


McKinsey
& Company

Superagency in the Workplace

Empowering people to unlock AI's full potential

Hannah Mayer
Lareina Yee
Michael Chui
Roger Roberts

January 2025



Contents

Introduction [2](#)

Chapters:


- 1. An innovation as powerful as the steam engine [5](#)**
- 2. Employees are ready for AI; now leaders must step up [11](#)**
- 3. Delivering speed and safety [18](#)**
- 4. Embracing bigger ambitions [26](#)**
- 5. Technology is not the barrier to scale [35](#)**

Conclusion: Meeting the AI future [40](#)

Acknowledgments [42](#)

Methodology [43](#)

Glossary [44](#)



Introduction

Almost all companies invest in AI, but just 1 percent believe they are at maturity. Our research finds the biggest barrier to scaling is not employees—who are ready—but leaders, who are not steering fast enough.

Artificial intelligence has arrived in the workplace and has the potential to be as transformative as the steam engine was to the 19th-century Industrial Revolution.¹ With powerful and capable large language models (LLMs) developed by Anthropic, Cohere, Google, Meta, Mistral, OpenAI, and others, we have entered a new information technology era. McKinsey research sizes the long-term AI opportunity at \$4.4 trillion in added productivity growth potential from corporate use cases.²

Therein lies the challenge: the long-term potential of AI is great, but the short-term returns are unclear. Over the next three years, 92 percent of companies plan to increase their AI investments. But while nearly all companies are investing in AI, only 1 percent of leaders call their companies “mature” on the deployment spectrum, meaning that AI is fully integrated into workflows and drives substantial business outcomes. The big question is how business leaders can deploy capital and steer their organizations closer to AI maturity.

This research report, prompted by Reid Hoffman’s book *Superagency: What Could Possibly Go Right with Our AI Future*,³ asks a similar question: How can companies harness AI to amplify human agency and unlock new levels of creativity and productivity in the workplace? AI could drive enormous positive and disruptive change. This transformation will take some time, but leaders must not be dissuaded. Instead, they must advance boldly today to avoid becoming uncompetitive tomorrow. The history of major economic and technological shifts shows that such moments can define the rise and fall of companies. Over 40 years ago, the internet was born. Since then, companies including Alphabet, Amazon, Apple, Meta, and Microsoft have attained trillion-dollar market capitalizations. Even more profoundly, the internet changed the anatomy of work and access to information. AI now is like the internet many years ago: The risk for business leaders is not thinking too big, but rather too small.


¹ [“Gen AI: A cognitive industrial revolution,”](#) McKinsey, June 7, 2024.

² [“The economic potential of generative AI: The next productivity frontier,”](#) McKinsey, June 14, 2023.

³ Reid Hoffman and Greg Beato, *Superagency: What Could Possibly Go Right with Our AI Future*, Authors Equity, January 2025.

Superagency: By the numbers

Employees are more ready for the change than their leaders imagine




3x


more employees are using gen AI for a third or more of their work than their leaders imagine; more than 70% of all employees believe that within 2 years gen AI will change 30% or more of their work

1.4x

more likely for millennials to report extensive familiarity with gen AI tools than peers in other age groups; they are also 1.2x more likely to expect workflows to change within a year



Companies need to move fast—employees trust their leaders to balance speed and safety



47% of the C-suite say their companies are developing gen AI tools too slowly, even though 69% started investing more than a year ago




1.3x more likely for employees to trust their own companies to get gen AI deployment right than they are to trust other institutions

Companies are investing in gen AI but have not yet achieved maturity


92%

of companies plan to invest more in gen AI over the next 3 years



1%

believe their investments have reached maturity



Leaders need to recognize their responsibility in driving gen AI transformation



2.4x

more likely for C-suite to cite employee readiness as a barrier to adoption vs their own issues with leadership alignment, despite employees currently using gen AI 3x more than leaders expect



48%

of employees rank training as the most important factor for gen AI adoption; yet nearly half feel they are receiving moderate or less support

This report explores companies' technology and business readiness for AI adoption (see sidebar "About the survey"). It concludes that employees are ready for AI. The biggest barrier to success is leadership.

Chapter 1 looks at the rapid advancement of technology over the past two years and its implications for business adoption of AI.

Chapter 2 delves into the attitudes and perceptions of employees and leaders. Our research shows that employees are more ready for AI than their leaders imagine. In fact, they are already using AI on a regular basis; are three times more likely than leaders realize to believe that AI will replace 30 percent of their work in the next year; and are eager to gain AI skills. Still, AI optimists are only a slight majority in the workplace; a large minority (41 percent) are more apprehensive and will need additional support. This is where millennials, who are the most familiar with AI and are often in managerial roles, can be strong advocates for change.

Chapter 3 looks at the need for speed and safety in AI deployment. While leaders and employees want to move faster, trust and safety are top concerns. About half of employees worry about AI inaccuracy and cybersecurity risks. That said, employees express greater confidence that their own companies, versus other organizations, will get AI right. The onus is on business leaders to prove them right, by making bold and responsible decisions.

Chapter 4 examines how companies risk losing ground in the AI race if leaders do not set bold goals. As the hype around AI subsides, companies should put a heightened focus on practical applications that empower employees in their daily jobs. These applications can create competitive moats and generate measurable ROI. Across industries, functions, and geographies, companies that invest strategically can go beyond using AI to drive incremental value and instead create transformative change.

Chapter 5 looks at what is required for leaders to set their teams up for success with AI. The challenge of AI in the workplace is not a technology challenge. It is a business challenge that calls upon leaders to align teams, address AI headwinds, and rewire their companies for change.

About the survey

To create our report, we surveyed 3,613 employees (managers and independent contributors) and 238 C-level executives in October and November 2024. Of these, 81 percent came from the United States, and the rest came from five other countries: Australia, India, New Zealand, Singapore, and the United Kingdom. The employees spanned many roles, including business development, finance, marketing, product management, sales, and technology.

All the survey findings discussed in the report, aside from two sidebars presenting international nuances, pertain solely to US workplaces. The findings are organized in this way because the responses from US employees and C-suite executives provide statistically significant conclusions about the US workplace. Analyzing global findings separately allows a comparison of differences between US responses and those from other regions.



1

An innovation as powerful as the steam engine

‘Scientific discoveries and technological innovations are stones in the cathedral of human progress.’

– Reid Hoffman, cofounder of LinkedIn and Inflection AI, partner at Greylock Partners, and author

Imagine a world where machines not only perform physical labor but also think, learn, and make autonomous decisions. This world includes humans in the loop, bringing people and machines together in a state of superagency that increases personal productivity and creativity (see sidebar “AI superagency”). This is the transformative potential of AI, a technology with a potential impact poised to surpass even the biggest innovations of the past, from the printing press to the automobile. AI does not just automate tasks but goes further by automating cognitive functions. Unlike any invention before, AI-powered software can adapt, plan, guide—and even make—decisions. That’s why AI can be a catalyst for unprecedented economic growth and societal change in virtually every aspect of life. It will reshape our interaction with technology and with one another.

Many breakthrough technologies, including the internet, smartphones, and cloud computing, have transformed the way we live and work. AI stands out from these inventions because it offers more than access to information. It can summarize, code, reason, engage in a dialogue, and make choices. AI can lower skill barriers, helping more people acquire proficiency in more fields, in any language and at any time. AI holds the potential to shift the way people access and use knowledge. The result will be more efficient and effective problem solving, enabling innovation that benefits everyone.

Over the past two years, AI has advanced in leaps and bounds, and enterprise-level adoption has accelerated due to lower costs and greater access to capabilities. Many notable AI innovations have emerged (Exhibit 1). For example, we have seen a rapid expansion of [context windows](#), or the short-term memory of LLMs. The larger a context window, the more information an LLM can process at once. To illustrate, Google’s Gemini 1.5 could process one million tokens in February 2024, while its Gemini 1.5 Pro could process two million tokens by June of that same year.⁴ Overall, we see five big innovations for business that are driving the next wave of impact: enhanced intelligence and reasoning capabilities, agentic AI, multimodality, improved hardware innovation and computational power, and increased transparency.

