# 4

# Economy

## Overview

In 2025, more money flowed into AI than ever before, and faster. Global corporate AI investment more than doubled, revenue at leading frontier companies grew at historically fast rates. Generative AI reached close to 53% population-level adoption within three years of its mass-market introduction, faster than the personal computer or the internet, and that rapid uptake is translating into real value. U.S. consumer surplus from generative AI reached an estimated $172 billion annually by early 2026. But the benefits of this expansion are not distributed evenly. Investment is heavily concentrated in a small number of countries, companies and deals. In labor markets, demand for AI skills is rising across sectors but the workforce impact is showing signs of falling disproportionately on the youngest workers in AI-exposed occupations. Productivity gains are measurable within narrow tasks, but the evidence at the macro level remains early and mixed. The AI economy is scaling quickly, but how widely and how fairly that growth translates into real economic value is still an open question.

## Contents

Chapter Highlights 173

**4.1 Year in Review: 2025 175**

**4.2 Investment and Infrastructure 178**

Corporate Investment Activity 178

Private Investment Activity 179

AI Company Economics 189

Revenue 189

Annual Compute Spend 190

Capital Expenditures 191

Highlight: What Is Generative AI Worth? 192

**4.3 Corporate AI Adoption 193**

Industry Usage 193

By Industry and Function 194

Deployment Stages 197

Highlight: Measuring Signals of AI Diffusion 199

**4.4 Jobs 204**

AI Labor Demand Across Geographies 204

AI Hiring in the United States 205

Skill Composition 205

By Sector 209

By State 210

AI Hiring 212

AI Talent Concentration 214

AI's Labor Impact 219

Productivity Trends 219

Workforce Impact 221

**4.5 Robot Deployments 227**

Global Installation Patterns 227

Geographic Patterns 229

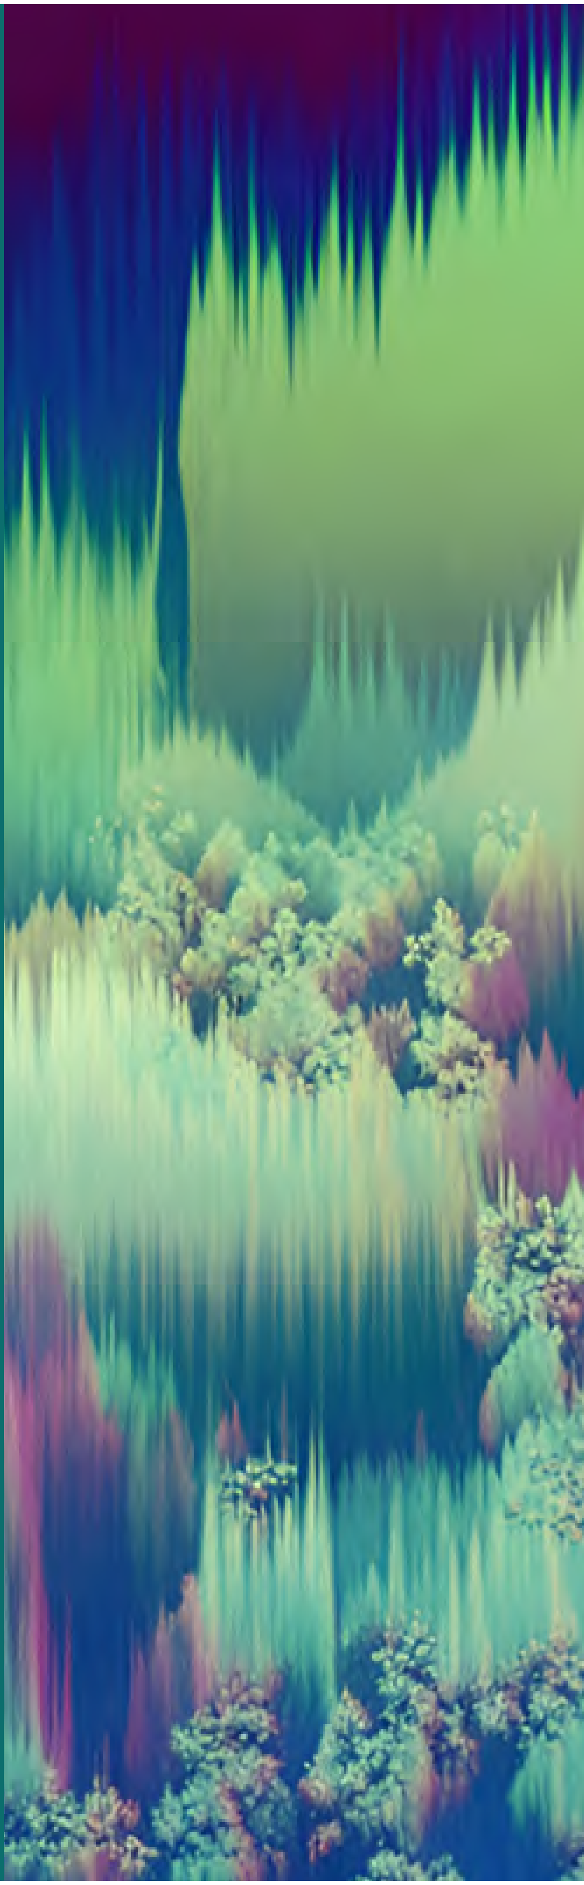

# Chapter Highlights

1 **Global corporate AI investment more than doubled in 2025.** Private investment grew fastest at 127.5% and now accounts for 60% of the total. Generative AI led the surge, growing more than 200% and capturing nearly half of all private AI funding. Newly funded AI companies rose 71%, and billion-dollar funding events nearly doubled.

2 **The United States continues to lead in global private AI investment, committing 23 times more than China.** In generative AI, U.S. investment exceeded the combined total of China and Europe by a wide margin. However, private investment figures likely understate China’s total AI spending, as government guidance funds have deployed an estimated $184 billion into AI firms between 2000 and 2023.

3 **AI company revenue is rising at historically fast rates, but compute costs and infrastructure spending are also reaching record levels.** Leading frontier companies are reaching meaningful revenue scale in a short period of time, but compute spend has increased significantly year-over-year. Major cloud providers have accelerated capital expenditures, with Google reporting more than $150 billion in annual capex in 2025.

4 **The value consumers get from generative AI grew 54% in a year.** Estimated U.S. consumer surplus reached $172 billion annually by early 2026, up from $112 billion a year earlier, with the median value per user tripling over the same period. Most of these tools remain free or close to it.

5 **Organizational AI adoption continued to rise in 2025, up to 88% of surveyed organizations, though AI agent use remains early.** Generative AI is now used in at least one business function at 70% of organizations, and China and Europe posted the highest year-over-year increases. AI agent deployment was in the single digits across nearly all business functions.

6 **Generative AI reached 53% adoption in three years, faster than the personal computer or the internet.** Adoption varies widely across countries and correlates strongly with GDP per capita, though some outpace what income would predict, including Singapore at 61% and the United Arab Emirates at 54%. Despite its lead in AI investment and model development, the United States ranks 24th at 28.3%.`

7 **AI’s labor market effects are showing up unevenly, concentrated in hiring pipelines and the youngest workers in exposed occupations.** Employment for software developers ages 22 to 25 has fallen nearly 20% from 2024. Employer surveys point to further change ahead, with one-third of respondents expecting workforce reductions over the coming year.

8 **One-third of organizations expect AI to reduce their workforce in the coming year, even though large-scale job losses have not yet shown up in overall employment data.** Almost half of organizations surveyed expected little to no change. Anticipated reductions are highest in service operations, supply chain, and software engineering. Across nearly all functions, anticipated decreases outpaces those already observed.

9 **Productivity gains from AI are largest in structured, measurable work where outputs are easy to monitor.** Studies report gains of 14% to 15% in customer support, 26% in software development, and 50% in marketing output. Gains are smaller in tasks requiring deeper reasoning, and recent evidence raises concerns that heavy AI reliance may carry long-term learning penalties that slow skill development over time.

10 **China continues to install more industrial robots than the rest of the world combined, and the gap widened in 2024.** China accounted for 54% of industrial robots installed globally, up from 51.1% in 2023. Global year-over-year growth was flat, and several major markets, including the United States, Germany, and Italy saw declines. Taiwan was an exception, recording the highest year-over-year growth at 33%.

# 4.1 Year in Review: 2025

2025

January 21

INVESTMENT

**$500B:** “Stargate Project” AI Infrastructure joint venture announced: OpenAI, SoftBank, Oracle, and MGX—supported by Nvidia and others—launch Stargate, a major AI infrastructure project announced at the White House. The venture plans to invest between $100 billion and $500 billion to build advanced AI data centers across the United States by 2029.

January 27

MILESTONE

**No. 1:** DeepSeek reaches No. 1 as the most downloaded free app on Apple’s U.S. App Store.

March 6

INVESTMENT

**$138B:** China announces a $138 billion state VC fund to invest in AI and other cutting-edge technologies.

March 10

ACQUISITION

**Agentic AI:** ServiceNow announces plans to acquire Moveworks to drive use of its agentic AI platform across key growth areas including CRM.

March 28

FUNDING

**$23B:** CoreWeave an AI data center company, has the largest U.S. tech IPO since 2021, raising $1.5 billion and valuing the company at $23 billion.

March 31

FUNDING

**$300B:** OpenAI raises $40 billion at a $300 billion post-money valuation.

May 13

INVESTMENT

**$5B:** AWS and HUMAIN announce a $5 billion AI infrastructure deal to accelerate AI adoption in Saudi Arabia and globally.

May 21

ACQUISITION

**$6.5B:** OpenAI acquires IO, the AI hardware startup founded by Jony Ive, for $6.5 billion to develop a new generation of consumer AI devices.

2025

June 2

ACQUISITION

**Watsonx AI:** IBM acquires the AI startup Seek AI to launch Watsonx AI Labs, an AI accelerator and innovation hub in New York City.

July 9

MILESTONE

**$4T:** Nvidia becomes the first public company worth $4 trillion.

July 11

DEAL

**$2.4B:** Google hires key staff from AI code-generation startup Windsurf and agrees to pay $2.4 billion in license fees to use some of Windsurf's technology on a nonexclusive basis.

July 15

FUNDING

**$12B:** Thinking Machines Lab an AI company founded by Mira Murati and other former OpenAI researchers, raises a $2 billion seed round at a $12 billion valuation.

September 2

FUNDING

**$183B:** Anthropic raises $13 billion in Series F funding at a $183 billion post-money valuation.

September 10

INVESTMENT
PARTNERSHIP

**$300B:** OpenAI signs a $300 billion, five-year cloud contract with Oracle, beginning in 2027. Oracle will provide 4.5 gigawatts of computing capacity for OpenAI's Stargate data center initiative.

October 27

FUNDING

**$10B:** Mercor which connects AI labs with domain experts for training their foundation AI models, raises $350 million Series C at a $10 billion valuation, making their founders, both 22 years old, the youngest ever self-made billionaires.

November 10

FUNDING

**$2.1B:** Gamma a startup that creates AI-generated presentations, websites, and social media posts, announces a $68 million Series B round at a $2.1 billion valuation led by Andreessen Horowitz.

2025

**November 11**

INVESTMENT

**$6.4B:** Google announces it will invest $6.4 billion in cloud infrastructure in Germany from 2026–29 to expand its data center capacity there.

**November 14**

FUNDING

**$29.3B:** Anysphere which sells the popular AI coding assistant Cursor, raises $2.3 billion at a $29.3 billion valuation.

**November 14**

INVESTMENT

**$40B:** Google announces a $40 billion investment in Texas data centers and AI infrastructure through 2027.

**November 20**

FUNDING

**$5.6B:** Physical Intelligence a robotics AI startup building general-purpose foundation-model "brains" for robots, raises $600 million led by CapitalG at a $5.6 billion valuation.

**December 9**

INVESTMENT

**$17.5B:** Microsoft pledges a $17.5 billion investment in India's AI Infra.

**December 10**

INVESTMENT

**1M:** Amazon announces it will invest over $35 billion in India by 2030 to expand AI and logistics, targeting 1 million new jobs and $80 billion in seller exports..

**December 22**

ACQUISITION

**$4.75B:** Alphabet says it will acquire Intersect for $4.75 billion in cash plus assumed debt to accelerate U.S. energy innovation and data center infrastructure build-out.

# 4.2 Investment and Infrastructure

The scale and direction of investment into AI provides a signal of the technology's broader economic trajectory. As AI systems become more capable and infrastructure-intensive, the capital required to develop and deploy them has expanded. Viewed alongside the broader trends discussed in other chapters of this report, these investment patterns capture not just market interest, but the rising cost of participating in the AI economy. This section examines those patterns across corporate infrastructure spending, startup funding activity, and the operational economics of AI companies themselves. The analysis draws primarily from Quid's database of AI-related investments, supplemented by publicly disclosed financial metrics from leading AI companies, as tracked by Epoch AI. The Quid investment data captures four categories of capital flowing into AI companies, including mergers and acquisitions, minority stake investments, private investment, and public offerings. Corporate investment, as used in this section, refers to the aggregate of all four. The private investment subsection that follows focuses more narrowly on private financing events, such as venture capital or private equity funding, directed at AI companies that have received over $1.5 million in funding since 2013. This subset represents a portion of total corporate investment.

## Corporate Investment Activity

Global corporate AI investment has grown over the past decade, accelerating within recent years, and further emphasizing how much AI has moved from an emerging technology to a key strategic priority. Across mergers and acquisitions, minority stakes, private investment, and public offerings, AI-related investment increased approximately fortyfold since 2013 (Figure 4.2.1). In 2025, total investment reached $581.69 billion, marking a 129.9% increase from the previous year. Private investments represented the largest share of activity with $344.66 billion, up 127.5% from 2024. Mergers and acquisitions showed similar signs of growth, rising 132.6% year over year. Though the composition of investment varies year to year, it is clear that organizations are committing growing sums of capital to strengthen their AI capabilities and position.

**Global corporate investment in AI by investment activity, 2013–25**

Source: Quid, 2025 | Chart: 2026 AI Index report

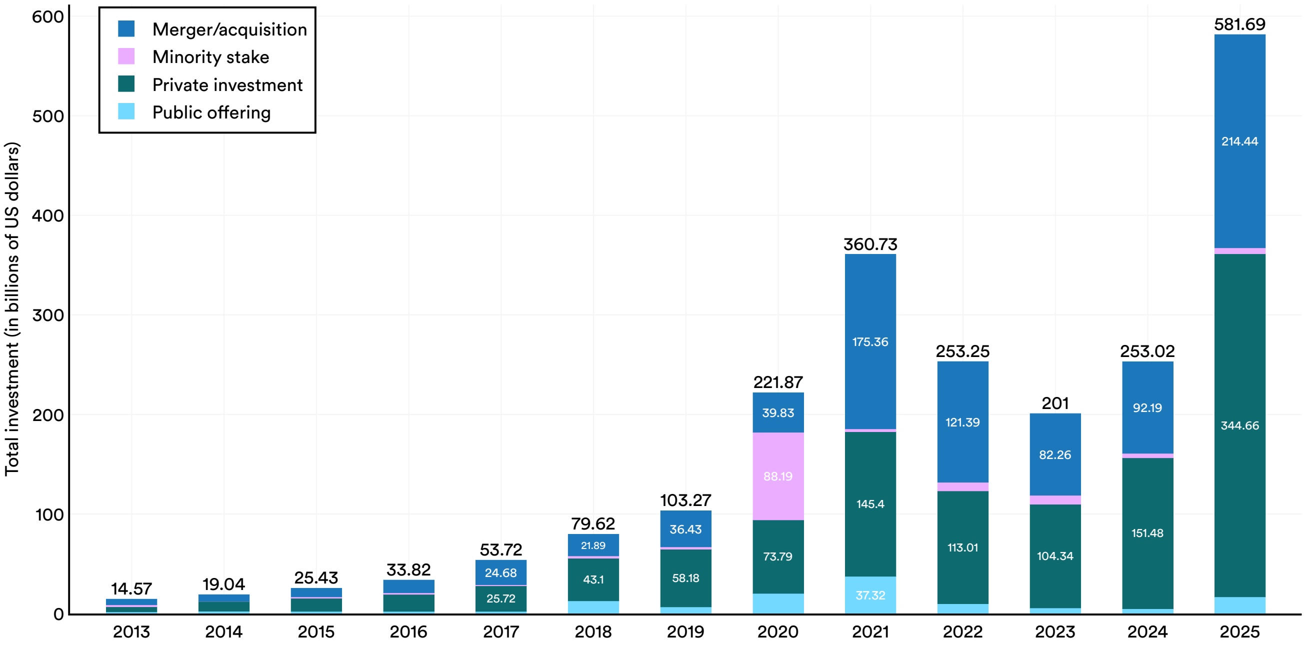

Figure 4.2.1

## Private Investment Activity

Within the broader investment landscape, private investment data, covering AI and ML companies with over $1.5 million in funding since 2013, offers a granular view into which firms are being funded and where that funding is concentrated. In 2025, global private investment in AI reached $344.7 billion, a 127.5% increase over the previous year (Figure 4.2.2). Generative AI companies accounted for $170.9 billion of that total, representing nearly half of all private investment and an increase of over 200% from 2024 (Figure 4.2.3).

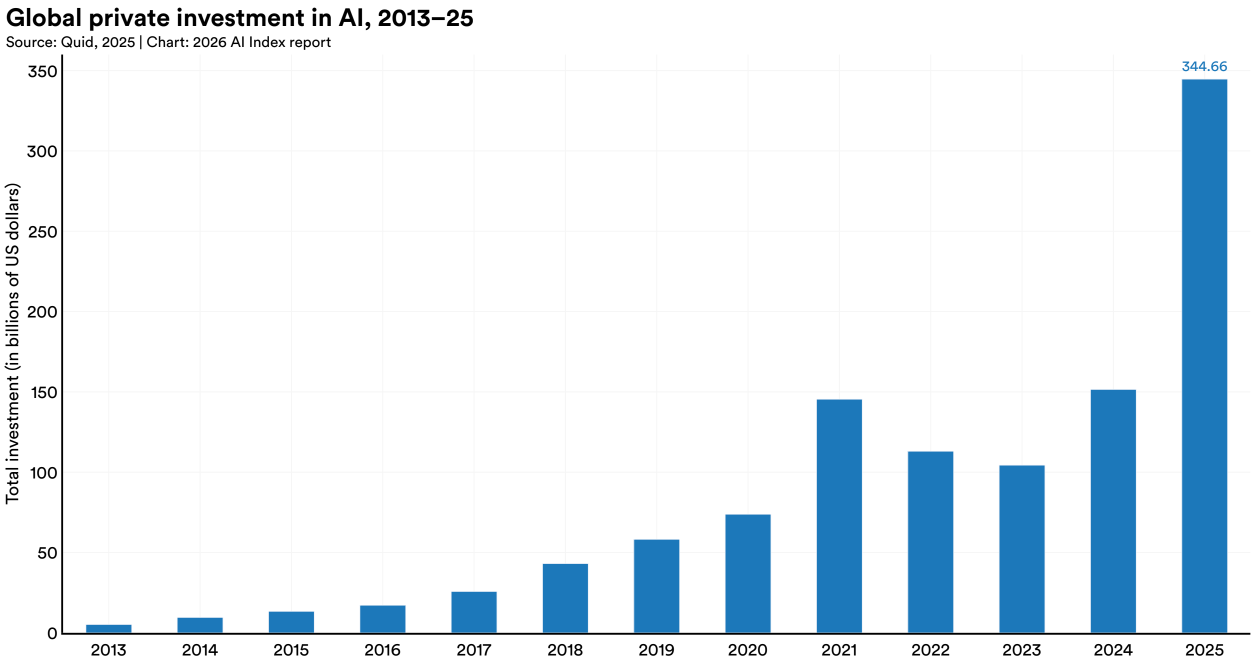

Figure 4.2.2

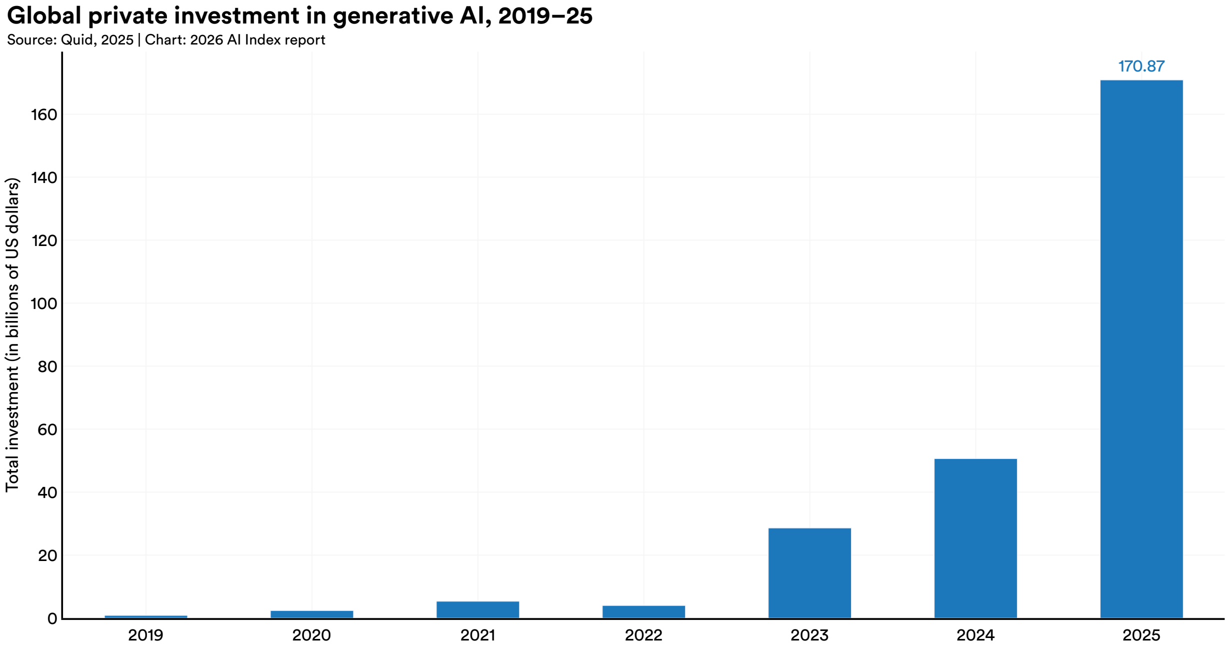

Figure 4.2.3

The private investment market is expanding in breadth but even more so in concentration. While the absolute number of newly funded AI companies has grown in 2025 (70.8% year-over-year increase), distribution of capital has dropped and the majority of investment dollars flow through a small number of deals (Figures 4.2.4-4.2.7). Compared to 2024, the average private AI investment event in 2025 increased 46% to $66.5 million. Investment activity increased across all funding sizes, but the strongest growth was at the upper end of that distribution, with 28 events exceeding $1 billion, up from 15 in 2024. The timeline of Section 4.1 notes several of these large funding rounds, including OpenAI's $40 billion raise and Anysphere's $2.3 billion round at a $29.3 billion valuation, highlighting the increasing skew in the funding landscape.

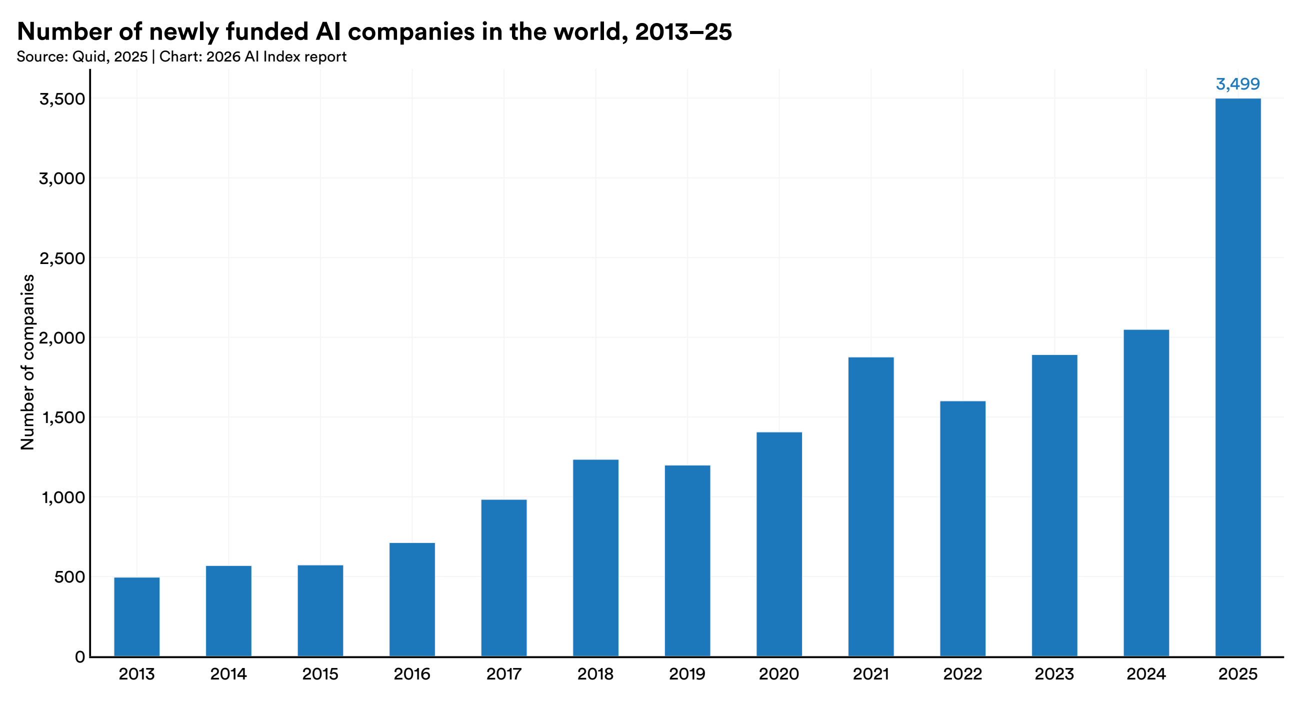
**Number of newly funded AI companies in the world, 2013–25**

Source: Quid, 2025 | Chart: 2026 AI Index report

Figure 4.2.4

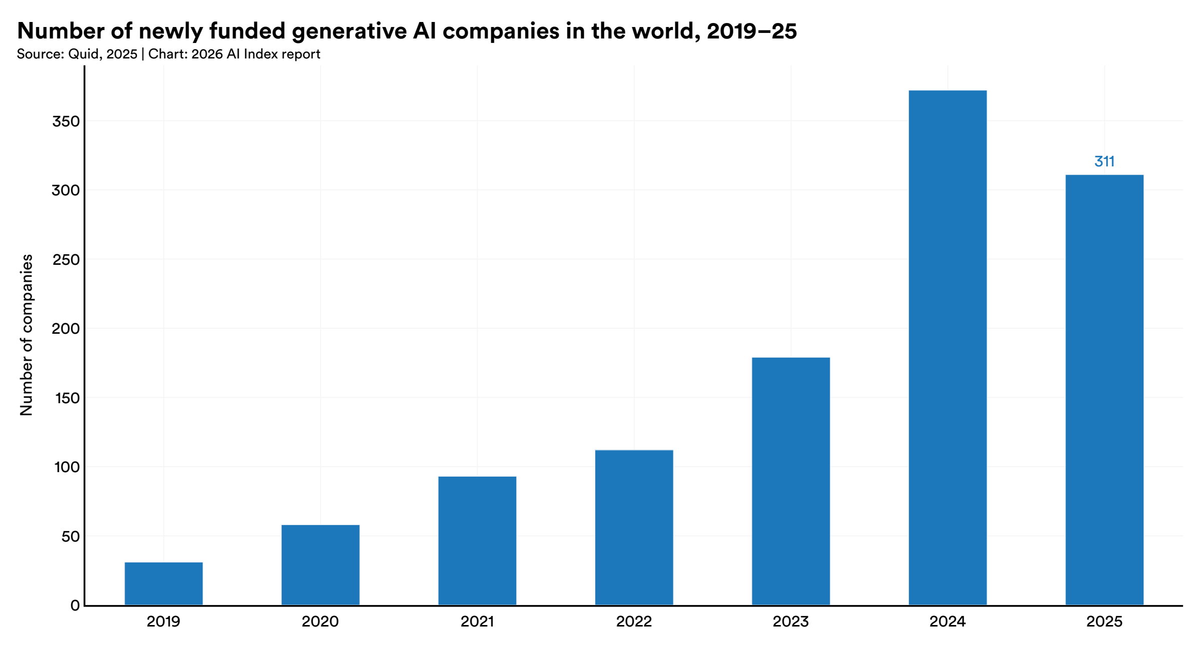
**Number of newly funded generative AI companies in the world, 2019–25**

Source: Quid, 2025 | Chart: 2026 AI Index report

Figure 4.2.5

## Average size of global AI private investment events, 2013–25

Source: Quid, 2025 | Chart: 2026 AI Index report

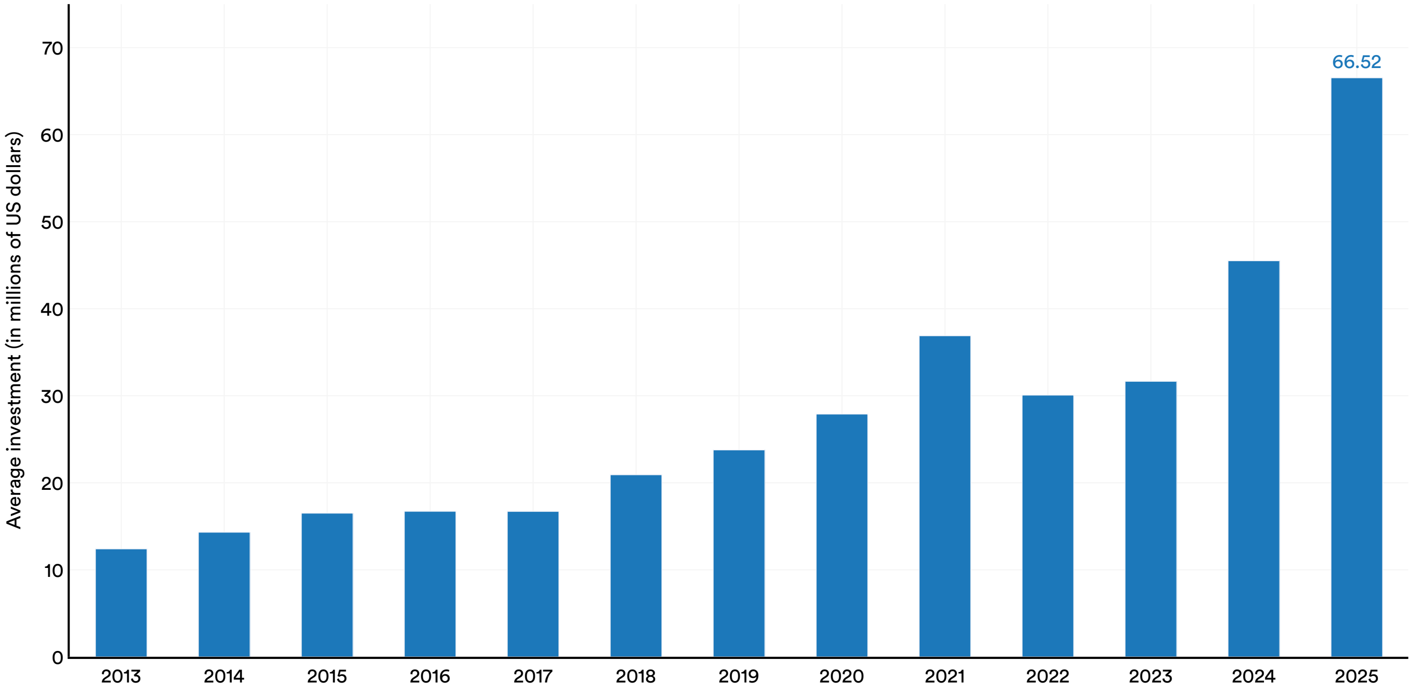

Figure 4.2.6

## Global AI private investment events by funding size, 2024 vs. 2025

Source: Quid, 2025 | Table: 2026 AI Index report

| Funding size (US dollars) | 2024 | 2025 |
|---|---|---|
| Over 1 billion | 15 | 28 |
| 500 million – 1 billion | 20 | 30 |
| 100 million – 500 million | 146 | 286 |
| 50 million – 100 million | 197 | 373 |
| Under 50 million | 2,951 | 4,464 |
| Undisclosed | 209 | 324 |
| Total | 3,538 | 5,505 |

Figure 4.2.7

## Geographic Distribution of Private Investment

As measured by both investment totals and the number of newly funded companies, private AI investment remains highly concentrated in a small number of countries. In 2025, the United States was the global leader with nearly $285.9 billion total invested, 23.1 times greater than the amount invested in the next highest country, China ($12.4 billion), and 48.5 times the amount invested in the United Kingdom ($5.9 billion) (Figure 4.2.8). This disparity was also seen in entrepreneurial activity, as the United States led with 1,953 newly funded AI companies in 2025, compared to 172 in the United Kingdom and 161 in China (Figure 4.2.9). In the United States, more than half of total private AI investment was generative AI-related ($163.6 billion), while the combined investment by China and Europe was $4.7 billion (Figure 4.2.10). Since 2024, private AI investment in the United States increased 160.2%, compared to an increase of 32.2% in China and 7.2% in Europe (Figure 4.2.11).

