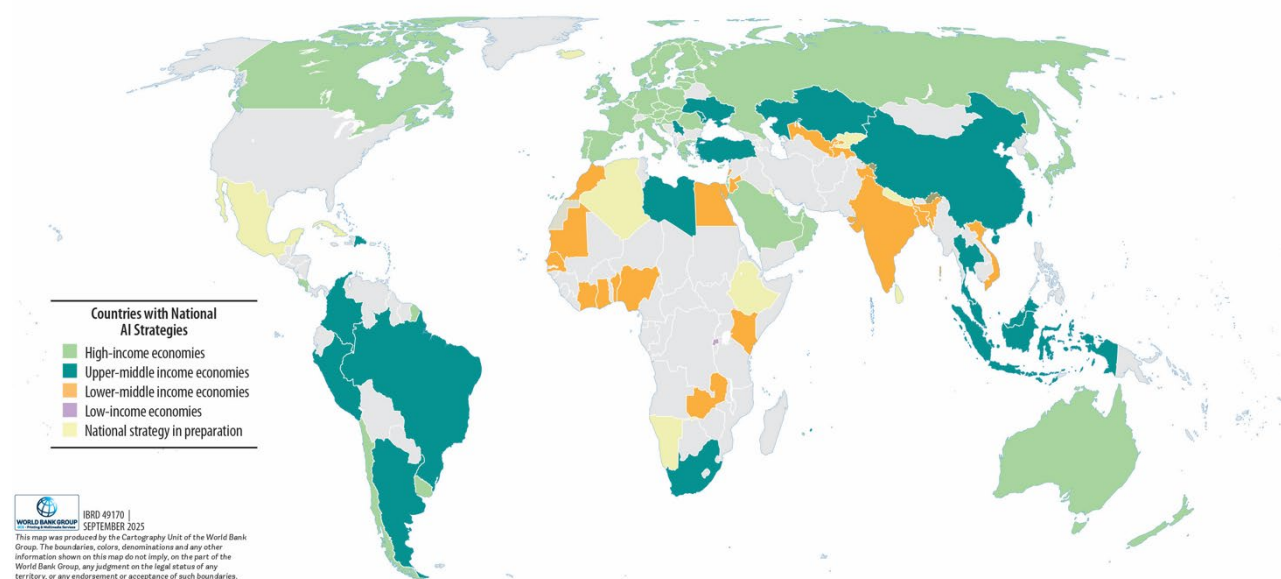
Third, the social, environmental, and political harm enabled by AI will undermine trust in AI and inhibit its diffusion to bring about welfare gains. Government policies will need to ensure trust in society around the deployment of AI to guard against the potential *misuse of AI*. This includes adopting good practices from around the world while incorporating safeguards to comply with local norms and regulations. This also requires stronger institutional capabilities that can help evaluate AI solutions and help countries adjust to changes in the underlying technology, market structures, and geopolitics that govern the deployment of AI.

In addressing barriers and bottlenecks to the *productive use of AI*, while addressing the risks of *unequal returns* and *misuse of AI*, countries need to assess the value of adopting off-the-shelf AI solutions on the one hand, and adapting or advancing AI solutions on the other hand. Their relative importance will vary by the economic, social, and political pillars of development. For the economic pillar around growth and jobs, the adaptation of AI solutions brings value to local entrepreneurs and startups and matters more than AI advancement (the development of frontier AI models) which depends on very large investments in computing infrastructure, training data, and talent. However, the adoption of ready AI solutions is likely to matter most to bring about widespread productivity benefits. For the social pillar around government services, the adoption and adaptation of AI solutions both matter equally to guard the possibility of social harm across sensitive use cases. For the political pillar around the AI value chain, AI advancement also matters due to concerns about over-dependence and AI sovereignty.

The role of governments is central to guide these choices. More than 80 countries have national AI strategies that are published, but there are stark differences between countries at different levels of per capita income (refer to map 1). For example, only Rwanda among 26 low-income countries had published a dedicated AI strategy compared to more than half of high-income countries, as of June 2025. Since 2021, countries around the world have increasingly turned to the goal of “AI sovereignty” in their national AI strategies (Radu 2025). This trend is motivated by concerns around the extraction of value from a country’s data,

people, and resources while profits accrue abroad; dependency on foreign technology firms; and a desire to retain control over key elements of the AI value chain.

Map 1 Global landscape of national AI strategies



Source: WDR 2026 team based on Radu 2025.

Note: The map includes published strategies between March and June 2025 for which the full text could be found online. AI = artificial intelligence.

This section of the Report maps the key policy priorities for maximizing the productive use of AI in developing countries while reducing unequal returns and guarding against possible misuse across three dimensions: empowering entrepreneurs to leverage the value of AI; educating people to enhance workers’ capabilities; and ensuring trust by building effective governance and institutional frameworks to manage AI’s broader societal impacts.

1. Empowering Entrepreneurs to Leverage the Value of AI

Government can create enabling conditions for businesses to power the use of AI solutions on a large scale. On the supply side, a necessary first step is establishing the foundational infrastructure that enables access to affordable electricity and internet connectivity—even if a country is just seeking to adopt existing off-the-shelf AI solutions. On the demand side, limited awareness about AI, financial constraints, and lack of management capabilities can constrain AI adoption across the production process. For many developing countries, the adaptation of AI solutions to the local context could allow businesses to accrue greater economic value, promote innovation, and guard against unhealthy dependences on foreign AI providers. This would require governments to establish an enabling environment for supporting local AI innovation, including the use of industrial policy such as government procurement to jumpstart domestic AI solutions.

Establishing the foundational digital and analog complements to spur the adoption of AI

On the supply side, access to electricity and the internet powers the use of AI. Such access is far from universal in developing countries, especially in remote areas. Investing in energy infrastructure to ensure universal, reliable, and affordable electricity access and use is therefore a priority. So is expanding and upgrading internet infrastructure (broadband from the first, to the middle, to the last mile) to improve coverage and quality; fast internet such as 4G, 5G, and fiber optics are particularly important for low-

latency AI applications. Such foundational infrastructure is the typical public good that is unlikely to be provided by the private sector, especially in remote areas. Yet, public investment in this infrastructure can be complemented by private investment, including by reducing restrictions on foreign investment. Establishing market competition in the provision of electricity and telecommunication services is also important to leverage the related investments in infrastructure.

If the use of smartphones and other digital devices is an important complement to power AI solutions, there might be merit in targeting support that improves access to such devices. At a minimum, reducing taxes on the purchase of low-cost digital devices that provide voice, internet, and mobile wallet services can improve access. It is estimated that the welfare costs of handset taxes are more than three times the revenue these taxes raise (Björkegren 2019). There might also be merit in considering innovative financing mechanisms for device ownership/access, such as the pay-as-you-go model, especially for low-income households/individuals.

On the demand side, AI adoption might be slowed by existing regulations that constrain entrepreneurship and the growth of productive businesses. Steps to make the business environment more conducive include simplifying the processes for business entry and exit, improving access to finance, streamlining labor regulations, and promoting market competition, including by liberalizing trade and promoting foreign direct investment, which have historically played an important role in facilitating technology diffusion and adoption. Similar to the adoption of other technologies, AI adoption might also be constrained by firm capabilities, including limited awareness about AI and weak management practices. Governments can intervene to plug in these information gaps.

Supporting competition and a local innovation ecosystem that spurs AI adaptation

Healthy competition in AI markets drives innovation, lowers costs, and enables local players to emerge (Martens 2024; Vipra and Korinek 2023). The role of competition policy depends on how AI solutions are monetized, and whether they are developed by domestic or foreign AI companies. If AI tools are monetized through subscription-based services, concentration of the technology in the hands of a few big technology companies based in advanced economies carries the risk of creating supplier lock-in and high prices for AI solutions, limiting adoption and use, especially in developing countries that lack suppliers of core elements of the technology (Athey and Morton 2025). Even if AI tools developed by the big technology companies are free for users, competition might still lag owing to the lack of a level-playing field in terms of key inputs, such as data sets and computing infrastructure (Tucker 2024). This lack of competition can limit innovation and concentrate the benefits of AI among a few large corporations, create unhealthy dependencies on foreign AI providers, and limit the ability of local stakeholders to shape AI in ways that benefit their communities. Further, the potential misalignment of AI models with local contexts can make AI solutions less effective in developing countries and slow diffusion.

The adaptation of AI models to the local context and the development of AI applications requires investments in data sets, computing infrastructure, and talent. Public investment (including through crowding in private investment to address prevailing market failures) can open up computing infrastructure and data sets to small and medium enterprises (SMEs), academia, and start-ups to be able to run and fine-tune AI models (Warso and Keller 2023). Here, the aim need not be to develop large-scale sovereign AI models using fully localized data and computers, as this may have the unintended consequences of stifling cross-border data flows, increasing geopolitical tensions, fragmenting global AI standards, increasing surveillance risks, and magnifying environmental impacts.²⁷ The focus could instead be on enabling “small AI” models that require smaller investments in computing infrastructure but are intentionally task-specific, more easily interpretable, and can be made accessible offline to consumers (Hadgu et al. 2023).

Governments can experiment with a range of “computing commons” approaches to make investments more viable (Davies and Vipra 2025), although results are still nascent. Approaches may include decentralized provision (such as the use of public funding to facilitate a network of smaller distributed facilities across a region, which can be publicly or privately owned, as is being considered in China) or market-based provision (using public funds to provide vouchers or subsidies for users to access compute from existing market providers, as is the case in the IndiaAI mission). Governments can similarly support the development of “data commons” approaches. For example, governments can lead open data initiatives to make local language data accessible to AI developers. Innovative licensing models can ensure that local communities that provide such data are appropriately compensated (Rajab et al. 2025). And data governance practices, such as data intermediaries and public data trusts, could ensure the responsible sharing of sensitive public sector data (such as health data) (Viljoen 2021). International cooperation can support these approaches, but this will depend on support for trusted flows of data across borders (World Bank 2021), as well as the adoption of international standards to promote interoperability and data portability (World Bank 2025c).