AI superagency

What impact will AI have on humanity? Reid Hoffman and Greg Beato’s book *Superagency: What Could Possibly Go Right with Our AI Future* (Authors Equity, January 2025) explores this question. The book highlights how AI could enhance human agency and heighten our potential. It envisions a human-led, future-forward approach to AI.

Superagency, a term coined by Hoffman, describes a state where individuals, empowered by AI, supercharge their creativity, productivity, and positive impact. Even those not directly engaging with AI can benefit from its broader effects on knowledge, efficiency, and innovation.

AI is the latest in a series of transformative supertools, including the steam engine, internet, and smartphone, that have reshaped our world by amplifying human capabilities. Like its predecessors, AI can democratize access to knowledge and automate tasks, assuming humans can develop and deploy it safely and equitably.

⁴ *The Keyword*, “Our next-generation model: Gemini 1.5,” blog entry by Sundar Pichai and Demis Hassabis, Google, February 15, 2024; *Google for Developers*, “Gemini 1.5 Pro 2M context window, code execution capabilities, and Gemma 2 are available today,” blog entry by Logan Kilpatrick, Shrestha Basu Mallick, and Ronen Kofman, June 27, 2024.

Exhibit 1

Gen AI capabilities have evolved rapidly over the past two years.

Illustrative capabilities of gen AI platforms from select frontier labs, nonexhaustive

| | 2022–23 ¹ | Jan 2025 ² |
|----------------------|---|---|
| Anthropic | Claude <ul style="list-style-type: none"> • Not multimodal (text only) • Limited contextual understanding (difficulty with complex conversations) • No tool usage | Claude 3.5 <ul style="list-style-type: none"> • Multimodal (text, audio, and images) • Enhanced contextual understanding and coherence during long interactions • Experimental computer usage capability for some users |
| Google Gemini | Google Bard <ul style="list-style-type: none"> • Not multimodal (text only) • Fair reasoning • Limited contextual understanding (difficulty with complex conversations) • Limited real-time data integration • Low personalization (limited adaptability) | Gemini 2.0 Flash <ul style="list-style-type: none"> • Multimodal (text, audio, and images) • Advanced reasoning (capable of multistep problem-solving and nuanced analysis) • Enhanced contextual understanding (maintains coherence in long dialogues) • Real-time data integration (from Google Search) • Advanced personalization (user context) |
| Meta | Llama 1 <ul style="list-style-type: none"> • Not multimodal (text only) • Fair reasoning • Limited contextual understanding (difficulty with complex conversations) • No API access | Llama 3.3 <ul style="list-style-type: none"> • Text-based (earlier versions were multimodal, LLaMa 3,2) • Advanced reasoning (capable of multistep problem-solving and nuanced analysis) • Enhanced contextual understanding (maintains coherence in long dialogues) • API access (tools for model and agent development) |
| Microsoft | Phi-1 <ul style="list-style-type: none"> • Not multimodal (text only) • Fair reasoning (ie, limited to coding tasks) • Focused training (smaller, coding-focused data set) | Phi-4 <ul style="list-style-type: none"> • Multimodal (text, audio, and images) • Advanced reasoning (capable of multistep problem-solving and nuanced analysis) • Comprehensive training (diverse data) |
| OpenAI | GPT-3.5 <ul style="list-style-type: none"> • Not multimodal (text only) • Fair reasoning ability (eg, scored high on SAT, but bottom 10% on bar examination) • Limited contextual understanding (difficulty with coherence in complex conversations) • Standard API access (for text generation) | OpenAI o1 <ul style="list-style-type: none"> • Multimodal (text and images) • Advanced reasoning (eg, top 10% on bar examination) • Enhanced contextual understanding (maintains coherence in long dialogues) • Advanced API access (supports multimodal inputs) |

Note: Exhibit is not intended as an evaluation or comparison but as an illustration of the rapid progress in capabilities.

¹Initial models released between Mar 2022 and Mar 2023.

²Latest models released between Nov and Dec 2024.

Source: Company websites and press releases; McKinsey analysis

McKinsey & Company

Intelligence and reasoning are improving

AI is becoming far more intelligent. One indicator is the performance of LLMs on standardized tests. OpenAI's Chat GPT-3.5, introduced in 2022, demonstrated strong performance on high-school-level exams (for example, scoring in the 70th percentile on the SAT math and the 87th percentile on the SAT verbal sections). However, it often struggled with broader reasoning. Today's models are near the intelligence level of people who hold advanced degrees. GPT-4 can so easily pass the Uniform Bar Examination that it would rank in the top 10 percent of test takers,⁵ and it can answer 90 percent of questions correctly on the US Medical Licensing Examination.⁶

The advent of reasoning capabilities represents the next big leap forward for AI. Reasoning enhances AI's capacity for complex decision making, allowing models to move beyond basic comprehension to nuanced understanding and the ability to create step-by-step plans to achieve goals. For businesses, this means they can fine-tune reasoning models and integrate them with domain-specific knowledge to deliver actionable insights with greater accuracy. Models such as OpenAI's o1 or Google's Gemini 2.0 Flash Thinking Mode are capable of reasoning in their responses, which gives users a human-like thought partner for their interactions, not just an information retrieval and synthesis engine.⁷

Agentic AI is acting autonomously

The ability to reason is growing more and more, allowing models to autonomously take actions and complete complex tasks across workflows. This is a profound step forward. As an example, in 2023, an *AI bot* could support call center representatives by synthesizing and summarizing large volumes of data—including voice messages, text, and technical specifications—to suggest responses to customer queries. In 2025, an *AI agent* can converse with a customer and plan the actions it will take afterward—for example, processing a payment, checking for fraud, and completing a shipping action.

'I've always thought of AI as the most profound technology humanity is working on . . . more profound than fire or electricity or anything that we've done in the past.'

– Sundar Pichai, CEO of Alphabet

⁵ GPT-4 technical report, OpenAI, March 27, 2023.

⁶ Dana Brin, Vera Sorin, Akhil Vaid, et al., "Comparing ChatGPT and GPT-4 performance in USMLE soft skill assessments," *Scientific Reports*, October 1, 2023.

⁷ "Learning to reason with LLMs," OpenAI, September 12, 2024; "Gemini 2.09 Flash Thinking Mode," Google, January 21, 2025.

‘AI, like most transformative technologies, grows gradually, then arrives suddenly.’

– Reid Hoffman, cofounder of LinkedIn and Inflection AI, partner at Greylock Partners, and author

Software companies are embedding agentic AI capabilities into their core products. For example, Salesforce’s Agentforce is a new layer on its existing platform that enables users to easily build and deploy autonomous AI agents to handle complex tasks across workflows, such as simulating product launches and orchestrating marketing campaigns.⁸ Marc Benioff, Salesforce cofounder, chair, and CEO, describes this as providing a “digital workforce” where humans and automated agents work together to achieve customer outcomes.⁹

Multimodality is bringing together text, audio, and video

Today’s AI models are evolving toward more advanced and diverse data processing capabilities across text, audio, and video. Over the last two years, we have seen improvements in the quality of each modality. For example, Google’s Gemini Live has improved audio quality and latency and can now deliver a human-like conversation with emotional nuance and expressiveness.¹⁰ Also, demonstrations of Sora by OpenAI show its ability to translate text to video.¹¹

Hardware innovation is enhancing performance

Hardware innovation and the resulting increase in compute power continue to enhance AI performance. Specialized chips allow faster, larger, and more versatile models. Enterprises can now adopt AI solutions that require high processing power, enabling real-time applications and opportunities for scalability. For example, an e-commerce company could significantly improve customer service by implementing AI-driven chatbots that leverage advanced graphics processing units (GPUs) and tensor processing units (TPUs). Using distributed cloud computing, the company could ensure optimal performance during peak traffic periods. Integrating edge hardware, the company could deploy models that analyze photos of damaged products to more accurately process insurance claims.

⁸ Sammy Spiegel, “The future of AI agents: Top predictions and trends to watch in 2025,” Salesforce, December 2024.