**Global private investment in AI by geographic area, 2025**

Source: Quid, 2025 | Chart: 2026 AI Index report

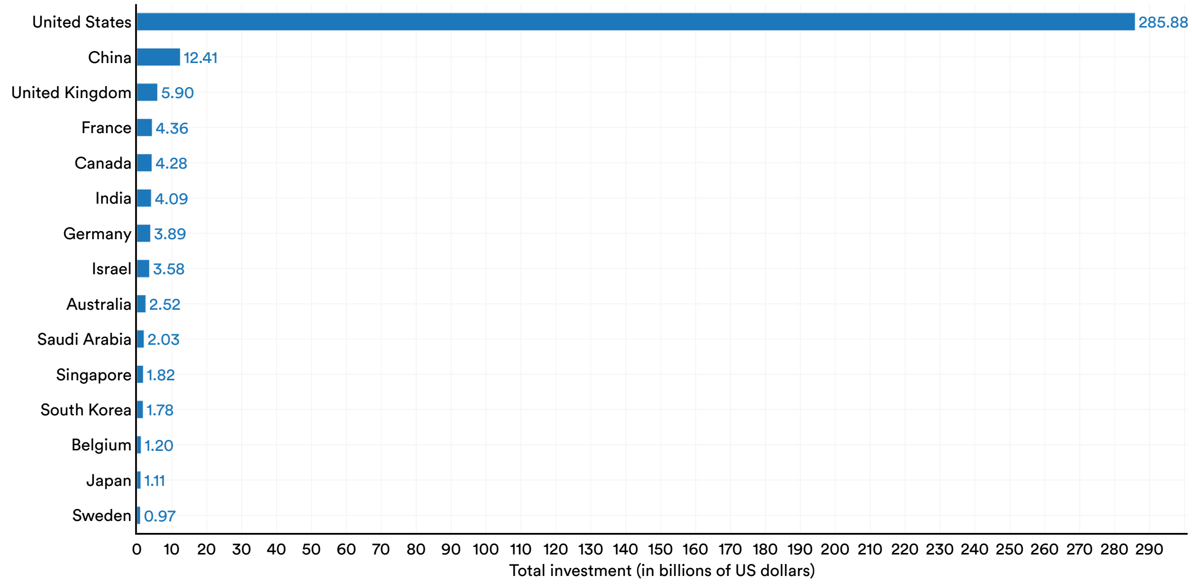

Figure 4.2.8

**Number of newly funded AI companies by geographic area, 2025**

Source: Quid, 2025 | Chart: 2026 AI Index report

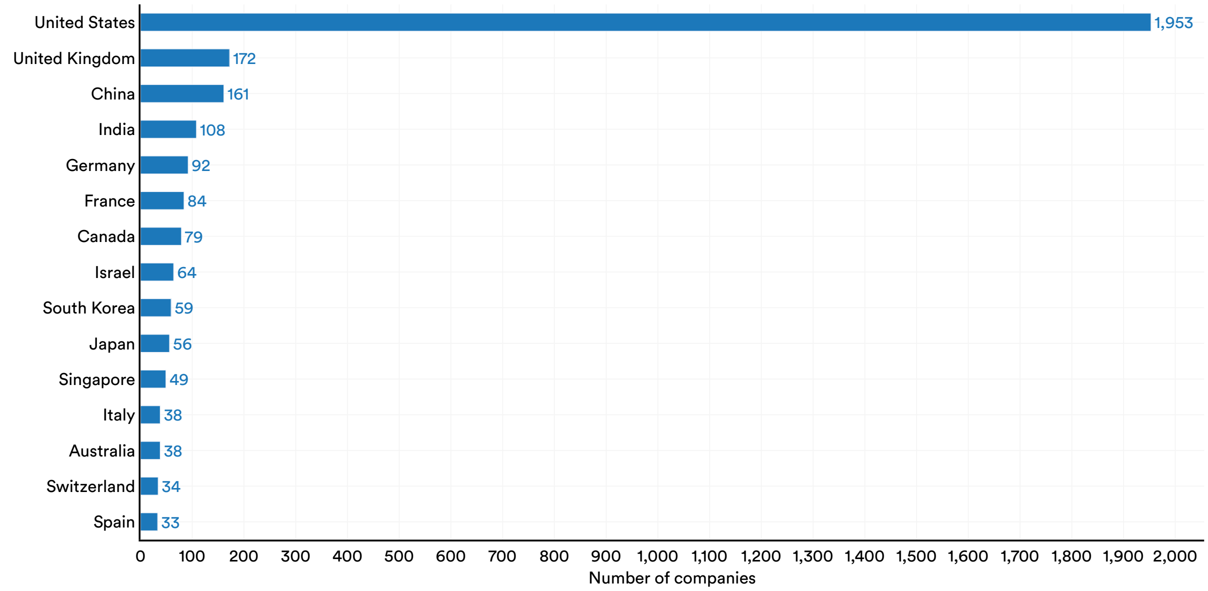

Figure 4.2.9

### Global private investment in generative AI by geographic area, 2019–25

Source: Quid, 2025 | Chart: 2026 AI Index report

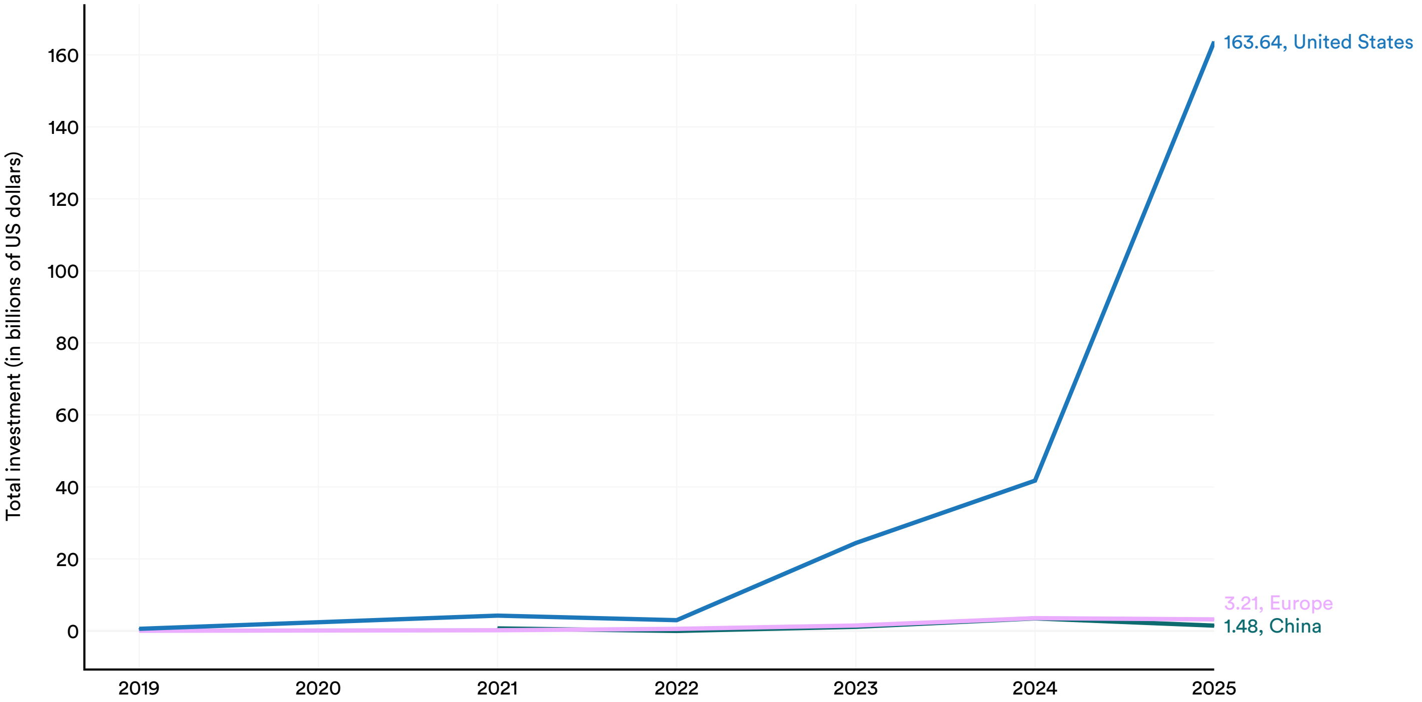

**Figure 4.2.10**

### Global private investment in AI by geographic area, 2013–25

Source: Quid, 2025 | Chart: 2026 AI Index report

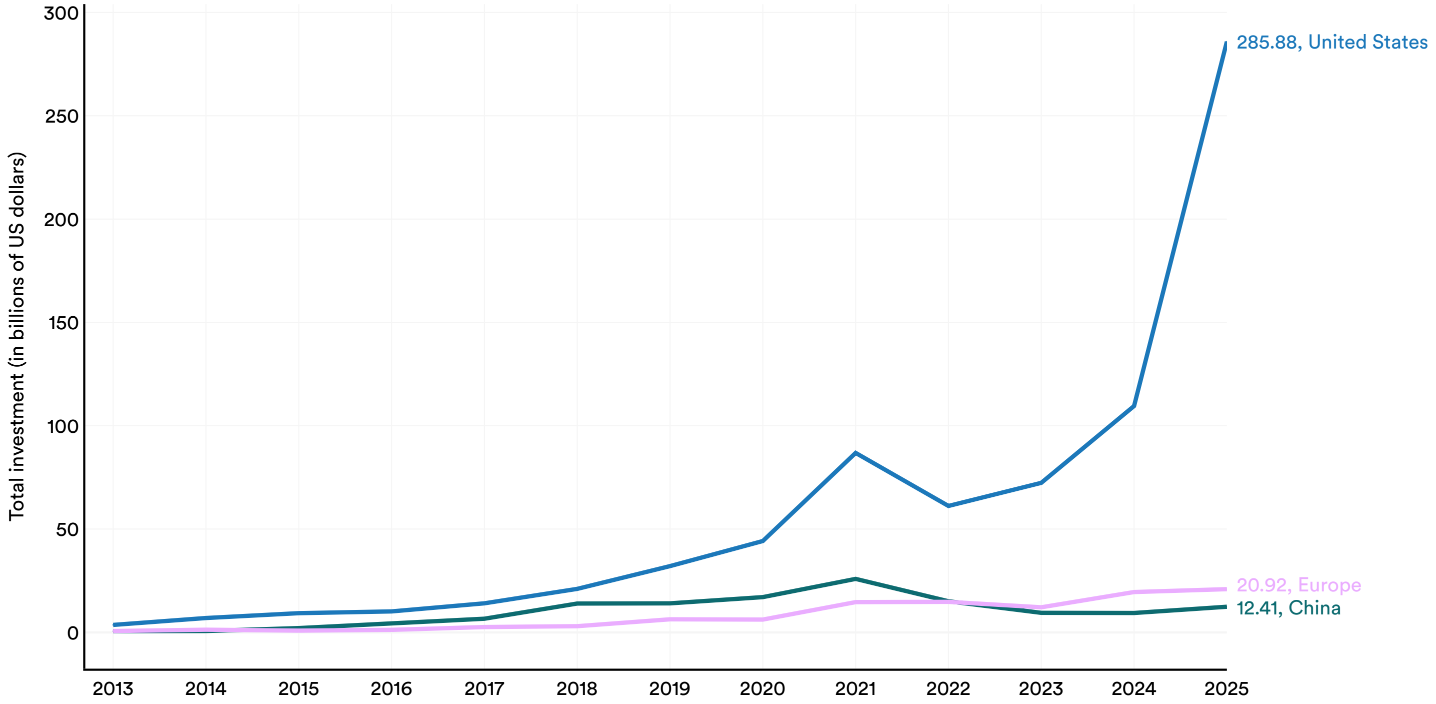

**Figure 4.2.11**

As noted earlier, the private investment figures in this section are drawn from Quid and do not account for government-backed funding in countries like China. For example, the Chinese government channels resources through government guidance funds, which are state-initiated investment funds that aim to both produce financial returns and further the government's strategic priorities (Beraja et al., 2024; Luong et al., 2021). Between 2000 and 2023, it was estimated that $912 billion of these funds were deployed across industries, with an estimated $184 billion allocated towards AI companies. Given this, comparisons based solely on private investment alone likely understate how much capital China is directing toward AI.

The trends in geographic concentration are also visible over a longer time horizon. Since 2013, the United States has attracted $757.3 billion in total private AI investment, far ahead of China at $131.8 billion (Figure 4.2.12). Other countries with notable cumulative investment totals include the United Kingdom ($34.1 billion), Canada ($19.6 billion), Israel ($18.5 billion), and Germany ($17.2 billion). Over the same period, the number of newly funded U.S. companies far exceeds other geographic areas, including five times that of China and 8.4 times the amount in the United Kingdom (Figures 4.2.13 and 4.2.14). The United States' growth rate continues to accelerate, with a 77.8% year-over-year increase in the number of funded AI startups.

**Global private investment in AI by geographic area, 2013–25 (sum)**

Source: Quid, 2025 | Chart: 2026 AI Index report

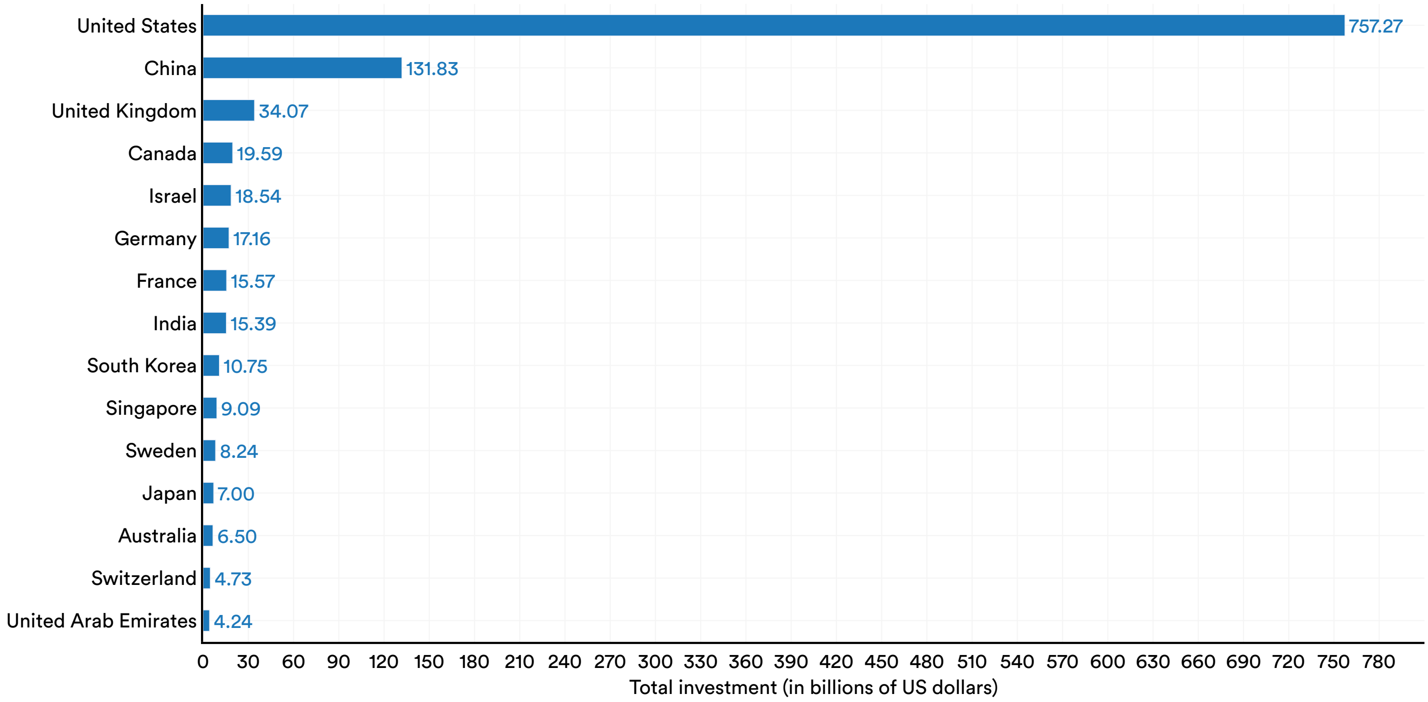

Figure 4.2.12

### Number of newly funded AI companies by geographic area, 2013–25 (sum)

Source: Quid, 2025 | Chart: 2026 AI Index report

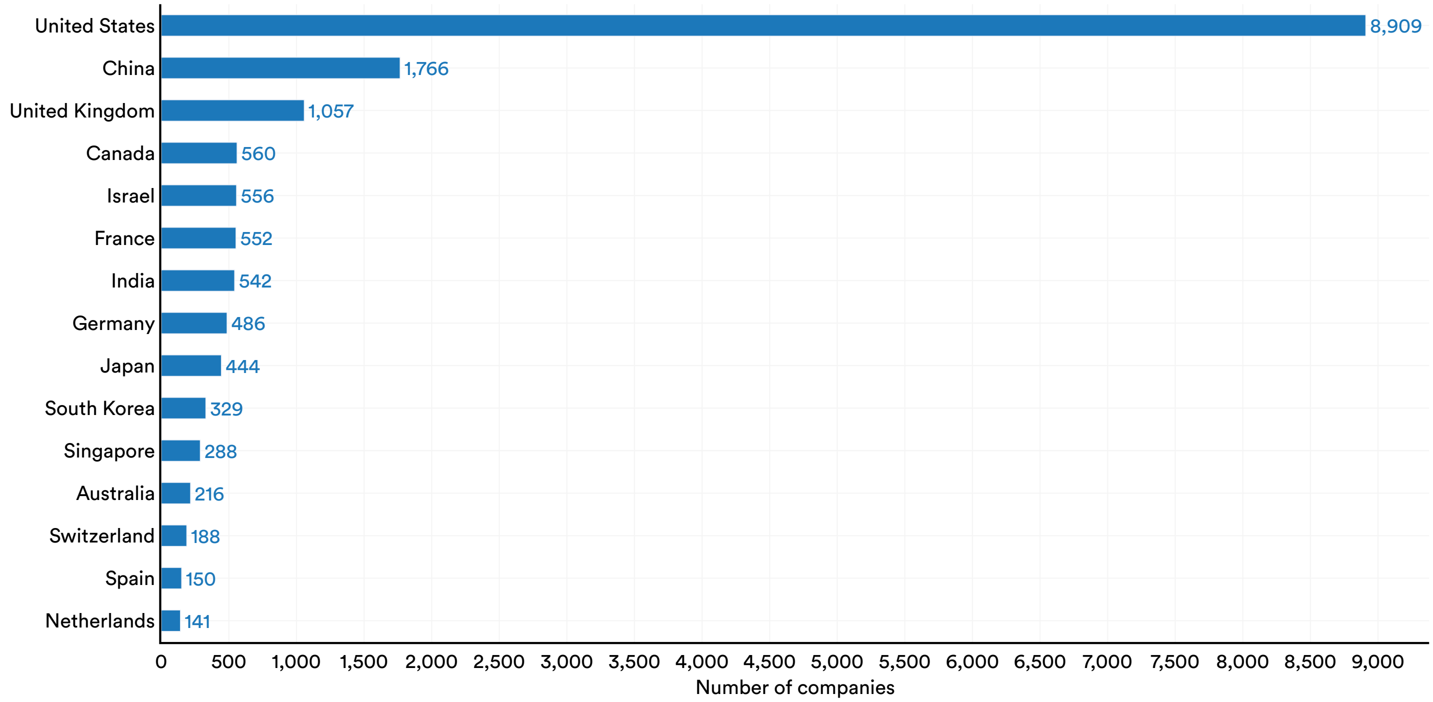

**Figure 4.2.13**

### Number of newly funded AI companies by geographic area, 2013–25

Source: Quid, 2025 | Chart: 2026 AI Index report

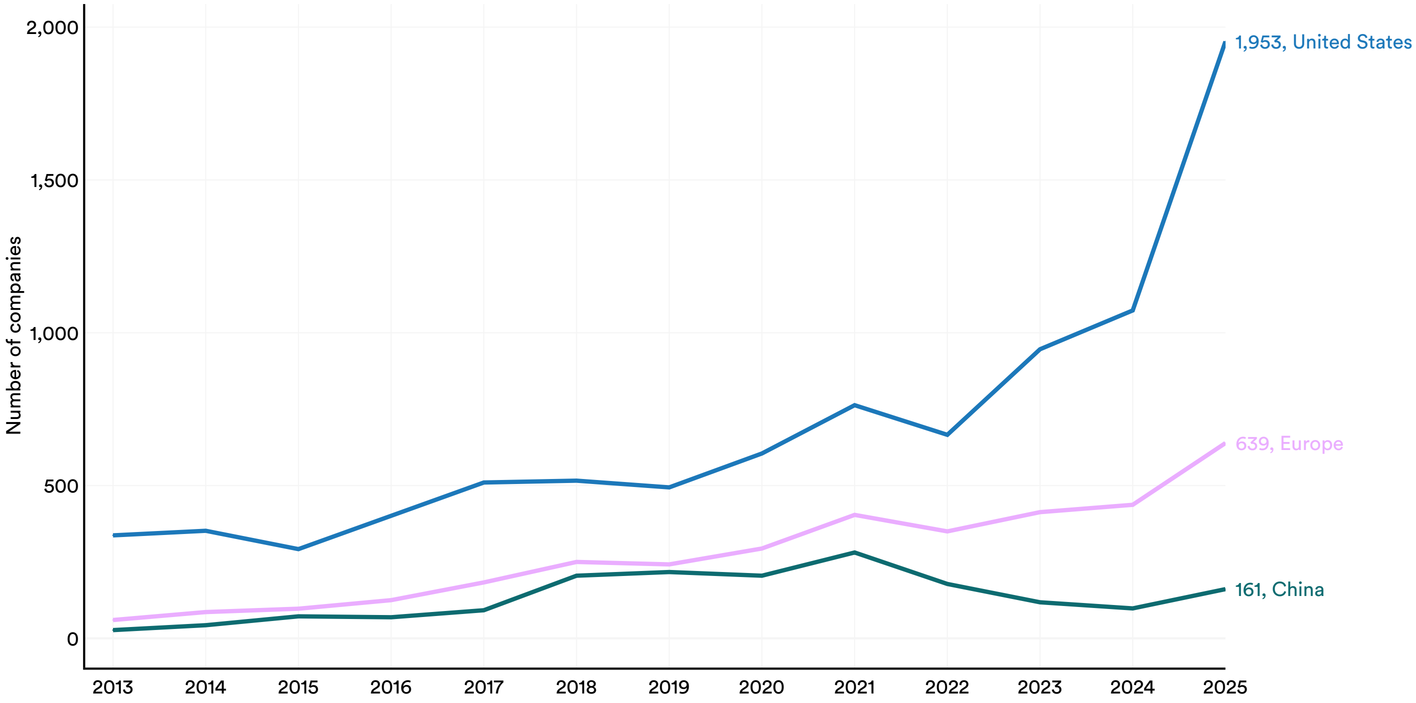

**Figure 4.2.14**

Within the United States, funding and entrepreneurial activity is heavily concentrated in a small number of states (Figure 4.2.15). California accounted for $218 billion in 2025, representing over 75% of the national total. Colorado ($19 billion), New York ($13 billion), and Florida ($6 billion) tracked the next largest investments. More than half of all U.S. states received less than $100 million in private AI investment, and a few, including South Dakota, Oklahoma, Arkansas, and West Virginia, reported no mapped investment activity. The underlying data for these state-level figures is not exhaustive but the overall pattern is clear.

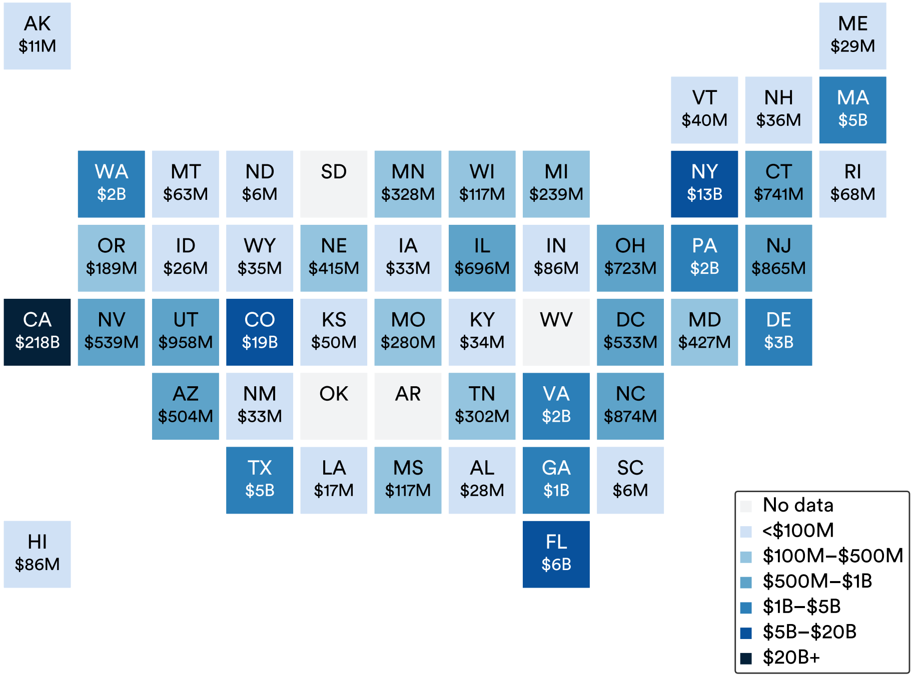

**US state private investment in AI, 2025**

Source: Quid, 2025 | Chart: 2026 AI Index report

Figure 4.2.15

### Focus Areas of Private Investment

In 2025, the breakdown of private AI startup investment by focus area suggests that capital was directed more heavily toward segments closest to building and scaling AI systems. The category of AI infrastructure, models, research, and governance attracted the largest share of funding, reaching $143.2 billion (Figures 4.2.16 and 4.2.17). As this category combines several types of priorities, the trend is best interpreted as evidence of growing investment in the foundational layers of the AI ecosystem rather than a precise measure of any single one of those areas. In recent years, this category has experienced the steepest growth in investment compared with all other areas. Other focus areas have expanded as well, including data management and processing and Internet of Things (IoT), yet none approaches the scale of foundational infrastructure. Alongside 2025 trends in technical performance (Chapter 2), and research and development (Chapter 1), this suggests that investment is tracking the infrastructure demands of deploying increasingly capable AI systems.

**Global private investment in AI by focus area, 2024 vs. 2025**

Source: Quid, 2025 | Chart: 2026 AI Index report

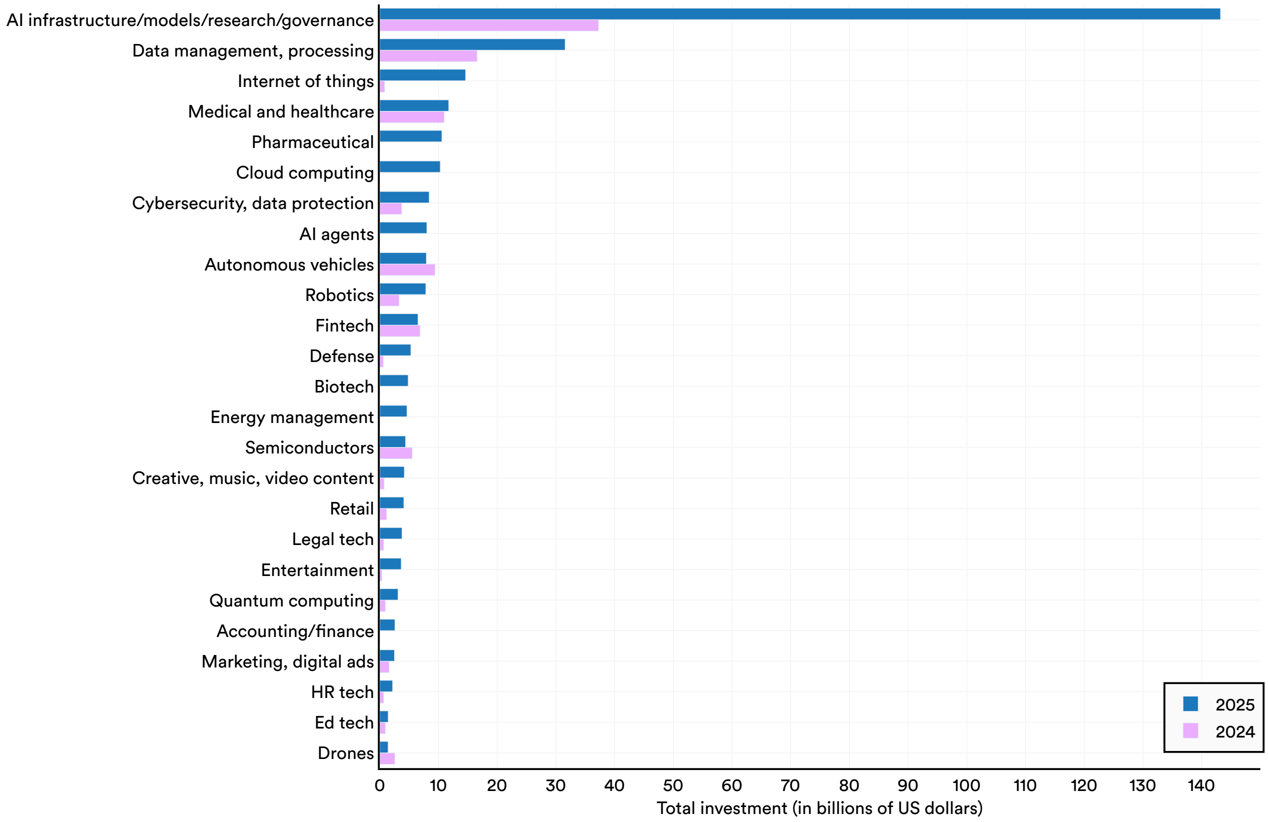

Figure 4.2.16

## Global private investment in AI by focus area, 2018–25

Source: Quid, 2025 | Chart: 2026 AI Index report

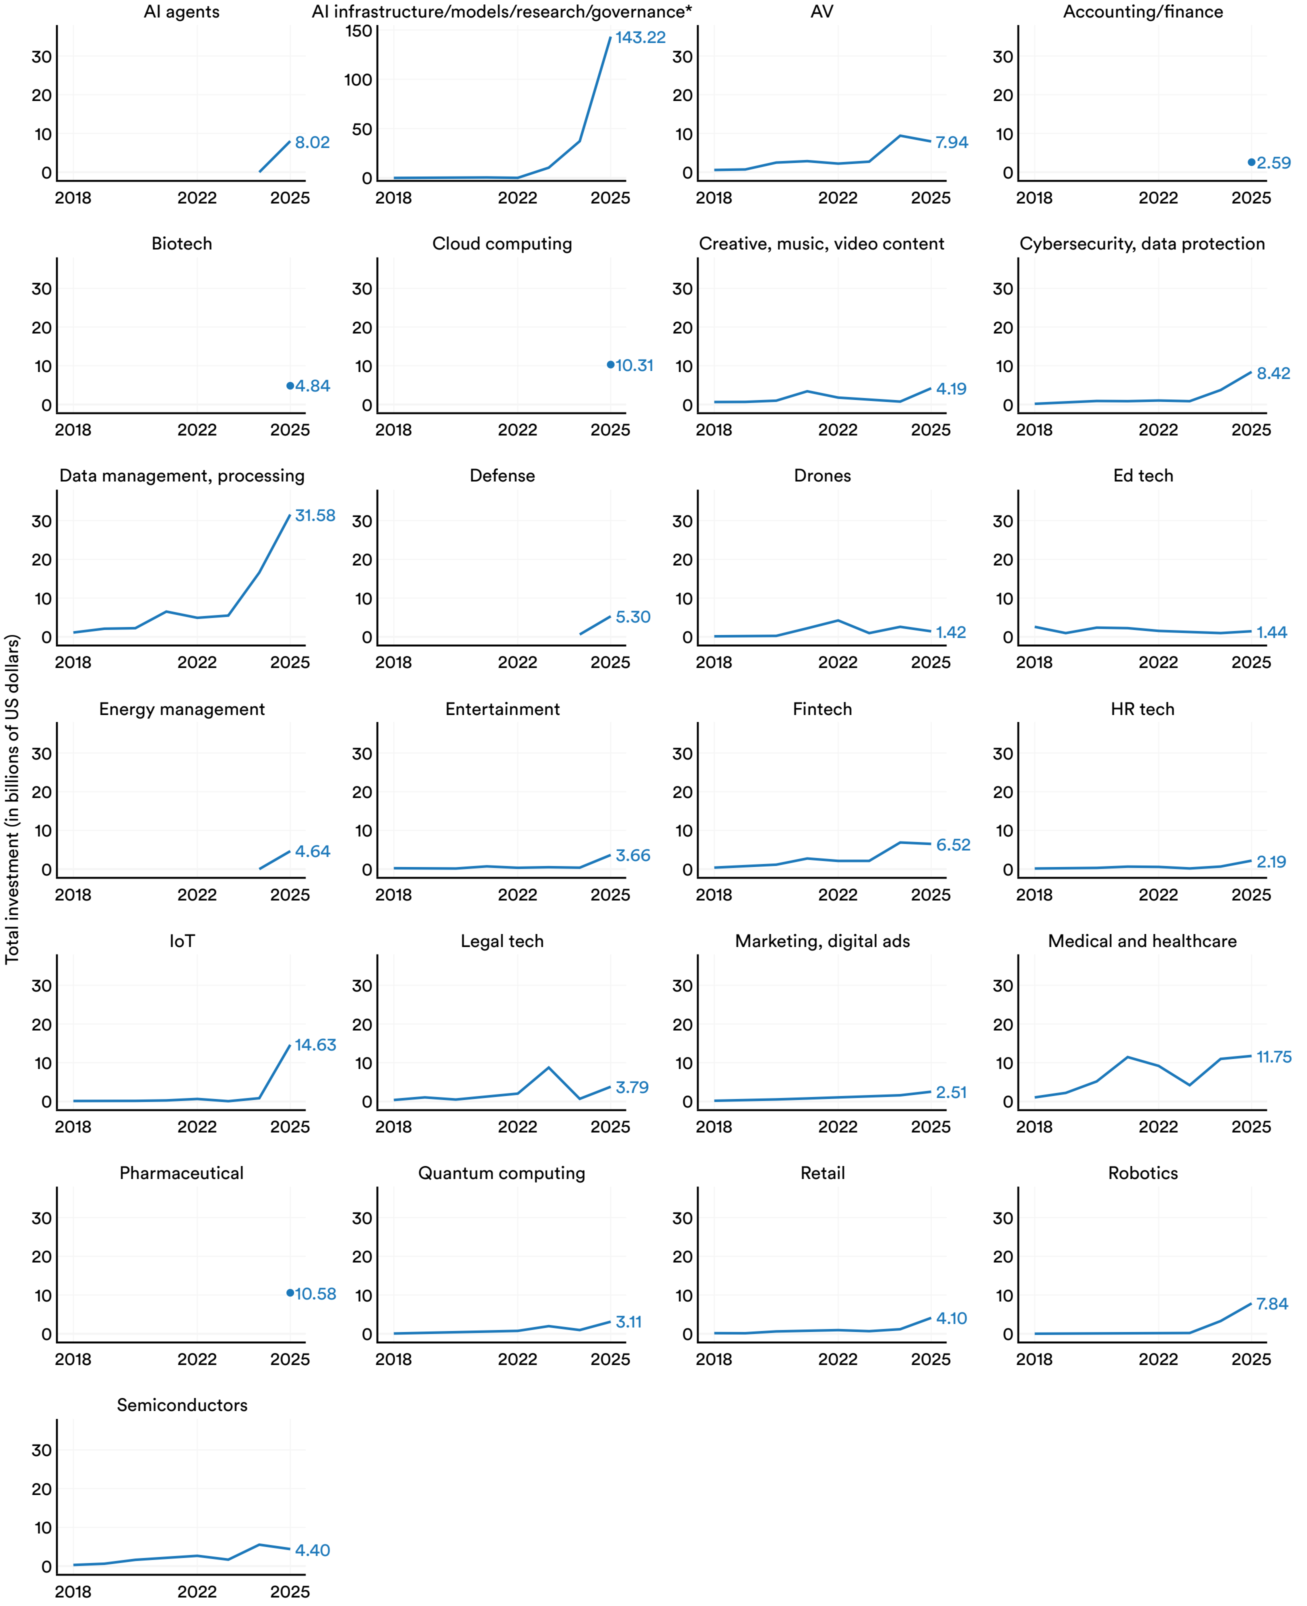

Figure 4.2.17

# AI Company Economics

Investment patterns show where capital is flowing across the AI ecosystem, but the operating economics of frontier AI companies reveal how advances in technical performance and deployment translate into commercial scale, compute demand, and infrastructure buildout. Using publicly disclosed data tracked by Epoch AI, this section examines those revenue trajectories, ongoing compute expenses, and the infrastructure costs to support AI development.