Supporting AI adaptation through government procurement and other fiscal tools

Government procurement of AI solutions, such as through the development of digital public infrastructure (DPI) is a powerful tool to shape AI ecosystems. Used well, these tools can incentivize responsible innovation and crowd in privately developed local AI solutions aligned with public goals. This approach is analogous to advance market commitments made in the context of vaccine development. However, poorly designed safeguards can inadvertently lock in harmful systems (Zick et al. 2024). Governments can also use financial incentives including subsidies, grants, and R&D credits to support the development of more bespoke AI solutions. This is particularly relevant for application areas, where either the risk of societal harm with misaligned AI models is high (for example, health, education, and justice delivery) or because they provide public goods (for example, teacher assistants and digital public infrastructure).

However, once embedded in public services or infrastructure, AI systems are difficult to unwind, creating path dependencies which might be especially costly because AI solutions need to be continuously upgraded.²⁸ Further, poorly designed rules can entrench incumbents or favor those with greater compliance capacity, and limit transparency. Therefore, procurement rules need to avoid vendor lock-in, maintain system interoperability, and mitigate dependencies on a single provider (Cohen et al. 2024). Embedding DPI principles (such as data interoperability, open application programming interfaces [APIs], and transparent audit trails) in procurement processes for AI tools can ensure that they integrate safely with existing public systems and remain vendor-neutral.

2. Educating People to Enhance Workers’ Capabilities

AI is expected to reshape the labor market profoundly in the coming years. However, the risk of immediate disruption in labor markets is relatively low in low- and middle-income countries owing to the lower incidence of cognitive tasks.²⁹ At the same time, AI is likely to transform occupations and industries, and create entirely new roles, leading to a mismatch between the skills that are being imparted today and the skills likely to be in demand in the future. Some adjustments and different points of emphasis in education policies will be necessary to guard against this mismatch. Broad-based skills development will spur AI adoption and help adjust to AI-related disruption, while AI-related technical skills will matter for AI adaptation. There will also be implications for reimagining pedagogy and curriculum design in formal education systems to guard against AI hollowing out cognitive abilities such as writing and problem-solving.

Emphasizing broad-based skills development that will enable AI adoption and help adjust to AI-related disruption

AI adoption is likely to be less widespread without literacy, numeracy, and basic digital skills, including know-how to effectively “prompting” AI on what to do. Formal education during early childhood and adolescence would be instrumental to impart these skills. For adults that have graduated from formal education, training programs or on-the-job learning could be a bridge. For the less-skilled, less-educated parts of the population, targeted information campaigns and training will likely make people more willing and able to adopt AI technologies. For managers in governments and small enterprises, upskilling programs might be the relevant policy lever to raise their capacity to absorb AI and complementary technologies.

Emphasizing the development of socio-emotional and other soft skills that are often shaped by early childhood education might be important for people to adjust to AI-related disruption. For example, tasks involving face-to-face interaction are less likely to be automated by AI and these are tasks where interpersonal skills are important. Similarly, social-emotional skills will allow people to cope better with the challenges that AI is bringing (Deming 2017; Dolev and Itzkovich 2020). The role of college education in imparting specialized skills, such as coding and entry-level programming, might need to be revisited considering the potential of AI to automate some of these specific knowledge-intensive tasks.

Investing in AI-related technical skills that matter for AI adaptation

The adaptation and advancement of AI requires investments across a range of specialized skills, including computer science, data science, cybersecurity, AI ethics, and systems engineering (Martins-Neto et al. 2025). Some of these technical skills will be important for R&D testing to detect model bias, risk of misinformation, and possible AI hallucinations. Entrepreneurial skills that help develop an understanding of expected market impacts of AI applications in specific use cases may also play a crucial role in successfully bringing customized AI products and services to the market. These aspects of human capital need to be integrated into formal general and technical education systems, vocational training, and lifelong learning on the job. AI-based personalized learning can be used to help people accumulate skills across these channels of learning. In addition, establishing industry-academia partnerships to develop industry-specific AI use skills would be important enablers.

Leveraging diaspora and foreign AI talent

There is the risk of brain drain associated with STEM educated workers and other AI talent migrating from low- and middle-income countries to high-income countries. Low- and middle-income countries can leverage the expertise of their diaspora and skilled foreign workers to strengthen local AI capabilities. By working alongside local workers, these skilled workers can share best practices and innovative ideas, leading to knowledge transfer and capacity building of the local workforce. Indeed, countries with successful diasporas often have a higher potential for knowledge transfer (World Bank 2024). Governments can tap into such talent that has migrated to high-income countries through diaspora engagement, research incentives, and establishing AI hubs. To attract foreign AI talent, streamlined visa applications systems could play a role.

Guarding against the possible erosion of human capital

AI has profound implications for pedagogy, curriculum design, and education systems more broadly. By delegating to AI, students may not acquire these fundamental cognitive abilities, such as critical thinking or writing, in school and/or further develop them at work. Further, education assessment might become less meaningful and there might be a breakdown of traditional signals to potential employers (for example, via

writing or communication tasks and test scores). Governments should ensure that the accumulation of core cognitive abilities is not hampered by the use of AI. This may include curriculums that teach students how to use AI tools in a way that avoids teaching, assessments that are conducted in the classroom rather than at home, and pedagogy that emphasizes AI use by teachers rather than students. There may be lessons that can be drawn from the ways in which the use of calculators and the use of the internet might have impacted learning outcomes.

3. Ensuring Trust by Building Effective Governance and Institutional Frameworks to Manage AI’s Broader Societal Impacts

Trust is the cornerstone of effective AI adoption. Without confidence that AI will not cause social or economic harm, individuals and institutions will resist its use, thereby undermining its potential to deliver public value. For example, while AI requires more and better data to improve accuracy, reduce bias, and benefit people, greater data-sharing also heightens violations of privacy rights and surveillance (D’Ignazio and Klein 2020; World Bank 2021). Similarly, there is the risk of extractive behavior where data and AI models are harvested for profit (Widder et al. 2023) without meaningful benefit-sharing, especially for communities that supplied the data in the first place (Abebe et al. 2021). To create and sustain trust against the risk of social harms linked to AI use, there is a need for countries to establish robust AI governance safeguards and move beyond principles toward an operational and enforceable trust infrastructure. Implementing regulatory frameworks that achieve this is particularly challenging given the rapid pace of change in AI technologies and institutional weaknesses in low-income countries. Governments can experiment with different regulatory approaches, adopt good practices, strengthen institutional preparedness, and invest in the evaluation of AI solutions.

Experimenting with different AI governance approaches

The establishment of robust AI governance frameworks can encompass both binding “hard” law and nonbinding tools like industry self-governance, “soft” law, and standards. Some of these frameworks involve cross-border agreements. The European Union (EU) AI Act and the Brazil AI bill are examples of hard law that create clear, enforceable guidelines and generate legal certainty. At the same time, such hard law is time-consuming and resource-intensive to design, implement, and enforce and requires significant regulatory capacity.

The OECD/G20 AI Principles and the Singapore Model AI Governance Framework for Generative AI are examples of soft law that set key normative standards while remaining adaptable and responsive to technological advances. Such soft law focuses on high-level principles rather than specific rights and responsibilities and can be “binding” to varying degrees.

The Colombia regulatory sandbox on privacy by design and default in AI projects and the Brazil regulatory sandbox pilot for AI and data protection are examples of experimental approaches that provide a controlled environment to test and evaluate new regulatory approaches, without exposing the public to unchecked risks.

The Microsoft Aether Committee and Responsible AI Standard Playbook, Google AI Principles, and IBM’s AI Ethics Board are examples of industry self-governance. They can be directly integrated into business models but are nonbinding and potentially inappropriate for high-risk use cases.

Given the degree of uncertainty around the future trajectory of AI, countries are increasingly combining these approaches to create hybrid AI governance models: for example, embedding soft-law principles or

technical standards into binding legislation, or coupling regulatory sandboxes with clear enforcement mandates. Taking a problem-driven, iterative approach to governance means that countries should foster active experimental iterations of different governance structures, identifying what works and what doesn't in terms of protecting against certain risks or harms before scaling these into country-wide solutions (Andrews et al. 2017).

Adopt good practices from around the world

With no established single approach for governing the development and use of AI, especially in developing countries, adopting good practices can be a useful first step. Countries can tailor their legal and policy frameworks by adopting good practices based on where they sit along the “Adopt, Adapt, Advance” spectrum, recognizing that AI governance lies on a continuum. Countries can be developers in some areas or sectors and primarily adopters in others. All countries need a baseline of AI governance good practices and independent oversight, but the scope and stringency of regulation should evolve with their level of adoption, institutional capacity, and the complexity of local AI use.

For AI adoption, governments can start by ensuring that AI solutions comply with local laws and ethical standards, through baseline testing, registration, and human-in-the-loop safeguards.³⁰ They can also identify clear red lines for unacceptable uses of AI, like social scoring, exploitative/harmful deepfake imagery, and mass surveillance, as does the EU AI Act list of unacceptable uses of AI. Governments can further prioritize the enforcement of sectoral regulations. For example, in the health sector, regulators can strengthen oversight of AI-enabled diagnostic and decision-support tools to ensure safety, accuracy, and respect for patient privacy. These considerations have shaped the regulation of algorithmic systems even before the latest advances in generative AI and will likely remain relevant, regardless of the direction of the technology in the future.