⁹ Marc Benioff, “How the rise of new digital workers will lead to an unlimited age,” *Time*, November 25, 2024.

¹⁰ Ivan Solovyev and Shrestha Basu Mallick, “Gemini 2.0: Level up your apps with real-time multimodal interactions,” Google, December 23, 2024.

¹¹ “OpenAI releases AI video generator Sora but limits how it depicts people,” Associated Press, December 10, 2024.

Transparency is increasing

AI is gradually becoming less risky, but it still lacks greater transparency and explainability. Both are critical for improving AI safety and reducing the potential for bias, which are imperative for widescale enterprise deployment. There is still a long way to go, but new models and iterations are rapidly improving. Stanford University's Center for Research on Foundation Models (CRFM) reports significant advances in model performance. Its Transparency Index, which uses a scale of 1 to 100, shows that Anthropic's transparency score increased by 15 points to 51 and Amazon's more than tripled to 41 between October 2023 and May 2024.¹²

Beyond LLMs, other forms of AI and machine learning (ML) are improving explainability, allowing the outputs of models that support consequential decisions (for example, credit risk assessment) to be traced back to the data that informed them. In this way, critical systems can be tested and monitored on a near-constant basis for bias and other everyday harms that arise from model drift and shifting data inputs, which happens even in systems that were well calibrated before deployment.

All of this is crucial for detecting errors and ensuring compliance with regulations and company policies. Companies have improved [explainability practices](#) and built necessary checks and balances, but they must be prepared to evolve continuously to keep up with growing model capabilities.

Achieving AI superagency in the workplace is not simply about mastering technology. It is every bit as much about supporting people, creating processes, and managing governance. The next chapters explore the nontechnological factors that will help shape the deployment of AI in the workplace.

¹² "The Foundation Model Transparency Index," Stanford Center for Research on Foundation Models, May 2024.



Employees are ready for AI; now leaders must step up

‘People are using [AI] to create amazing things. If we could see what each of us can do 10 or 20 years in the future, it would astonish us today.’

– Sam Altman, cofounder and CEO of OpenAI

Employees will be the ones to make their organizations AI powerhouses. They are more ready to embrace AI in the workplace than business leaders imagine. They are more familiar with AI tools, they want more support and training, and they are more likely to believe AI will replace at least a third of their work in the near future. Now it's imperative that leaders step up. They have more permission space than they realize, so it's on them to be bold and capture the value of AI. Now.

Beyond the tipping point

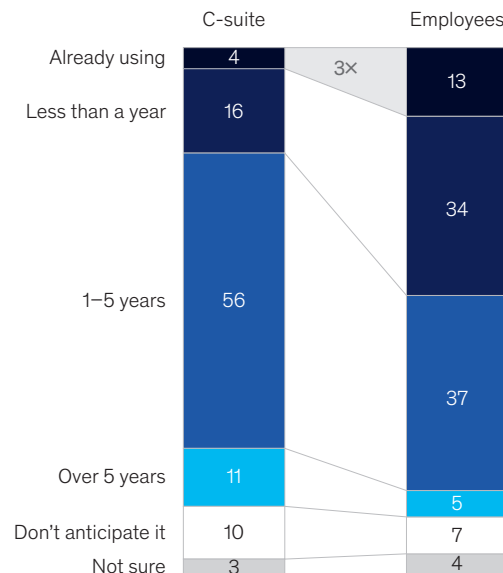
In our survey, nearly all employees (94 percent) and C-suite leaders (99 percent) report having some level of familiarity with gen AI tools. Nevertheless, business leaders underestimate how extensively their employees are using gen AI. C-suite leaders estimate that only 4 percent of employees use gen AI for at least 30 percent of their daily work, when in fact that percentage is three times greater, as self-reported by employees (Exhibit 2). And while only a total of 20 percent of leaders believe employees will use gen AI for more than 30 percent of their daily tasks within a year, employees are twice as likely (47 percent) to believe they will (see sidebar “Who is using AI at work? Nearly everyone, even skeptical employees”).

The good news is that our survey suggests three ways companies can accelerate AI adoption and move toward AI maturity.

Exhibit 2

Employees are three times more likely to be using gen AI today than their leaders expect.

US employees' and C-suite's timeline for employees using gen AI for >30% of daily tasks, % of respondents



Note: Figures may not sum to 100%, because of rounding.
 Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118); McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

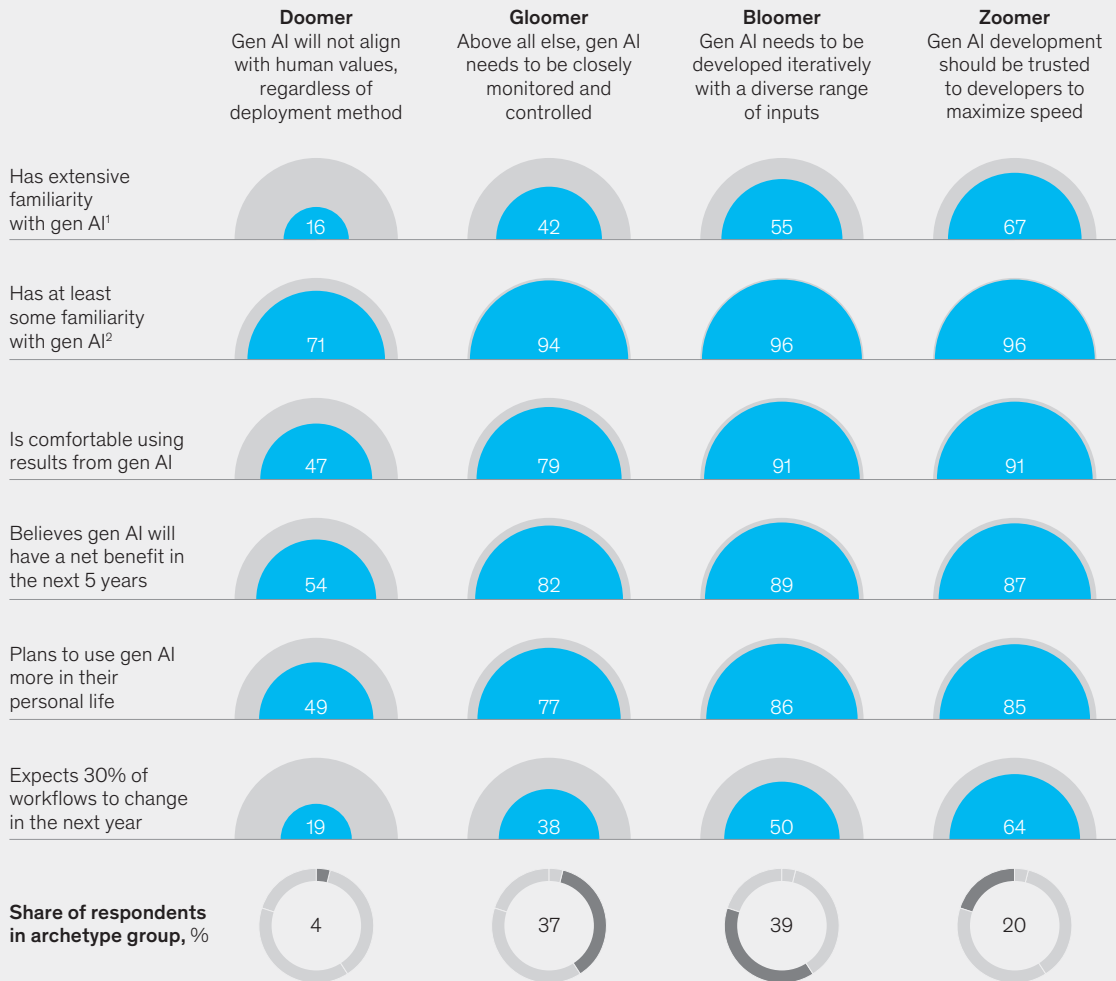
Our research looked at people who self-identify as “Zoomers,” “Bloomers,” “Gloomers,” and “Doomers” in their attitudes toward AI—a set of archetypes introduced in *Superagency*. We find that 39 percent of employees identify as Bloomers, who are AI optimists that want to collaborate with their companies to create responsible solutions. Meanwhile, 37 percent identify as Gloomers, who are more skeptical about AI and want extensive top-down AI regulations; 20 percent identify as Zoomers, who want AI to be quickly deployed with few guardrails; and just 4 percent identify as Doomers, who have a fundamentally negative view of AI (exhibit).