## Revenue

Annualized revenue estimates for leading AI companies, including OpenAI, Anthropic, xAI, and Mistral AI, have grown quickly in recent years (Figure 4.2.18). These estimates are drawn from direct company statements or established media reporting from 2023 through 2025. They may differ from annual recurring revenue calculations, and the underlying data varies in reliability depending on source credibility and accounting practices. Therefore, these figures should be interpreted as directional rather than precise. The chart uses a logarithmic scale to accommodate the exponential growth pattern, meaning that a straight line represents consistent percentage growth rather than absolute growth. With those notes in mind, the overall dynamic points to a small set of frontier AI companies reaching meaningful revenue scale in a relatively short amount of time.

To contextualize this growth, a separate comparison places OpenAI's revenue trajectory alongside those of other high-growth companies in the years after crossing $1 billion in annual revenue (Figure 4.2.19). While Google remains the only company in the comparison set to be scaling toward $100 billion in annual revenue, OpenAI's early revenue growth outpaces that of Uber, Cheniere Energy, and Moderna over comparable time periods.

**AI company annualized revenue**

Source: Epoch AI, 2026 | Chart: 2026 AI Index report

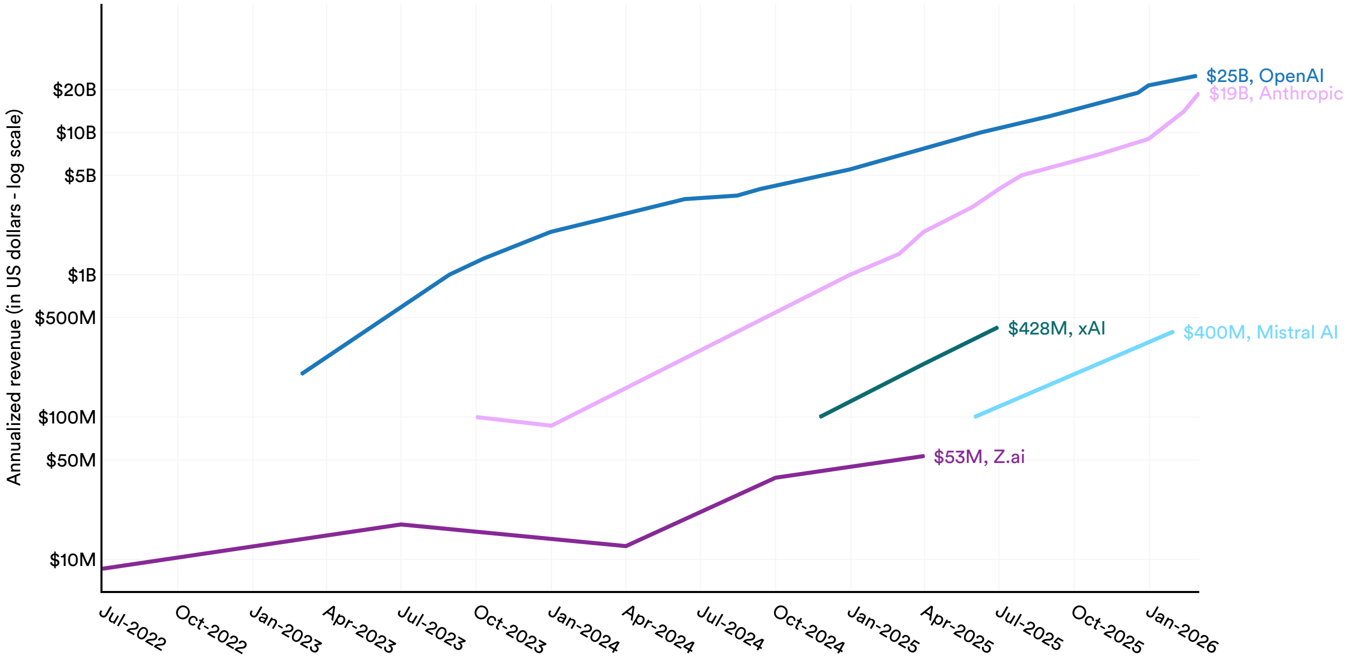

Figure 4.2.18

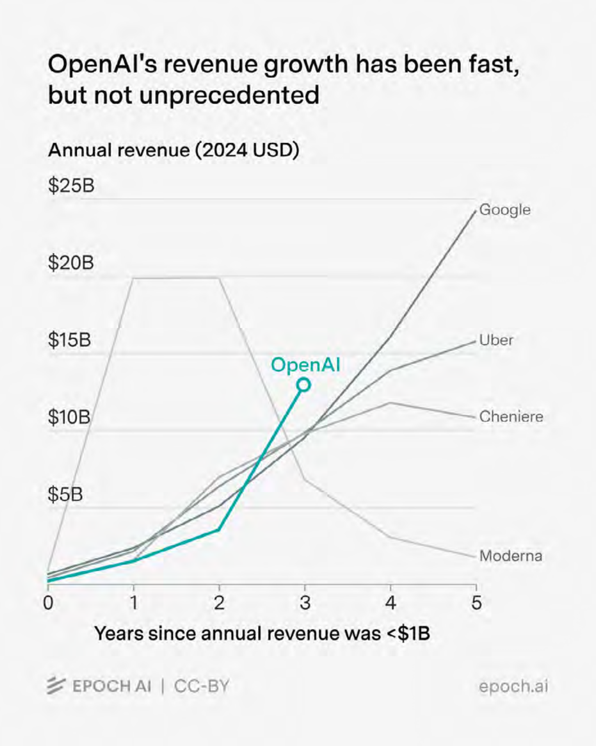

Figure 4.2.19

## Annual Compute Spend

The rapid revenue growth of leading AI companies has come with increasing compute costs. Reported annual compute spend, which largely reflects rented cloud capacity rather than owned data centers, offers a proxy for how much compute these companies procure each year to train and operate models at scale (Figure 4.2.20). OpenAI's reported compute spend increased significantly from 2024 to 2025, as did Anthropic's. The drive to meet growing commercial demand with increasingly capable systems means that the economics of frontier AI is tied to large-scale compute and its associated costs.

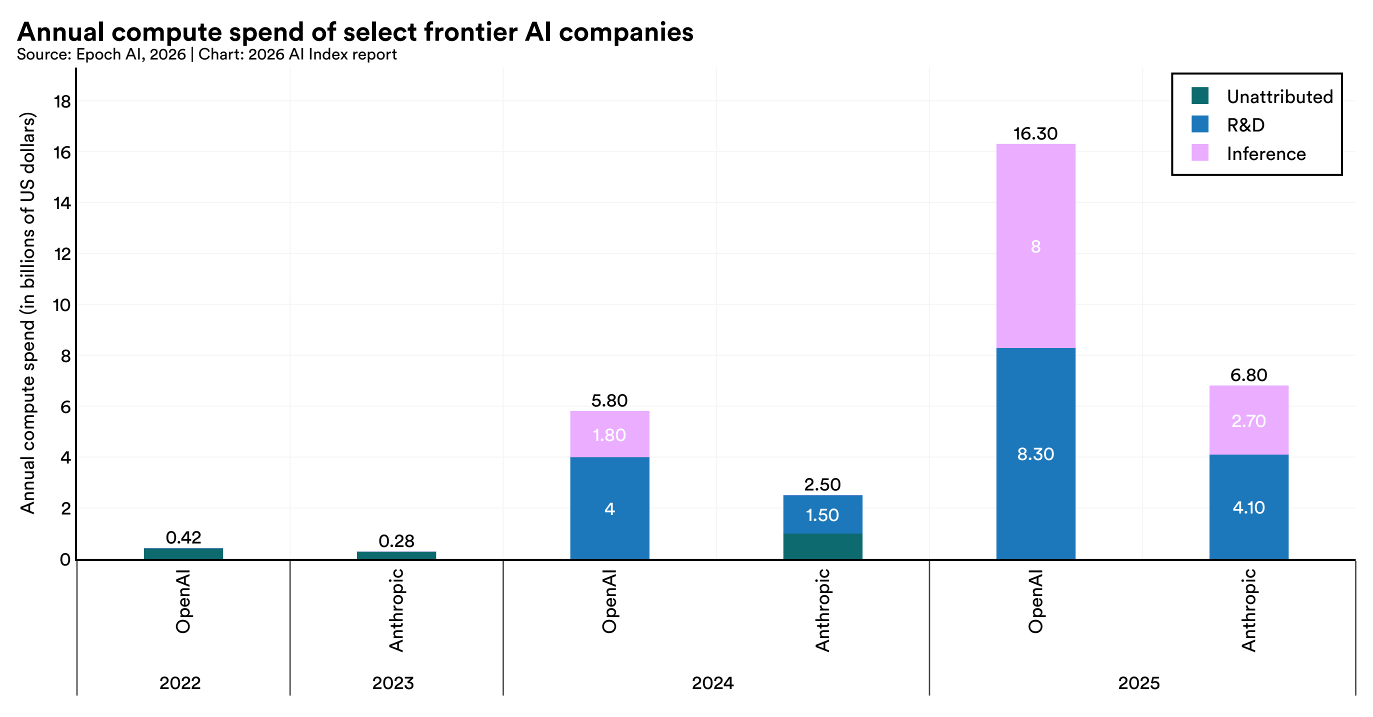

Figure 4.2.20

## Capital Expenditures

The infrastructure needed to support frontier AI is being financed not only by AI companies but also by the major cloud providers that lease them compute capacity. These providers have accelerated their own infrastructure investments to support the increasingly advanced AI models (Figure 4.2.21). In 2025, Google and Amazon led in total annual capital expenditures (capex), with Google reporting more than $150 billion in capex. This infrastructure investment is seen in the chapter's 2025 timeline, including the $100–$500 billion Stargate Project announced by OpenAI, SoftBank, Oracle, and others, as well as Google's $40 billion commitment to Texas data centers and Microsoft's $17.5 billion investment in AI infrastructure in India.

Hyperscaler's annual capex has more than doubled since ChatGPT's release

Hyperscaler annual capex (2025 and 2026 reflect estimates)

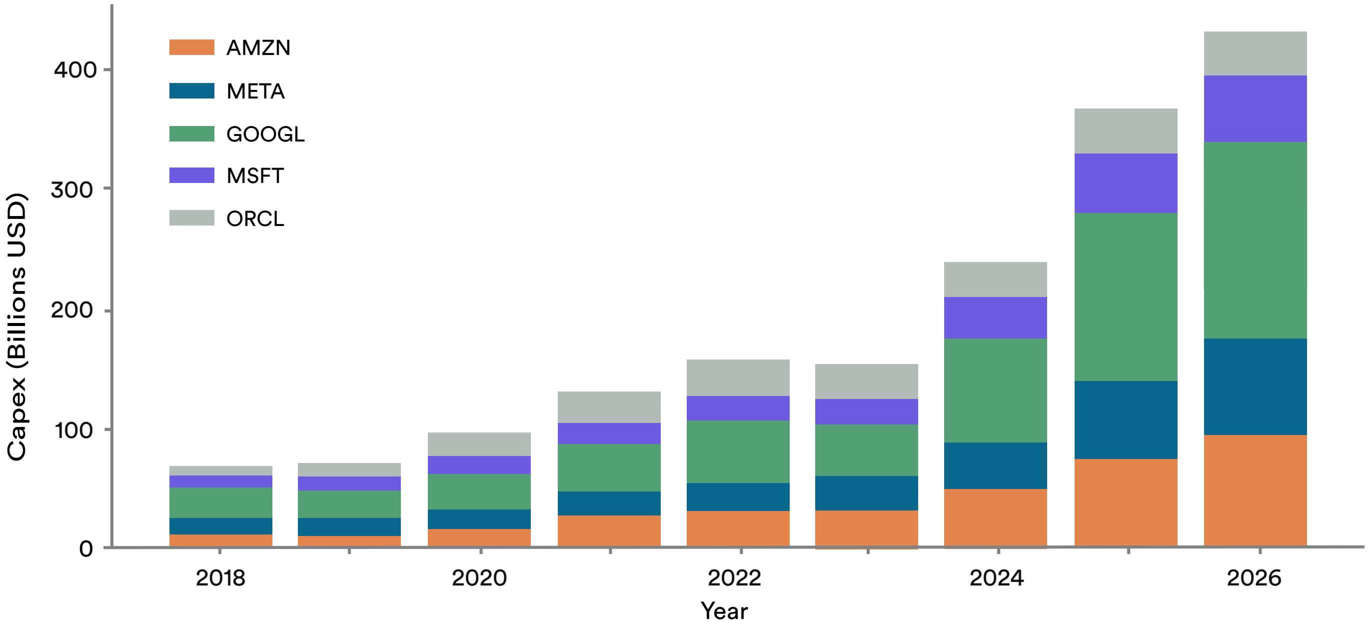

Exponential view. Source: Citi Research

Figure 4.2.21

**HIGHLIGHT:**

## What Is Generative AI Worth?

Investment, revenue and compute costs all measure the value of AI to the companies building and deploying it. They do not capture what the technology is worth to people using it. For most people, generative AI tools are free or close to it, making their economic value easy to undercount. Brynjolfsson et al. (2026) provide the first longitudinal estimates of that value, using online choice experiments conducted in 2025 (N=1,400) and early 2026 (N=2,000). Rather than measuring productivity effects, the study directly asked users' how much compensation they would accept to give up access to all generative AI tools for one month. This measure of "consumer surplus" is theoretically appropriate for goods that are largely free and already in the consumers' possession.

The study finds that total consumer surplus is estimated to have grown from $112 billion to $172 billion annually in the United States (4.2.22). This reflects the growing share of U.S. adults using generative AI, which increased from 48% to 56% (Bick et al., 2026) as well as a higher value per user. In particular, the average consumer surplus among U.S. generative AI users increased by 27% from $98 in 2025 to $125 by March 2026, while the median value per user tripled, from $3.40 to $11.40 over the same period. This increase in both adoption and per-user value is plausibly driven by a broadening and deepening of the capabilities of AI models.

This consumer surplus figure dwarfs estimated U.S. generative AI revenues, suggesting that the social returns from the technology far exceed the private returns captured by producers. This pattern is consistent with findings by Nordhaus (2004) that innovators historically capture only ~3% of total social returns from major technologies. The authors also find that usage frequency is the strongest individual-level predictor of surplus, followed by work use, number of different products used, and paid subscription status. Usage of generative AI for practical guidance, technical help, or information seeking are all associated with higher surplus.

**Generative AI consumer surplus in the United States, 2025 vs. 2026**

Source: Brynjolfsson et al., 2026 | Chart: 2026 AI Index report

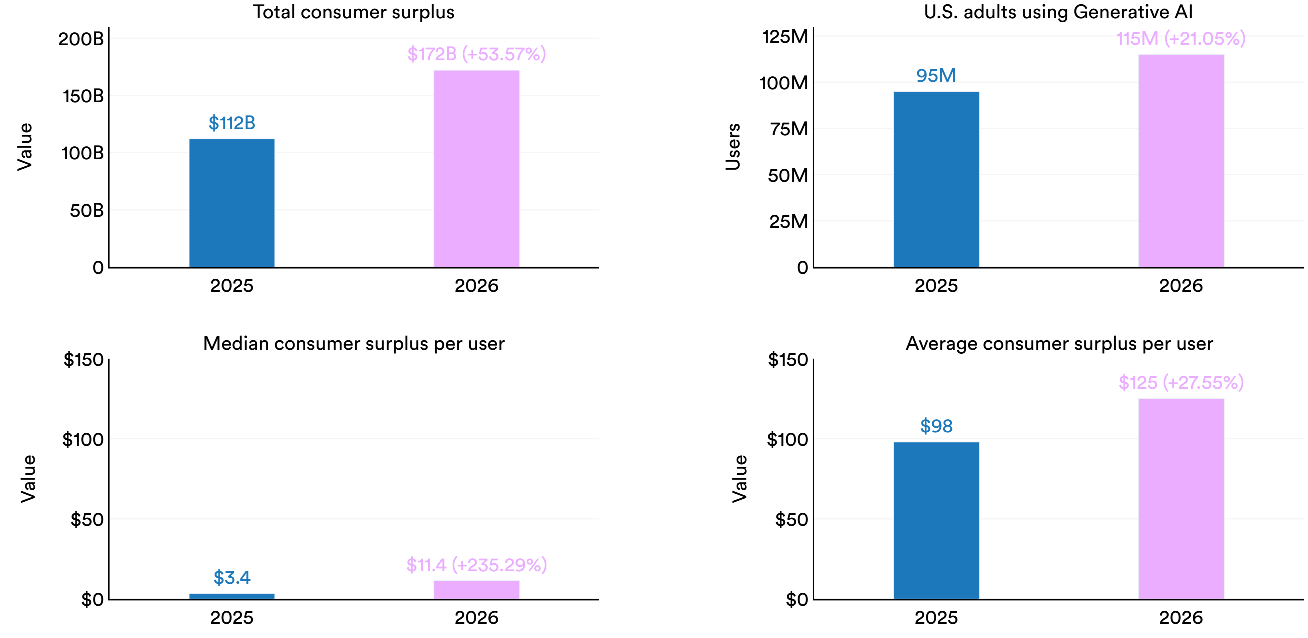

Figure 4.2.22

# 4.3 Corporate AI Adoption

The investment and infrastructure activity described earlier in this chapter establish the scale of resources being directed toward AI as well as the capacity being built to support it. This section examines how those investments translate into organizational use and reported business outcomes. Drawing on McKinsey & Company's annual State of AI surveys and other enterprise measures, the analysis traces the breadth and depth of adoption, and its associated benefits. As with other survey-based data in this chapter, the results are self-reported and should be viewed as directional rather than comprehensive.

## Industry Usage

In 2025, organizational adoption of AI continued to expand in both usage and function. A large majority of respondents reported that their organization uses AI in at least one business function, up to 88% in 2025 from 78% in 2024 (Figure 4.3.1). Over half of respondents reported three or more business functions leveraging AI. Use of generative AI mirrored that growth, with 79% of respondents reporting that their organizations regularly use generative AI in at least one business function, compared to 71% in 2024. This expanded adoption of AI was seen across all regions, though at different rates (Figure 4.3.2). China and Europe experienced higher year-over-year increases, with reported organizational AI use growing 13 and 11 percentage points, respectively.

**Share of respondents who say their organization uses AI in at least one function, 2017–25**

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

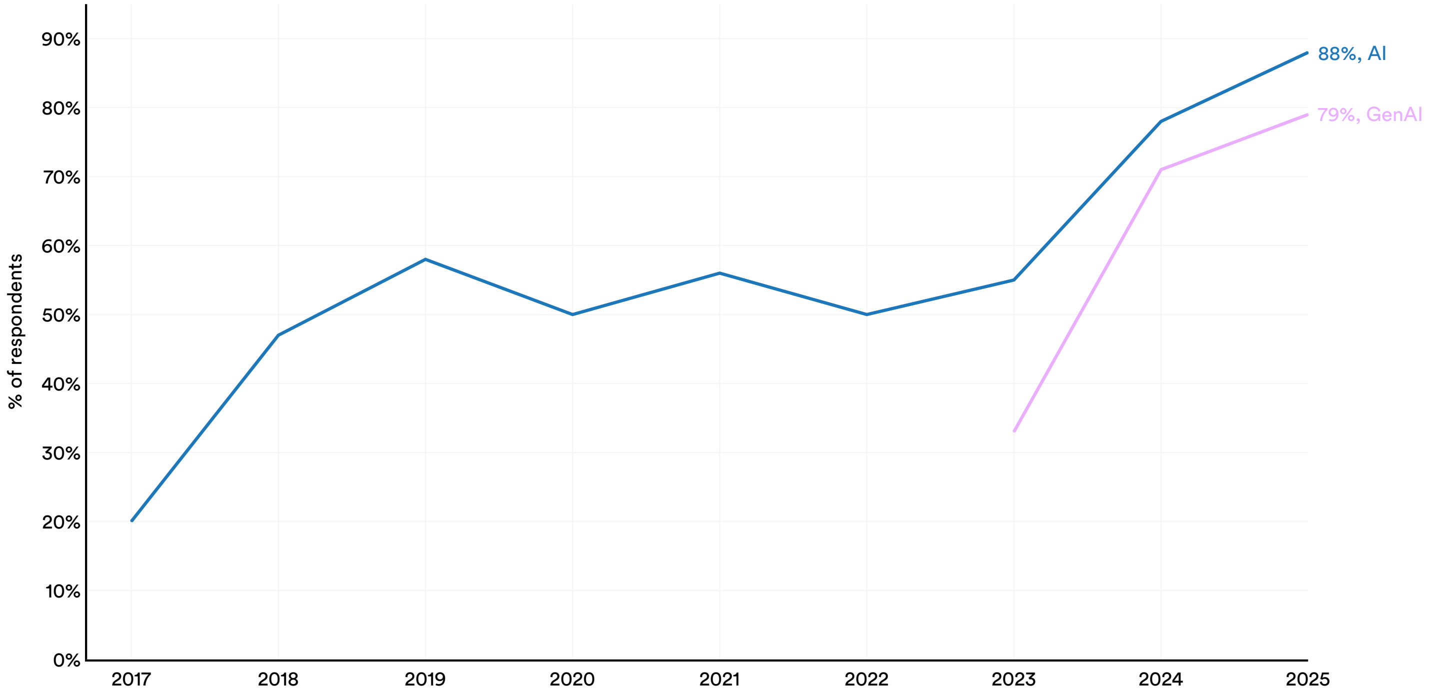

Figure 4.3.1

**AI use by organizations in the world, 2023–25**

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

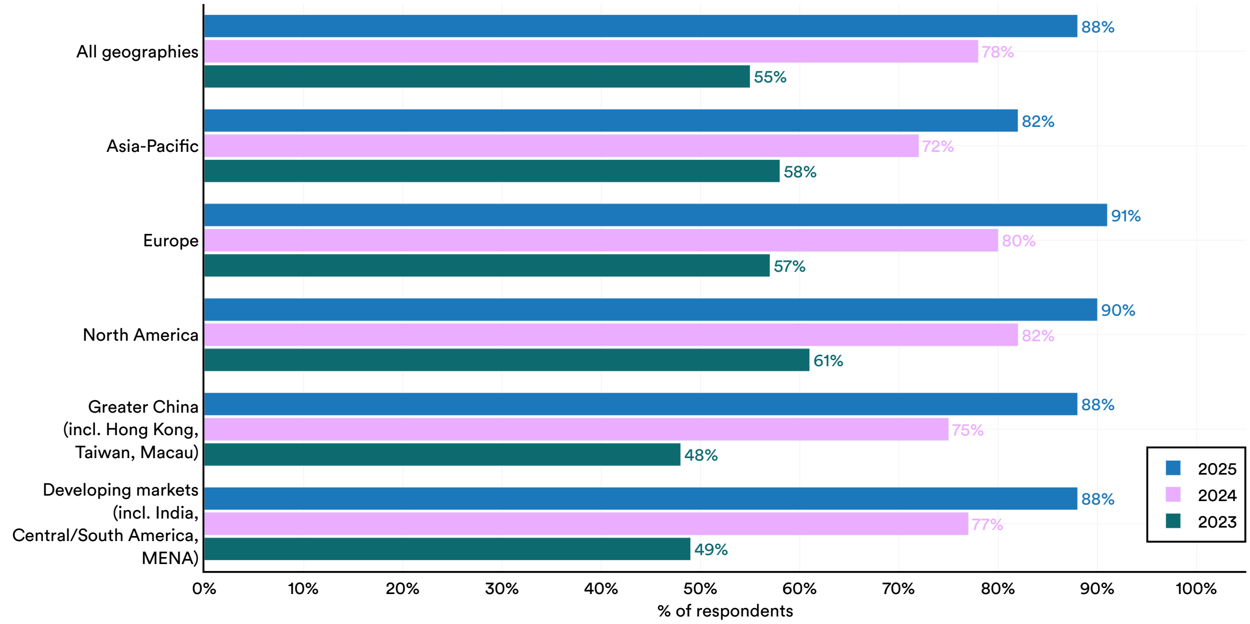

Figure 4.3.2

## By Industry and Function

Adoption patterns varied across industry and function, with some industry/function pairings showing higher rates of diffusion than others (Figure 4.3.3). The highest reported AI usage was in knowledge management for business, legal, and professional services (58%) and in software engineering and IT in the technology sector (58% and 56%, respectively). This was closely followed by marketing and sales for consumer goods and retail (51%). More broadly, functions tied to information processing, software, customer engagement, and internal knowledge work reported higher adoption than areas such as strategy and corporate finance and risk and compliance, where uptake remains low across most sectors. Financial services were an exception; they reported high use in risk and compliance functions, which are more central to their core operations.

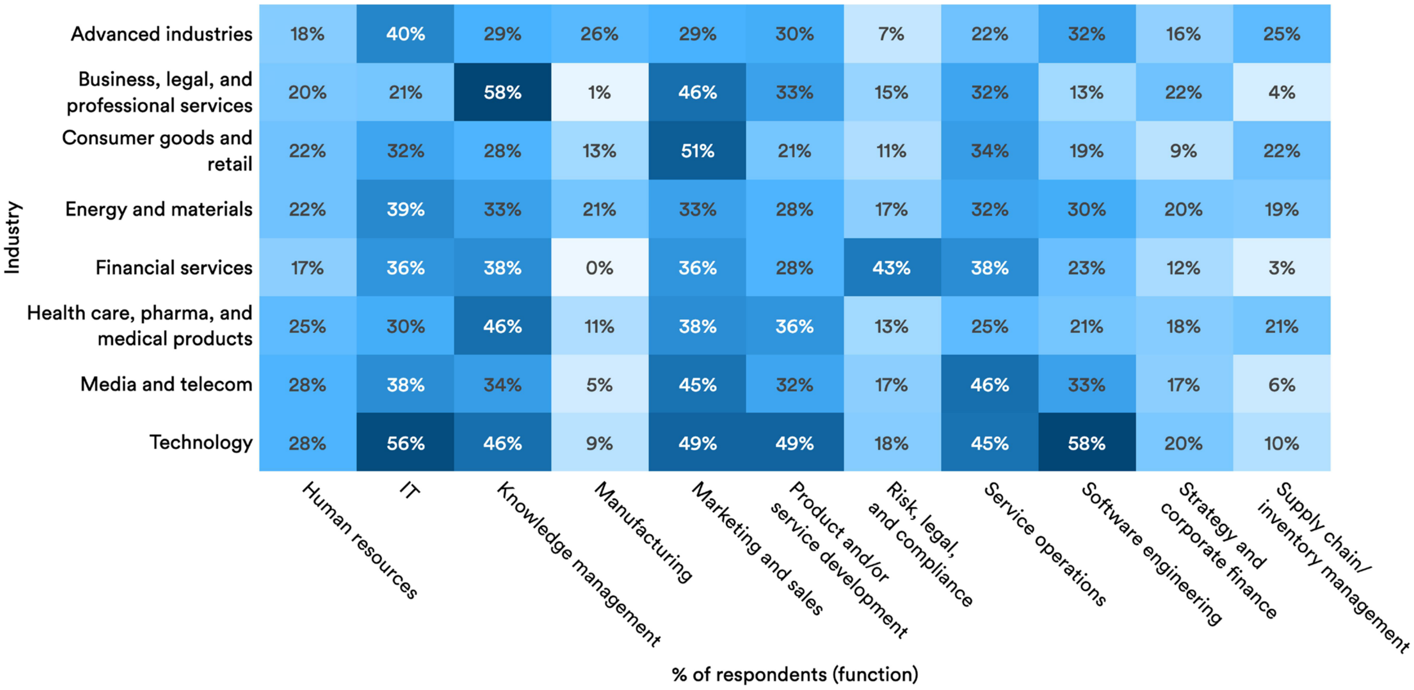

**AI use by industry and function, 2025**

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

| Industry | Human resources | IT | Knowledge management | Manufacturing | Marketing and sales | Product and/or service development | Risk, legal, and compliance | Service operations | Software engineering | Strategy and corporate finance | Supply chain/ inventory management |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Advanced industries | 18% | 40% | 29% | 26% | 29% | 30% | 7% | 22% | 32% | 16% | 25% |
| Business, legal, and professional services | 20% | 21% | 58% | 1% | 46% | 33% | 15% | 32% | 13% | 22% | 4% |
| Consumer goods and retail | 22% | 32% | 28% | 13% | 51% | 21% | 11% | 34% | 19% | 9% | 22% |
| Energy and materials | 22% | 39% | 33% | 21% | 33% | 28% | 17% | 32% | 30% | 20% | 19% |
| Financial services | 17% | 36% | 38% | 0% | 36% | 28% | 43% | 38% | 23% | 12% | 3% |
| Health care, pharma, and medical products | 25% | 30% | 46% | 11% | 38% | 36% | 13% | 25% | 21% | 18% | 21% |
| Media and telecom | 28% | 38% | 34% | 5% | 45% | 32% | 17% | 46% | 33% | 17% | 6% |
| Technology | 28% | 56% | 46% | 9% | 49% | 49% | 18% | 45% | 58% | 20% | 10% |

Figure 4.3.3[1]

Respondents more often associated AI with the highest cost savings in software engineering and manufacturing functions (56%), while revenue gains were cited with marketing and sales (67%), strategy and corporate finance (65%), and product and/or service development (62%) (Figure 4.3.4). Across broader organizational outcomes, 64% of respondents reported that AI usage had improved innovation, and 45% reported improvements in employee and customer satisfaction (Figure 4.3.5). Often, the number of respondents who believed AI usage had improved various organizational measures was similar to the number who did not believe it had any effect. Overall, negative effects were reported less frequently, with no more than 7% believing AI usage had worsened cost metrics.

---

[1] "Advanced industries" comprises respondents from sectors such as advanced electronics, aerospace and defense, automotive and assembly, and semiconductors. "Energy and materials" encompasses respondents from agriculture, chemicals, electric power and natural gas, metals and mining, oil and gas, as well as paper, forest products, and packaging.

## Cost decrease and revenue increase from analytical AI use by function, 2025

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

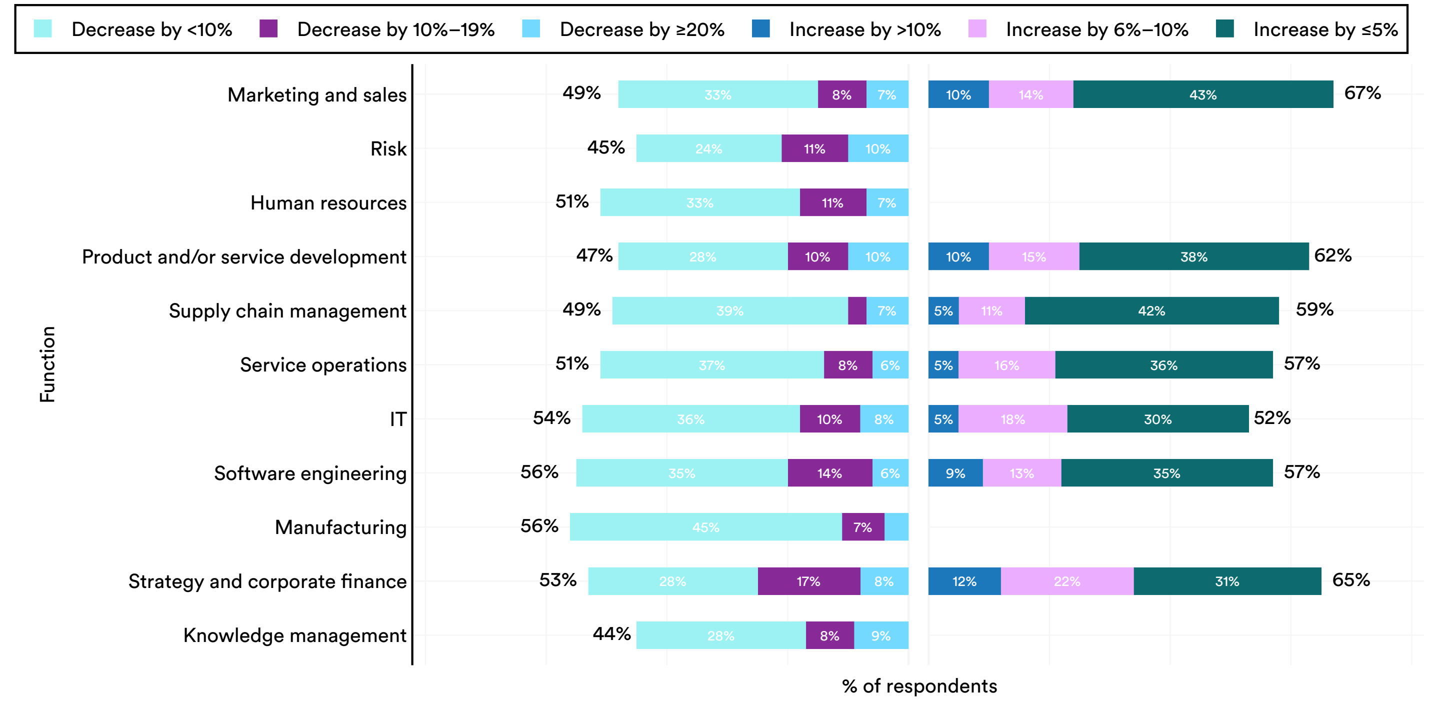

Figure 4.3.4

## AI impact on organizational measures over the past year, 2025

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

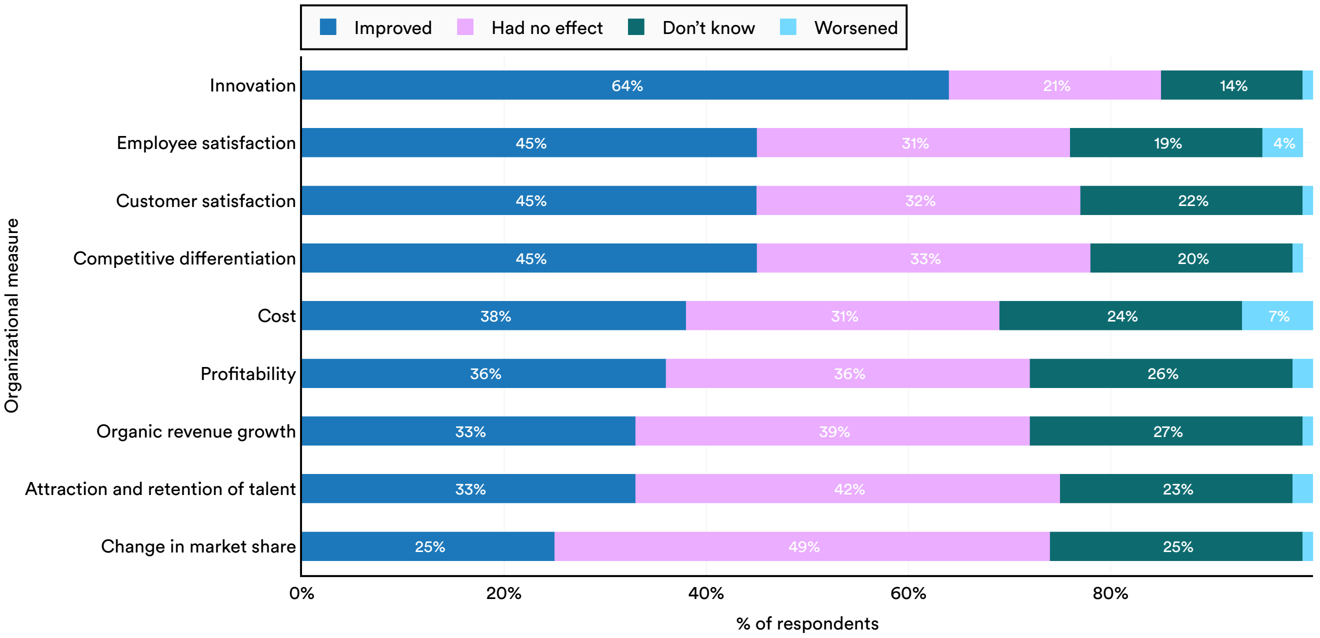

Figure 4.3.5[2]

---

[2] This includes only those respondents whose organizations regularly use AI in at least one business function. Figures may not add up to 100% because of rounding.