For AI adaptation, governments can go further to adopt best practices that help evaluate, contextualize, and improve imported general-purpose AI systems before deployment. For instance, governments can establish transparency requirements, such as in the California Generative Artificial Intelligence: Training Data Transparency Act 2024 that mandates developers of generative AI systems or services to disclose detailed information about the data sets they used to train their models. Similarly, technical standards (such as the AI risk management frameworks ISO/IEC 23053:2022) can help local AI developers comply with intellectual property regulations, especially when embedded in procurement rules, or certification schemes. The adoption of such technical standards should be paired with mechanisms to bridge participation gaps in international standard-setting for developing countries (World Bank 2025a).

Where governments choose to invest in local computing infrastructure, policy must ensure environmental sustainability. This includes mandatory environmental impact disclosures for data centers, or market-based mechanisms like tiered electricity and water pricing that discourage excessive energy and water use by charging higher rates during peak demand periods, alleviating stress on the electricity grid and water systems.

Investing in the evaluation of AI solutions

Governments in low- and middle-income countries should invest in frameworks and implementation plans to evaluate AI's contribution to human capital development where context matters greatly, but where the evidence base is thin. These frameworks could seek to verify the output being generated by the AI solution, study whether the AI solution is user friendly or not, test the impact of the AI solution on relevant outcomes, and analyze whether the AI solution is cost effective or not (The Agency Fund 2025).

Investing in the development of institutional agility

Countries will require stronger and more agile institutions, especially if they seek to extract value from AI and ensure guardrails against potential social harms, given the uncertain direction of AI development and potential conflicts of interest where governments serve as regulators, users, and in some cases, investors in AI. How can governments be better prepared? First, governments can engage multiple stakeholders such as end-users, affected communities, businesses, and civil society to align AI policies with the new “social contract” on AI—ensuring they are inclusive, responsive to local needs, and receive broader support.³¹

Second, governments can invest in raising the capabilities of bureaucrats by promoting contestability in AI regulatory offices, jobs, and ideas. Contestability in positions of power can foster responsive leadership and prevent incumbents from becoming entrenched in their roles (refer, for example, to Acemoglu et al. 2014). Contestability in government jobs through open recruitment and competitive hiring can promote merit and attract qualified candidates (Aneja and Xu 2024). Contestability in the space of ideas through active policy debates that include academia and civil society can promote evidence-based solutions.

Third, governments can foster global cooperation to help bridge capability gaps, reduce global asymmetries in AI access, and establish trust frameworks that allow AI to operate safely across borders.

Consultations, Timeline, and WDR 2026 Team

Consultations and Timeline

Since the topic for the *World Development Report 2026* was announced, the WDR 2026 team has held internal consultations within the World Bank Group with global practices, verticals, corporate units, and technical colleagues, and has had preliminary engagements with a range of external stakeholders (discussed later in this section).

The WDR 2026 will also benefit from the guidance and advice of an Academic Advisory Committee and consultations with other academics, as well as a High-Level Advisory Panel (discussed later in this section).

WDR 2026 will be published in late 2026.

The WDR 2026 Team

Gaurav Nayyar is the Director of the Report. Susan Athey from the Stanford Graduate School of Business and Ufuk Akcigit from the University of Chicago are the Lead Academics. Joyce Ibrahim is the Task Team Leader. The core team comprises S. Anukriti, Sharmista Appaya, Elwyn Davies, Lelys Dinarte-Diaz, Samuel Paul Fraiberger, Yi Jie Gwee, Yan Liu, Jeremy Ng, Manuel Ramos Maqueda, Tiago Carneiro Peixoto, Anja Sautmann, Katherine Stapleton, and Shu Yu. Deepak Mishra serves as an advisor to the WDR 2026 team. Aya Aboul Hosn, Alegria Burneo, Sonam Tshering Gurung, Sharif Kazemi, Bolor Erdene Munkhbayar, Shyam Jayanti Patel, and Uzma Sahar Rauf serve as research assistants. Sandi Soe Lwin provided administrative support through December 2025. Gabriela Calderon Motta has provided administrative support since January 2026.

The Report is co-sponsored by the Development Economics Vice-Presidency and the Digital and AI Vertical. Overall guidance for the preparation of the Report is provided by Indermit Gill (Senior Vice

President and Chief Economist of the World Bank Group) and Somik Lall (Director of Strategy, Development Economics Vice-Presidency).

External Consultations

The team has held consultations with the following organizations: Accel, Adalat AI, AI4Bharat, AI for the People, Africa Practice, African Development Bank (AfDB), AfriLabs, Alibaba Group, Alpha Intelligence Fund, Amazon, Anthropic, Arkam Ventures, ASEAN Foundation, Asian Development Bank (ADB), Aspen Institute, Atlas Computing, Bank of Mongolia, BharathGen, Cambodia Academy of Digital Technology, Carnegie Endowment for International Peace, Center for Analysis of Economic Reforms and Communication, Centre for Intellectual Property and Information Technology Law, Centre for Open Societal Systems, Cognitive Insights for Artificial Intelligence, Consumer Reports, Cyber Sierra, Deloitte, Department of Communications and Digital Technologies (South Africa), Digital Government Committee (Cambodia), EkStep Foundation, E-Mongolia, European Bank for Reconstruction and Development (EBRD), European Commission, European Investment Bank (EIB), Foreign Correspondents Club of Japan, Gates Foundation, GitHub, Google, Google I/O partnerships, Google Research, Grab, Homai, Humane Intelligence, IDB Invest (Inter-American Development Bank Group), Independent Communications Authority of South Africa (ICASA), Infosys, Innovation and Digital Development Agency (Azerbaijan), Inter-American Development Bank (IDB), International Centre for Theoretical Sciences, International Economics Consulting, International Monetary Fund (IMF), International Telecommunication Union (ITU), ISEAS – Yusof Ishak Institute, Jus Mundi, Kalaari Capital, Karnataka Digital Economy Mission, KPMG, Leo Capital, LinkedIn, Masakhane African Languages Hub, Meta, Microsoft, Ministry of Electronics and Information Technology (India), Ministry of Industry and New Technologies (Artificial Intelligence Council) (Tajikistan), Ministry of Industry and Trade (E-Commerce and Digital Economy Agency) (Viet Nam), Ministry of National Development Planning (Indonesia), Ministry of Science, Technology and Innovation (Malaysia), MOSIP, National Artificial Intelligence Center (Azerbaijan), National Artificial Intelligence Expert Council (Mongolia), National Association of Software and Service Companies (NASSCOM), National Institution for Transforming India (NITI Aayog), Nikkei, Nirmai.ai, NVIDIA, Office of the Council of State (Thailand), Omidyar, OpenAI, Open Forum for AI, Organization for Economic Co-operation and Development (OECD), Otoritas Jasa Keuangan (Indonesia Financial Services Authority), PayGlocal, People + ai, Qure.ai, RBI Innovation Hub, RevolutionAize, Rocket Learning, Rookie Systems LLC, Sarvam AI, Shutterstock, Smart Government II Project Implementation Unit (Mongolia), Snowmountain AI, STelligence, Swiss AI, Takshashila Institution, Tencent Group, TROTRO Tractor Ltd, United Nations Children’s Fund (UNICEF), United Nations Development Programme (UNDP), upGrad Enterprise, and zypl.ai.

The team consulted with many of these organizations as part of a series of consultations that took place and are being planned across several regions. As of this publication, the team held: (i) a regional consultation in Tokyo (Japan) in December 2025 co-sponsored with the Asian Development Bank Institute (ADBI), with stakeholders from Azerbaijan, Bangladesh, Cambodia, China, Japan, India, Indonesia, Republic of Korea, Malaysia, Mongolia, the Philippines, Singapore, Tajikistan, Thailand, and Viet Nam; and (ii) consultations in Bangalore (co-hosted by InfoSys) and Delhi (India) in December 2025. These consultations brought together a diverse group of stakeholders—including government entities, research institutions/think tanks, and private sector actors—to share regional experiences, identify pressing policy questions, and ensure the report’s recommendations are grounded in real-world development challenges and opportunities. Further planned consultations include: (i) a regional consultation in Pretoria, South Africa in March 2026 co-sponsored with the University of Pretoria and focused on several Sub-Saharan African countries; (ii) a regional consultation in Rome in March 2026 focused on countries from Europe and Central Asia; and (iii) a regional consultation in Saudi Arabia in April 2026 co-sponsored with the World Bank’s Knowledge Hub focused on countries from the Middle East, North Africa, Afghanistan, and Pakistan.

Regional Initiatives

There are several ongoing regional initiatives, coordinated by the regional chief economist offices, that are complementary to this Report and are, indeed, informing the Report. This includes two recent regional updates on East Asia and the Pacific and South Asia that discussed the implications of AI for jobs. Further, the analysis in a World Policy Research Working Paper, entitled “Buffer or Bottleneck? Employment Exposure to Generative AI and the Digital Divide in Latin America” and authored by colleagues in the Latin America and Caribbean Region, is being extended to larger set of low- and middle-income countries to better understand the potential labor market impacts of AI. Going forward, the Europe and Central Asia Chief Economist’s Office is working on the topic of AI and aging for a forthcoming regional update/report. And the Middle East, North Africa, Afghanistan, and Pakistan Chief Economist’s Office is working on a report on labor markets that will include a discussion on AI.