Even those with a skeptical take on AI are familiar with it; 94 percent of Gloomers and 71 percent of Doomers say they have some familiarity with gen AI tools. Furthermore, approximately 80 percent of Gloomers and about half of Doomers say they are comfortable using gen AI at work.

Exhibit

Employee segments differ, but all indicate a high familiarity with gen AI.

US employee sentiment on gen AI, by archetype, % of respondents



¹Defined as those who have “extensive experience (use several tools for complex tasks)” and “experts.”

²Defined as those who have “some familiarity (use 1–2 tools a few times)” and “extensive experience (use several tools for complex tasks)” and “experts.”

Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

Leaders can invest more in their employees


As noted at the beginning of this chapter, employees anticipate AI will have a dramatic impact on their work. Now they would like their companies to invest in the training that will help them succeed. Nearly half of employees in our survey say they want more formal training and believe it is the best way to boost AI adoption. They also would like access to AI tools in the form of betas or pilots, and they indicate that incentives such as financial rewards and recognition can improve uptake.

Yet employees are not getting the training and support they need. More than a fifth report that they have received minimal to no support (Exhibit 3). Outside the United States, employees also want more training (see sidebar “Global perspectives on training”).


Exhibit 3

Employees long for more support and training on gen AI.

Share of US employees agreeing that a company initiative would make them more likely to increase day-to-day usage of gen AI tools, %



US employees’ perceived level of support for gen AI capability building at their organizations, % of respondents



Note: Figures do not sum to 100%, because of rounding.
¹Objectives and key results.
 Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

McKinsey & Company

Global perspectives on training


To get a clearer picture of global AI adoption trends, we looked at trends across five countries: Australia, India, New Zealand, Singapore, and the United Kingdom. Broadly speaking, these employees and C-suite leaders—the “international” group in this report—have similar views of AI as their US peers. In some key areas, however, including the topic of training, their experiences differ.

Many international employees are concerned about insufficient training, even though they report receiving far more support than US employees. Some 84 percent of international employees say they receive significant or full organizational support to learn AI skills, versus just over half of US employees. International employees also have more opportunities to participate in developing gen AI tools at work than their US counterparts, with differences of at least ten percentage points in activities such as providing feedback, beta testing, and requesting specific features (exhibit).


Exhibit

International employees get more encouragement to use gen AI tools.

Sources encouraging employees’ use of gen AI tools at work, % of respondents reporting practice in place at their organization



Employee involvement in developing gen AI tools, % of respondents



Source: McKinsey international employee survey, Oct–Nov 2024 (Australia and New Zealand, n = 139; India, n = 134; Singapore, n = 140; UK, n = 201); McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

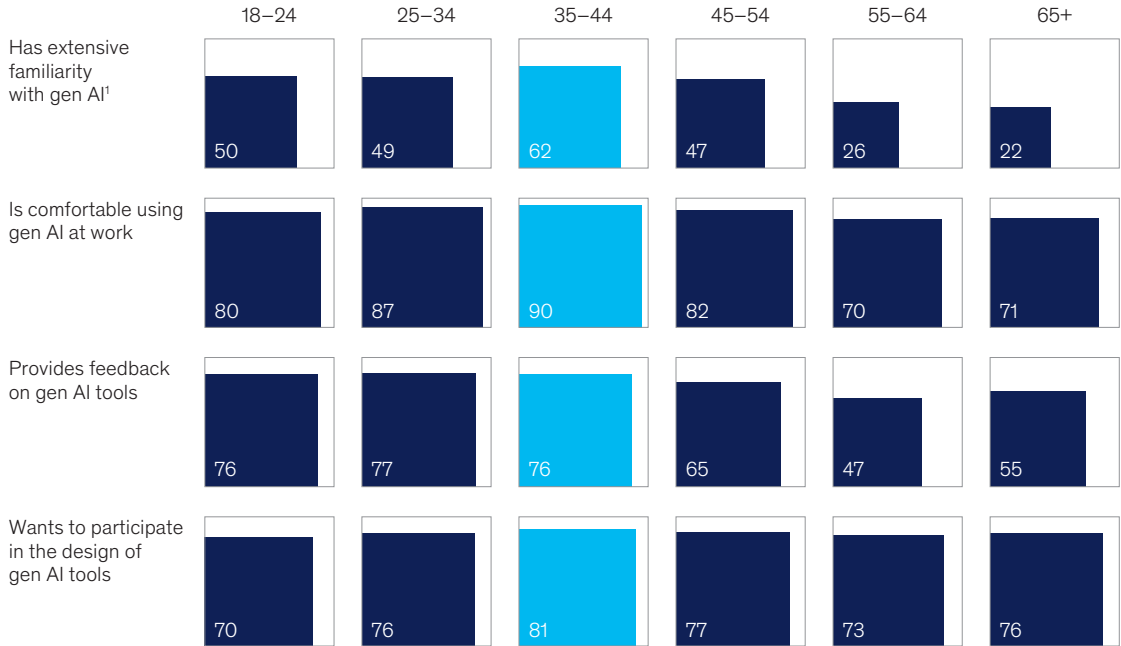
C-suite leaders can help millennials lead the way

Many millennials aged 35 to 44 are managers and team leaders in their companies. In our survey, they self-report having the most experience and enthusiasm about AI, making them natural champions of transformational change. Millennials are the most active generation of AI users. Some 62 percent of 35- to 44-year-old employees report high levels of expertise with AI, compared with 50 percent of 18- to 24-year-old Gen Zers and 22 percent of baby boomers over 65 (Exhibit 4). By tapping into that enthusiasm and expertise, leaders can help millennials play a crucial role in AI adoption.

Exhibit 4

Millennials aged 35 to 44 are AI optimists, with 90 percent indicating confidence in their gen AI abilities.

US employee sentiment on gen AI, by age group, % of respondents




¹Defined as those who have "extensive experience (use several tools for complex tasks)" and "experts."
Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

McKinsey & Company

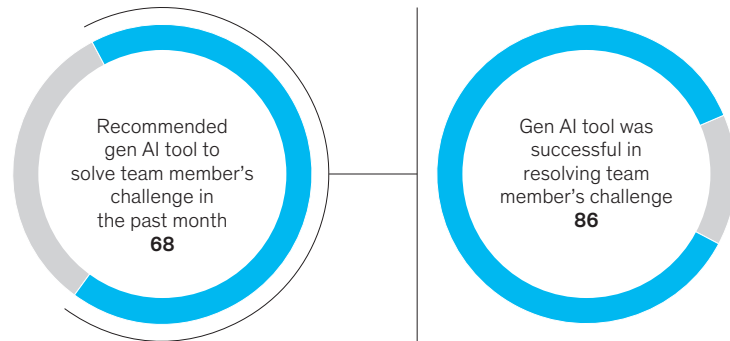
Exhibit 5

Two-thirds of managers regularly act as sounding boards for their teams on gen AI.

Frequency of team inquiries about using new gen AI tools at work, % of US manager respondents (n = 1,440)



Use of gen AI tools to resolve a team member's challenge, % of US manager respondents (n = 1,440)



Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

McKinsey & Company

Since many millennials are managers, they can support their teams to become more adept AI users. This helps push their companies toward AI maturity. Two-thirds of managers say they field questions from their team about how to use AI tools at least once a week, and a similar percentage say they recommend AI tools to their teams to solve problems (Exhibit 5).

Since leaders have the permission space, they can be bolder

In many transformations, employees are not ready for change, but AI is different. Employee readiness and familiarity are high, which gives business leaders the permission space to act. Leaders can listen to employees describe how they are using AI today and how they envision their work being transformed. They also can provide employees with much-needed training and empower managers to move AI use cases from pilot to scale.