## Deployment Stages

The McKinsey survey also captured how deeply AI had been integrated into an organization's operations by looking at different stages of the deployment life cycle (Figure 4.3.6). As expected, given the resource and investment demands of integration, larger companies were the most likely to report that their AI programs had reached a scaling phase.

**Stage of AI deployment by organization revenue, 2025**

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

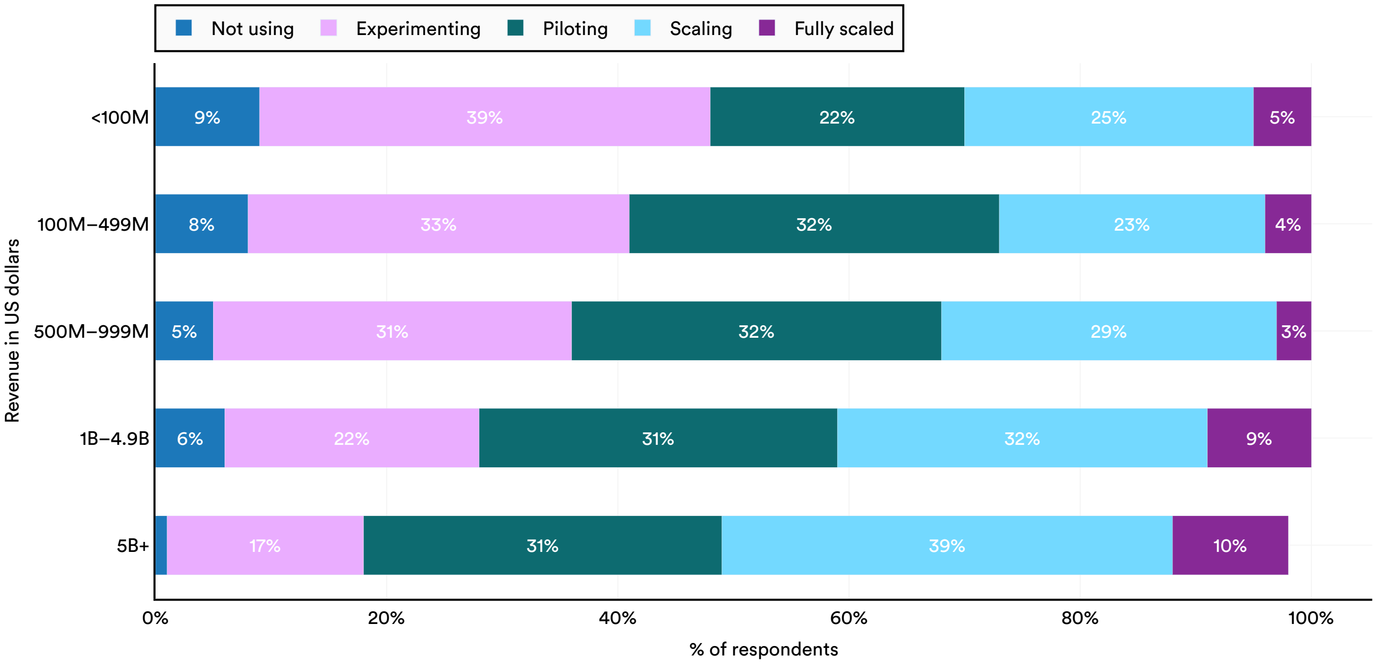

Figure 4.3.6[3]

Early indicators on AI agent adoption show that diffusion is still at an early stage. Across most business functions, a majority of respondents reported no agent use at all (Figure 4.3.7). Scaled use was in the single digits for nearly all functions. Even in functions with the most activity, including IT and knowledge management, about two-thirds or more of respondents reported no use. At the industry level, the technology sector had comparatively higher rates of scaled agent use in software engineering (24%), IT (22%), and service operations (21%) (Figure 4.3.8). The business functions reporting the highest rates of AI agent use tend to be the same as those with broader, more established AI adoption.

---

[3] Figures may not add up to 100% because of rounding; respondents who said "I don't know" were not shown but represent <1% of the total, which could also cause bars to not add up to 100%.

## AI agents in scaled use by industry and function, 2025

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

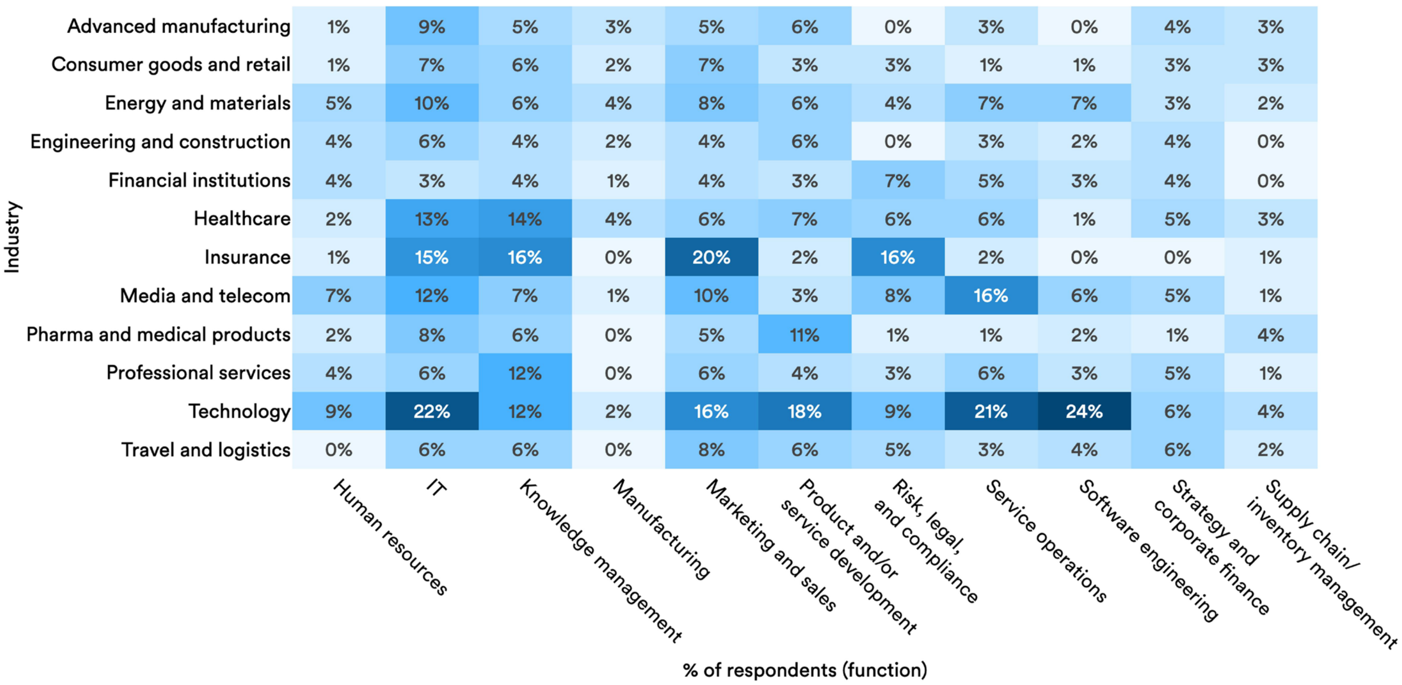

| Industry | Human resources | IT | Knowledge management | Manufacturing | Marketing and sales | Product and/or service development | Risk, legal, and compliance | Service operations | Software engineering | Strategy and corporate finance | Supply chain/ inventory management |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Advanced manufacturing | 1% | 9% | 5% | 3% | 5% | 6% | 0% | 3% | 0% | 4% | 3% |
| Consumer goods and retail | 1% | 7% | 6% | 2% | 7% | 3% | 3% | 1% | 1% | 3% | 3% |
| Energy and materials | 5% | 10% | 6% | 4% | 8% | 6% | 4% | 7% | 7% | 3% | 2% |
| Engineering and construction | 4% | 6% | 4% | 2% | 4% | 6% | 0% | 3% | 2% | 4% | 0% |
| Financial institutions | 4% | 3% | 4% | 1% | 4% | 3% | 7% | 5% | 3% | 4% | 0% |
| Healthcare | 2% | 13% | 14% | 4% | 6% | 7% | 6% | 6% | 1% | 5% | 3% |
| Insurance | 1% | 15% | 16% | 0% | 20% | 2% | 16% | 2% | 0% | 0% | 1% |
| Media and telecom | 7% | 12% | 7% | 1% | 10% | 3% | 8% | 16% | 6% | 5% | 1% |
| Pharma and medical products | 2% | 8% | 6% | 0% | 5% | 11% | 1% | 1% | 2% | 1% | 4% |
| Professional services | 4% | 6% | 12% | 0% | 6% | 4% | 3% | 6% | 3% | 5% | 1% |
| Technology | 9% | 22% | 12% | 2% | 16% | 18% | 9% | 21% | 24% | 6% | 4% |
| Travel and logistics | 0% | 6% | 6% | 0% | 8% | 6% | 5% | 3% | 4% | 6% | 2% |

Figure 4.3.7

## Stage of AI agent use by business function, 2025

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

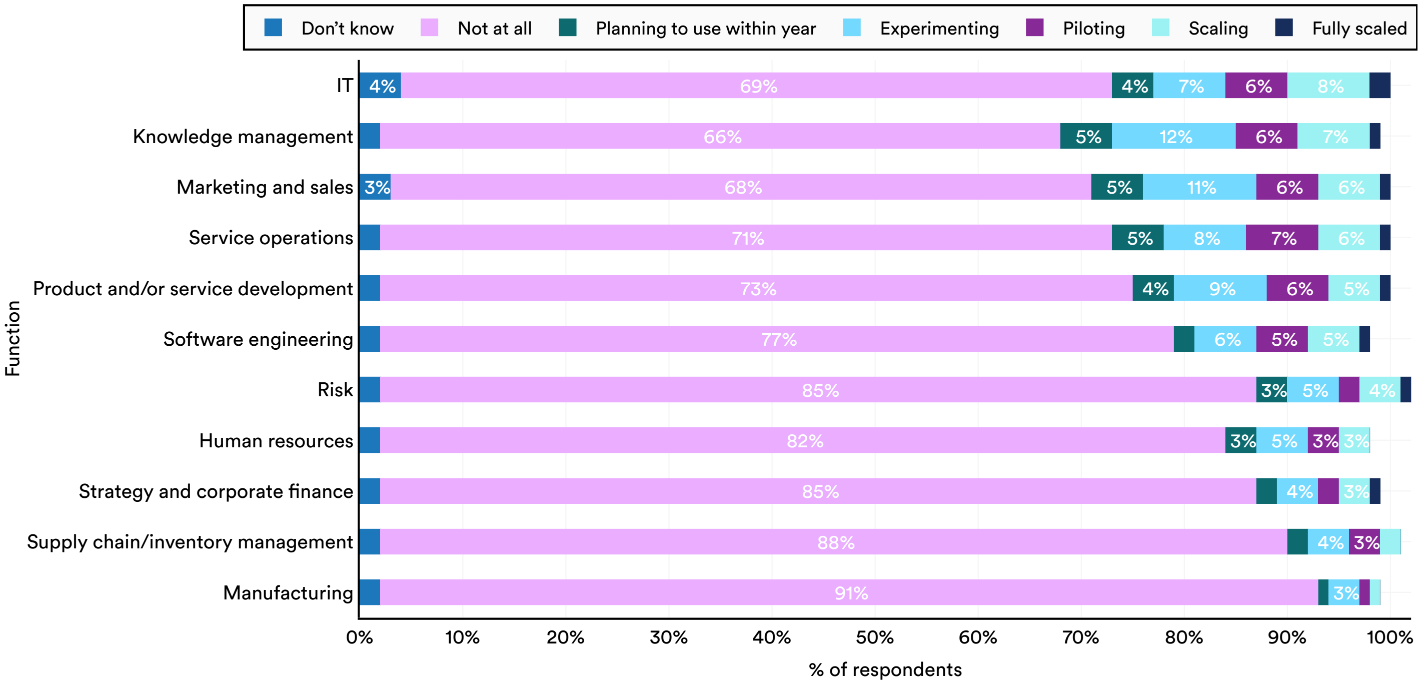

Figure 4.3.8[4]

---

[4] Figures may not add up to 100% because of rounding.

**HIGHLIGHT:**

## Measuring Signals of AI Diffusion

The full economic impact of AI is difficult to assess through investment patterns or organizational adoption alone. The assessment also requires tracking diffusion, or how widely AI tools are being adopted across populations, countries, occupations, and everyday tasks. This section brings together several complementary signals of AI diffusion, including population-level survey estimates, cross-country comparisons, historical adoption benchmarks, and platform-level usage data. Once combined, these measures offer a comprehensive view of how quickly AI is being integrated into work and daily life.

Compared with earlier transformative technologies, generative AI's adoption has been rapid in the years after its mass market introduction (Figure 4.3.9). Measured from the release of each technology's first widely available product, generative AI reached approximately 53% adoption within three years, well above the initial trajectories of the personal computer and the internet over comparable time frames (Bick et al., 2024). The sharp uptake is also reflected in the revenue trajectories of leading AI companies, as seen in the company-level revenue analysis in section 4.2, where commercial scale was reached in comparably shorter time frames (Figure 4.3.9).

**Speed of AI adoption by technology**

Source: The Project on Workforce at Harvard, 2025 | Chart: 2026 AI Index report

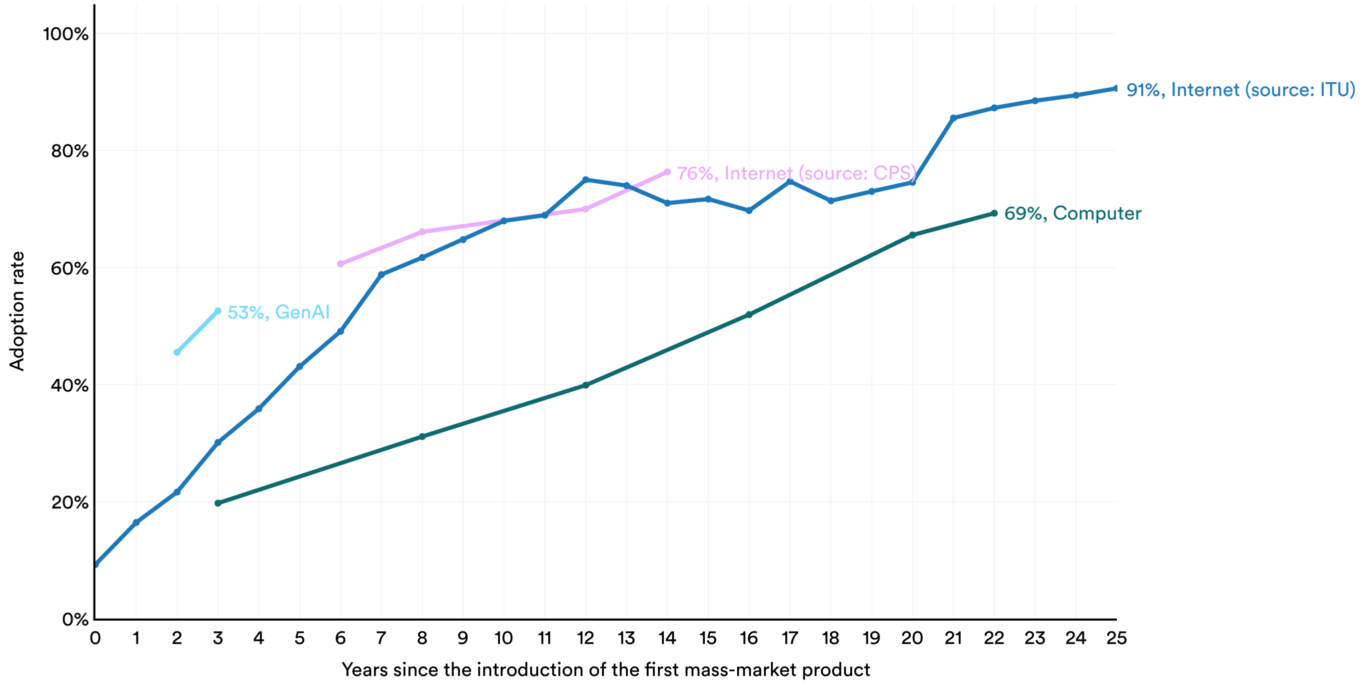

**Figure 4.3.9**[5]

---

[5] Source: https://www.genaiadoptiontracker.com. The figure shows overall usage rates for three technologies: generative AI, computers, and the internet. The horizontal axis represents years since the introduction of the first mass-market product for each technology. We use 1981 as the introduction year for computers, which was the year the IBM PC was released. We use 1995 as the introduction year for the internet, which was the year that the NSF decommissioned NSFNet and allowed the internet to carry commercial traffic. We use 2022 as the introduction year for generative AI, which was the year ChatGPT was released. The data source for computers is the 1984–2003 Computer and Internet Use Supplement of the CPS. We plot two estimates of internet use: the 2001–2009 Computer and Internet Use Supplement of the CPS and the ITU. The sample for the RPS and CPS is all individuals ages 18–64. The sample for the ITU is individuals of all ages. We pool RPS waves by year.

**HIGHLIGHT:**

Another broad signal comes from survey-based estimates of AI usage across countries (Figure 4.3.10). Adoption varies widely, and shows a strong, statistically significant positive correlation with GDP per capita (Misra et al., 2025) (Figure 4.3.11). Most high-income economies cluster between 25% and 45% adoption, with European and North American averages reaching approximately 27% and 22%, respectively. Lower usage is reported in South Asia and sub-Saharan Africa, where GDP per capita is also lower. However, there are exceptions to the relationship between GDP and AI adoption. The United Arab Emirates and Singapore report adoption levels above 54% and 61%, respectively, well above what their GDP per capita would predict. Some wealthy economies, such as the United States and Denmark, fall below the trend.

**AI diffusion by geographic area, second half 2025**

Source: Microsoft AI Economy Institute, 2025 | Chart: 2026 AI Index report

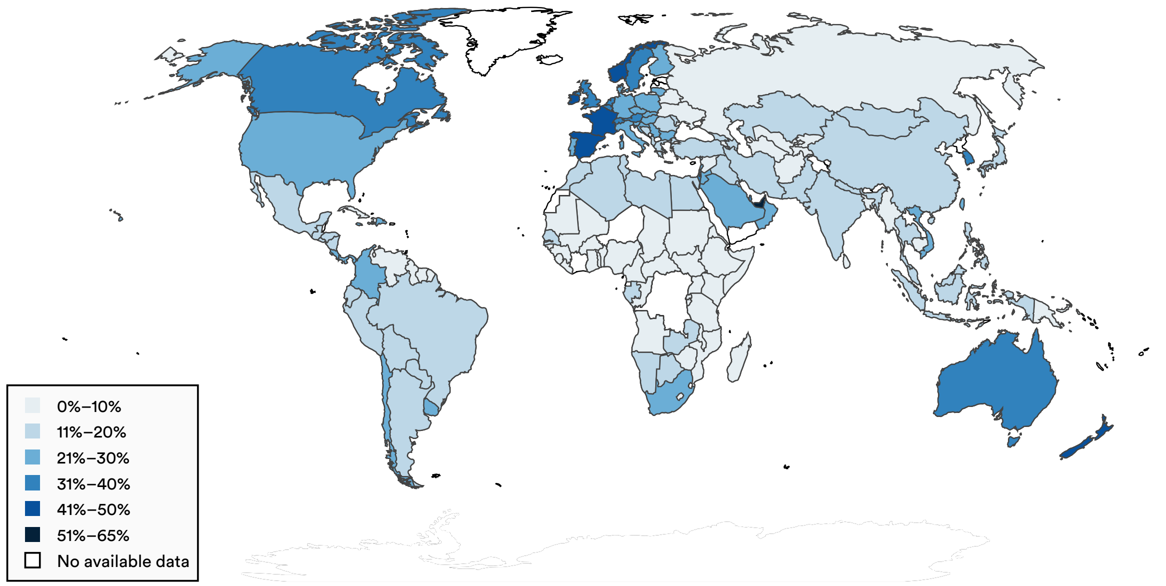

Figure 4.3.10

**AI diffusion relative to GDP per capita by geographic area, 2025**

Source: Microsoft AI Economy Institute, 2025; World Bank, 2025 | Chart: 2026 AI Index report

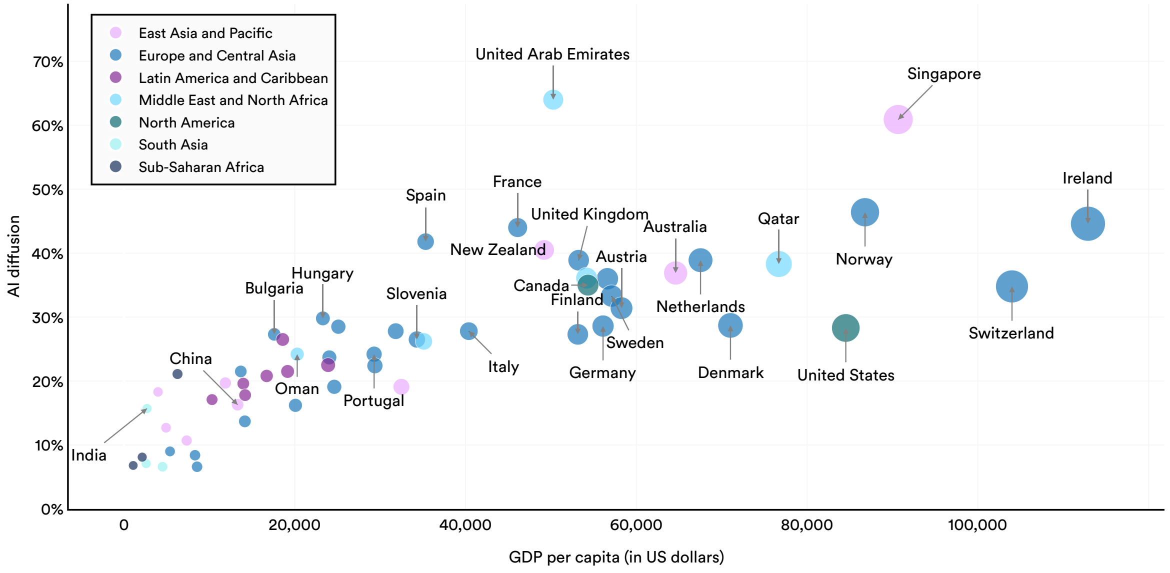

Figure 4.3.11

**HIGHLIGHT:**

Between the first and second half of 2025, AI adoption grew across the majority of the top 30 economies (Figure 4.3.12). South Korea posted the largest gain of 4.8%, climbing the rankings from 25th to 18th. The United States, despite its leading position in AI investment and model development, dropped to 24th place with a population-level adoption rate of 28.3%. Even as usage grows, the United States remains in the lower half of the global adoption ranking, in line with the more cautious public mood toward AI explored in Chapter 9.

**AI diffusion by top 30 geographic areas, first vs. second half 2025**

Source: Microsoft AI Economy Institute, 2025 | Chart: 2026 AI Index report

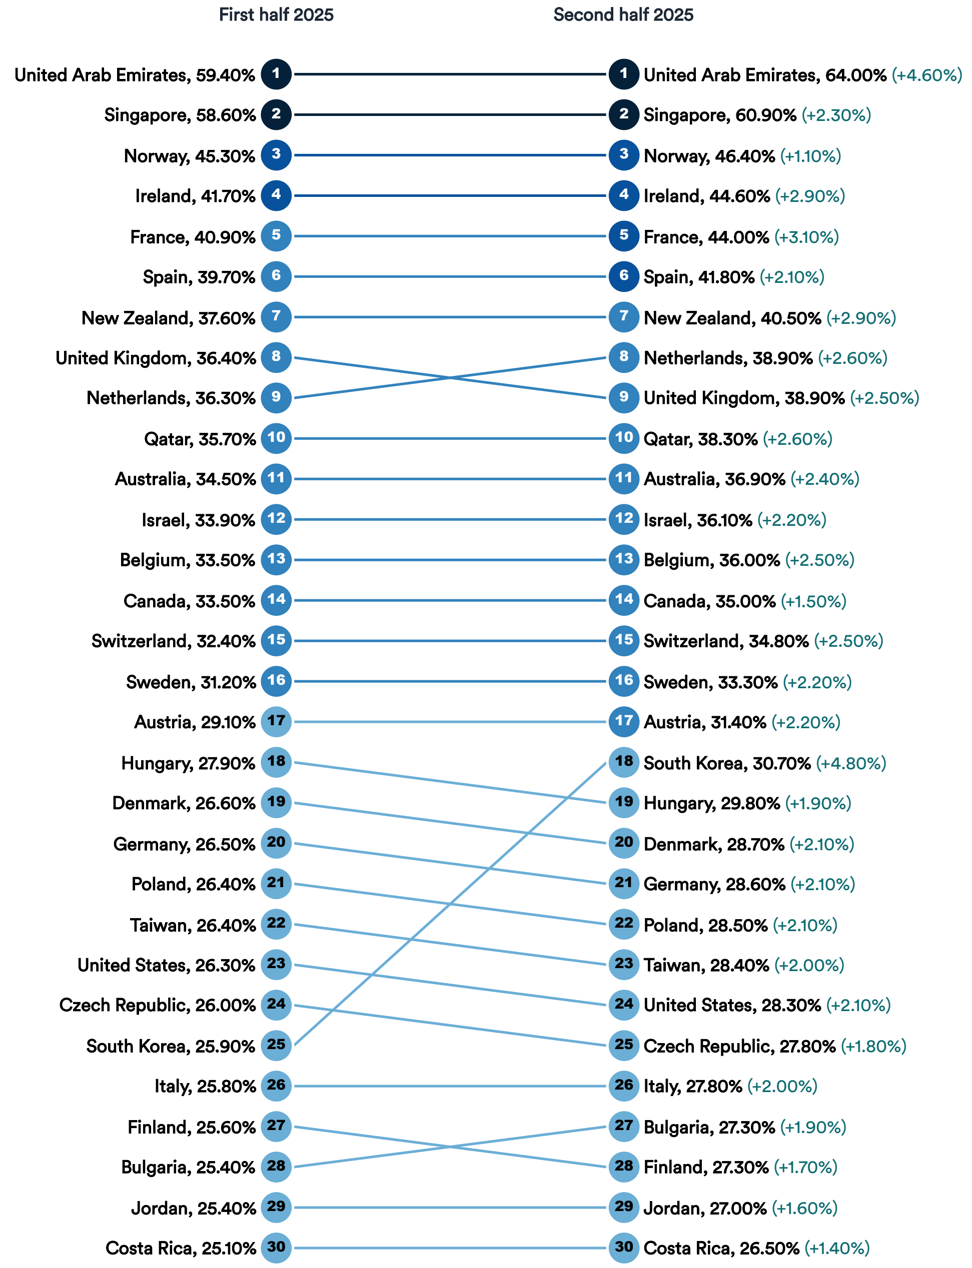

Figure 4.3.12

**HIGHLIGHT:**

On a more granular level, platform-level data from Anthropic's AI Usage Index[6] provides a view of adoption across occupations and tasks (Massenkof et. al, 2026). Throughout 2025, computer and mathematical tasks accounted for the largest share of overall usage, consistently representing close to 40% of activity (Figure 4.3.13). Educational instruction and library tasks showed the most significant growth, rising from 9% early in the year to approximately 14% by late 2025. This growth in educational settings is worth noting alongside Chapter 7, which explores how institutional guidance and readiness still lag behind adoption. In addition, an analysis of the conversation patterns reveals a shift in how users interact with the tools (Figure 4.3.14). The share of automation-oriented conversations, where users instruct the tool to complete a task autonomously, rose from 41% at the start of 2025 to 49% in August. This surpassed augmentation-style interactions for the first time. However, by November, augmentation had moved ahead, representing 52% of the share of conversations. The fluctuation over the course of the year suggests that automation-oriented use is becoming more prevalent, which is consistent with the early-stage AI agent adoption patterns for organizations, described in the previous section (Figure 4.3.8).

**Task usage share by occupation group, V1–V4 2025**

Source: Anthropic Economic Index, 2026 | Chart: 2026 AI Index report

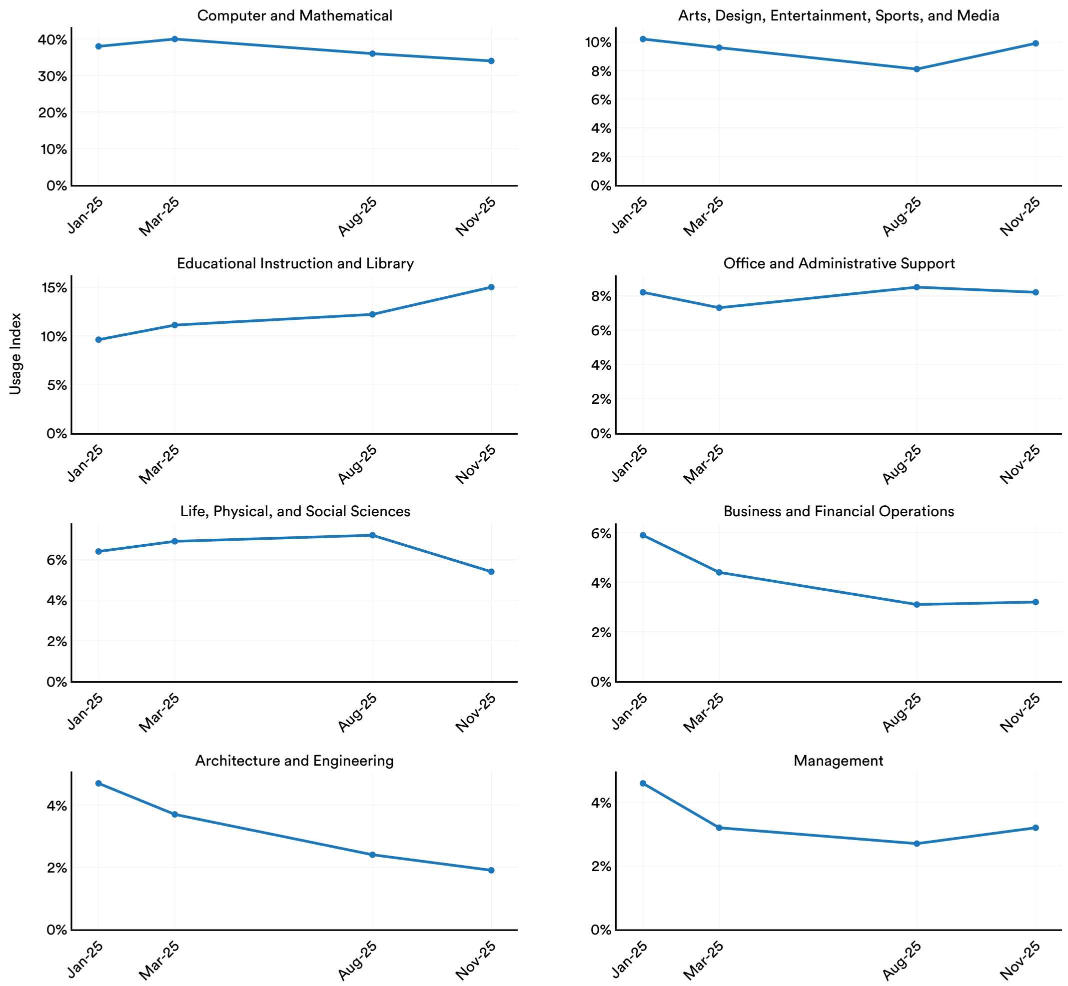

Figure 4.3.13[7]

---

[6] The Anthropic AI Usage Index (AUI) measures Claude usage relative to the working-age population by calculating each geography's share of Claude usage divided by its share of the working-age population (ages 15–64). Countries with an AUI greater than 1 use Claude more often than expected based on their working-age population alone, while those with an AUI less than 1 use it less.