Academic Advisory Committee and Other Academic Consultations

The Academic Advisory Committee is comprised of the following members: David H. Autor (Massachusetts Institute of Technology), Daniel Björkegren (Columbia University), Nicholas Bloom (Stanford University), Erik Brynjolfsson (Stanford University), Bhaskar Chakravorti (Tufts University), Julie Cohen (University of Chicago), Avi Goldfarb (University of Toronto), Aziz Z. Huq (University of Chicago), Michael Kremer (University of Chicago), David Leslie (Queen Mary London), Arvind Narayanan (Princeton University), Ziad Obermeyer (University of California, Berkeley), Benjamin Rosman (University of the Witwatersrand), Catherine Tucker (Massachusetts Institute of Technology), and Yi Zeng (Institute of Automation, Chinese Academy of Sciences, University of Chinese Academy of Sciences).

In addition, the WDR 2026 team has consulted with the following academics: Rachel Adams (African Observatory on Responsible AI), Philippe Aghion (College de France and INSEAD), Khondaker Abdullah Al Mamun (United International University, Bangladesh), Asrat Mulatu Beyene (Addis Ababa Science and Technology University), Altan Cakir (Istanbul Technical University), Maggie Chen (George Washington University), Zhao Chen (China Center for Economic Studies, Fudan University), David Eaves (University College London), James Evans (University of Chicago), Nick Feamster (University of Chicago), Robert Grossman (University of Chicago), Sivaramakrishnan R. Guruvayur (Amritapuri Vishwa Vidyapeetham, India), Yuki Higuchi (Sophia University), Jingbo Huang (United Nations University Institute in Macau SAR, China), Seiro Ito (JETRO Institute of Developing Economies), Zhiyuan Jiang (Shanghai University), Ben Jones (Northwestern University), Taro Katsurai (JICA Ogata Sadako Research Institute for Peace and Development), Rembrand M. Koning (Harvard Business School), Anton Korinek (University of Virginia), Ponnurangam Kumaraguru (International Institute of Information Technology (IIIT) Hyderabad) Lixing Li (National School of Development, Peking University), Zhikuo Liu (School of Economics, Fudan University), Joel Mokyr (Northwestern University), Jabu Mtsweni (Council for Scientific and Industrial Research [CSIR] South Africa), Deeparghya Mukherjee (Indian Institute of Management Nagpur), Atsuko Munemura (Chiba University of Commerce), Yadati Narahari (Indian Institute of Science (IISc) Bangalore), Izumi Ohno (GRIPS [National Graduate Institute for Policy Studies]), Nohyoung Park (Korea University Law School), Alan Ross (Global AI Solutions), Jun Saito (Japan Center for Economic Research), Aduwati Sali (Institute for Mathematical Research and Intelligent Systems [INSPEM], Universiti Putra Malaysia), Sunita Sarawagi (Indian Institute of Technology (IIT) Bombay), Jonathan Shock (University of Cape Town), Ram Singh (Delhi School of Economics), Tetushi Sonobe (GRIPS [National Graduate Institute for Policy Studies]), Milind Tambe (Harvard University), Neil Thompson (Massachusetts Institute of Technology), Yasushi Ueki (JETRO Institute of Developing Economies), Lan Xue (National School of Development, Peking University), Hiroyuki Yamada (Keio University), Naoyuki Yoshino (Keio

University), Dandan Zhang (National School of Development, Peking University), and Yuan Zhang (China Center for Economic Studies, Fudan University).

High-Level Advisory Panel

The High-Level Advisory Panel is comprised of the following members: Abdullah Alswaha (Minister of Communications and Information Technology, Saudi Arabia), Karan Bhatia (Global Head, Global Affairs and Public Policy, Google), Rumman Chowdhury (Founder, Humane Intelligence), Amal El Fallah Seghrouchni (Minister of Digital Transition and Administrative Reform, Morocco), Ryu Je-Myung (Vice Minister of Science and ICT, Republic of Korea), Kate Kallot (CEO and Co-Founder, Amini), S. Krishnan (Secretary, Ministry of Electronics and Information Technology, India), Albert Park (Chief Economist, Asian Development Bank), Amandeep Singh Gill (Under-Secretary-General and Secretary-General's Envoy on Technology, United Nations), and Bosun Tijani (Minister of Communications, Innovation, and Digital Economy, Nigeria).

Notes

¹ A general-purpose technology (GPT) is characterized by widespread use, the potential for continuous innovation, and innovation in application industries.

² A General purpose AI chatbots are AI applications that can perform a wide variety of tasks, such as answering questions, generating content, and assisting with coding. Examples include OpenAI's ChatGPT, Google's Gemini, Microsoft's Copilot, and Anthropic's Claude.

³ This uses job posting measures of widespread use, potential for innovation, and innovation in application industries.

⁴ The continued pace of improvement in AI performance across some of these functions is uncertain because underlying prediction algorithms may not be able to generalize to sophisticated cognitive achievements (Shojaee et al. 2025).

⁵ GitHub is a software developer platform that allows developers to create, store, manage, and share their code.

⁶ More broadly, AI is also an input to other "frontier technologies" such as biotechnology and nanotechnology, which implies that progress in AI will be correlated with advancement in these other domains (United Nations 2018).

⁷ For example, training of AI models that cost less than US\$1,000 in 2017 increased to US\$4million to US\$6 million by 2020, and now routinely exceed US\$170 million (Cottier et al. 2025; Sevilla and Roldán 2024). The estimated training costs for OpenAI's GPT-4 and Meta's LLaMA 3.1 are approximately US\$80 million and US\$170 million, respectively.

⁸ Recent developments show that smaller open-source AI models trained efficiently on high-quality data can match or even exceed the performance of much larger proprietary systems at a fraction of the cost (Sevilla and Roldán 2024). For example, DeepSeek-V3 was trained for just US\$5.6 million, compared to the tens or hundreds of millions spent on models like GPT-4 (DeepSeek 2024).

⁹ The latest GenAI models, such as Claude AI, also show more and more robust performance in languages with modest existing resources, such as Mali's 13 official languages (Dembele et al. 2025).

¹⁰ However, recent data from Anthropic's Claude.ai chatbot show that users in countries with relatively low per capita AI usage prefer delegating tasks to AI (automation), whereas users in economies with high per capita AI usage prefer more collaborative or learning-based interactions with AI (augmentation) (Anthropic 2025).

¹¹ Evidence from the United States shows that 40 percent of US adults adopted generative AI tools within two years of the first mass-market product release compared to 20 percent within three years for personal computers.

¹² Open-source foundation AI models are models whose architecture is publicly available for others to use, modify, and build upon.

¹³ The gap between open and closed AI models is narrowing. On the Chatbot Arena leaderboard, for example, the top open-weight model lagged the leading closed model by 8.0 percent in early 2024; by early 2025, that gap had shrunk to just 1.7 percent (LMSYS 2025).

¹⁴ The lack of sufficient technology diffusion among poor countries in each GPT wave is reflected in the fact that the intensity of use of adopted technologies in lower-income countries compared to high-income countries has diverged with each GPT wave (Comin and Mestieri 2018).

¹⁵ Historical examples of technological lock-in, such as the "QWERTY" keyboard layout (David 1986), the VHS system in videocassette recording (Arthur 1990), and the "light-water" design for civilian nuclear power reactors (Cowan 1990), show how the initial adoption of specific technologies can persist for decades even if they could be suboptimal in some instances.

¹⁶ While evidence is scarce, there is evidence in developed countries that the polarization of views can lead to hate crimes (Chan et al. 2016; Müller and Schwarz 2021).

¹⁷ The target sample is around 120 technical officers with expertise in various types of digital system or activity in government

¹⁸ The target sample is around 2,500-3,000 business establishments in India, Jordan, Kenya, Malaysia, Mexico, Nigeria, Türkiye, and the United States.

¹⁹ The survey instrument is adapted from existing surveys around technology adoption, including the World Bank's Firm-Level Adoption of Technology (FAT) survey (Comin et al. 2022), the Eurostat ICT and E-commerce Surveys, and the US Business Trends and Outlook Survey (BTOS). The module on expectations is based on the methodology developed by the Federal Reserve Bank of Atlanta to measure business uncertainty (Altig et al. 2022) and used by the World Bank during COVID-19 to measure uncertainty around the world (Avalos et al. 2025).

²⁰ The impact of technological change on jobs has typically been assessed through a task-based framework (Acemoglu and Restrepo 2019). First, there is a “displacement” effect whereby the technology replaces labor in certain tasks. Second, there is a “productivity” effect whereby automation raises labor productivity and increases the demand for labor in non-automated tasks. Third, there is a “reinstatement” effect whereby technologies can create new tasks.

²¹ Whether these productivity gains translate into employment losses or gains depends on a large extent on whether scale effects (increased demand due to productivity gains) outweighs the substitution effect of needing fewer workers to provide the same.

²² In 2019, there were 2.9 physicians per 10,000 people in Sub-Saharan Africa, far below the global average of 16.7; these physician-patient densities were also low (less than 10 per 10,000) in South Asia and the Middle East and North Africa (Haakenstad et al. 2022). Similarly, the average student-teacher ratio in primary schools across low-income countries is approximately 40:1, compared to about 15:1 in high-income countries (UNESCO Institute for Statistics, <https://databrowser.uis.unesco.org/>).

²³ This capacity of generative AI builds on the success of previous digital technologies in expanding access to government services in remote areas by enabling less qualified workers to complete more expert tasks across sectors including health, education, environmental management, and social protection (Mitchell et al. 2013; Muralidharan et al. 2016; Muralidharan et al. 2019; Slough et al. 2021; Tan et al. 2024).

²⁴ For example, evidence from India, Kenya, Nepal, the Philippines, and Uganda shows that phone call tutorials have improved learning outcomes through targeted instruction tailored to students’ needs (Angrist et al. 2023). Similarly, digital technology interventions have improved learning in contexts where the regular school system was unavailable (Angrist et al. 2022).