It's critical that leaders meet this moment. It's the only way to accelerate the probability that their companies will reach AI maturity. But they must move with alacrity, or they will fall behind.



Delivering speed and safety

‘Soon after the first automobiles were on the road, there was the first car crash. But we didn’t ban cars—we adopted speed limits, safety standards, licensing requirements, drunk-driving laws, and other rules of the road.’

– Bill Gates, cofounder of Microsoft

A I technology is advancing at record speed. ChatGPT was released about two years ago; OpenAI reports that usage now exceeds 300 million weekly users¹³ and that over 90 percent of Fortune 500 companies employ its technology.¹⁴ The internet did not reach this level of usage until the early 2000s, nearly a decade after its inception.

The majority of employees describe themselves as AI optimists; Zoomers and Boomers make up 59 percent of the workplace. Even Gloomers, who are one of the two less-optimistic segments in our analysis, report high levels of gen AI familiarity, with over a quarter saying they plan to use AI more next year.

Business leaders need to embrace this speed and optimism to ensure that their companies don't get left behind. Yet despite all the excitement and early experimentation, 47 percent of C-suite leaders say their organizations are developing and releasing gen AI tools too slowly, citing talent skill gaps as a key reason for the delay (Exhibit 6).

Exhibit 6

Half of business leaders believe the development and release of gen AI tools is too slow in their organizations.

US C-suite categorization of the pace at which their organizations are developing and releasing gen AI tools, % of respondents



US C-suite's perceived top reason organizations are developing and releasing gen AI tools too slowly, % of respondents reporting development pace as "too slow"



Note: Figures may not sum to 100%, because of rounding.
Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118)

McKinsey & Company

¹³ Hayden Field, "OpenAI's active user count soars to 300 million people per week," CNBC, December 4, 2024.

¹⁴ Krystal Hu and Dawn Chmielewski, "OpenAI's Altman pitches ChatGPT Enterprise to large firms, including some Microsoft customers," Reuters, April 12, 2024.


Business leaders are trying to meet the need for speed by increasing investments in AI. Of the executives surveyed, 92 percent say they expect to boost spending on AI in the next three years, with 55 percent expecting investments to increase by at least 10 percent from current levels. But they can no longer just spend on AI without expecting results. As companies move on from the initial thrill of gen AI, business leaders face increasing pressure to generate ROI from their gen AI deployments.

We are at a turning point. The initial AI excitement may be waning, but the technology is accelerating. Bold and purposeful strategies are needed to set the stage for future success. Leaders are taking the first step: One quarter of those executives we surveyed have defined a gen AI road map, while just over half have a draft that is being refined (Exhibit 7). With technology changing this fast, all road maps and plans will evolve constantly. For leaders, the key is to make some clear choices about what valuable opportunities they choose to pursue first—and how they will work together with peers, teams, and partners to deliver that value.

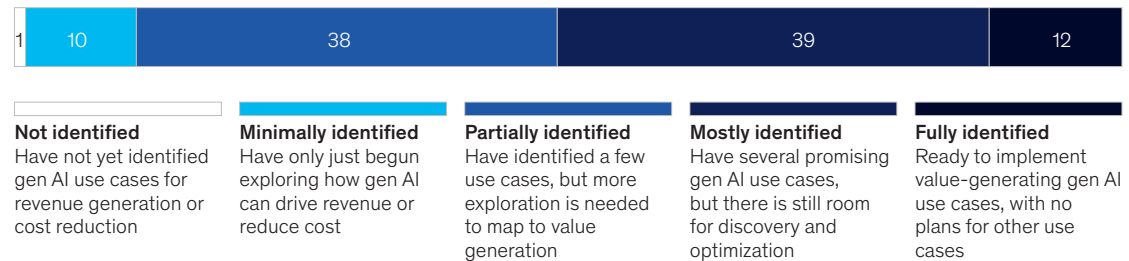
Exhibit 7

Most C-suite respondents have road maps to guide their gen AI strategies and have begun identifying use cases.

Presence of a defined gen AI road map, % of US C-suite respondents



Level of identifying revenue-generating use cases for gen AI, % of US C-suite respondents



Note: Figures may not sum to 100%, because of rounding.
Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118)

McKinsey & Company

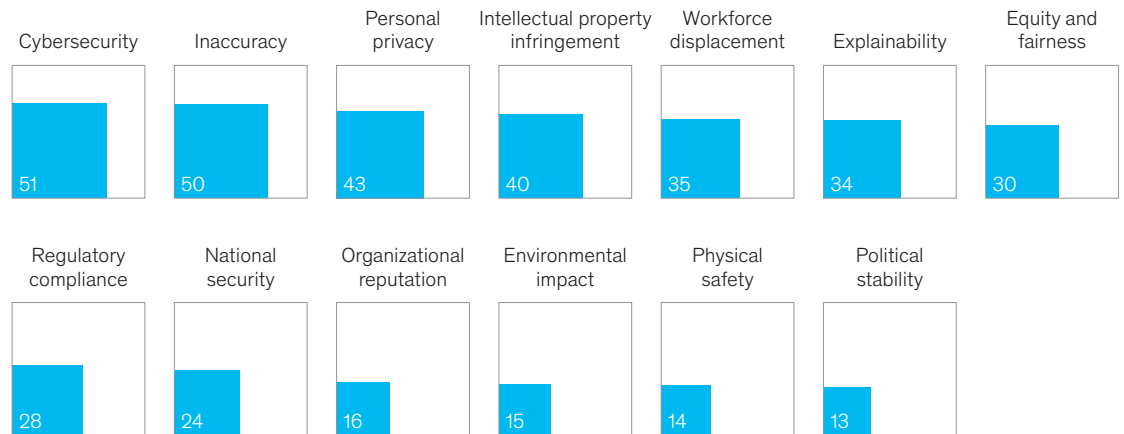
The dilemma of speed versus safety

There's a spanner in the works: Regulation and safety often continue to be seen as insurmountable challenges rather than opportunities. Leaders want to increase AI investments and accelerate development, but they wrestle with how to make AI safe in the workplace. Data security, hallucinations, biased outputs, and misuse (for example, creating harmful content or enabling fraud) are challenges that cannot be ignored. Employees are well aware of AI's safety challenges. Their top concerns are cybersecurity, privacy, and accuracy (Exhibit 8). But what will it take for leaders to address these concerns while also moving ahead at light speed?

Exhibit 8

Employees have concerns about gen AI—namely cybersecurity risks, inaccuracies, and data leaks.

Share of US employees citing risk as a concern associated with gen AI, %



Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

McKinsey & Company

Employees trust business leaders to get it right

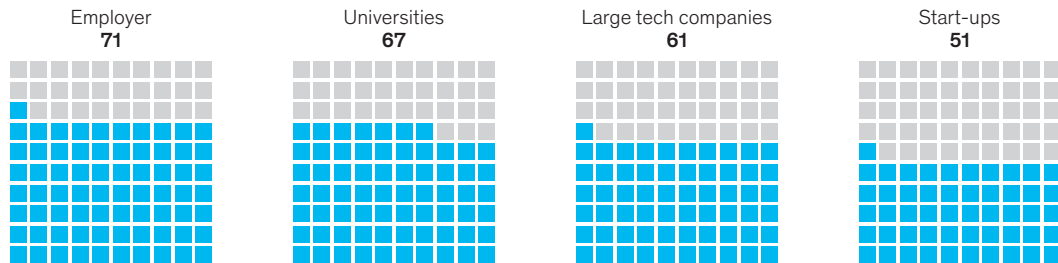
While employees acknowledge the risks and even the likelihood that AI may replace a considerable portion of their work, they place high trust in their own employers to deploy AI safely and ethically. Notably, 71 percent of employees trust their employers to act ethically as they develop AI. In fact, they trust their employers more than universities, large technology companies, and tech start-ups (Exhibit 9).

According to our research, this is in line with a broader trend in which employees show higher trust in their employers to do the right thing in general (73 percent) than in other institutions, including the government (45 percent). This trust should help leaders act with confidence as they tackle the speed-versus-safety dilemma. That confidence also applies outside the United States, even though employees in other regions may have more desire for regulation (see sidebar “Global perspectives on regulation”).