[7] V1–V4 refer to the four successive releases of the Anthropic Economic Index in 2025, corresponding to January, March, August, and November 2025, respectively.

**HIGHLIGHT:**

Claude.ai collaboration mode share, 2025

Source: Anthropic Economic Index, 2026 | Chart: 2026 AI Index report

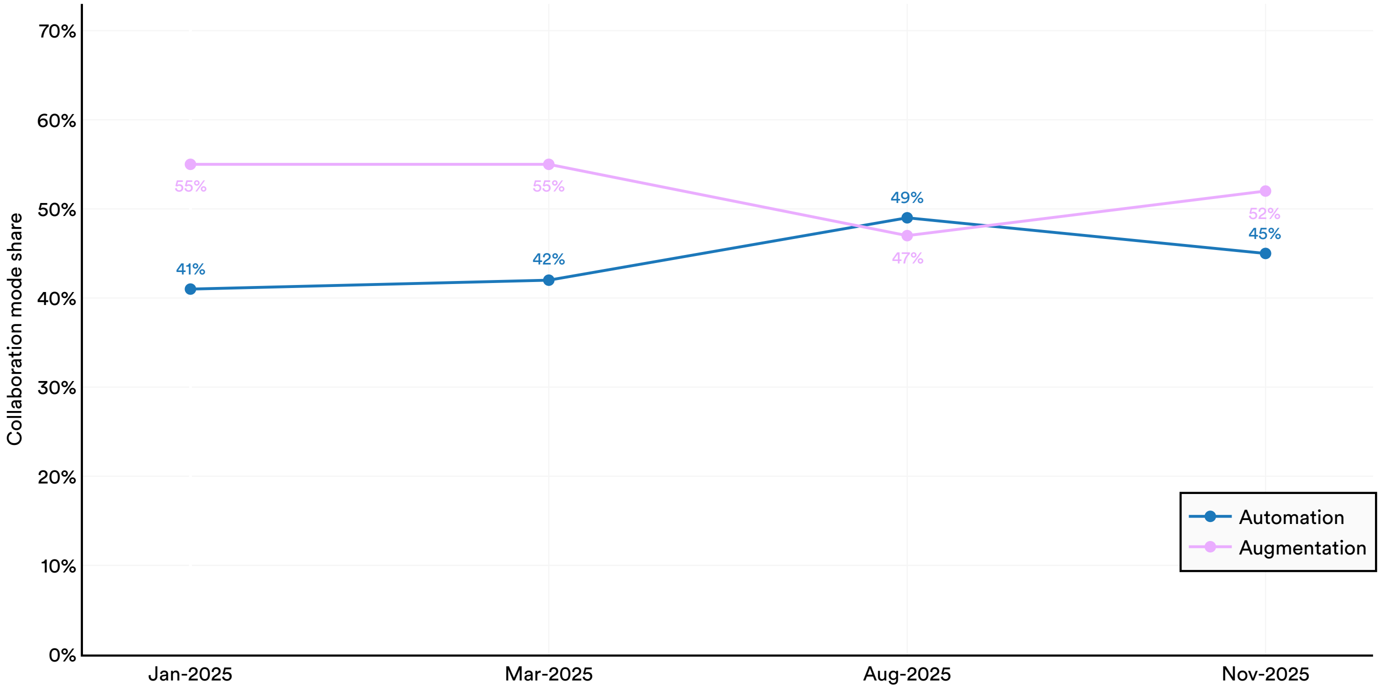

**Figure 4.3.14**

AI diffusion is also shaped by broader societal attitudes, including public trust and optimism about the technology. Chapter 9 studies these trends to determine how excitement for and exposure to AI vary across countries and what they suggest about the societal experience of increasing adoption.

# 4.4 Jobs

Labor markets provide signals of how investment, technical progress, and organizational adoption are changing workforce dynamics. This section tracks both the demand side of the labor market, through job postings and skill requirements, and the supply side, through talent flows, before examining the impact on employment outcomes and employee sentiment. The analysis draws from Lightcast's job posting database, LinkedIn's talent and hiring metrics, and recent research on AI's effects on the labor market.

## AI Labor Demand Across Geographies

Across the countries tracked by Lightcast, demand for AI-related talent continued to increase in 2025[8] as job listings that require AI skills continue to make up a growing share of overall postings (Figures 4.4.1 and 4.4.2). While most countries are hitting new peaks of demand, the intensity varies across countries. In 2025, Singapore led with 4.69% of all job postings that required AI skills, followed by Hong Kong (3.5%), Luxembourg (3.4%), and Spain (3.3%). The United States reached 2.6%, followed by Chile (2.4%) and the United Kingdom (1.9%).

**AI job postings (% of all job postings) by select geographic areas, 2014–25 (part 1)**

Source: Lightcast, 2025 | Chart: 2026 AI Index report

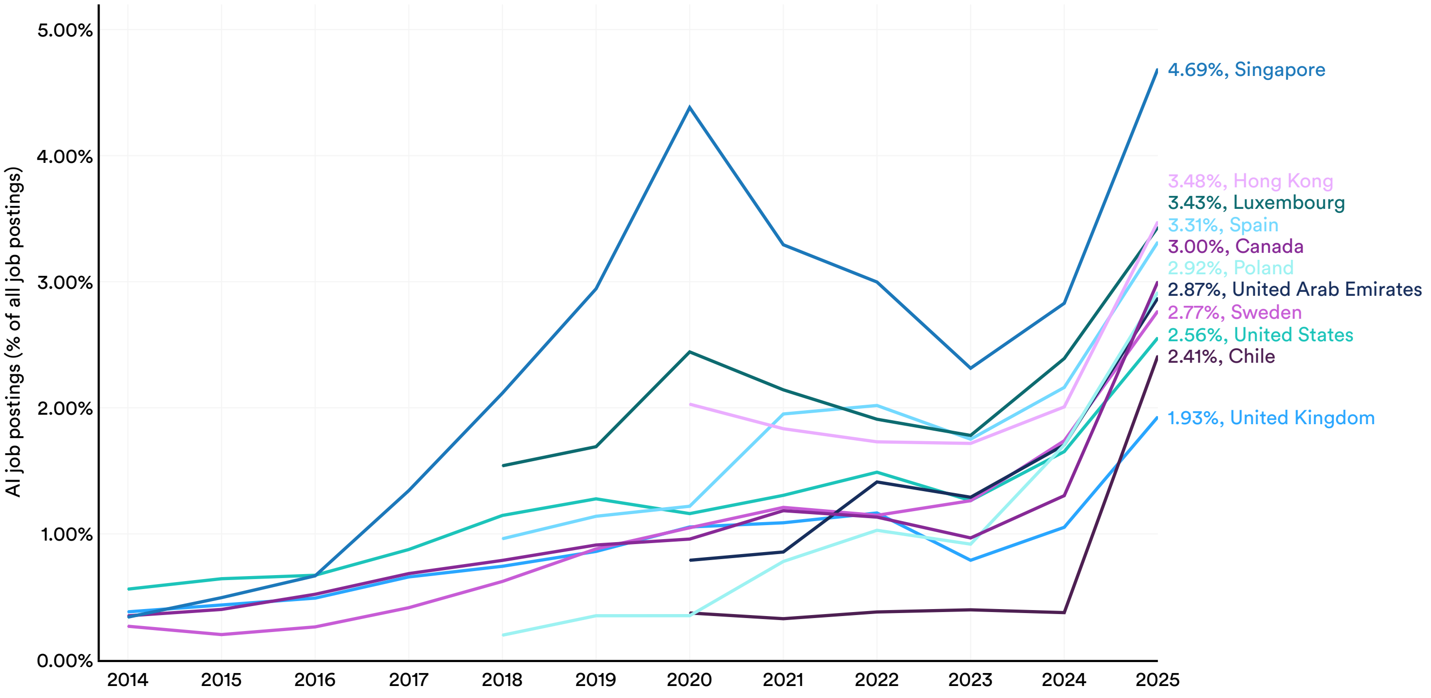

Figure 4.4.1

---

[8] Historical posting counts may differ from previously published versions due to data updates and revisions. However, year-over-year trends remain consistent with earlier analyses. See here for more details.

**AI job postings (% of all job postings) by select geographic areas, 2014–25 (part 2)**
Source: Lightcast, 2025 | Chart: 2026 AI Index report

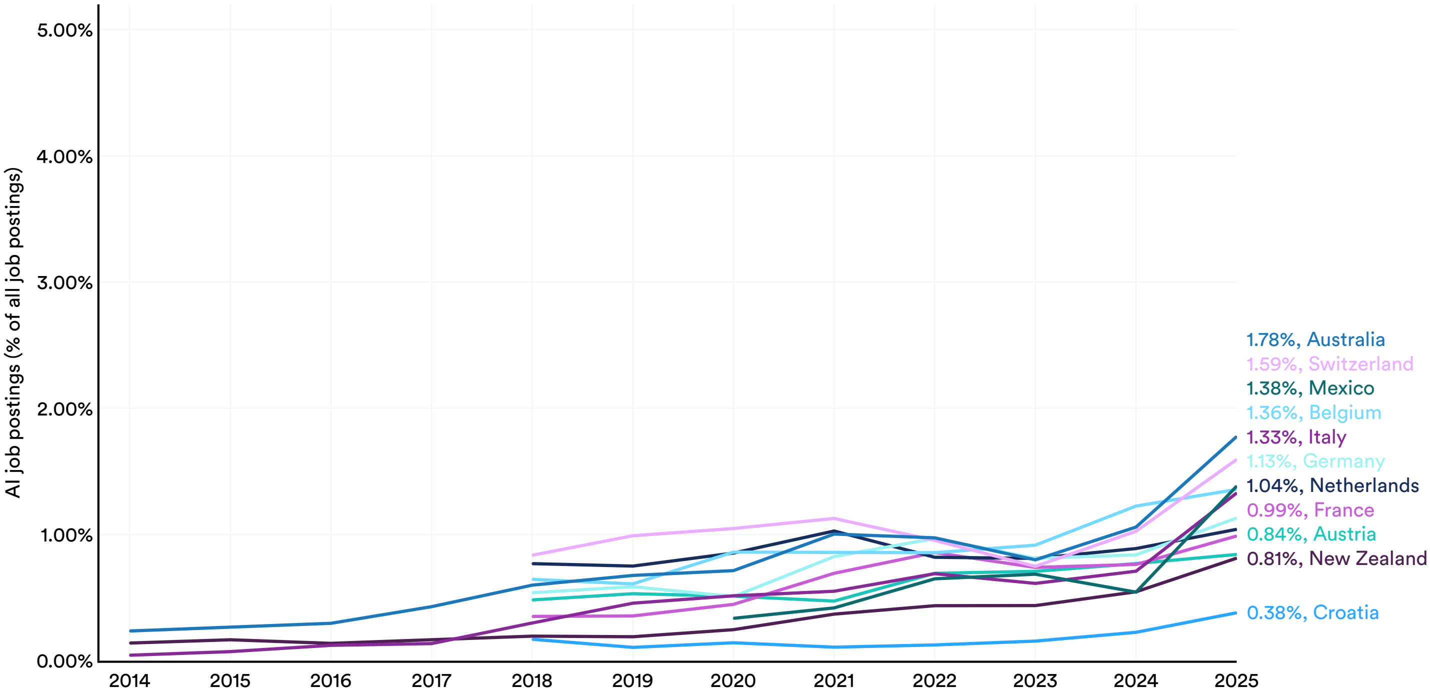

Figure 4.4.2

## AI Hiring in the United States

### Skill Composition

Within the United States, AI labor demand can be disaggregated by skillset to reveal how the workforce footprint is evolving (Figures 4.4.3–4.4.8). In 2025, broad AI and machine learning skill clusters remain the most frequently cited categories in AI job posting, accounting for 1.7% and 1.0% of all job postings. Among the top specialized skills, Python appeared the most often, in 258,674 posts, a 391% increase compared to the 2013–15 time period and a near 30% increase from 2024. The fastest growth appears in skills needed to build and operate systems at scale, with employer demand mirroring the broader investment shift toward AI infrastructure and deployment capacity. Amazon Web Services expanded significantly compared to a decade ago (+1,358%) alongside an increasing emphasis on scalability (+733%) and workflow management (+818%).

Mentions of generative AI skills in AI job postings grew 111% from 2024 to 2025, though their share of total AI job postings decreased by 5%. With overall AI labor demand rising, a newer skill cluster tied to AI agents emerged. From 2024 to 2025, postings referencing agentic AI, AI agents, or agentic systems exponentially increased. The share of AI job postings that mentioned ChatGPT, chatbot, or conversational AI declined, while posts referencing agentic terms or orchestration frameworks such as LangGraph increased. Job demand appears to be shifting from general familiarity with chat-based tools toward skills required to coordinate and operationalize task-oriented systems.

### AI job postings (% of all job postings) in the United States by skill cluster, 2010–25

Source: Lightcast, 2025 | Chart: 2026 AI Index report

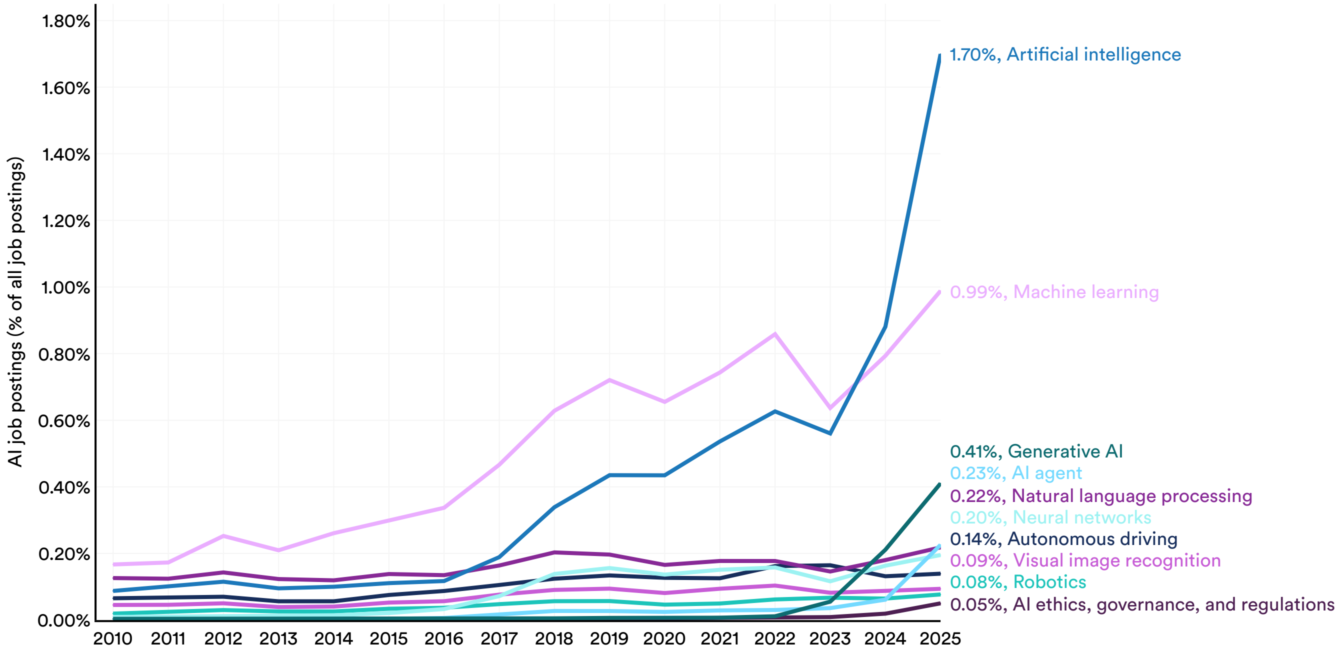

Figure 4.4.3[9]

### Top 10 specialized skills in 2025 AI job postings in the United States, 2013–15 vs. 2025

Source: Lightcast, 2025 | Chart: 2026 AI Index report

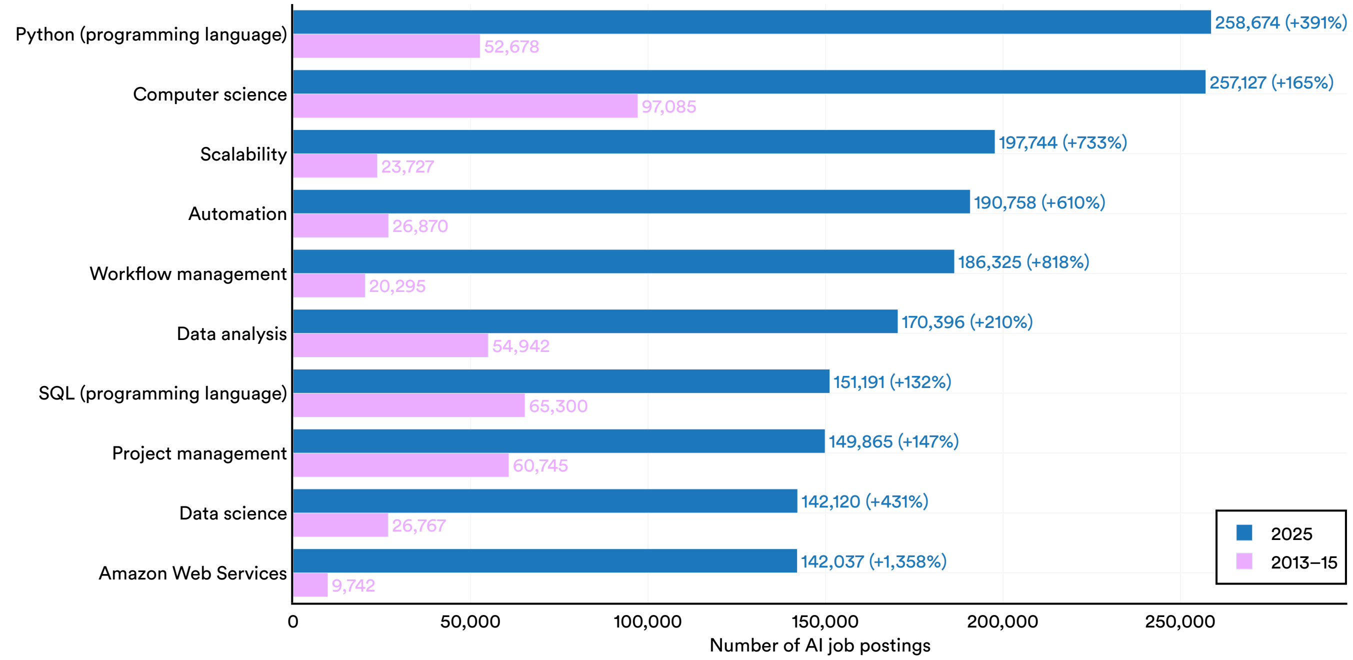

Figure 4.4.4

---

9 A single job posting can list multiple AI skills.

## Generative AI skills in AI job postings in the United States, 2024 vs. 2025

Source: Lightcast, 2025 | Chart: 2026 AI Index report

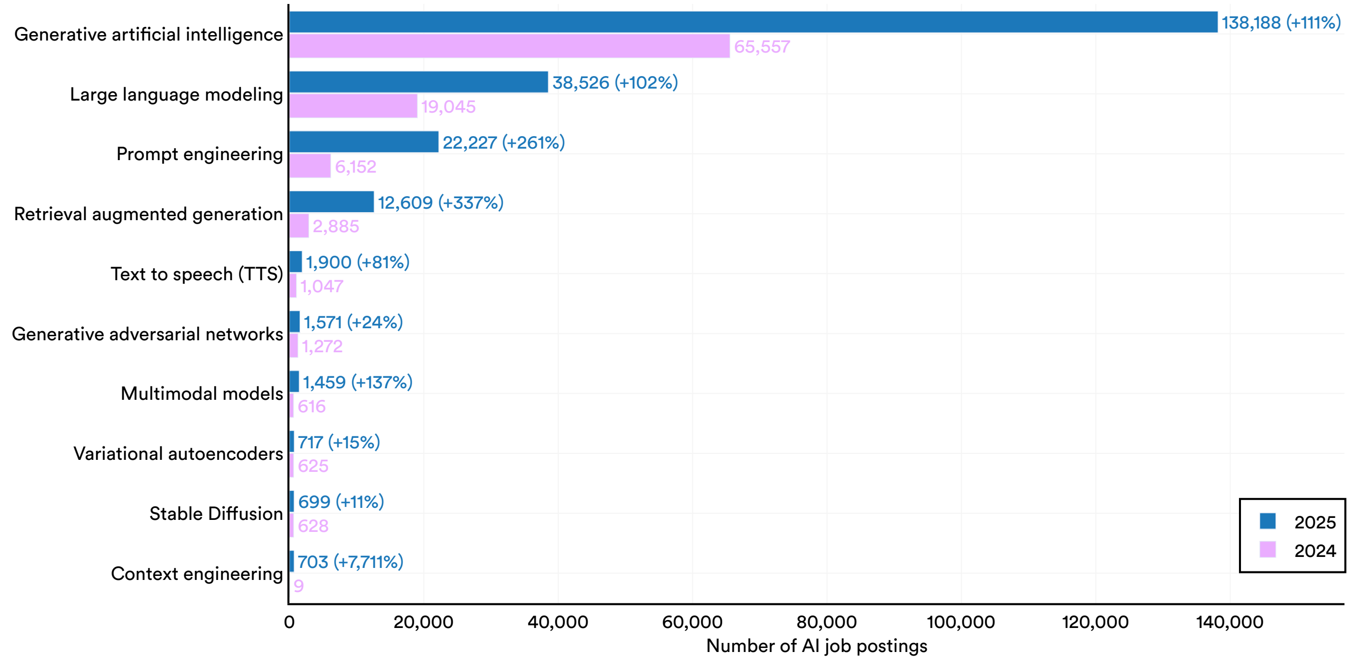

Figure 4.4.5

## Share of generative AI skills in AI job postings in the United States, 2024 vs. 2025

Source: Lightcast, 2025 | Chart: 2026 AI Index report

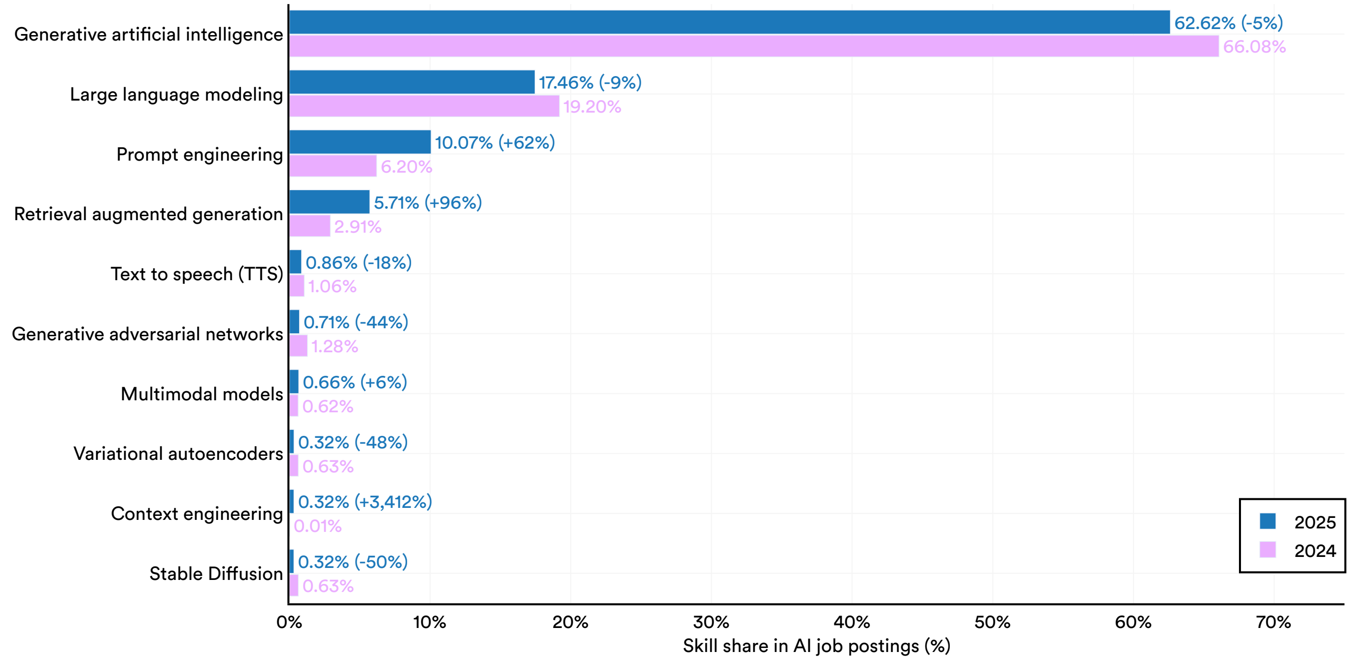

Figure 4.4.6

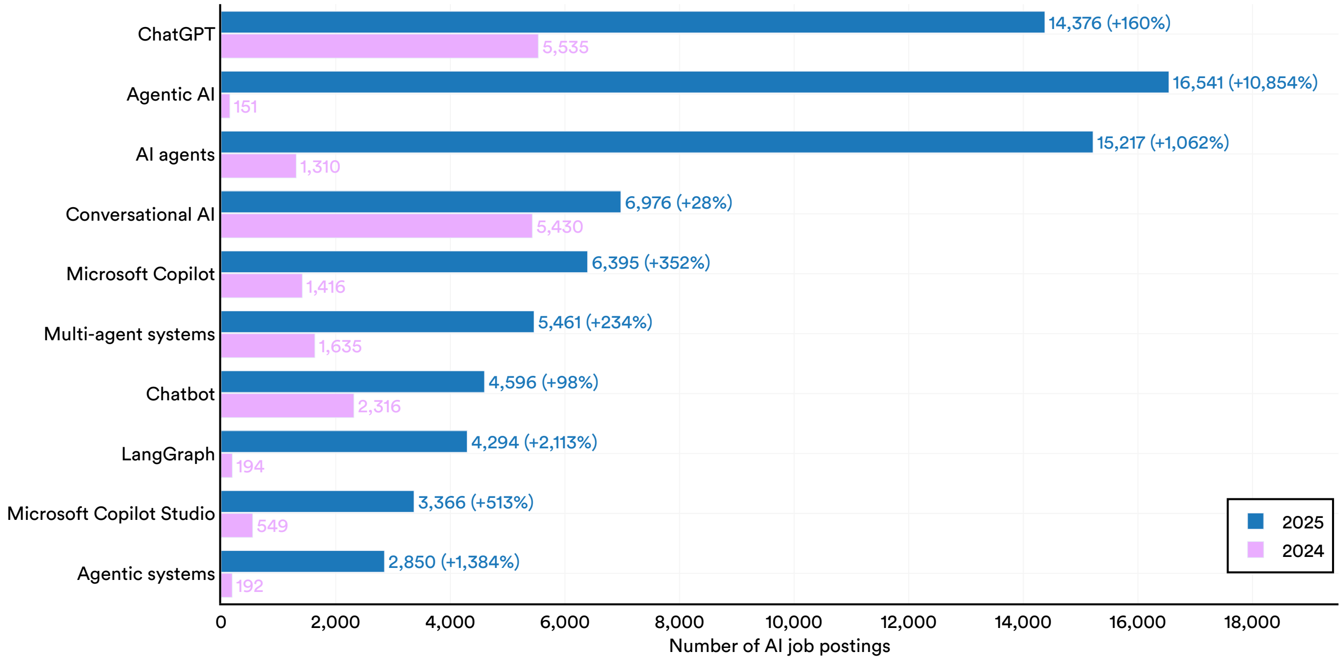

**AI agent skills in AI job postings in the United States, 2025**

Source: Lightcast, 2025 | Chart: 2026 AI Index report

Figure 4.4.7[10]

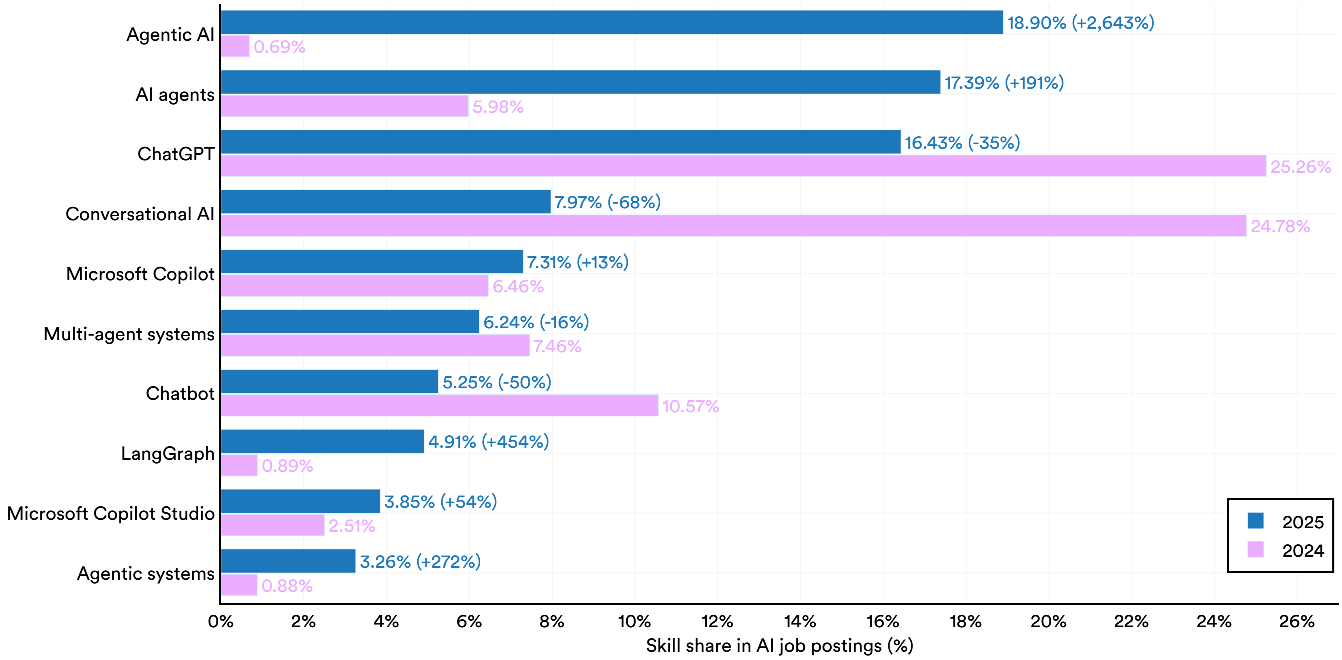

**Share of AI agent skills in AI job postings in the United States, 2024 vs. 2025**

Source: Lightcast, 2025 | Chart: 2026 AI Index report

Figure 4.4.8

---

[10] For definitions of each skill category, see the Lightcast taxonomy at https://lightcast.io/open-skills or the Appendix.

## By Sector

Demand for AI talent increased across all economic sectors in 2025, though the pace of growth varied (Figure 4.4.9). The information sector leads, with AI skills appearing in a 13.2% share of its job postings, up from 7.8% in 2024. Other sectors with relatively high AI posting shares include professional, scientific, and technical services (6.5%), finance and insurance (5.3%), and manufacturing (4.7%). In 2025, AI hiring also expanded in sectors with historically low adoption rates. Transportation and warehousing, real estate, and education showed year-over-year increases, evidence that the diffusion is reaching beyond traditional technology-driven industries.

**AI job postings (% of all job postings) in the United States by sector, 2024 vs. 2025**

Source: Lightcast, 2025 | Chart: 2026 AI Index report

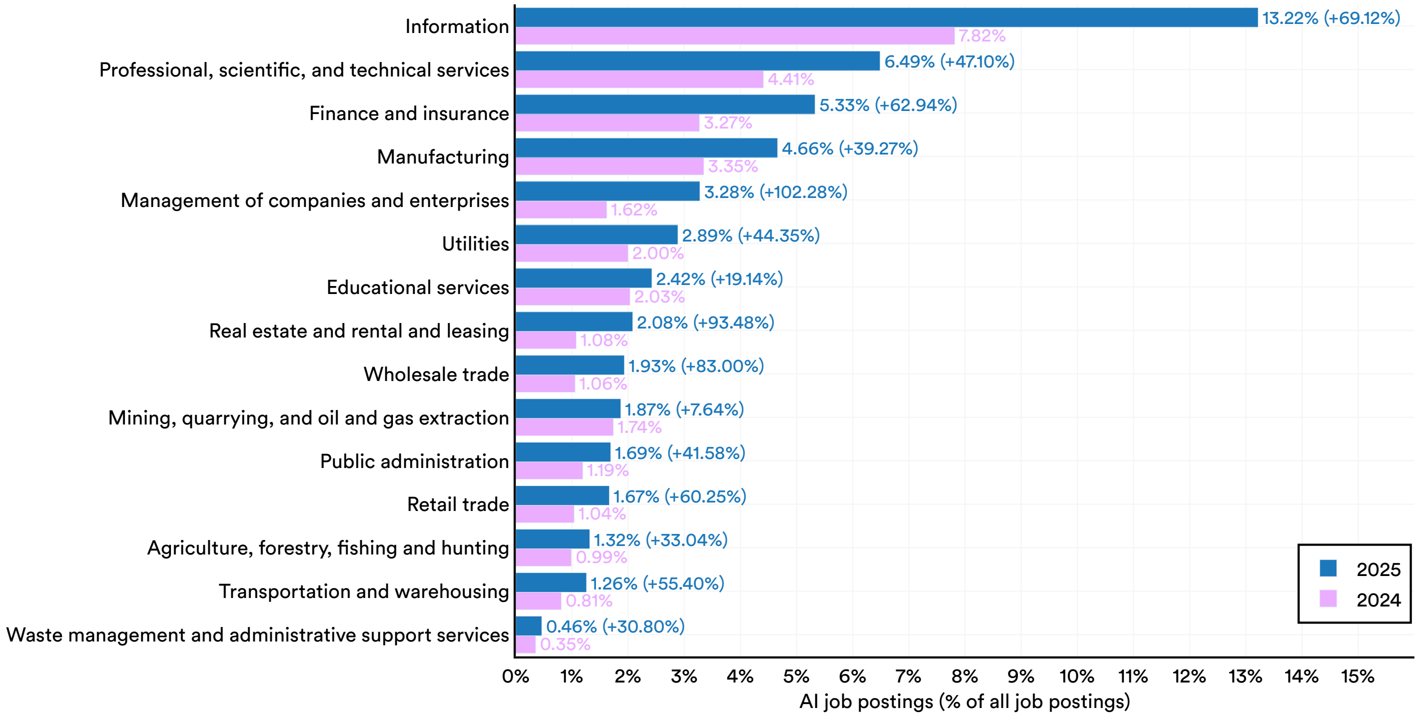

**Figure 4.4.9**[11]

---

[11] The sector classifications in Figure 4.2.9 are based on two-digit NAICS codes. For more information on the Bureau of Labor Statistics' supersector and NAICS classifications, see the following reference.