²⁵ LLMs do not have network effects whereby a product or service becomes more valuable as more people use it and have higher marginal cost. This could weaken the market position of providers of LLMs, with the most gains accruing to providers of AI chips (Nvidia), and cloud computing services (AWS, GCP, Azure).

²⁶ AI development may also be skewed toward certain sectors because AI investment (usually) flows toward high-profit sectors (finance, advertising, defense) rather than public goods (health, education, agriculture).

²⁷ This approach contrasts with that of many advanced economies and large emerging markets that are actively pursuing the creation of sovereign large-scale AI models. Driven by strategic competition, concerns over dependency, and the geopolitical value of AI capabilities, the aim is to secure control over critical components of the AI supply chain, including semiconductors, computing infrastructure, data sets, and talent (Belli 2023; Jiang and Belli 2025; Mügge 2024; Nothias 2025).

²⁸ For example, an AI model trained on data collected in violation of data subject rights cannot be easily “un-trained” on that data, short of deleting the model entirely (refer to Shreve 2024).

²⁹ To the extent that there is disruption, governments may consider social safety nets and re-skilling training programs to facilitate the reallocation of workers who are displaced by AI to occupations that are either less affected by AI or those that are expanding due to AI.

³⁰ Human-in-the-loop safeguards are mechanisms that use a human to actively review, validate, or override AI-driven decisions, often by handling complex, ambiguous, or high-stakes situations that automated systems are not equipped to manage alone.

³¹ For example, Bangladesh plans to create an independent National Artificial Intelligence Center of Excellence (NAICE) that will be responsible for constituting a multistakeholder High-Level National AI Advisory Council including official agencies, civil society experts, and representatives from industry and academia to guide the implementation of AI initiatives in various sectors (refer to Baree 2024).

References

- Abdin, Marah, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Hassan Awadalla, Nguyen Bach, et al. 2024. “Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone.” *arXiv* (April 23). <https://doi.org/10.48550/arXiv.2404.14219>.
- Abebe, Rediet, Kehinde Aruleba, Abeba Birhane, Sara Kingsley, George Obaido, Sekou L. Remy, and Swathi Sadagopan. 2021. “Narratives and Counternarratives on Data Sharing in Africa.” *arXiv preprint arXiv:2103.01168*. <https://arxiv.org/abs/2103.01168>.
- Acemoglu, Daron. 2024. “Harms of AI.” In *The Oxford Handbook of AI Governance*, edited by Justin B. Bullock, Yu-Che Chen, Johannes Himmelreich, Valerie M. Hudson, Anton Korinek, Matthew M. Young, and Baobao Zhang, 660–706. Oxford University Press.
- Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” *Handbook of Labor Economics* 4B: 1043–1171. Elsevier. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).
- Acemoglu, Daron, Suresh Naidu, Pascual Restrepo, and James A. Robinson. 2014. “Democracy Does Cause Growth.” *Journal of Political Economy* 122 (1): 47–91.
- Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.
- Adapa, Karthik, Ashu Gupta, Sandeep Singh, Hitinder Kaur, Abhinav Trikha, Ajoy Sharma, and Kumar Rahul. 2025. “A Real World Evaluation of an Innovative Artificial Intelligence Tool for Population-Level Breast Cancer Screening.” *NPJ Digital Medicine* 8 (1): 2.
- Agarwal, Sumit, Shashwat Alok, Pulak Ghosh, and Sudip Gupta. 2019. “Financial Inclusion and Alternate Credit Scoring: Role of Big Data and Machine Learning in Fintech.” Indian School of Business, Hyderabad. December. SSRN: [https://ssrn.com/abstract=3507827\[1\]](https://ssrn.com/abstract=3507827[1]).
- Agrawal, Ajay K., Joshua S. Gans, and Avi Goldfarb. 2023. “Similarities and Differences in the Adoption of General Purpose Technologies.” In *Technology, Productivity, and Economic Growth*, edited by Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner, 57–73. National Bureau of Economic Research. <https://www.nber.org/chapters/c14751>.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2024. “Generative AI Is Still Just a Prediction Machine.” *Harvard Business Review*, November 18, 2024. <https://hbr.org/2024/11/generative-ai-is-still-just-a-prediction-machine>.
- Albanesi, Stefania, António Dias da Silva, Juan F. Jimeno, Ana Lamo, and Alena Wabitsch. 2023. “New Technologies and Jobs in Europe.” NBER Working Paper 31357, National Bureau of Economic Research.
- Altig, David, Jose Maria Barrero, Nicholas Bloom, Steven J. Davis, Brent Meyer, and Nicholas Parker. 2022. “Surveying Business Uncertainty.” *Journal of Econometrics* 231 (1): 282–303.
- Andrews, Matt, Lant Pritchett, and Michael Woolcock. 2017. *Building State Capability: Evidence, Analysis, Action*. Oxford University Press.
- Aneja, Abhay, and Guo Xu. 2024. “Strengthening State Capacity: Civil Service Reform and Public Sector Performance during the Gilded Age.” *American Economic Review* 114 (8): 2352–87. <https://doi.org/10.1257/aer.20221712>.
- Angrist, Noam, Micheal Ainomugisha, Sai Pramod Bathena, Peter Bergman, Colin Crossley, Claire Cullen, Thato Letsomo, Moitshepi Matsheng, Rene Marlon Panti, Shwetlena Sabarwal, and Tim Sullivan. 2023. “Building Resilient Education Systems: Evidence from Large-Scale Randomized Trials in Five Countries.” NBER Working Paper 31208, National Bureau of Economic Research, Cambridge, MA.
- Angrist, Noam, Peter Bergman, and Moitshepi Matsheng. 2022. “Experimental Evidence on Learning Using Low-Tech When School Is Out.” *Nature Human Behaviour* 6 (6): 941–50.
- Anthropic. 2025. *Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption*. <https://www.anthropic.com/research/anthropic-economic-index-september-2025-report>.

- Arthur, W. Brian. 1990. "Positive Feedbacks in the Economy." *Scientific American* 262 (2): 92–99.
- Asad, Sher Afghan, Benjamin Olken, Mahvish Shaukat, and Adnan Khan. 2019. "Increasing Revenue Collection with Computer Vision: Experiments in Pakistan." International Growth Centre. <https://www.theigc.org/collections/increasing-revenue-collection-computer-vision-experiments-pakistan>.
- Athey, Susan, and Fiona Scott Morton. 2025. "Artificial Intelligence, Competition, and Welfare." Presented at Economics of Transformative AI Workshop, September 18. <https://www.nber.org/research/videos/2025-economics-transformative-ai-workshop-artificial-intelligence-competition-and-welfare>.
- Attai, Kingsley, Moses Ekpenyong, Constance Amannah, Daniel Asuquo, Peterben Ajuga, Okure Obot, Ekemini Johnson, Anietie John, Omosivie Maduka, Christie Akwaowo, and Faith-Michael Uzoka. 2024. "Enhancing the Interpretability of Malaria and Typhoid Diagnosis with Explainable AI and Large Language Models." *Tropical Medicine and Infectious Disease* 9 (9): 216.
- Autor, David H. 2024. "Applying AI to Rebuild Middle Class Jobs." NBER Working Paper No. 32140, National Bureau of Economic Research.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–97. <https://doi.org/10.1257/aer.103.5.1553>.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96 (2): 189–94. <https://doi.org/10.1257/000282806777212620>.
- Autor, David, and Neil Thompson. 2025. "Expertise." NBER Working Paper 33941, National Bureau of Economic Research. doi: 10.3386/w33941.
- Avalos, Edgar, Jose Maria Barrero, Elwyn Davies, Leo Iacovone, and Jesica Torres. 2025. "Business Uncertainty in Developing and Emerging Economies." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4327860.
- Balashankar, Ananth, Lakshminarayanan Subramanian, and Samuel P. Fraiberger. 2023. "Predicting Food Crises using News Streams." *Science Advances* 9 (9). <https://www.science.org/doi/10.1126/sciadv.abm3449>.
- Baree, M. Nafis Manjoor. 2024. "An Overview of Bangladesh National Artificial Intelligence Policy 2024." *Daily Star*, April 19, 2024. <https://www.thedailystar.net/law-our-rights/law-vision/news/overview-bangladesh-national-artificial-intelligence-policy-2024-3590351>.
- Bastos, Paulo, and Gaurav Nayyar. 2025. "Artificial Intelligence and International Trade in Services: Evidence from the United States." Unpublished manuscript.
- Belli, Luca. 2023. "To Get Its AI Foothold, Brazil Needs to Apply the Key AI Sovereignty Enablers (KASE)." In *New Digital Dilemmas: Resisting Autocrats, Navigating Geopolitics, Confronting Platforms*, edited by Steven Feldstein. Carnegie Endowment for International Peace.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Margaret Mitchell. 2021. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, 610–23. Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
- Bento, Pedro, and Diego Restuccia. 2017. "Misallocation, Establishment Size, and Productivity." *American Economic Journal: Macroeconomics* 9 (3): 267–303.
- Beraja, Martin, Andrew Kao, David Yang, and Noam Yuchtman. 2023. "Exporting the Surveillance State via Trade in AI." NBER Working Paper 31676, National Bureau of Economic Research.
- Beraja, Martin, and Noam Yuchtman. 2025. "Generalized Disruption: Society, Work, and Property Rights in the Age of AI." In *The Political Economy of Artificial Intelligence*, edited by Ajay Agrawal, Joshua Gans, Avi Goldfarb, and Catherine Tucker. University of Chicago Press.
- Betai, N., and M. Chen. 2024. "Generative AI and International Service Outsourcing: Evidence from Online Labor Marketplaces." Unpublished manuscript.
- Bhuller, Manudeep, Tarjei Havnes, Edwin Leuven, and Magne Mogstad. 2013. "Broadband Internet: An Information Superhighway to Sex Crime?" *Review of Economic Studies* 80 (4): 1237–66.