Exhibit 9

Employees trust their employers most for a safe rollout of gen AI.

Share of US employees who highly trust institution to deploy gen AI tools responsibly, safely, and ethically,¹%



¹High trust is defined as a 4 or 5 on a scale of 1 to 5.
Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

McKinsey & Company

Global perspectives on regulation


A high percentage of international C-suite leaders we surveyed across five regions (Australia, India, New Zealand, Singapore, and the United Kingdom) are Gloomers, who favor greater regulatory oversight. Between 37 to 50 percent of international C-suite leaders self-identify as Gloomers, versus 31 percent in the United States. This may be because top-down regulation is more accepted in many countries outside the United States. Of the global C-suite leaders surveyed, half or more worry that ethical use and data privacy issues are holding back their employees from adopting gen AI.

However, our research shows that attitudes about regulation are not inhibiting the economic expectations of business leaders outside the United States. More than half of the international executives (versus 41 percent of US executives) indicate they want their companies to be among the first adopters of AI, with those in India and Singapore being especially bullish (exhibit). The desire of international business leaders to be AI first movers can be explained by the revenue they expect from their AI deployments. Some 31 percent of international C-suite leaders say they expect AI to deliver a revenue uplift of more than 10 percent in the next three years, versus just 17 percent of US leaders. Indian executives are the most optimistic, with 55 percent expecting a revenue uplift of 10 percent or more over the next three years.

Exhibit

Half of international C-suite respondents want to be early adopters.

C-suite sentiment on adoption of new technologies, % of respondents



Source: McKinsey international employee survey, Oct–Nov 2024 (Australia and New Zealand, n = 139; India, n = 134; Singapore, n = 140; UK, n = 201); McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

McKinsey & Company

Risk management for gen AI

In *Superagency*, Hoffman argues that new risks naturally accompany new capabilities—meaning they should be managed but not necessarily eliminated.¹⁵ Leaders need to contend with external threats, such as infringement on intellectual property (IP), AI-enabled malware, and internal threats that arise from the AI adoption process. The first step in [building fit-for-purpose risk management](#) is to launch a comprehensive assessment to identify potential vulnerabilities in each of a company’s businesses. Leaders can then establish a robust governance structure, implement real-time monitoring and control mechanisms, and ensure continuous training and adherence to regulatory requirements.

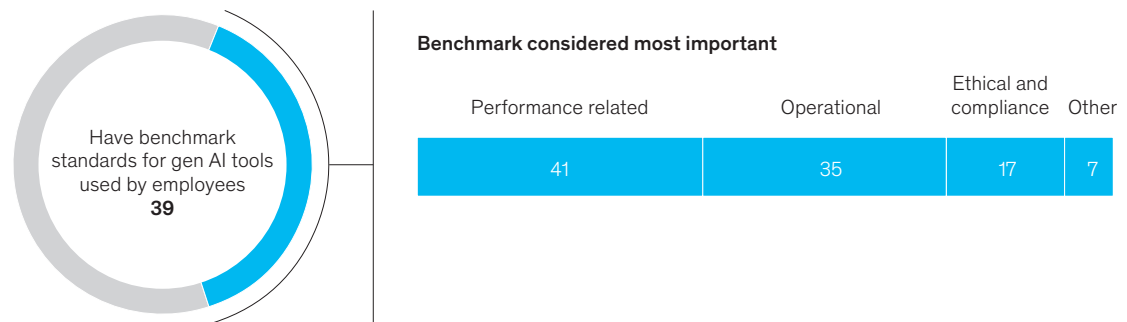
One powerful control mechanism is respected third-party benchmarking that can increase AI safety and trust. Examples include Stanford CRFM’s Holistic Evaluation of Language Models (HELM) initiative—which offers comprehensive benchmarks to assess the fairness, accountability, transparency, and broader societal impact of a company’s AI systems—as well as MLCommons’s AILuminate tool kit on which researchers from Stanford collaborated.¹⁶ Other organizations such as the Data & Trust Alliance unite large companies to create cross-industry metadata standards that aim to bring more transparency to enterprise AI models.

While benchmarks have significant potential to build trust, our survey shows that only 39 percent of C-suite leaders use them to evaluate their AI systems. Furthermore, when leaders do use benchmarks, they opt to measure operational metrics (for example, scalability, reliability, robustness, and cost efficiency) and performance-related metrics (including accuracy, precision, F1 score, latency, and throughput). These benchmarking efforts tend to be less focused on ethical and compliance concerns: Only 17 percent of C-suite leaders who benchmark say it’s most important to measure fairness, bias, transparency, privacy, and regulatory issues (Exhibit 10).

Exhibit 10

More than a third of C-suite respondents use benchmarks for gen AI, but with less focus on ethical metrics.

US C-suite use of benchmarking for gen AI tools, % of respondents



Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118)

McKinsey & Company

¹⁵ Reid Hoffman and Greg Beato, *Superagency: What Could Possibly Go Right with Our AI Future*, Authors Equity, January 28, 2025.

¹⁶ “MLCommons launches AILuminate, first-of-its-kind benchmark to measure the safety of large language models,” Business Wire, press release, December 04, 2024.

The focus on operational and performance metrics reflects the understandable desire to prioritize immediate technical and business outcomes. But ignoring ethical considerations can come back to haunt leaders. When employees don't trust AI systems, they are less likely to accept them. Although benchmarks are not a panacea to eliminate all risk and can't ensure that AI systems are fully efficient, ethical, and safe, they are a useful tool.

Even companies that excel at all three categories of AI readiness—technology, employees, and safety—are not necessarily scaling or delivering the value expected. Nevertheless, leaders can harness the power of big ambitions to transform their companies with AI. The next chapter examines how.

4

Embracing bigger ambitions

‘This is a time when you should be getting benefits [from AI] and hope that your competitors are just playing around and experimenting.’

– Erik Brynjolfsson, professor at Stanford University and director of the Digital Economy Lab at the Stanford Institute for Human-Centered Artificial Intelligence (HAI)

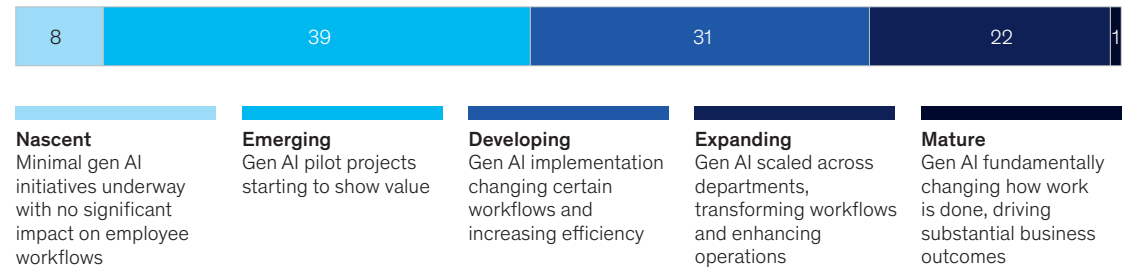
Most organizations that have invested in AI are not getting the returns they had hoped. They are not winning the full economic potential of AI. About half of C-suite leaders at companies that have deployed AI describe their initiatives as still developing or expanding (Exhibit 11). They have had the time to move further. Our research shows that more than two-thirds of leaders launched their first gen AI use cases over a year ago.

Pilots fail to scale for many reasons. Common culprits are poorly designed or executed strategies, but a lack of bold ambitions can be just as crippling. This chapter looks at patterns governing today’s investments in AI across industries and suggests the potential awaiting those who can dream bigger.

Exhibit 11

Only 1 percent of C-suite respondents describe their gen AI rollouts as mature.