## By State

Within the United States, AI labor demand remains highly concentrated within certain states (Figures 4.4.10 and 4.4.11). California leads with 170,881 postings and accounts for a disproportionate share of total 2025 U.S. AI job postings (17.2%). Texas follows with 80,547 postings (8.1% of total) and New York with 66,029 (6.6%). These three states represent approximately a third of all AI job postings nationally. This mirrors the state-level investment data described earlier, where AI funding was also concentrated in a small number of states, led by California. Over time, however, California's share of the national total has declined, from over 25% in 2012 to 17.9% in 2025, even as the state continues to lead in AI investment concentration (Figure 4.4.12). However, looking at the density of AI labor demand within each state, there are instances of above-average AI penetration relative to the particular total job market (Figure 4.4.13). For example, despite having smaller total numbers, Washington, D.C., accounts for a comparatively high 6.2% share of those postings followed by Delaware at 4.4%. From 2024 to 2025, California, Washington state, New York, and Texas all continued to see growth in AI job postings within their labor markets (Figure 4.4.14).

Number of AI job postings in the United States by state, 2025

Source: Lightcast, 2025 | Chart: 2026 AI Index report

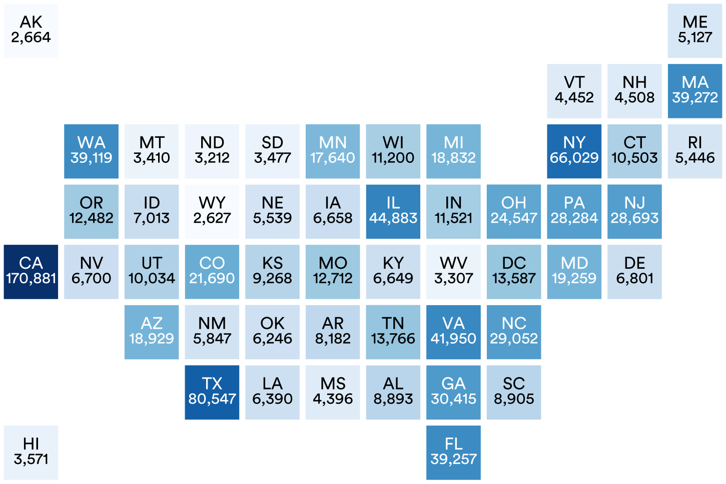

Figure 4.4.10

Percentage of US AI job postings by state, 2025

Source: Lightcast, 2025 | Chart: 2026 AI Index report

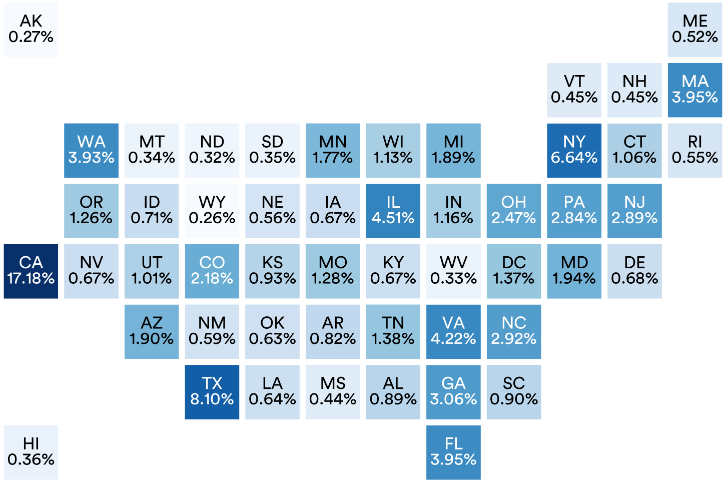

Figure 4.4.11

### Percentage of US AI job postings by select US state, 2010–25

Source: Lightcast, 2025 | Chart: 2026 AI Index report

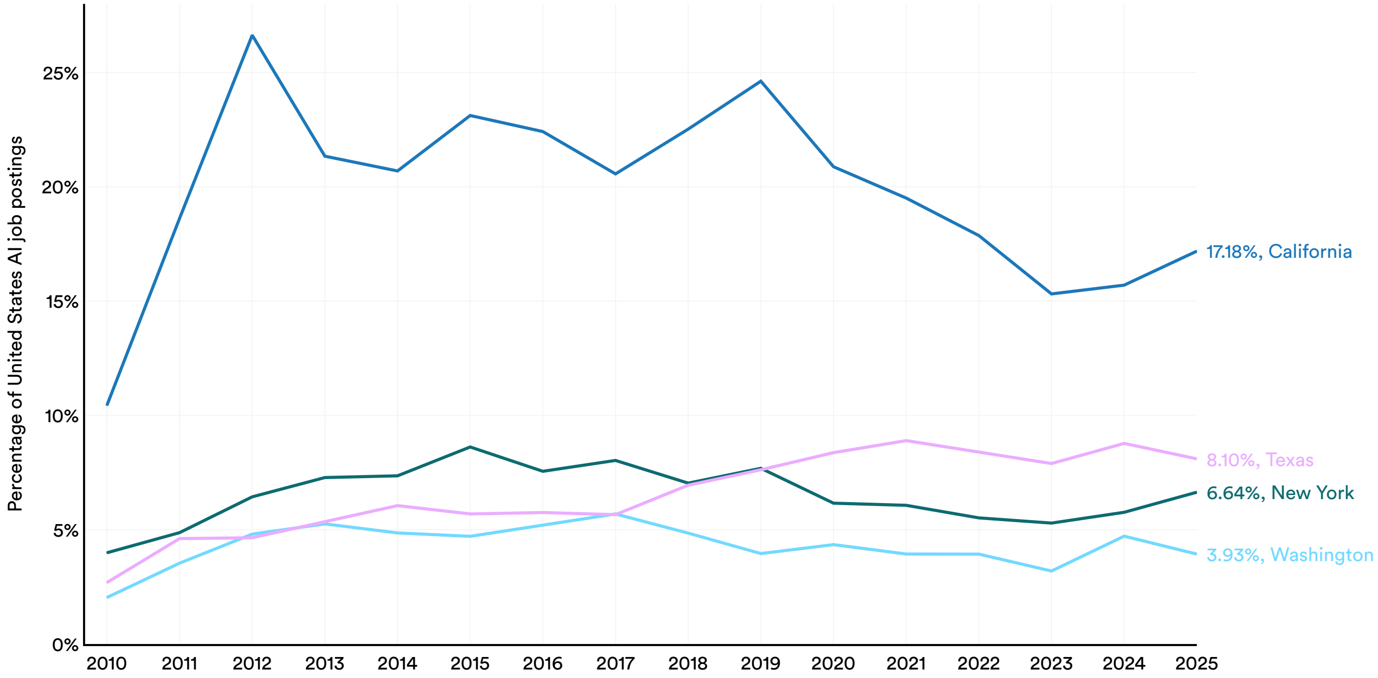

**Figure 4.4.12**

### Percentage of US states' job postings in AI, 2025

Source: Lightcast, 2025 | Chart: 2026 AI Index report

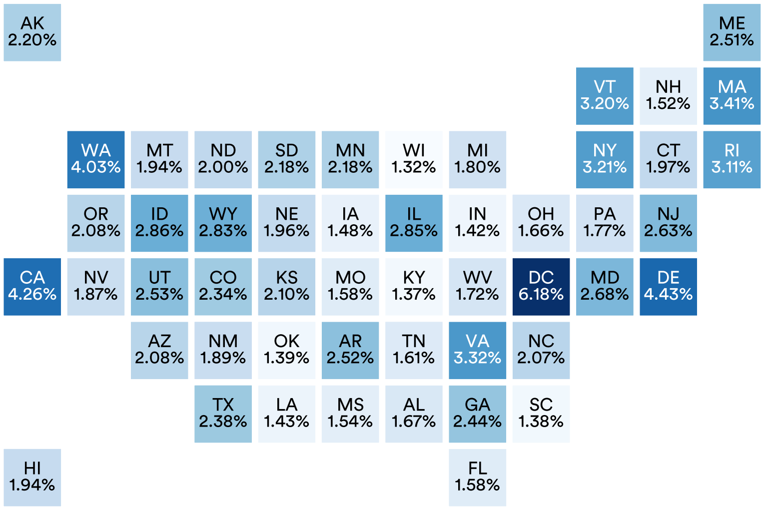

**Figure 4.4.13**

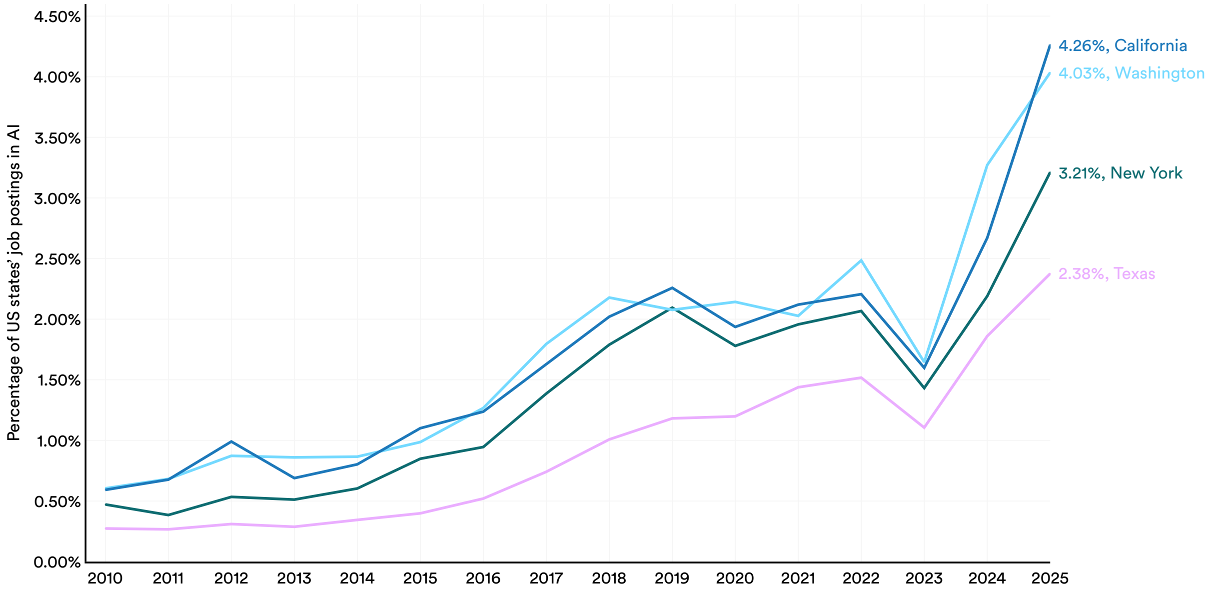

**Percentage of US states' job postings in AI by select US state, 2010–25**
Source: Lightcast, 2025 | Chart: 2026 AI Index report

**Figure 4.4.14**

## AI Hiring

LinkedIn's hiring and talent data provides a view of how AI labor demand is changing the actual workforce in practice. In most countries, AI hiring rates outpaced overall hiring growth in 2025 (Figures 4.4.15 and 4.4.16). Indonesia recorded the highest relative AI hiring growth at 31.7%, followed by Croatia (27.8%) and Belgium (21.5%). Since 2018, many countries show a sustained pattern of AI hiring rates that exceed general labor market growth. However, there are a few exceptions, such as Iceland and Sweden, where AI hiring growth lagged behind the broader market.

**AI vs. overall hiring rate growth by geographic area, 2025**
Source: LinkedIn, 2025 | Chart: 2026 AI Index report

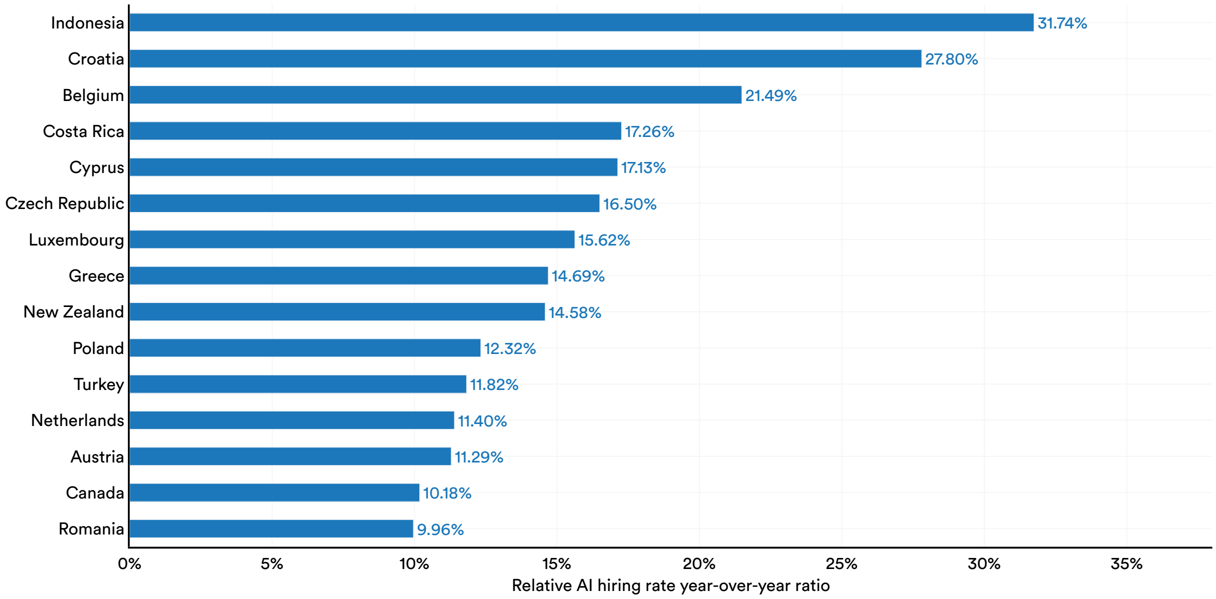

**Figure 4.4.15**[12]

---

[12] For the sake of brevity, the visualization includes only the top 15 countries for this metric.

### AI vs. overall hiring rate growth by geographic area, 2018–25

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

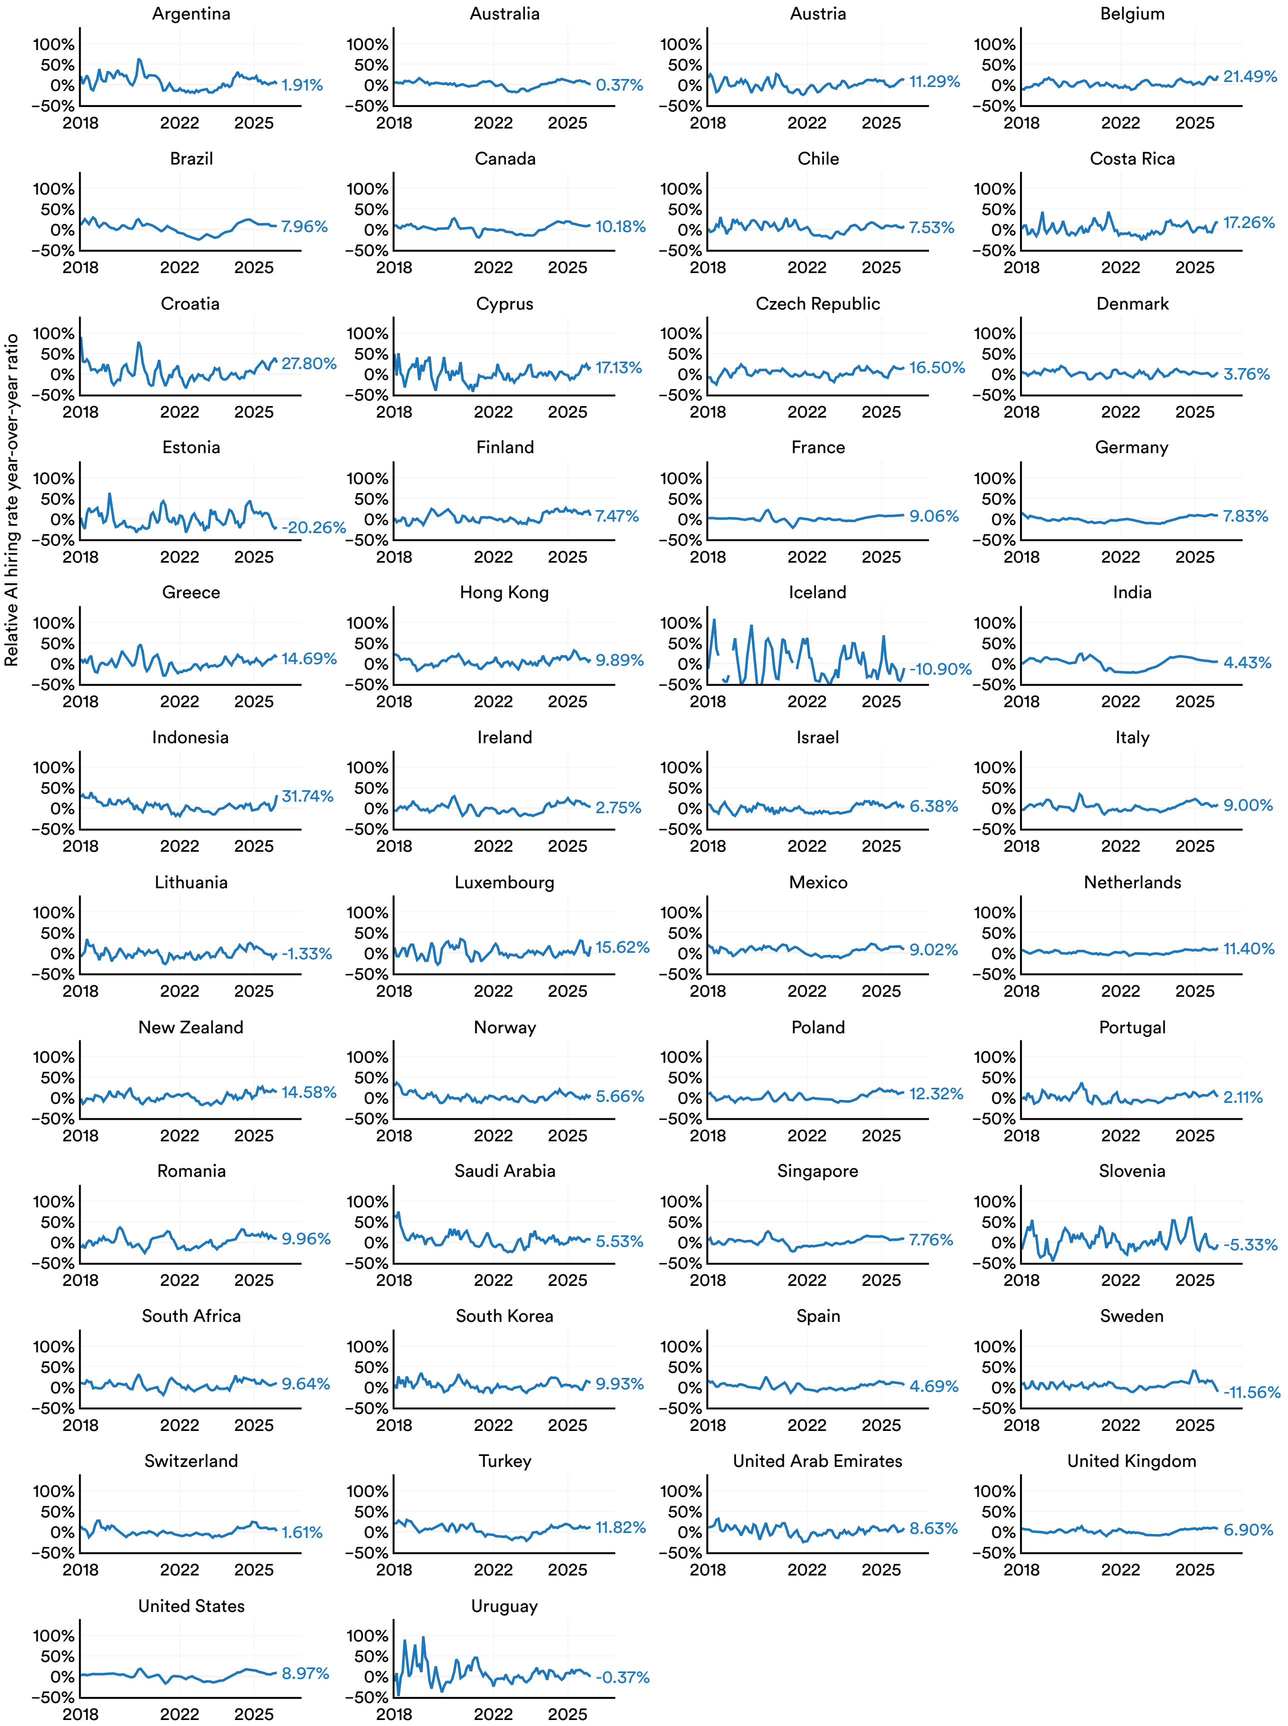

**Figure 4.4.16**

## AI Talent Concentration

Talent data helps show where AI capabilities are accumulating and how the AI workforce is distributed globally. This section reviews LinkedIn's measures on the concentration of AI talent within countries and the movement of that talent across borders. In 2025, Israel had the highest concentration of AI talent among LinkedIn members (2.1%), followed by Singapore (1.8%) and Luxembourg (1.6%) (Figures 4.4.21 and 4.4.22). However, the United Arab Emirates, India, and Saudi Arabia showed the fastest growth in their share of AI talent, each increasing over 100% between 2019 and 2025. Over the same time period, all countries in the sample grew by at least 75% in AI talent concentration.

**AI talent concentration by geographic area, 2025**

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

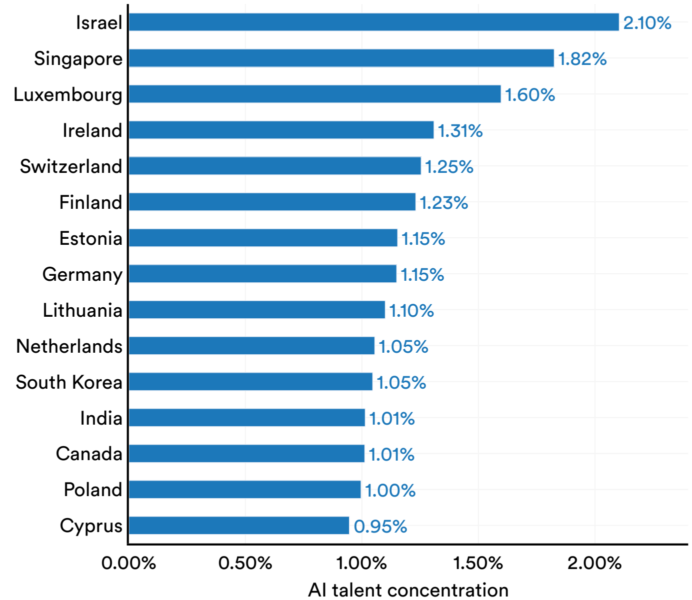

Figure 4.4.21[13]

**Percentage change in AI talent concentration by geographic area, 2019 vs. 2025**

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

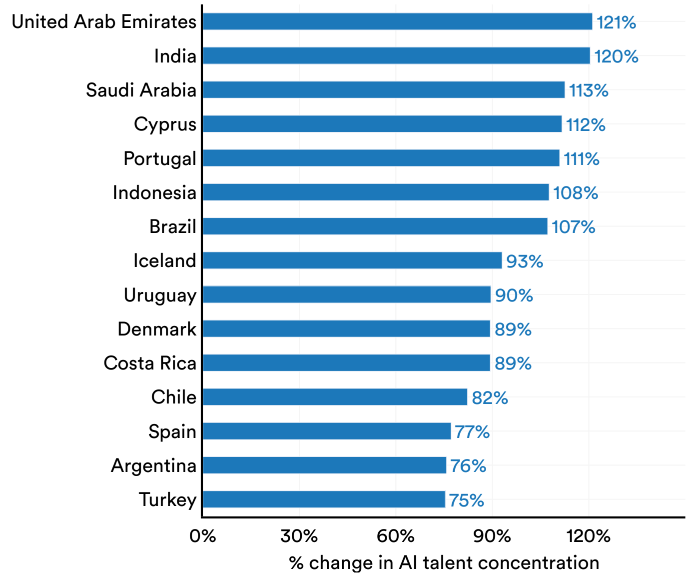

Figure 4.4.22[14]

---

[13] For the sake of brevity, the visualization includes only the top 15 countries for this metric.

[14] For the sake of brevity, the visualization includes only the top 15 countries for this metric.

Migration patterns show the dynamic global redistribution of AI talent (Figure 4.4.23 and 4.4.24). In 2025, Luxembourg recorded the highest net inflow relative to other tracked countries, with 5.23 per 10,000 LinkedIn members, as smaller economies actively try to attract more AI workers. The United States is a net importer of AI talent at 1.2 per 10,000 LinkedIn members. Similar to skills penetration, gender representation within AI talent continues to be uneven. Across the countries measured, men still account for the majority of AI talent, typically between 65% and 75% (Figures 4.4.25 and 4.4.26). Gender ratios have for the most part stayed flat since 2016, despite an expanding AI workforce. In the United States, women represent a 34.3% share of AI talent compared to men's 65.7% share. Other major labor markets, including the United Kingdom, Canada, France, and Singapore, show similar disparities.

**Net AI talent migration per 10,000 LinkedIn members by geographic area, 2025**

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

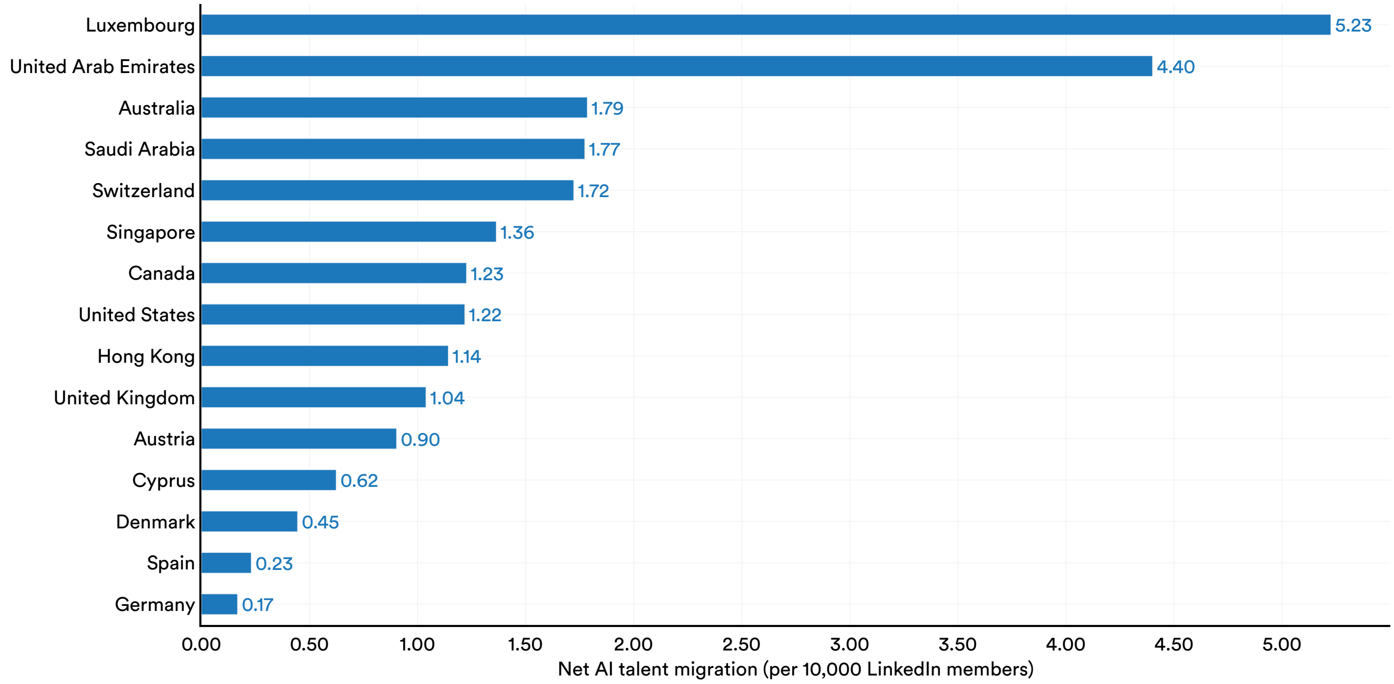

Figure 4.4.23[15]

---

[15] For the sake of brevity, the visualization includes only the top 15 countries for this metric.

## Net AI talent migration per 10,000 LinkedIn members by geographic area, 2021–25

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

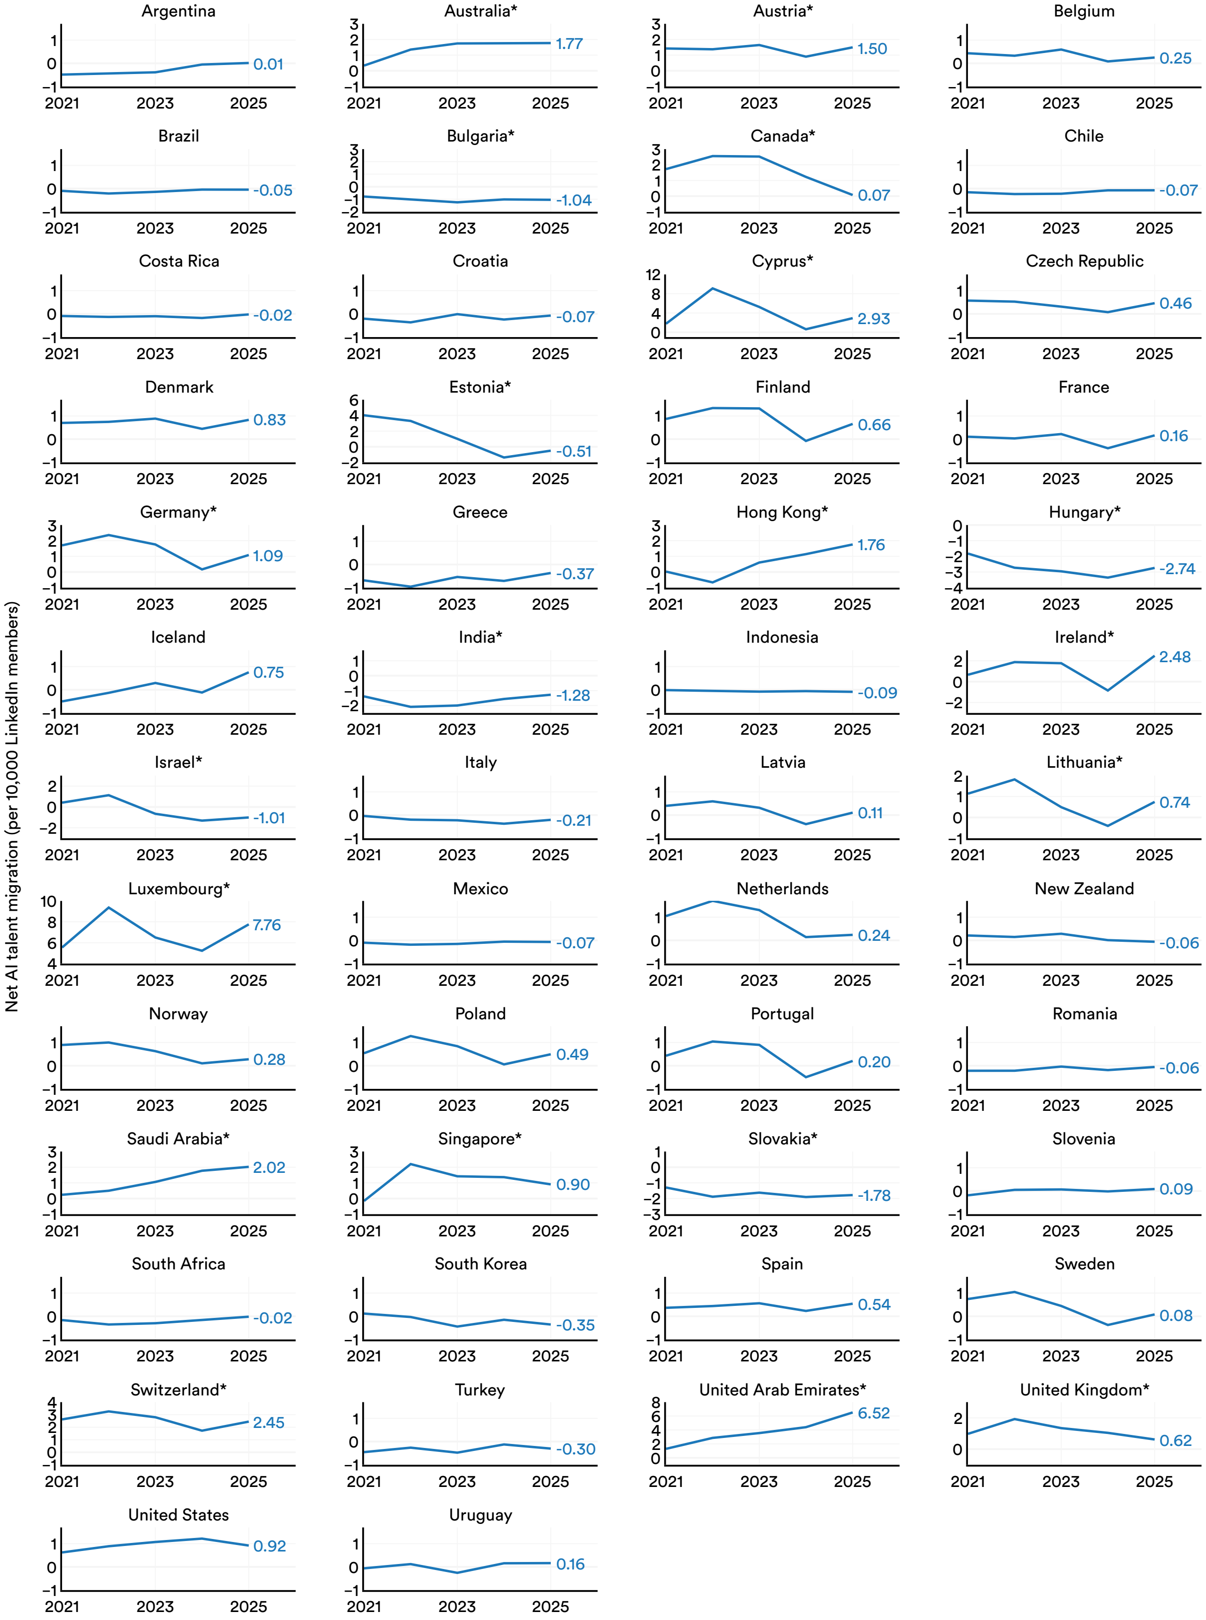

**Figure 4.4.24**[16]

---

[16] Asterisks indicate that a country's y-axis label is scaled differently than the y-axis label for the other countries.