- Bick, Alexander, Adam Blandin, and David J. Deming. 2024. “The Rapid Adoption of Generative AI.” NBER Working Paper 32966, National Bureau of Economic Research. https://www.nber.org/system/files/working_papers/w32966/w32966.pdf.
- Björkegren, Daniel. 2019. “The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda Purchased.” *Review of Economic Studies* 86 (3): 1033–60. <https://doi.org/10.1093/restud/rdy024>.
- Björkegren, Daniel, Jun Ho Choi, Divya Panchaksharappa Budihal, Dominic Sobhani, Oliver Garrod, and Paul Atherton. 2025. “Could AI Leapfrog the Web? Evidence from Teachers in Sierra Leone.” Preprint, arXiv:2502.12397.
- Björkegren, Daniel, and Darrell Grissen. 2020. “Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment Available for Purchase.” *World Bank Economic Review* 34 (3): 618–34. <https://doi.org/10.1093/wber/lhz006>.
- Blanco-Gonzalez, Alejandro, Alexandre Cabezas, Tara Sigdel, Rafael Romero, and Alejandro Pazos. 2023. “The Role of AI in Drug Discovery: Challenges, Opportunities, and Strategies.” *Pharmaceuticals* 16 (6): 891.
- Bloom, Nicholas, Christos Genakos, Raffaella Sadun, and John Van Reenen. 2012. “Management Practices Across Firms and Countries.” NBER Working Paper 17850, National Bureau of Economic Research. https://www.nber.org/system/files/working_papers/w17850/w17850.pdf.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. “Are Ideas Getting Harder to Find?” *American Economic Review* 110 (4): 1104–44.
- Bommasani, Rishi, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. “On the Opportunities and Risks of Foundation Models.” Stanford Institute for Human-Centered Artificial Intelligence. <https://doi.org/10.48550/arXiv.2108.07258>.
- Bonfiglioli, Alessandra, Rosario Crinò, Gino Gancia, and Ioannis Papadakis. 2025. “Artificial Intelligence and Jobs: Evidence from US Commuting Zones.” *Economic Policy* 40 (121): 145–94.
- Brachman, Michelle, Amina El-Ashry, Casey Dugan, and Werner Geyer. 2025. “Current and Future Use of Large Language Models for Knowledge Work.” Preprint, arXiv:2503.16774, arch 21. <https://doi.org/10.48550/arXiv.2503.16774>.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence.” *Quarterly Journal of Economics* 117 (1): 339–76.
- Bresnahan, Timothy F., and Manuel Trajtenberg. 1995. “General Purpose Technologies ‘Engines of Growth’?” *Journal of Econometrics* 65 (1): 83–108. [https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T).
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. “Language Models are Few-Shot Learners.” *Advances in Neural Information Processing Systems* 33 (NeurIPS 2020): 1877–1901. <https://arxiv.org/abs/2005.14165>.
- Brynjolfsson, Erik. 2022. “The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence.” *Daedalus* 151 (2): 272–87.
- Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2025. “Generative AI at Work.” *Quarterly Journal of Economics* 140 (2): 889–942.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.” *American Economic Journal: Macroeconomics* 13 (1): 333–72. doi:10.1257/mac.20180386.
- Caunedo, Julieta, Elisa Keller, and Yongseok Shin. 2023. “Technology and the Task Content of Jobs across the Development Spectrum.” *World Bank Economic Review* 37 (3): 479–93.
- Cazzaniga, Mauro, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J. Panton, Carlo Pizzinelli, Emma J. Rockall, and Marina Mendes Tavares. 2024. “Gen-AI: Artificial Intelligence and the Future

- of Work. IMF Staff Discussion Note 2024/001, International Monetary Fund, Washington, DC. <https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379>.
- Chan, Jason, Anindya Ghose, and Robert Seamans. 2016. “The Internet and Racial Hate Crime: Offline Spillovers from Online Access.” *MIS Quarterly* 40 (2): 381–404.
- Chatterji, Aaron, Thomas Cunningham, David J. Deming, Zoe Hitzig, Christopher Ong, Carl Yan Shan, and Kevin Wadman. 2025. “How People Use ChatGPT.” NBER Working Paper 34255, National Bureau of Economic Research.
- Chen, Zhisheng. 2023. “Ethics and Discrimination in Artificial Intelligence-Enabled Recruitment Practices.” *Humanities and Social Sciences Communications* 10 (567): 1–12.
- Christian, Brian. 2020. *The Alignment Problem: Machine Learning and Human Values*. W. W. Norton.
- Cohen, Julie E., Paul Ohm, Meg Leta Jones, Brenda Dvoskin, Smitha Krishna Prasad, and Nina-Simone Edwards. 2024. “Provisioning Digital Tools and Systems for Government Use.” Redesigning the Governance Stack Project at Georgetown Law, September 16. <https://ssrn.com/abstract=4958214>.
- Comin, Diego, Xavier Cirera, and Marcio Cruz. 2022. *Bridging the Technological Divide: Technology Adoption by Firms in Developing Countries*. World Bank Productivity Project, World Bank.
- Comin, Diego, and Martí Mestieri. 2018. “If Technology Has Arrived Everywhere, Why Has Income Diverged?” *American Economic Journal: Macroeconomics* 10 (3): 137–78. <https://doi.org/10.1257/mac.20150175>.
- Copestake, Alexander, Ashley Pople, Max Marczonek, and Katherine Stapleton. 2024. “AI and Services-Led Growth: Evidence from Indian Job Adverts.” <https://copestake.info/workingpaper/akai/AKAI.pdf>.
- Čorba, Juraj, Audrey Plonk, Karine Perset, Yoichi Iida. 2024. “Evolving with Innovation: The 2024 OECD AI Principles Update.” OECD.AI, May 20, 2024. <https://oecd.ai/en/wonk/evolving-with-innovation-the-2024-oecd-ai-principles-update>.
- Cottier, Ben, Robi Rahman, Loredana Fattorini, Nestor Maslej, and David Owen. 2025. “How Much Does It Cost Train Frontier AI Models?” Epoch AI blog, updated January 13, 2025. <https://epoch.ai/blog/how-much-does-it-cost-to-train-frontier-ai-models>.
- Cowan, Robin. 1990. “Nuclear Power Reactors: A Study in Technological Lock-In.” *Journal of Economic History* 50 (3): 541–67.
- Cross, James L., Michael A. Choma, and John A. Onofrey. 2024. “Bias in Medical AI: Implications for Clinical Decision-Making.” *PLOS Digital Health* 3 (11): e0000651.
- Cui, Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2024. “The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers.” SSRN: <https://ssrn.com/abstract=4945566>.
- David, Paul. 1986. “Understanding the Economics of QWERTY: The Necessity of History.” In *Economic History and the Modern Economist*, edited by William Parker. Basil Blackwell.
- David, Paul A. 1990. “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox.” *American Economic Review* 80 (2): 355–61.
- Davies, Matt, and Jai Vipra. 2025. *Computing Commons: Designing Public Compute for People and Society*. Ada Lovelace Institute. <https://www.adalovelaceinstitute.org/report/computing-commons/>.
- Dembele, Alou, Nouhoum Souleymane Coulibaly, and Michael Leventhal. 2025. “The Serendipity of Claude AI: Case of the 13 Low-Resource National Languages of Mali.” arXiv. <https://arxiv.org/abs/2503.03380>.
- Deming, David J. 2017. “The Growing Importance of Social Skills in the Labor Market.” *Quarterly Journal of Economics* 132 (4): 1593–1640.
- De Simone, Martín E., María Barron, Wuraola Mosuro, Federico Manolio, Federico Tiberti, and Eliot Dikoru. 2025. “From Chalkboards to Chatbots: Evaluating the Impact of Generative AI on Learning Outcomes in Nigeria.” Policy Research Working Paper 11125, World Bank. <https://doi.org/10.1596/1813-9450-11125>.
- DeepSeek. 2024. “DeepSeek-V3 Technical Report.” arXiv. <https://arxiv.org/abs/2412.19437>.