US C-suite’s categorization of organizations’ level of gen AI maturity, % of respondents



Note: Figures may not sum to 100%, because of rounding.
Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118)


McKinsey & Company

AI investments vary by industry

Different industries have different AI investment patterns. Within the top 25 percent of spenders, companies in healthcare, technology, media and telecom, advanced industries, and agriculture are ahead of the pack (Exhibit 12). Companies in financial services, energy and materials, consumer goods and retail, hardware engineering and construction, and travel, transport, and logistics are spending less. The consumer industry—despite boasting the second-highest potential for value realization from AI—seems least willing to invest, with only 7 percent of respondents qualifying in the top quartile, based on self-reported percentage of revenue spend on gen AI. That hesitation may be explained by the industry’s low average net margins in mass-market categories and thus higher confidence thresholds for adopting costly organization-wide technology upgrades.

Companies' gen AI spend does not match the economic potential in their industries.

Representation in top quartile of gen AI spending vs overall survey representation



Note: Not all industries represented in survey are shown; only industries in top quartile based on self-reported gen AI spend as % of revenue are represented. Economic potential was determined using the following mapping of industries listed in the McKinsey report *The economic potential of generative AI*, (advanced industries: advanced manufacturing), (technology: high tech), (financial services: banking), (hardware engineering and construction: advanced electronics and semiconductors and construction).

¹Includes biotechnology, healthcare equipment and services, life sciences, and pharmaceuticals.

²Includes aerospace and defense and automotive and assembly.

Source: *The economic potential of generative AI: The next productivity frontier*, McKinsey, June 14, 2023; McKinsey US CxO survey, Oct–Nov 2024 (n = 118)

McKinsey & Company

In some industries, employees are cautious


Employees in the public sector, as well as the aerospace and defense and semiconductor industries, are largely skeptical about the development of AI's future. In the public sector and aerospace and defense, only 20 percent of employees anticipate that AI will have a significant impact on their daily tasks in the next year, versus roughly two-thirds in media and entertainment (65 percent) and telecom, at 67 percent (Exhibit 13). What's more, our survey shows that just 31 percent of social sector employees trust that their employers will develop AI safely. That's the least confidence in any industry; the cross-industry average is 71 percent.

Exhibit 13

Employees in the social sector, aerospace and defense, and public sector are least optimistic about gen AI.

US employee sentiment on gen AI, by sector, % of respondents

0% 100%



Note: Level of familiarity is defined as those who have “extensive experience (use several tools for complex tasks)” and “experts.” High trust is “Level 4” and “Level 5” on a scale of 1 to 5. Perceived accuracy is based on past gen AI usage in a workplace setting.
 Source: McKinsey US employee survey, Oct–Nov 2024 (n = 3,002); *The economic potential of generative AI: The next productivity frontier*, McKinsey, June 14, 2023

Employees' relative caution about AI in these sectors likely reflects near-term challenges posed by external constraints such as rigorous regulatory oversight, outdated IT systems, and lengthy approval processes.

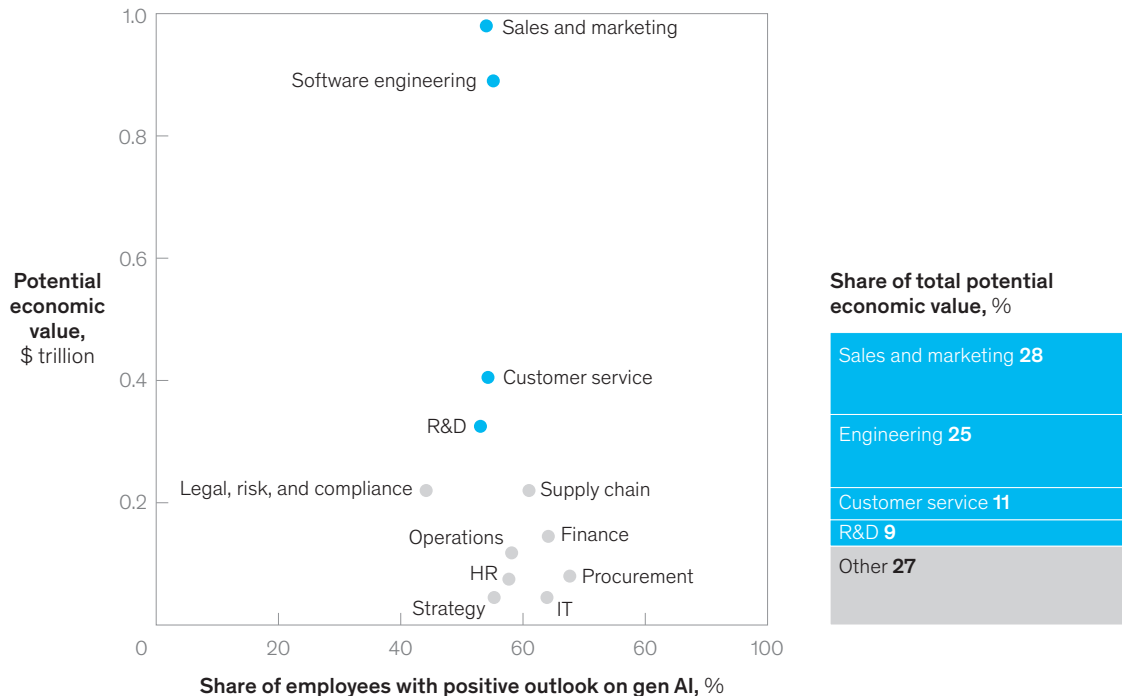
There's a lot of headroom in some functions

Our research finds that the functional areas where AI presents the greatest economic potential are also those where employee outlook is lukewarm. Employees in sales and marketing, software engineering, customer service, and R&D contribute roughly three-quarters of AI's total economic potential, but the self-reported optimism of employees in these functions is middling (Exhibit 14). It may be the case that these functions have piloted AI projects, leading employees to be more realistic about AI's benefits and limitations. Or perhaps the economic potential has made them worry that AI could replace their jobs. Whatever the reasons, leaders in these functions might consider investing more in employee support and elevating the change champions who can improve that sentiment.

Exhibit 14

The employees most optimistic about gen AI do not represent the most economic value potential.

Potential economic value from gen AI and share of employees with a positive outlook on gen AI,¹ by function



Note: For AI outlook data from survey, strategy function includes business development. Economic potential of functions was determined by grouping corporate use cases listed in the McKinsey report *The economic potential of generative AI* (eg, sales and marketing, legal and risk and compliance) others were aligned as follows: (customer service: customer operations), (procurement: procurement and management), (HR: talent and organization, including HR), (IT: corporate IT).
¹Optimism defined as share of Bloomer (believes in gen AI's potential and supports iterative deployment) and Zoomer (believes in speedy gen AI deployment by technologists) respondents.
 Source: *The economic potential of generative AI: The next productivity frontier*, McKinsey, June 14, 2023; McKinsey US employee survey, Oct–Nov 2024 (n = 3,002)

Gen AI has not delivered enterprise-wide ROI, but that can change


Across all industries, surveyed C-level executives report limited returns on enterprise-wide AI investments. Only 19 percent say revenues have increased more than 5 percent, with another 39 percent seeing a moderate increase of 1 to 5 percent, and 36 percent reporting no change (Exhibit 15). And only 23 percent see AI delivering any favorable change in costs.

Despite this, company leaders are optimistic about the value they can capture in the coming years. A full 87 percent of executives expect revenue growth from gen AI within the next three years, and about half say it could boost revenues by more than 5 percent in that time frame (Exhibit 16). That suggests quite a lot could change for the better over the next few years.

Exhibit 15

Gen AI has not yet delivered significant return on investment for enterprises.

US C-suite's perception on how gen AI has affected revenues and costs, % of respondents




Note: Figures may not sum to 100%, because of rounding.
Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118)

McKinsey & Company

Exhibit 16

Half of C-suite respondents expect gen AI to deliver more than 5 percent revenue growth in the next three years.

US C-suite's perception on how gen AI will affected revenue, next 3 years, % of respondents



Source: McKinsey US CxO survey, Oct–Nov 2024 (n = 118)

McKinsey & Company

Big ambitions can help solve big problems

To drive revenue growth and improve ROI, business leaders may need to commit to transformative AI possibilities. As the hype around AI subsidies and the focus shifts to value, there is a heightened attention on practical applications that can create competitive moats.