## AI talent representation by gender and geographic area, 2016–25

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

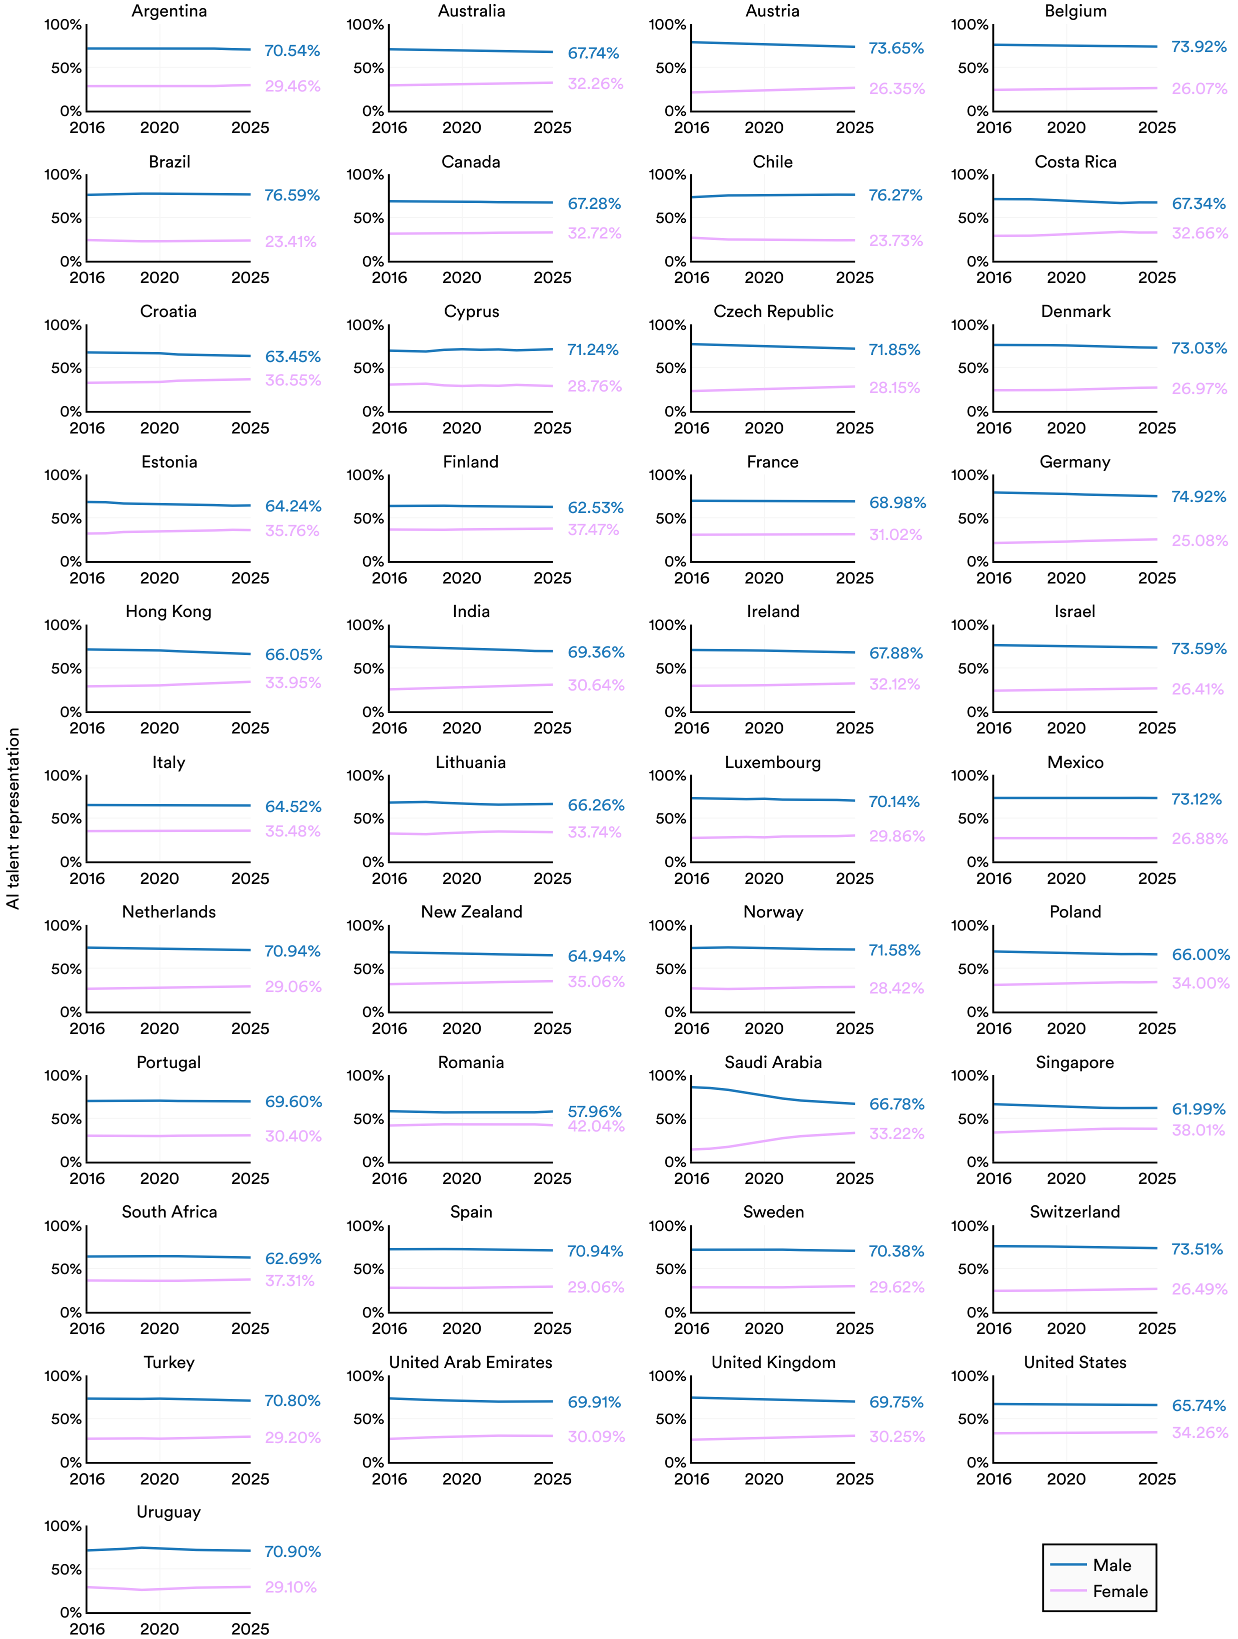

Figure 4.4.25

### AI talent concentration by gender and geographic area, 2016–25

Source: LinkedIn, 2025 | Chart: 2026 AI Index report

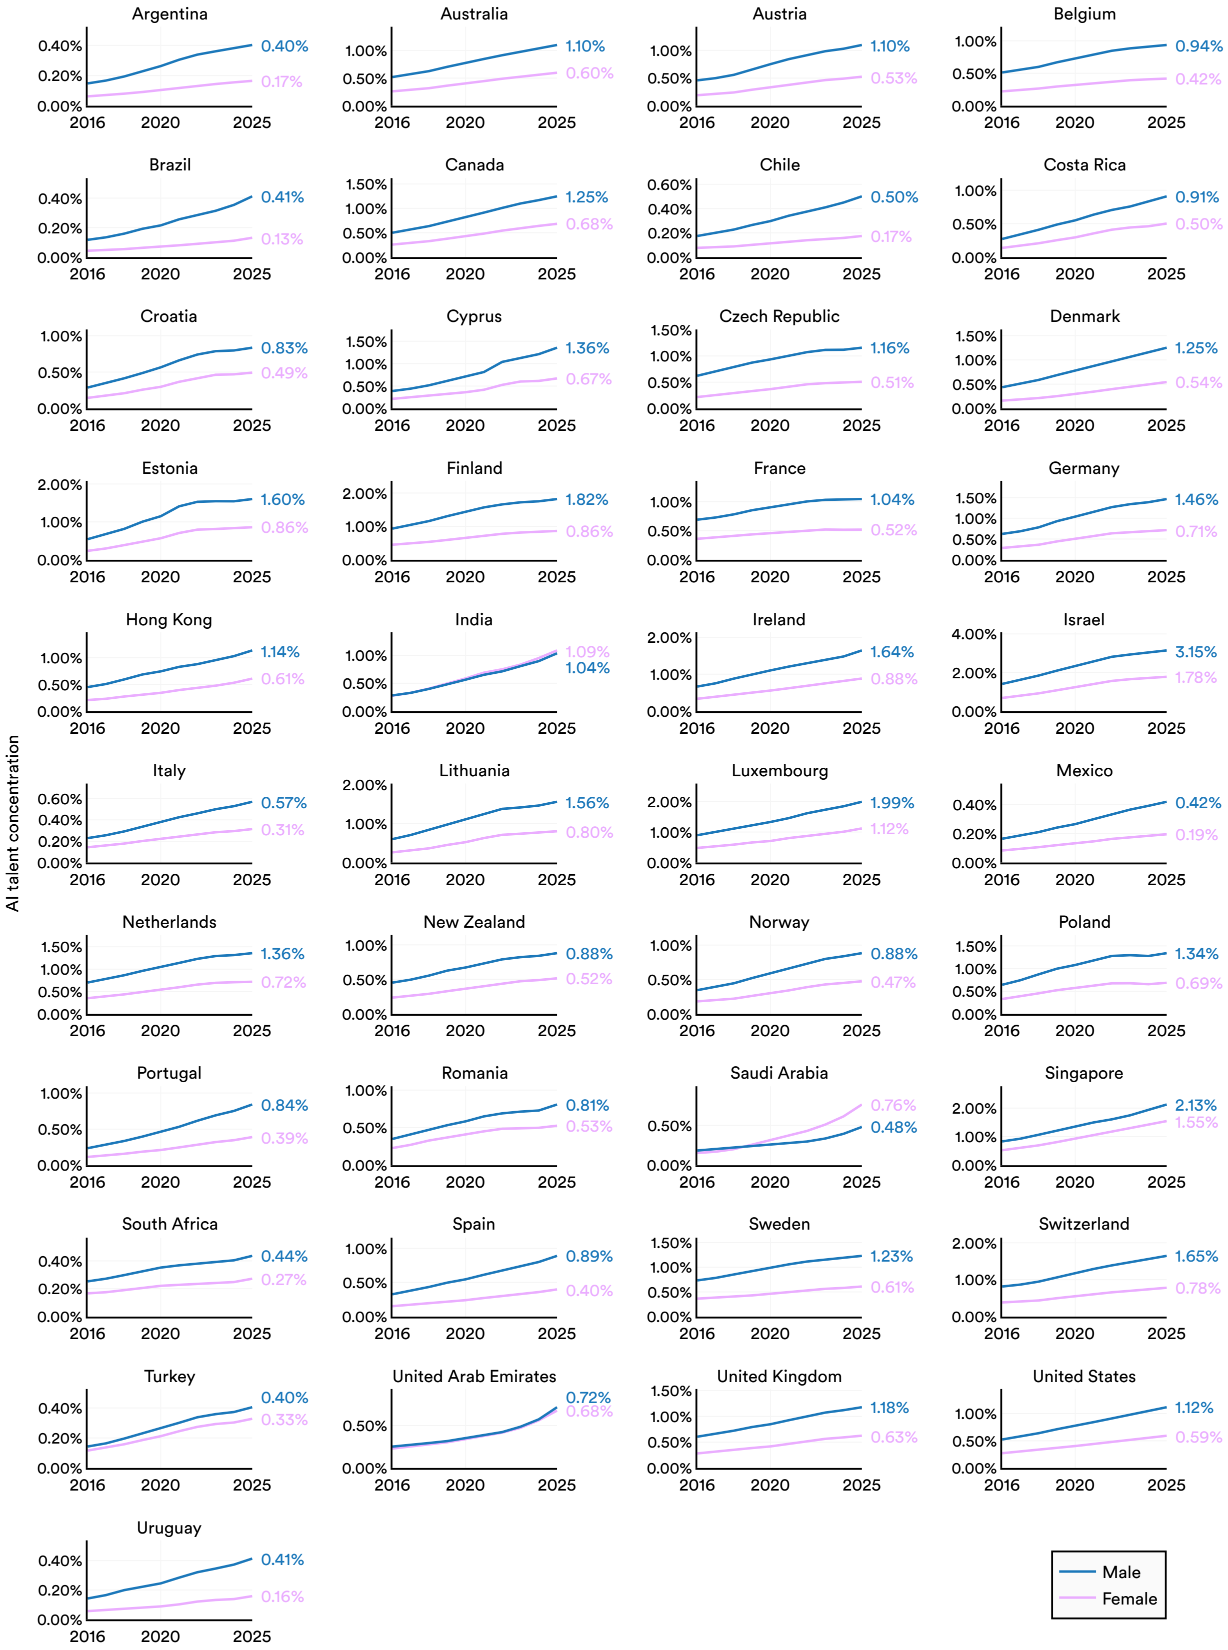

Figure 4.4.26

## AI's Labor Impact

Shifts in skill demand and talent flows are reshaping how work is organized by driving changes in productivity and hiring. A growing body of academic research has begun to measure AI's impact both at the micro level, examining how individual workers perform their jobs using AI tools, and at the macro level, examining how AI adoption impacts aggregate productivity and employment figures. In some settings, AI improves productivity, particularly for tasks that are structured, language heavy, or supported by clear feedback loops. In others, gains are marginal or even negative when tools are poorly matched to the task. Early macro-level evidence indicates that productivity gains may take longer to materialize and that the labor market costs may fall disproportionately on junior and entry-level workers. Findings from several notable studies are summarized in the tables below; as with any emerging research area, results vary in methodology, scope, and context.

## Productivity Trends

### Micro-level Studies

A growing number of studies have looked at how AI tools affect individual worker productivity across occupations (Figure 4.4.27). The results have been generally positive, but the size and distribution of the gains varies. The clearest impact is on support work, software development, and marketing. Customer support agents using a conversational AI assistant resolved 14%–15% more issues per hour (Brynjolfsson et al., 2025), software developers using GitHub Copilot completed 26% more pull requests (Cui et al., 2025), and marketing teams using multimodal AI for ad creation saw a 50% increase in output per worker (Ju and Aral, 2025). One consistent finding across several of these studies is that the less experienced workers tended to benefit the most, suggesting that AI tools may help close existing skill gaps.

Results are not uniformly positive, and for work that requires deeper reasoning or judgment, some studies have found AI tools produced little benefit or even slowed workers down. The most widely cited example comes from Model Evaluation & Threat Research (METR), which found that experienced open-source developers became 19 percent slower when using AI assistance, with a disconnect between how helpful the developers thought the tools were and how they actually performed (Becker et al., 2025). However, the METR team has not been able to replicate the results in a later study, primarily due to a growing reluctance among developers to work without AI, and that developers in late 2025 were likely sped up by AI relative to the original study period. When looking at the longer-term effects on skill development, research shows mixed results. Software engineers who relied heavily on AI for learning showed no measurable speed improvement and faced what researchers call "learning penalties" (Shen and Tamkin, 2025). Overall, AI's productivity effects are highly context dependent. The gains are strongest when work can be divided into well-defined, repeatable tasks with clear quality monitoring.

| Study | Occupation | AI application | Change in productivity | Who benefited most? |
|---|---|---|---|---|
| Reimers and Waldfogel (2026) | Authors | LLMs for content | +200%<br>(Output volume; releases tripled) | New entrants (drove quantity); pre-AI authors (maintained quality) |
| Shen and Tamkin (2025) | Software engineers | Learning new libraries | 0%<br>(Statistically insignificant speed change) | High scorers (65%+) who used AI for conceptual inquiry, avoiding "learning penalties" |
| Becker et al. (2025) | Developers | Open-source tools | -19%<br>(Speed; developers became slower with AI) | None (significant gap between perceived help and actual performance) |
| Brynjolfsson et al. (2025) | Support agents | Conversational assistant | +14%–15%<br>(Efficiency; issues resolved per hour) | Less experienced/skilled agents (30%–35% gains) |
| Cui et al. (2025) | Software developers | GitHub Copilot | +26%<br>(Output; completed pull requests) | Junior and less-experienced workers |
| Ju & Aral (2025) | Marketing teams | Multimodal ad creation | +50%<br>(Output; productivity per worker) | Human-AI teams (shifted focus from social coordination to task execution) |
| Choi & Xie (2025) | Accountants | AI-based accounting | +55%<br>(Throughput; weekly client support) | Experienced accountants (used AI confidence scores to target oversight) |

Figure 4.4.27

## Macro-level Studies

At the macro level, the evidence is earlier and less conclusive, but there are indicators that AI is starting to register in aggregate productivity data (Figure 4.2). A study of 12,000 European firms found that AI adoption boosted labor productivity by 4%, with training strengthening the outcome (Aldasoro et al., 2026). In the United States, productivity growth reached 2.7% in 2025, nearly double the 1.4% average of the previous decade. Brynjolfsson (2026) explains this may reflect the early stages of a "J-curve," where organizations absorb the costs of adopting AI before the larger productivity gains show up in the numbers. Similarly, OECD projections for G7 economies estimate annual productivity gains of 0.2 to 1.3 percentage points over the next decade (Filippucci et al., 2025). As mentioned earlier, the evidence is not conclusive nor is it all positive. A survey of 6,000 executives across four countries found widespread adoption but minimal realized productivity gains and a projected 0.7% reduction in employment over the next three years (Yotzov et al., 2026). The gap between adoption and measurable impact could be because AI is still early in its organizational integration, as illustrated earlier in this chapter through the deployment stage data. The pace at which these returns materialize will continue to be an important indicator to track.

| Study | Scope | Insight | Productivity / employment impact |
|---|---|---|---|
| Aldasoro et al. (2026) | 12,000 European firms (2019–2024) | AI adoption increases efficiency without reducing short-run employment; training significantly boosts gains. | +4% increase in labor productivity; +5.9 percentage point gain for every 1% spent on training. |
| Yotzov et al. (2026) | 6,000 executives (US, UK, DE, AU) | Representative survey showing high adoption but minimal realized impact on productivity to date. | +1.4% projected productivity boost; +0.8% projected output increase; -0.7% projected employment reduction (over next 3 years). |
| Brynjolfsson (2026) | United States economy | A "decoupling" of output from labor input is visible; framed through the "J-curve" hypothesis. | 2.7% US productivity growth in 2025 (nearly double the 1.4% annual average in the previous decade). |
| Filippucci et al. (OECD, 2025) | G7 economies (10-year horizon) | Projected annual gains based on sectoral specialization (e.g., high in finance/ICT, low in manufacturing). | +0.4 to +1.3 pp (US/UK) vs. +0.2 to +0.8 percentage points (Italy/Japan) annual labor productivity growth. |
| Frank et al. (2026) | LinkedIn profiles and US unemployment insurance records | Deterioration in AI-exposed labor markets (unemployment risk) began in early 2022, pre-ChatGPT. | Negative entry rates for AI-exposed roles. |
| Brynjolfsson et al. (2025) | US payroll data (ADP) through 2025 | "Canaries in the coal mine": large employment declines for junior workers in exposed fields. | -15% to -16% employment for early-career workers. |
| Hosseini Maasoum and Lichtinger (2025) | 62 million workers/ 285,000 US firms | "Seniority-biased technological change"; AI substitutes for junior labor while leaving senior roles intact. | Sharp decline in junior employment driven by slower hiring. |
| St. Louis Fed (2025) | US general labor market | Back-of-envelope calculations based on self-reported time savings. | +1.1% to +1.3% labor productivity increase. |
| Penn Wharton Budget Model (2025) | US economy | Projects AI's current contribution to total factor productivity (TFP). | +0.01 percentage points contribution to TFP (negligible). |

Figure 4.4.28

## Workforce Impact

It is challenging to measure AI's impact on employment, particularly because the technology is still in the early stages of widespread deployment. So far, effects on the workforce appear to be uneven, initially showing up in hiring pipelines, among younger workers, and within specific business functions. The evidence does not point to broad, uniform displacement. Firm-level survey data does suggest that many organizations expect the pace of workforce change to accelerate over the next year. According to McKinsey's survey, respondents also anticipated headcount reductions for the coming year to exceed those reported in the past year. (Figures 4.4.30 and Figure 4.4.31).

Recent employment data for software developers and customer service roles reveals a generational pattern (Brynjolfsson et al., 2025). In the United States, normalized headcount trends for both occupations show that employment among the youngest workers (ages 22–25) has declined since 2022, even as headcount for older age groups continues to grow (Figure 4.4.29). By September 2025, employment for software developers ages 22–25 had fallen close to 20% from its 2022 peak.

**Normalized headcount trends by age group for software developers and customer service agents, 2021–25**

Source: Brynjolfsson et al., 2025 | Chart: 2026 AI Index report

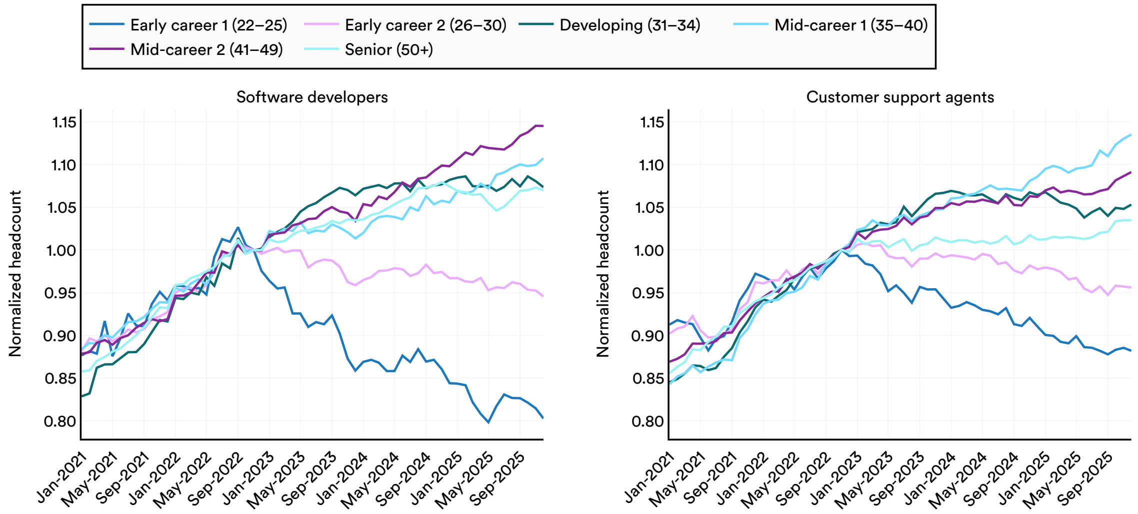

Figure 4.4.29

When occupations are grouped by their exposure to AI, the age-based pattern holds (Figure 4.4.30). Among workers ages 22–25, employment in the most AI-exposed occupations has fallen roughly 16% relative to the least-exposed, after controlling for firm-type effects, which isolate AI exposure from broader shocks like interest rate pressure or sector slowdowns. The gap began widening in mid-2024 and has grown steadily since.

**Headcount trends in high AI-exposure jobs (early career 22–25), 2021–25**

Source: Brynjolfsson et al., 2025 | Chart: 2026 AI Index report

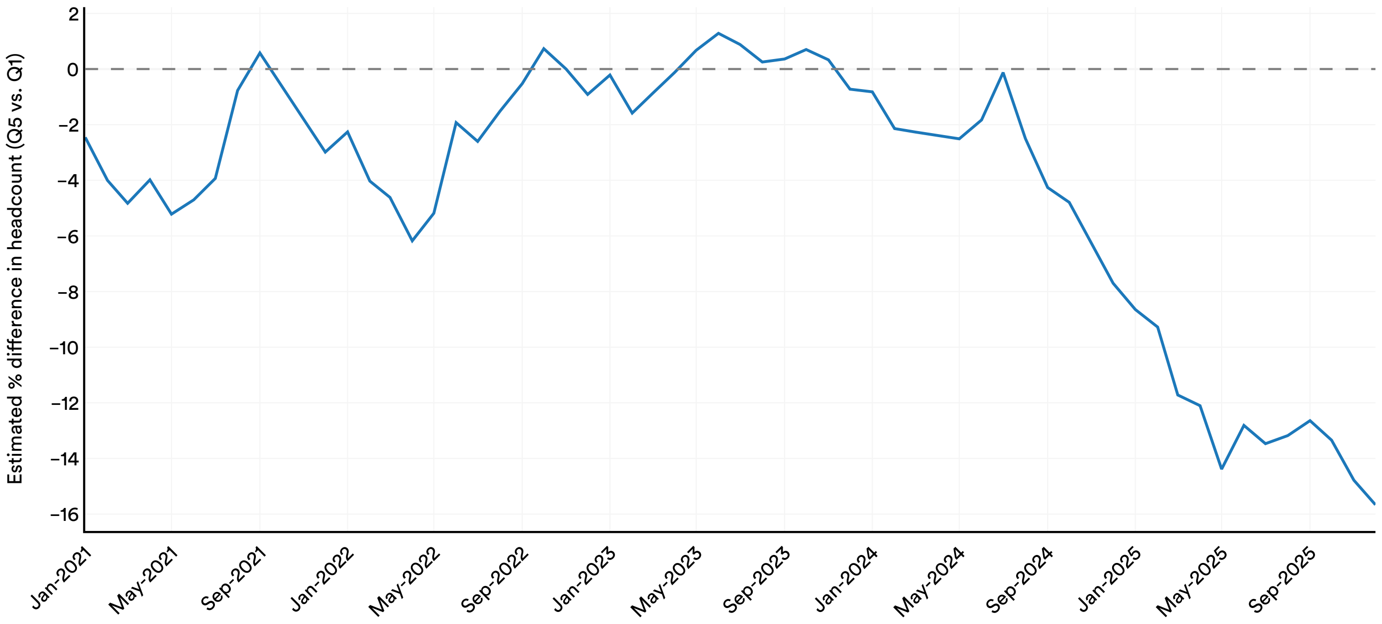

Figure 4.4.30[17]

---

[17] This figure includes firm-by-time fixed effects (within-firm, same-period comparisons), accounting for firm/industry hiring swings.

Unemployment data suggests an even more complicated dynamic (Felten et al., 2021; Eckhardt and Goldschlag, 2025) (Figure 4.4.31). From 2022 to early 2025, unemployment rose across all occupation groups regardless of AI exposure level. While the unemployment rate for the most AI-exposed workers (quintile 5) increased by 0.30 percentage points, it rose even more for the least exposed workers (quintile 1), climbing by 0.94 percentage points. AI exposure alone does not seem to be driving recent unemployment trends, but it appears to play a part in broader macroeconomic conditions and organizational changes.

### Unemployment rate by AI exposure quintile

Source: Eckhardt and Goldschlag, 2025

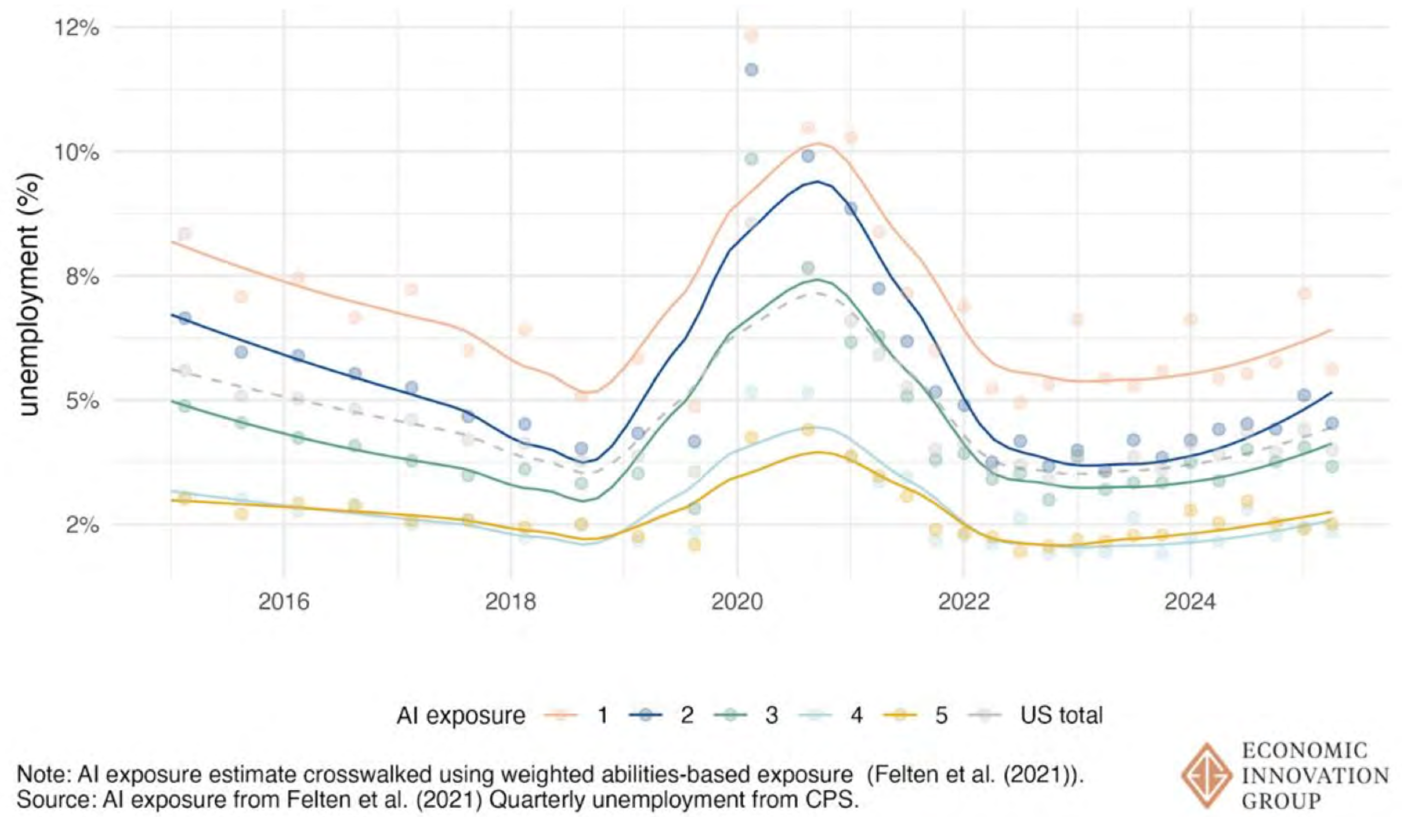

Note: AI exposure estimate crosswalked using weighted abilities-based exposure (Felten et al. (2021)).
Source: AI exposure from Felten et al. (2021) Quarterly unemployment from CPS.

Figure 4.4.31

A different view emerges from looking at occupational churn trends after the introduction of major technologies (Figure 4.4.32). Over comparable time frames, the occupational mix in the United States has shifted faster since the introduction of generative AI than the shift that followed the introduction of computers or the internet (Gimbel et al., 2025).

### Changes in the occupational mix from recent baselines

Source: Gimbel et al., 2025

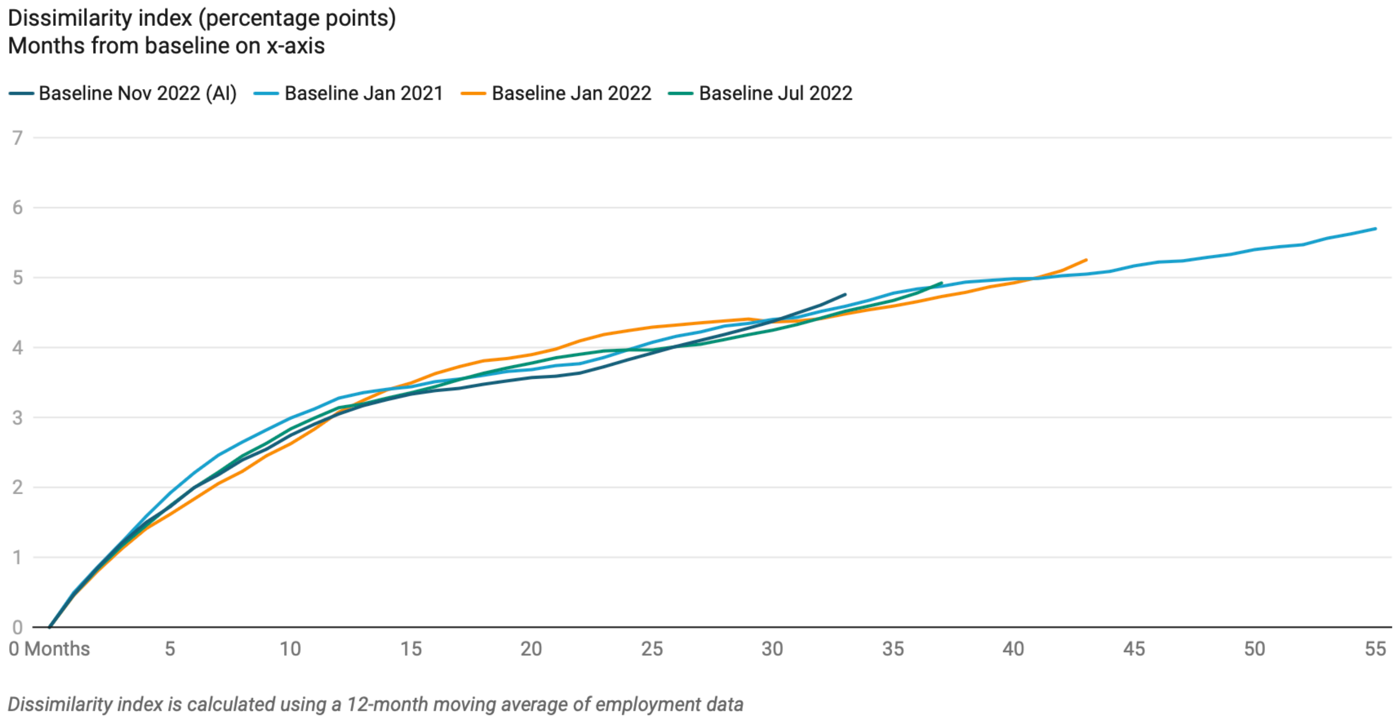

Figure 4.4.32

However, employer expectations seem to indicate that the pace may accelerate (Figure 4.4.33). According to McKinsey’s survey (2025), one-third of respondents anticipate a decrease in workforce size, a percentage that is higher at larger organizations (35% at organizations with ≥$1 billion in revenue) compared to smaller firms (30% at organizations with <$1 billion in revenue). Most respondents (43% overall) expect little or no change, while only a minority foresee an increase in workforce size. Even compared to workforce changes that have already taken place, the general sentiment leans toward headcount reductions (Figure 4.4.34). In nearly all functions, respondents anticipate a greater impact from AI on headcount next year than was observed in the past year, with expected decreases outpacing observed decreases. This trend is particularly pronounced in service operations, supply chain/inventory management, marketing and sales, and software engineering, where the expected decrease in employees for the next year significantly exceeds the actual decrease reported over the past year. Conversely, expectations for workforce increases remain relatively modest across business functions.