- Dell'Acqua, Fabrizio, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Kraye, François Candelon, and Karim R. Lakhani. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." Harvard Business School Technology and Operations Management Unit Working Paper 24-013.
- Demombynes, Gabriel, Jörg Langbein, and Michael Weber. 2025. "The Exposure of Workers to Artificial Intelligence in Low- and Middle-Income Countries." Policy Research Working Paper 11057, World Bank. <http://hdl.handle.net/10986/42765>.
- Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. "ImageNet: A Large-Scale Hierarchical Image Database." In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–55. Institute of Electrical and Electronics Engineers.
- D'Ignazio, Catherine, and Lauren F. Klein. 2020. *Data Feminism*. MIT Press. <https://data-feminism.mitpress.mit.edu/>
- Dolev, Niva, and Yariv Itzkovich. 2020. "In the AI Era, Soft Skills are the New Hard Skills." In *Management and Business Education in the Time of Artificial Intelligence: The Need to Rethink, Retrain, and Redesign (Research in Management Education and Development)*, edited by Agata Stachowicz-Stanusch and Wolfgang Amann. Information Age Publishing.
- Drozd, Lukasz, and Marina Tavares. 2024. "Generative AI: A Turning Point for Labor's Share?" Federal Reserve Bank of Philadelphia, Research Department. <https://www.philadelphiafed.org/the-economy/generative-ai-a-turning-point-for-labors-share>.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–08. <https://www.science.org/doi/10.1126/science.adj0998>.
- Epoch AI. 2025. *AI Model Releases Dataset, 1960–2025*. Epoch AI (accessed September 16, 2025). <https://epoch.ai>.
- Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses." *Strategic Management Journal* 42 (11): 2085–2111. <https://doi.org/10.1002/smj.3286>.
- Felten, Edward, Manav Raj, and Robert Seamans. 2023. "Occupational Heterogeneity in Exposure to Generative AI." arXiv [Preprint]. <https://arxiv.org/abs/2303.01157>.
- Finan, Frederico, Benjamin A. Olken, and Rohini Pande. 2017. "The Personnel Economics of the State." In *Handbook of Economic Field Experiments*, edited by Esther Duflo and Abhijit Banerjee, 467–514. North-Holland.
- Funk, Allie, Adrian Shahbaz, and Kian Vesteinsson. 2023. *Freedom on the Net 2023: The Repressive Power of Artificial Intelligence*. Freedom House. <https://freedomhouse.org/report/freedom-net/2023/repressive-power-artificial-intelligence>.
- GitHub. 2024. "Octoverse 2024: The State of Open Source" (accessed September 16, 2025). <https://octoverse.github.com/>.
- Gmyrek, Pawel, Mariana Viollaz, and Hernan Winkler. 2025. "Automation or Augmentation? How Development Constraints Split the Global Impact of GenAI." Background paper prepared for *World Development Report 2026: Artificial Intelligence for Development*.
- Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis. 2023. "Could Machine Learning Be a General Purpose Technology? A Comparison of Emerging Technologies Using Data from Online Job Postings." *Research Policy* 52 (1): 104653. <https://doi.org/10.1016/j.respol.2022.104653>.
- Goldin, Claudia, and Lawrence F. Katz. 2008. *The Race between Education and Technology*. Belknap Press of Harvard University Press.
- Goos, Maarten, Alan Manning, and Marcel Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104 (8): 2509–26. <https://doi.org/10.1257/aer.104.8.2509>.
- Gordin, Michael D. 2015. *Scientific Babel: How Science Was Done Before and After Global English*. University of Chicago Press.

- Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *Review of Economics and Statistics* 100 (5): 753–68. https://doi.org/10.1162/rest_a_00754.
- Greenbaum, Dov, and Mark Gerstein. 2025. "Hidden Human Costs of AI." *Science* 387 (6729): 32.
- Haakenstad, Annie, Caleb Mackay Salpeter Irvine, Megan Knight, Corinne Bintz, Aleksandr Y. Aravkin, Peng Zheng, Vin Gupta, Michael R. M. Abrigo, Abdelrahman I. Abushouk, Oladimeji M. Adebayo, et al. 2022. "Measuring the Availability of Human Resources for Health and its Relationship to Universal Health Coverage for 204 Countries and Territories from 1990 to 2019: A Systematic Analysis for the Global Burden of Disease Study 2019." *The Lancet* 399 (10341): 2129–54.
- Hadgu, Asmelash Teka, Paul Azunre, and Timnit Gebru. 2023. "Combating Harmful Hype in Natural Language Processing." Conference Paper, International Conference on Learning Representations (ICLR), Kigali, Rwanda. https://pml4dc.github.io/iclr2023/pdf/PML4DC_ICLR2023_39.pdf.
- Harari, Yuval N. 2024. *Nexus: A Brief History of Information Networks from the Stone Age to AI*. Random House.
- Henkel, Owen, Hannah Horne-Robinson, Nessie Kozhakhmetova, and Amanda Lee. 2024. "Effective and Scalable Math Support: Evidence on the Impact of an AI Tutor on Math Achievement in Ghana." *arXiv*. <https://arxiv.org/abs/2402.09809>.
- Hinton, Geoffrey, Oriol Vinyals, and Jeff Dean. 2015. "Distilling the Knowledge in a Neural Network." *arXiv preprint arXiv:1503.02531*.
- HubSpot. 2025. "AI Stats Every Startup Should Know" (accessed September 16, 2025). <https://www.hubspot.com/startups/ai/ai-stats-for-startups>.
- Hui, Xiang, Oren Reshef, and Luofeng Zhou. 2024. "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market." *Organization Science* 35 (6): 1977–1989.
- Jiang, Min, and Luca Belli. 2025. *Digital Sovereignty in the BRICS Countries: How the Global South and Emerging Power Alliances Are Reshaping Digital Governance*. Cambridge University Press.
- Juhász, Réka, Mara P. Squicciarini, and Nico Voigtländer. 2024. "Technology Adoption and Productivity Growth: Evidence from Industrialization in France." *Journal of Political Economy* 132 (10): 3215–59.
- Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. "Scaling Laws for Neural Language Models." *arXiv*. <https://arxiv.org/abs/2001.08361>.
- Khan, Misaal, Mahapara Khurshid, Mayank Vatsa, Richa Singh, Mona Duggal, and Kuldeep Singh. 2022. "On AI Approaches for Promoting Maternal and Neonatal Health in Low Resource Settings: A Review." *Frontiers in Public Health* 10: 880034.
- Kleiner, Morris M., Robert T. Marier, Rick Park, and Ian Wing. 2016. "Estimating the Impact of Scope-of-Practice Laws on Nurse Practitioners' Earnings." *Health Services Research* 51 (5): 1927–45.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2012. "ImageNet Classification with Deep Convolutional Neural Networks." *Advances in Neural Information Processing Systems* 25. https://proceedings.neurips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html.
- Kwa, Thomas, Ben West, Joel Becker, Amy Deng, Katharyn Garcia, Max Hasin, Sami Jawhar, Megan Kinniment, Nate Rush, Sydney Von Arx, et al. 2025. "Measuring AI Ability to Complete Long Tasks." <https://arxiv.org/pdf/2503.14499>.
- Lawson, Cina, Morlé Koudeka, Ana Lucía Cárdenas Martínez, and Luis Alberro. 2023. *Novissi Togo: Harnessing Artificial Intelligence to Deliver Shock-Responsive Social Protection*. Social Protection and Jobs Discussion Paper 2306, World Bank.
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–44.
- Liu, Yan. 2024. "Generative AI: Catalyst for Growth or Harbinger of Premature De-Professionalization?" Policy Research Working Paper (Digital), World Bank.
- Liu, Yan, and He Wang. 2024. "Who on Earth Is Using Generative AI?" Policy Research Working Paper 10870, World Bank.

- LMSYS. 2025. “Chatbot Arena Leaderboard.” LMSYS (accessed February 18, 2025). <https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>.
- Longpre, Shayne, Nikhil Singh, Manuel Cherep, Kushagra Tiwary, Joanna Materzynska, William Brannon, Robert Mahari, Naana Obeng-Marnu, Manan Dey, Mohammed Hamdy, et al. 2024. “Bridging the Data Provenance Gap across Text, Speech, and Video.” Data Provenance Initiative. <https://arxiv.org/pdf/2412.17847>.
- Ma, Michelle. 2025. “Thirsty AI Creates Another Climate Risk.” *Bloomberg*, May 8, 2025. <https://www.bloomberg.com/news/newsletters/2025-05-08/thirsty-ai-creates-another-climate-risk>.
- Marsal, Ales, and Patryk Perkowski. 2025. “Generative AI as Routine-Biased Technical Change? Evidence from a Field Experiment in Central Banking.” NBS Working Paper, National Bank of Slovakia. SSRN 5228176.
- Martens, Bertin. 2024. “Why Artificial Intelligence Is Creating Fundamental Challenges for Competition Policy.” Policy Brief 16/2024, Bruegel. <https://www.bruegel.org/policy-brief/why-artificial-intelligence-creating-fundamental-challenges-competition-policy>.
- Martins-Neto, Antonio, Yan Liu, Saloni Khurana, and Juan Manuel Porras Lopez. 2025. “Click, Code, Earn: The Returns to Digital Skills.” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5529280.
- Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, et al. 2024. *Artificial Intelligence Index Report 2025*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. https://hai.stanford.edu/assets/files/hai_ai-index-report-2024-smaller2.pdf.
- Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, et al. 2025. *Artificial Intelligence Index Report 2025*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf.
- McPeak, Grady, Anja Sautmann, Ohia George, Adham Hallal, Eduardo Arancón Simal, Aaron L. Schwartz, Jason Abaluck, Nirmal Ravi, and Robert Pless. 2024. “An LLM’s Medical Testing Recommendations in a Nigerian Clinic: Potential and Limits of Prompt Engineering for Clinical Decision Support.” *Proceedings of the 2024 IEEE 12th International Conference on Healthcare Informatics (ICHI)*. <https://doi.org/10.1109/ICHI61247.2024.00094>.
- Minniti, Antonio, Klaus Prettnner, and Francesco Venturini. 2025. “AI Innovation and the Labor Share in European Regions.” *European Economic Review* 177: 105043.
- Mitchell, Michelle, Bethany L. Hedt-Gauthier, Daniel Msellemu, Misheck Nkaka, and Neal Lesh. 2013. “Using Electronic Technology to Improve Clinical Care—Results from a Before-After Cluster Trial to Evaluate Assessment and Classification of Sick Children According to Integrated Management of Childhood Illness (IMCI) Protocol in Tanzania.” *BMC Medical Informatics and Decision Making* 13 (1): 95. <https://doi.org/10.1186/1472-6947-13-95>.
- Moore, Tyler, Richard Clayton, and Ross Anderson. 2009. “The Economics of Online Crime.” *Journal of Economic Perspectives* 23 (3): 3–20. <https://doi.org/10.1257/jep.23.3.3>.
- Mügge, Daniel. 2024. “EU AI Sovereignty: For Whom, to What End, and to Whose Benefit?” *Journal of European Public Policy* 31 (8): 2200–25.
- Müller, Karsten, and Carlo Schwarz. 2021. “Fanning the Flames of Hate: Social Media and Hate Crime.” *Journal of the European Economic Association* 19 (4): 2131–67.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar. 2016. “Building State Capacity: Evidence from Biometric Smartcards in India.” *American Economic Review* 106 (10): 2895–2929. <https://doi.org/10.1257/aer.20141346>.
- Muralidharan, Karthik, Abhijeet Singh, and Alejandro J. Ganimian. 2019. “Disrupting Education? Experimental Evidence on Technology-Aided Instruction in India.” *American Economic Review* 109 (4): 1426–60. <https://doi.org/10.1257/aer.20171112>.
- Narayanan, Arvind, and Sayash Kapoor. 2024. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton University Press.

- Nayyar, Gaurav, Mary Hallward-Driemeier, and Elwyn Davies. 2021. *At Your Service?: The Promise of Services-Led Development*. World Bank.
- Nayyar, Gaurav, Regina Pleninger, Dana Vorisek, and Shu Yu. 2024. “Digitalization and Inclusive Growth: A Review of the Evidence.” Policy Research Working Paper 10941, World Bank.
- Nothias, Toussaint. 2025. “An Intellectual History of Digital Colonialism.” *Journal of Communication* jqaf003.
- Noy, Shakked, and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381 (6657): 187–92. DOI: [10.1126/science.adh2](https://doi.org/10.1126/science.adh2).
- Oakden-Rayner, Luke, Jared Dunnmon, Gustavo Carneiro, and Christopher Ré. 2019. “Hidden Stratification Causes Clinically Meaningful Failures in Machine Learning for Medical Imaging.” Preprint, arXiv:1909.12475.
- O’Donnell, James, and Casey Crownhart. 2025. “We Did the Math on AI’s Energy Footprint. Here’s the Story You Haven’t Heard.” *MIT Technology Review*, May 20, 2025. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>.
- OECD (Organisation for Economic Co-operation and Development). 2024. “Robot Alice—Bid, Contract and Notice Analyser.” <https://oecd-opsi.org/innovations/robot-alice-bid-contract-and-notice-analyser/>.
- Otis, Nicholas G., Rowan Clarke, Solène Delecourt, David Holtz, and Rembrand Koning. 2024. “The Uneven Impact of Generative AI on Entrepreneurial Performance.” <https://ssrn.com/abstract=4671369>.
- Perrigo, Billy. 2023. “Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic.” *TIME*, January 18, 2023. <https://time.com/6247678/openai-chatgpt-kenya-workers/>.
- Quinlan, Ross J. 1986. “Induction of Decision Trees.” *Machine Learning* 1 (1): 81–106. <https://doi.org/10.1007/BF00116251>.
- Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. “Improving Language Understanding by Generative Pre-Training.” OpenAI. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf
- Radu, Roxana. 2025. “The Global Landscape of National AI Strategies.” Background paper prepared for *World Development Report 2026: Artificial Intelligence for Development*. World Bank.
- Rajab, Jenalea, Anuoluwapo Aremu, Everlyn Asiko Chimoto, Dale Dunbar, Graham Morrissey, Fadel Thior, Luandrie Potgieter, Jessico Ojo, Atnafu Lambebo Tonja, Maushami Chetty, Onyothi Nekoto, Pelonomi Moilola, Jade Abbott, Vukosi Marivate, and Benjamin Rosman. 2025. *The Esethu Framework: Reimagining Sustainable Dataset Governance and Curation for Low-Resource Languages*. arXiv preprint arXiv:2502.15916. [arXiv+2export.arxiv.org+2](https://arxiv.org/abs/2502.15916)
- Roberts, Sarah T. 2019. *Behind the Screen*. Yale University Press.
- Sevilla, Jaime, and Eduardo Roldán. 2024. *Training Compute of Frontier AI Models Grows by 4–5x per Year*. Epoch AI, May 28, 2024. <https://epoch.ai/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year>.
- Shojaee, Parshin, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. 2025. “The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity.” arXiv:2506.06941. Preprint, arXiv, July 18. <https://doi.org/10.48550/arXiv.2506.06941>.
- Shreve, James. 2024. “The Disgorgement Hammer: The FTC and Equitable Remedies.” JDSUPRA, November 14. <https://www.jdsupra.com/legalnews/the-disgorgement-hammer-the-ftc-and-7205380/>.
- Slough, Tara, Mark T. Buntaine, Cody Hickey, Patrick Hunnicut, Benjamin Leo, and Michael Wehner. 2021. “Testing the Accuracy of Non-Experts for Remote Monitoring of Water Services Using Satellite Imagery and Machine Learning.” *Proceedings of the National Academy of Sciences* 118 (49): e2109912118. <https://doi.org/10.1073/pnas.2109912118>.
- Stapleton, Katherine, and Layla O’Kane. 2021. “Artificial Intelligence and Services Offshoring.” Unpublished manuscript.

- Tan, Ruo Lin, Zhen Chee Kwan, Muhammad Abdul Bujang, Mohd Azahadi bin Omar, Siti Salwana binti Kamsani, Norizzati Bukhary, and Chee Ken Chan. 2024. “The Effectiveness of a Telemedicine Program for Managing Adults with Coronary Artery Disease Risk Factors in Underserved Communities: A Preliminary Randomized Controlled Trial.” *PLOS ONE* 19 (4): e0301340. <https://doi.org/10.1371/journal.pone.0301340>.
- The Agency Fund. 2025. “Rethinking AI Evaluation for the Social Sector: Launching the Living Playbook for AI Evaluation in the Social Sector.” The Agency Fund, August 19. <https://theagencyfund.substack.com/p/rethinking-ai-evaluation-for-the>.
- Tirole, Jean. 2021. “Digital Dystopia.” *American Economic Review* 111 (6): 2007–48.
- Tucker, Catherine. 2024. “How Does Competition Policy Need to Change in a World of Artificial Intelligence?” *Oxford Review of Economic Policy* 40: 834–42.
- Turing, Alan M. 1950. “Computing Machinery and Intelligence.” *Mind* 59 (236): 433–60. <https://doi.org/10.1093/mind/LIX.236.433>.
- United Nations. 2018. “Frontier Technologies for a Sustainable Future.” United Nations, September 28. <https://www.un.org/en/desa/frontier-technologies-sustainable-future>.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. “Attention Is All You Need.” <https://arxiv.org/pdf/1706.03762>.
- Viljoen, Salomé. 2021. “A Relational Theory of Data Governance.” *The Yale Law Journal*. https://www.yalelawjournal.org/pdf/131.2_Viljoen_1n12myx5.pdf.
- Vipra, Jai, and Anton Korinek. 2023. “Market Concentration Implications of Foundation Models: The Invisible Hand of ChatGPT.” Center on Regulation and Markets Working Paper No. 9, Brookings Institution. <https://www.brookings.edu/wp-content/uploads/2023/09/Market-concentration-implications-of-foundation-models-FINAL-1.pdf>
- Wang, Hanchen, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. 2023. “Scientific Discovery in the Age of Artificial Intelligence.” *Nature* 620: 47–60.
- Wang, Peng, Ling-Yu Zhang, Asaf Tzachor, and Wei-Qiang Chen. 2024. “E-Waste Challenges of Generative Artificial Intelligence.” *Nature Computational Science* 4: 818–23.
- Warso, Zuzanna, and Paul Keller. 2023. “Open Source, AI and the Paradox of Open.” OpenFuture. <https://openfuture.eu/blog/open-source-ai-and-the-paradox-of-open/>.
- Webb, Michael. 2020. “The Impact of Artificial Intelligence on the Labor Market.” https://www.michaelwebb.co/webb_ai.pdf.
- Weizenbaum, Joseph. 1966. “ELIZA—A Computer Program for the Study of Natural Language Communication between Man and Machine.” *Communications of the ACM* 9 (1): 36–45. <https://doi.org/10.1145/365153.365168>.
- Widder, David Gray, Sarah West, and Meredith Whittaker. 2023. “Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI.” Accepted to appear in *Nature*. SSRN. <https://ssrn.com/abstract=4543807>.
- Wild, Sarah Emily. 2025. “AI Models are Neglecting African Languages—Scientists Want to Change That.” *Nature*, July 29, 2025. <https://www.nature.com/articles/d41586-025-02292-5>.
- World Bank. 2021. *World Development Report 2021: Data for Better Lives*. World Bank.
- World Bank. 2024. *World Development Report 2024: The Middle-Income Trap*. World Bank.
- World Bank. 2025a. *Digital Progress and Trends Report 2025: Strengthening AI Foundations*. World Bank.
- World Bank. 2025b. *Harnessing Artificial Intelligence for Agricultural Transformation*. World Bank.
- World Bank. 2025c. *World Development Report 2025: Standards for Development*. World Bank.
- Yang, Jenny, Nguyen Thanh Dung, Pham Thach, and David A. Clifton. 2024. “Generalizability Assessment of AI Models across Hospitals in a Low-Middle and High-Income Country.” *Nature Communications* 15 (1): 1234. <https://doi.org/10.1038/s41467-024-52618-6>.
- Yao, Shunyu, Dian Yu, and Daniel S. Weld. 2023. “Tree of Thoughts: Deliberate Problem Solving with Large Language Models.” *arXiv*. <https://arxiv.org/abs/2305.10601>.

Zick, Tom, Mason Kortz, David Eaves, and Finale Doshi-Velez. 2024. "AI Procurement Checklists: Revisiting Implementation in the Age of AI Governance." *arXiv*. <https://arxiv.org/abs/2404.14660>