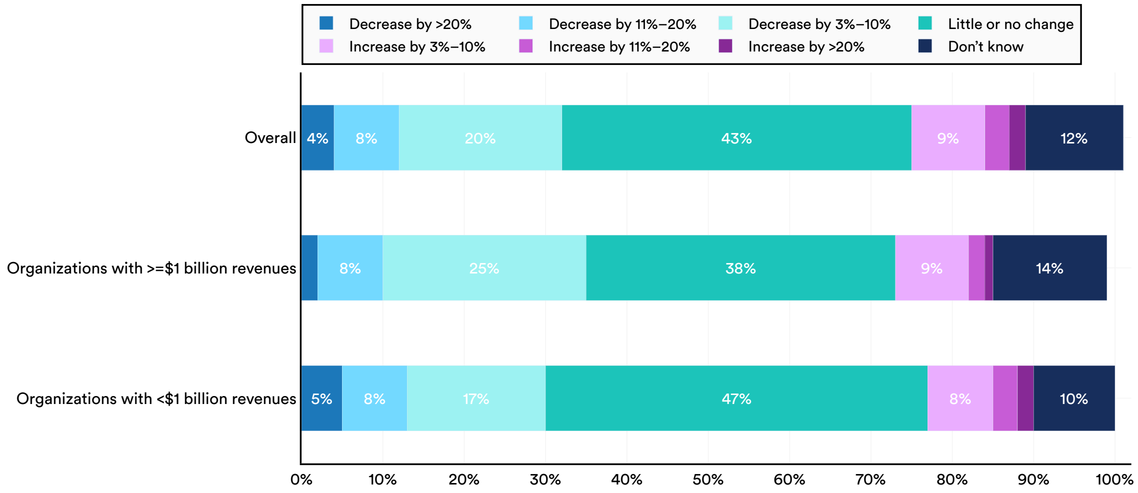

**Expected change in workforce size as a result of AI in the next year**

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

Figure 4.4.33

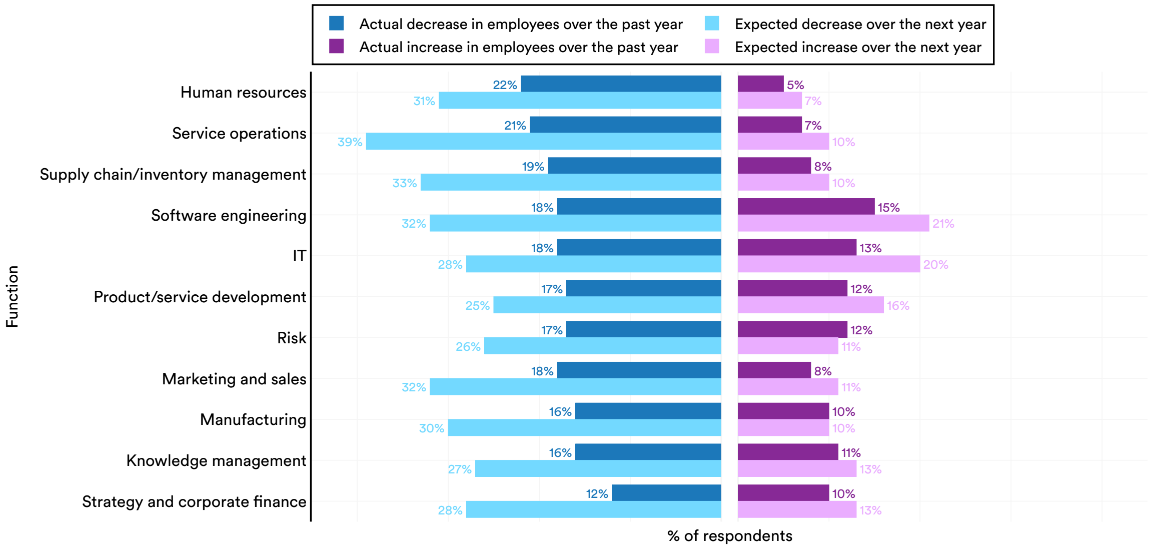

**Actual vs. expected change in workforce size as a result of AI by function**

Source: McKinsey & Company Survey, 2025 | Chart: 2026 AI Index report

Figure 4.4.34

A surprising finding from Shao et al. (2026) shows that many workers are not wholly resistant to automation (Figure 4.4.35). A survey of 844 occupational tasks found that 46.1% of workers actively want AI to take over those tasks. Support was especially strong in areas where workers believed automation would free up time for higher-value tasks, reduce repetitiveness, or improve quality of output. However, actual usage patterns do not necessarily reflect these preferences. Occupational tasks with the highest average automation scores account for only 1.3% of Claude.AI usage. A related framework maps these tasks across four zones based on worker desire and technical feasibility (Figure 4.4.36). Looked at this way, AI's labor impact will likely register in how specific tasks are redesigned rather than blunt automation at the occupation level, with the reorganization of work unfolding gradually.

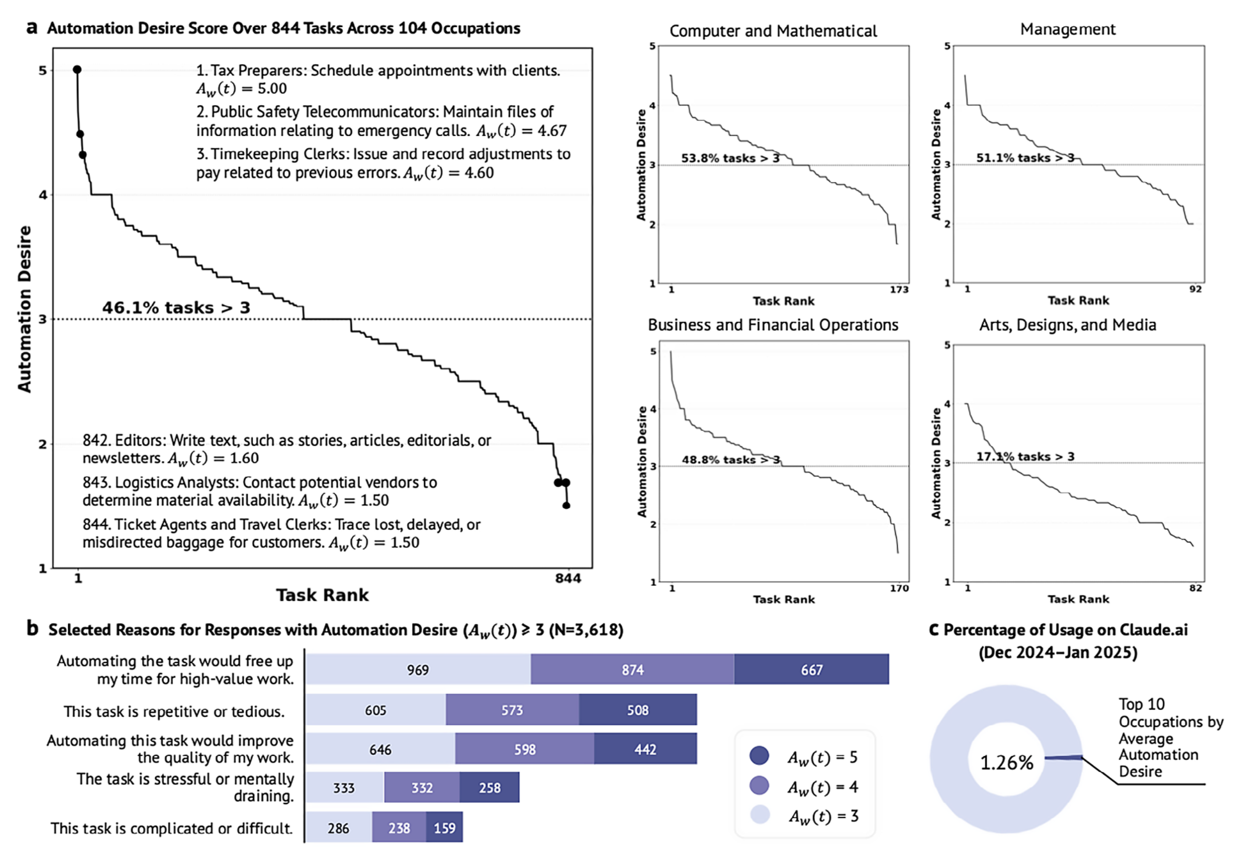

Source: [Shao et al. (2026)](#)

Figure 4.4.35

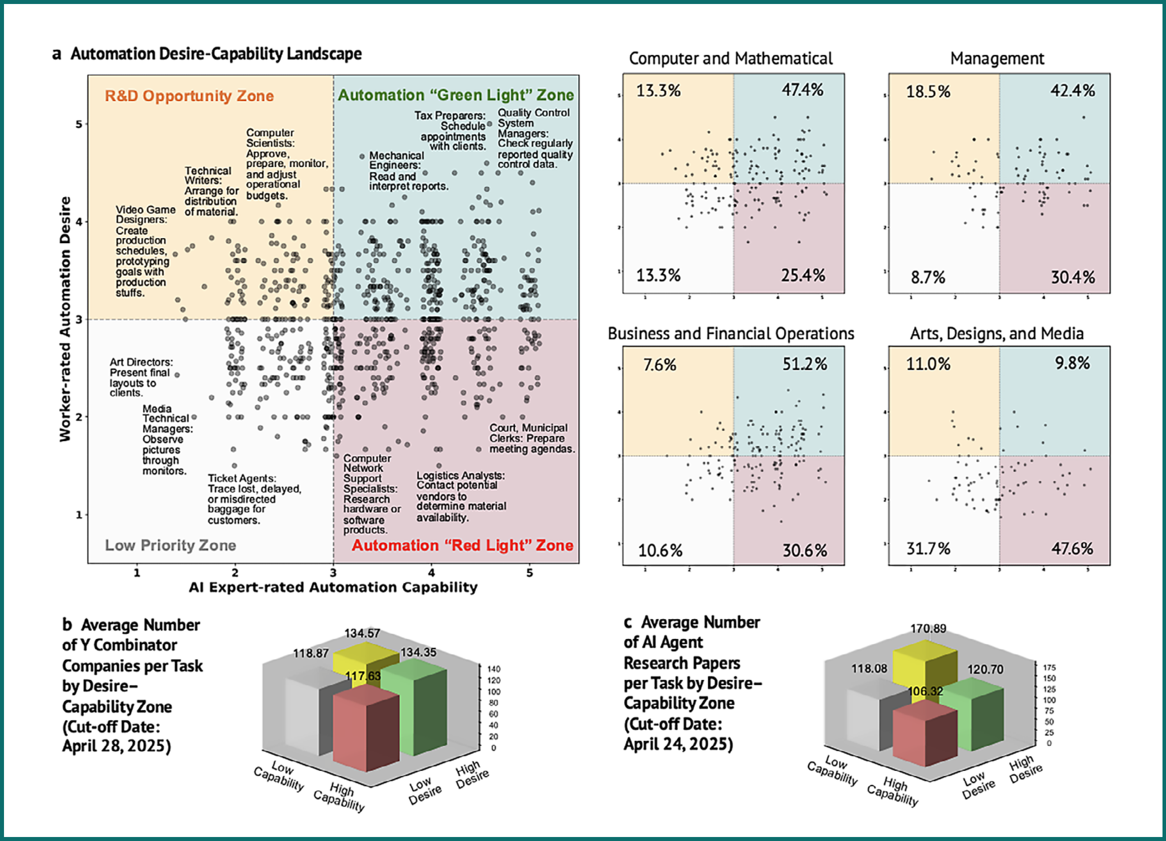

Source: [Shao et al. (2026)](#)

Figure 4.4.36

# 4.5 Robot Deployments

Physical automation through robotics represents one form of AI's economic integration in industrial environments, particularly in production settings such as manufacturing lines or warehouses. To track the trends, the AI Index uses data from the International Federation of Robotics, IFR 2025, a nonprofit that publishes annual World Robotics reports on global installation patterns and operational stock. Industrial robots are defined by the IFR, and in this reporting,[18] as "automatically controlled, reprogrammable, multipurpose manipulators, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications."

## Global Installation Patterns

Global industrial robot activity continues to rise, though year-over-year growth has flattened. In 2024, 542,000 industrial robots were installed globally, a slight increase (0.2%) from the previous year (Figure 4.5.1). The composition of those robots has also shifted over time. Collaborative robots, which are designed to work alongside human operators, continue to gain market share over traditional robots. In 2017, collaborative robots accounted for just 2.8% of all new industrial robot installations, compared to 13.6% in 2024. The total operational stock in 2024 grew to 4,664,000, up from 4,282,000 in 2023 (Figure 4.5.2). Overall, industrial automation capacity has shown a consistent upward trajectory, with the global fleet of industrial robots quadrupling since 2012.

**Number of industrial robots installed in the world, 2012–24**

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

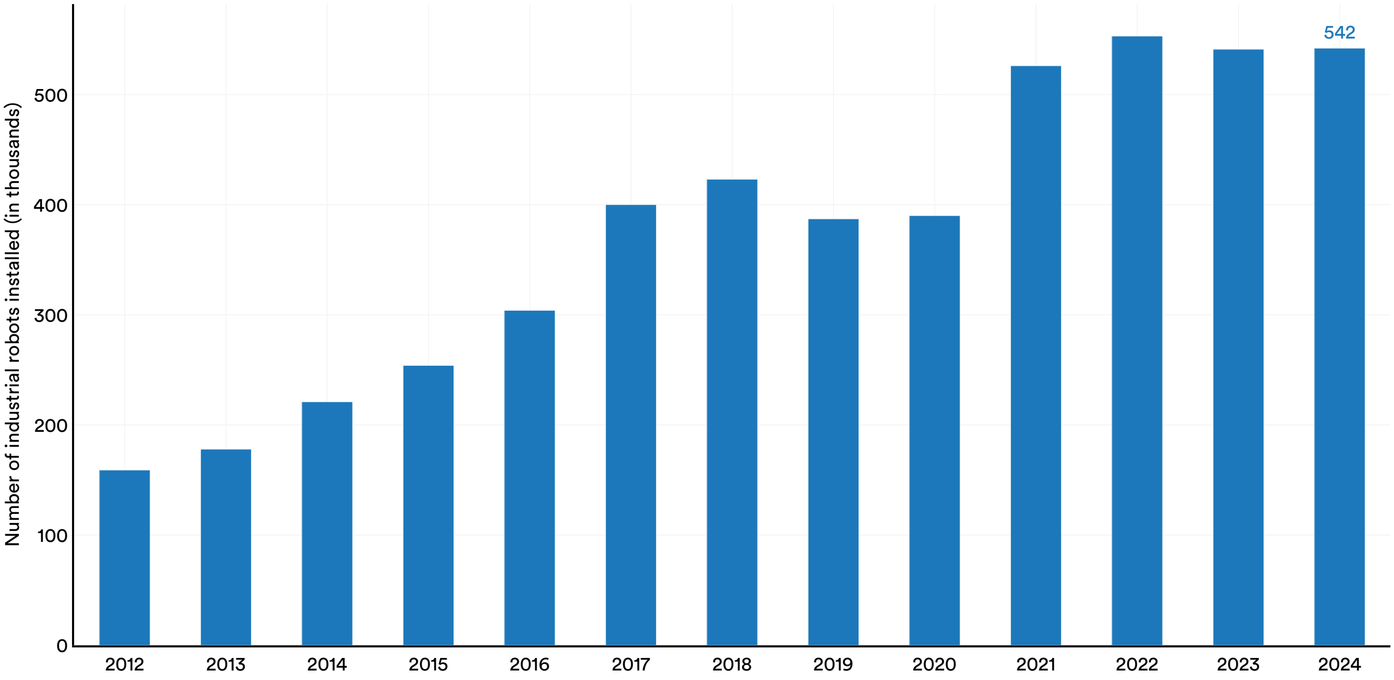

Figure 4.5.1

---

[18] Due to the timing of the IFR report, the most recent data is from 2024. Every year, the IFR revisits data collected for previous years and will occasionally update the data if more accurate figures become available. Therefore, some of the data reported in this year's report might differ slightly from data reported in previous years.

## Operational stock of industrial robots in the world, 2012–24

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

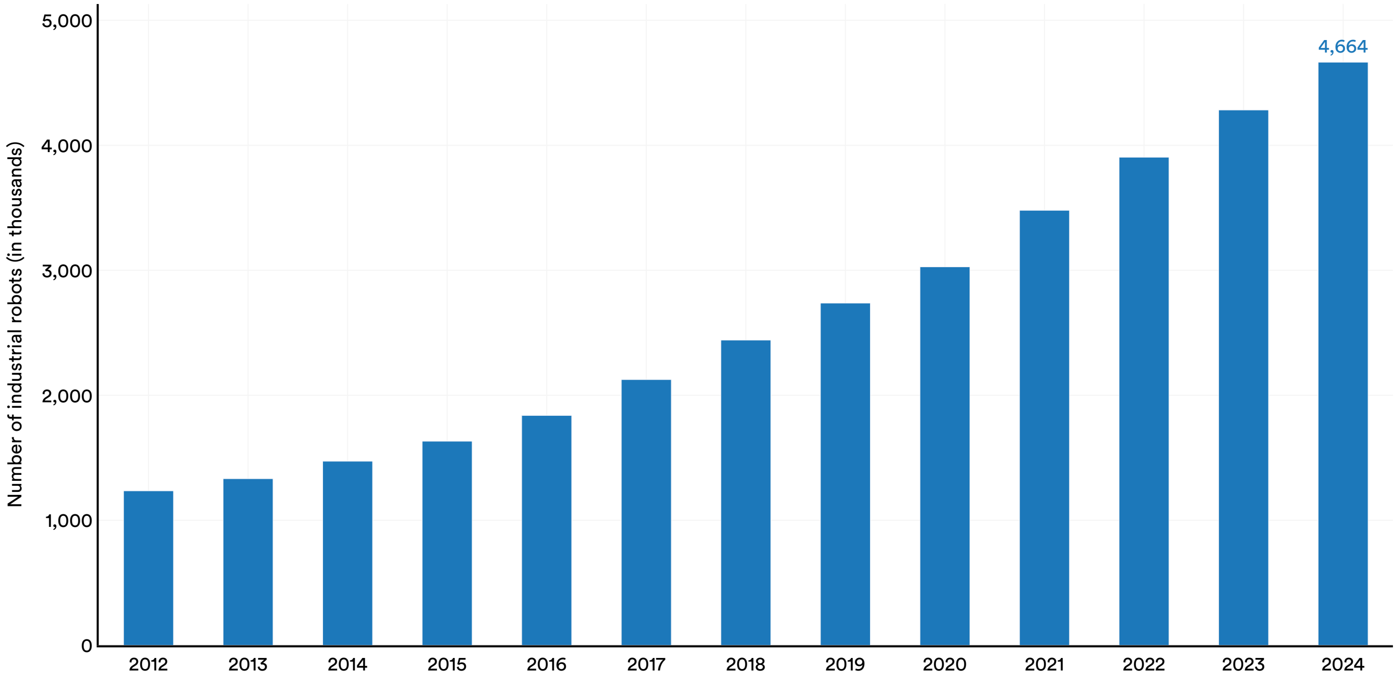

Figure 4.5.2

## Number of industrial robots installed in the world by type, 2017–24

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

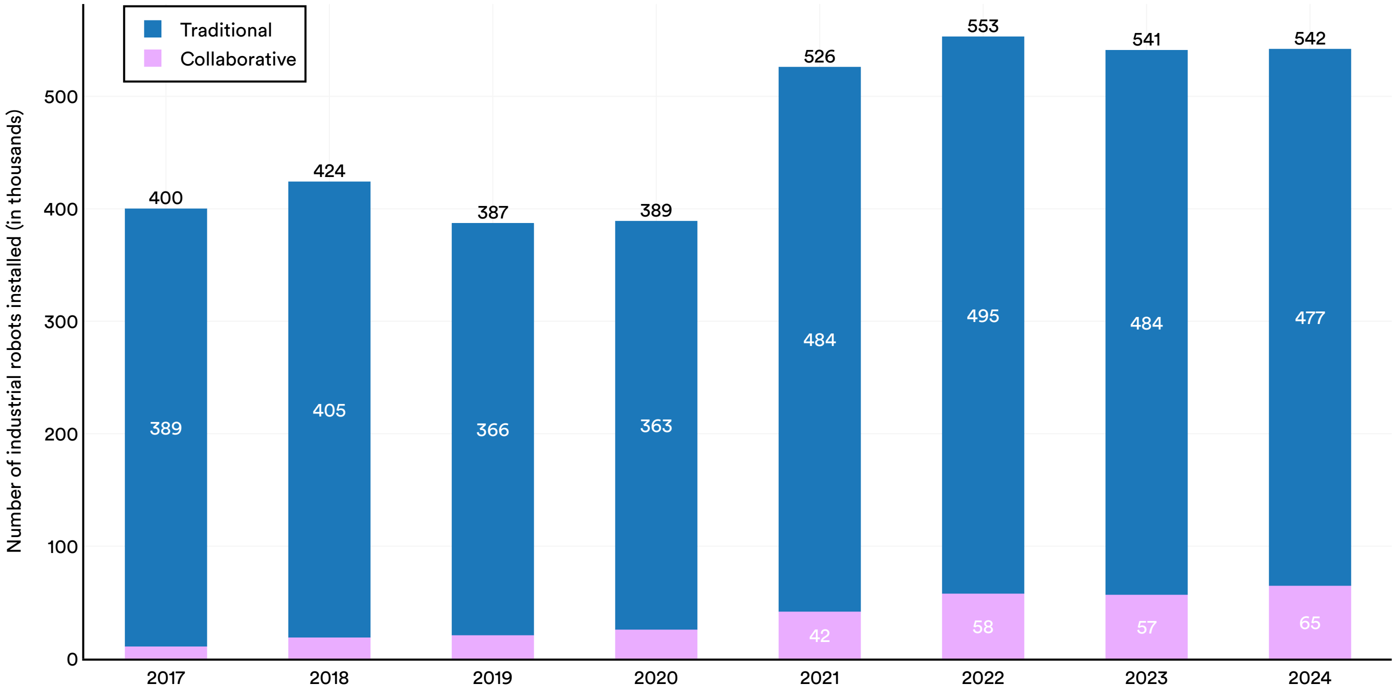

Figure 4.5.3

## Geographic Patterns

Industrial robot installation follows patterns similar to the investment and talent trends discussed above, although its geographic distribution is relatively narrow. In 2024, China led the world with 295,000 industrial robot installations, six times more than Japan's 44,500 and 8.6 times more than the United States' 34,200 (Figure 4.5.4). South Korea and Germany followed with 30,600 and 27,000 installations, respectively. China's share of global installations has increased substantially from 20.8% in 2013 to 54.4% in 2024 (Figure 4.5.5).

**Number of industrial robots installed by geographic area, 2024**

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

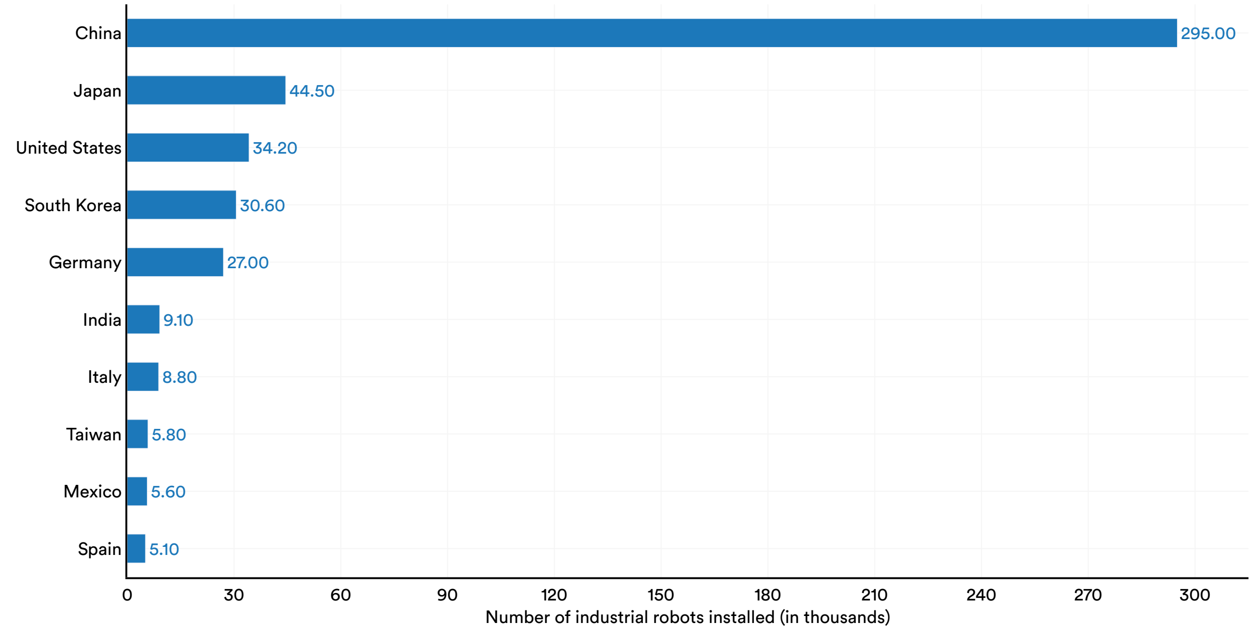

Figure 4.5.4

**Number of new industrial robots installed in top 5 countries, 2011–24**

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

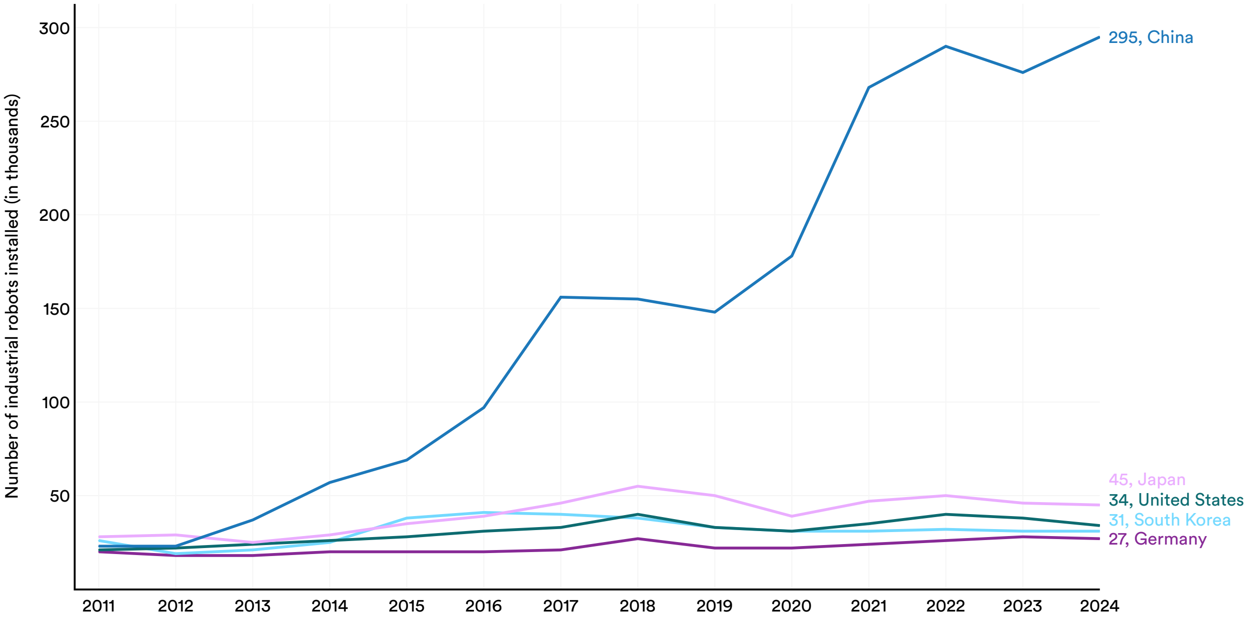

Figure 4.5.5

**Annual growth rate of industrial robots installed by geographic area, 2023 vs. 2024**

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

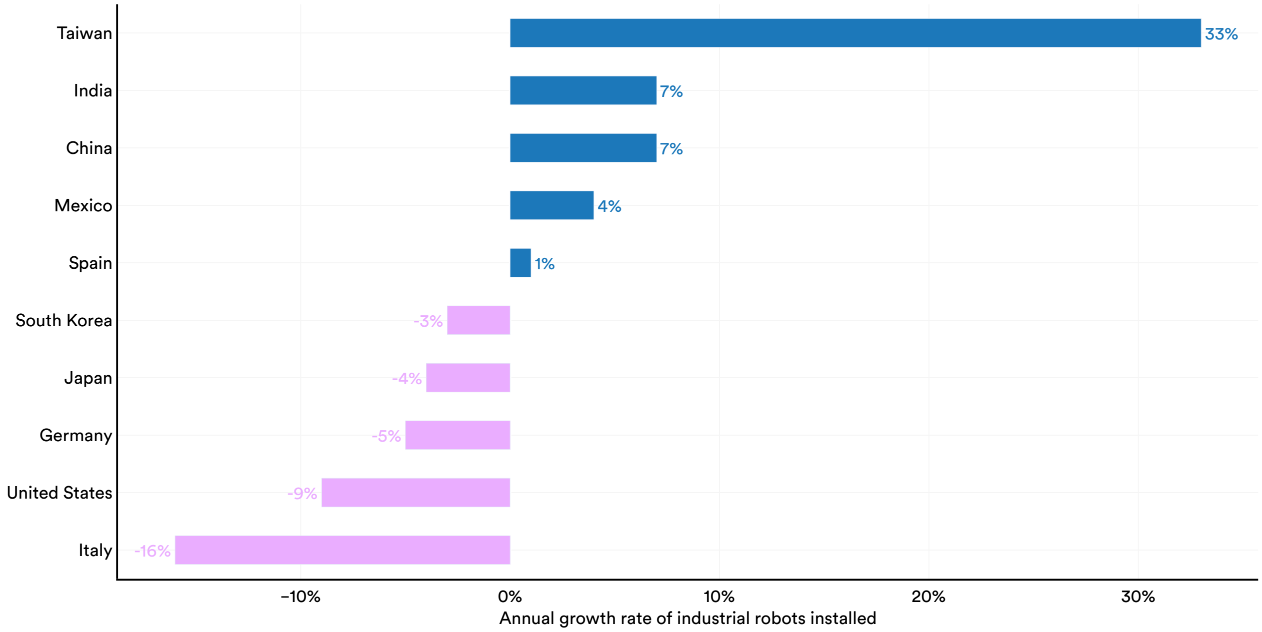

Figure 4.5.6

## Service Robotics

Nonindustrial or service robots designed for tasks such as logistics, hospitality, and agriculture showed growth in 2024 (Figure 4.5.7). Service robot installations increased across most application areas compared to 2023, though agriculture saw particularly strong adoption. The number of service robots deployed in an agricultural setting increased 2.5-fold. Only the hospitality category saw a year-over-year decline.

**Number of service robots installed in the world by application area, 2021–24**

Source: International Federation of Robotics (IFR), 2025 | Chart: 2026 AI Index report

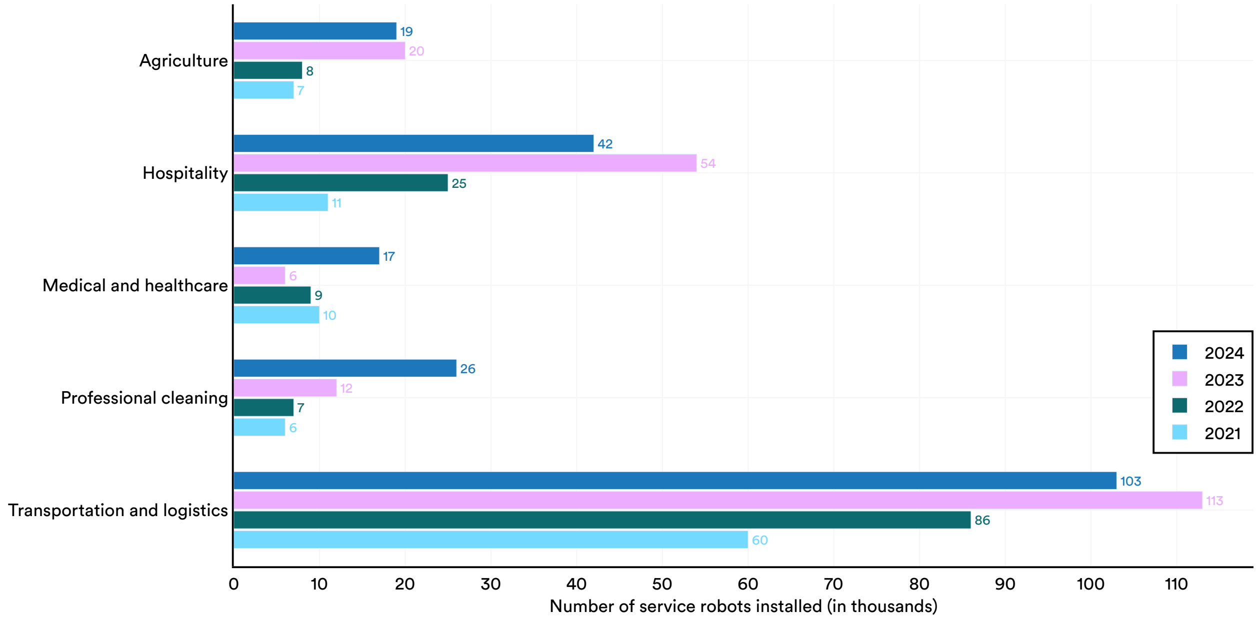

Figure 4.5.7