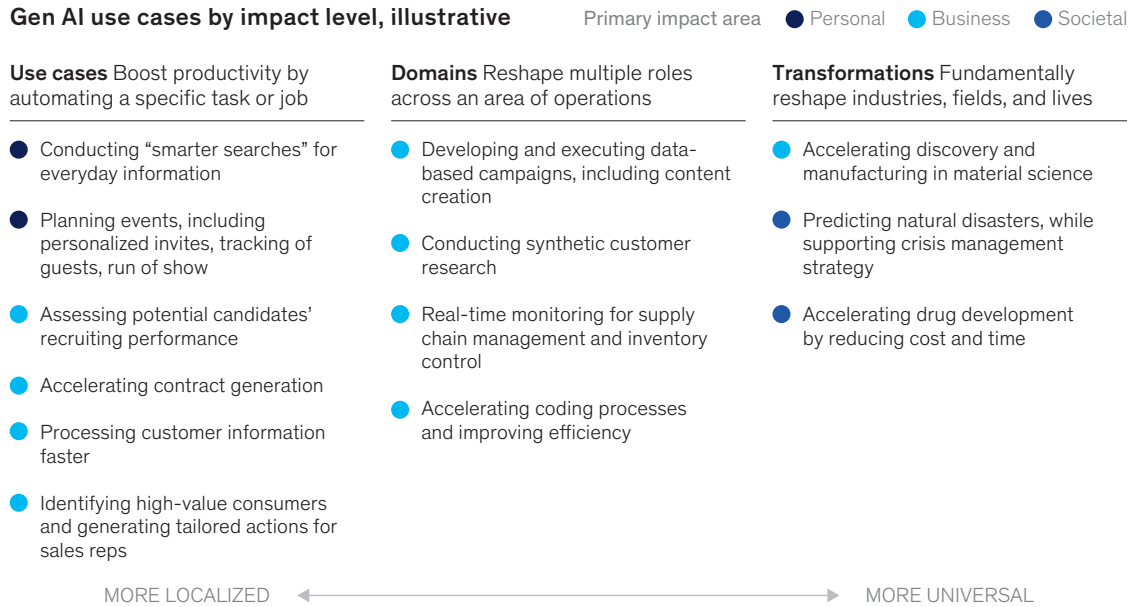
To assess how far along companies are in this shift, we examined three categories of AI applications: personal use, business use, and societal use (see sidebar “AI’s potential to enhance our personal lives”). We mapped over 250 applications from our work and publicly shared examples to understand the spectrum of impact levels, from localized use cases to transformations with more universal impact. Our conclusion? Given that most companies are early in their AI journeys, most AI applications are localized use cases still in the pilot stages (Exhibit 17).

‘[It] is critical to have a genuinely inspiring vision of the future [with AI] and not just a plan to fight fires.’

– Dario Amodei, cofounder and CEO of Anthropic

Exhibit 17

Over the past two years, personal and business gen AI applications have often focused on localized impact.



McKinsey & Company

‘It is in [the] collaboration between people and algorithms that incredible scientific progress lies over the next few decades.’

– Demis Hassabis, cofounder and CEO of Google DeepMind

In many cases, that’s perfectly appropriate. But creating AI applications that can revolutionize industries and create transformative value requires something more. Robotics in manufacturing, predictive AI in renewable energy, drug development in life sciences, and personalized AI tutors in education—these are the kinds of transformative efforts that can drive the greatest returns.¹⁷ These weren’t created from a reactive mindset. They are the result of inspirational leadership, a unique concept of the future, and a commitment to transformational impact. This is the kind of courage needed to develop AI applications that can revolutionize industries.

¹⁷ Dario Amodei, “Machines of loving grace: How AI could transform the world for the better,” Dario Amodei website, October 2024.

To truly harness the potential of AI, companies must challenge themselves to envision and implement more breakthrough initiatives. Success in the era of AI hinges not just on technology deployment or employee willingness but also on visionary leadership. The ingredients are here. The technology is already highly capable and rapidly advancing, and employees are more ready than leaders think. Leaders have more permission space than they realize to deploy AI quickly in the workplace. To do so, leaders need to stretch their ambitions toward systematic change, laying the foundation for real competitive differentiation. If they want to be more ambitious about AI, companies must increase the proportion of transformational initiatives in their portfolios. The next chapter examines the headwinds that leaders must overcome—and how they can do so.

AI's potential to enhance our personal lives

Outside of the business context, individuals are increasingly using AI in their personal lives. In previous research, we analyzed the potential impact of AI across 77 personal activities and across age, gender, and working status in the United States. While individuals have limited desire to automate certain personal activities, including leisure, sleeping, and fitness, the data shows significant opportunity for AI combined with other technologies to help with chores or labor-intensive tasks. Already in 2024, our research identified about an hour of such daily activities with the technical potential to be automated. By 2030, expansion of use cases and continued improvements in AI safety could increase automation potential up to three hours per day. When people use AI-enabled tools—say, an autonomous vehicle for transportation or an interactive personal finance bot—they can repurpose time for personal fulfillment activities or being productive in other ways.

Using human-centric design and tapping into gen AI's potential for “emotional intelligence” are unlocking new personal AI applications that go beyond basic efficiencies. Individuals are beginning to use conversational and reasoning AI models for counseling, coaching, and creative expression. For example, people are using conversational AI for advice and emotional support or to bring their artistic visions to life with only verbal cues. Further, to the notion that AI superagency will advance society, AI has potential to become a democratizing force, making experiences that were previously expensive or exclusive—such as animation generation, career coaching, or tax advice—available to much wider audiences.

5

Technology is not the barrier to scale

‘As we build this next generation of AI, we made a conscious design choice to put human agency both at a premium and at the center of the product. For the first time, we have the access to AI that is as empowering as it is powerful.’

– Satya Nadella, chairman and CEO of Microsoft

There is no question: AI offers a rare and phenomenal opportunity. Almost 90 percent of leaders anticipate that deploying AI will drive revenue growth in the next three years. But securing that growth entails corporate transformation, and businesses have a poor track record in this area. Nearly [70 percent of transformations fail](#).

To make their companies part of the minority that succeed, C-level executives must turn the mirror on themselves. They need to embrace the vital role their leadership plays. C-suite leaders participating in our survey are more than twice as likely to say employee readiness is a barrier to adoption as they are to blame their own role. But as previously noted, employees indicate that they are quite ready.

This chapter looks at how leaders can take the reins, recognizing and owning the fact that the AI opportunity requires more than technology implementation. It demands a strategic transformation. There is no denying that companies face a set of AI headwinds. To tackle these challenges, leadership teams will need to commit to [rewiring their enterprises](#).

The operational headwinds that slow execution

Business adoption of AI faces several operational headwinds. Our interviews and research surfaced five that are most challenging: aligning leadership, addressing cost uncertainty, workforce planning, managing supply chain dependencies, and meeting the demand for explainability.

Leadership alignment is a challenging but critical first step

Securing consensus from senior leaders on a strategy-led gen AI road map is no simple task. The key to meeting this challenge is first recognizing that leadership alignment cannot be oversimplified or assumed. The process requires ongoing engagement from senior leaders across business domains, each of which may have distinct objectives and risk appetites. Together, leaders must clearly define where value lies, how AI will drive this value, and how risk will be mitigated. They must collectively establish metrics for performance evaluation and investment recalibration. To facilitate alignment, they may want to appoint a gen AI value and risk leader or institute an enterprise-wide leadership and orchestration function. These actions can enhance collaboration among business, technology, and risk teams. Although challenging, aligning leadership is a crucial step to ensure that AI projects are not disparate, avoid liability, and deliver transformative business outcomes.

Cost uncertainty makes it difficult for enterprises to predict ROI

Many companies are still determining if they can “take” AI solutions off the shelf from tech vendors or if they need to “shape” and customize them, which can be more costly but brings the potential for greater differentiation from competitors. Additionally, while leaders can budget for AI pilots, the full cost of building and managing AI applications at scale remains uncertain. Planning for a limited pilot is very different from assessing the costs of a mature solution that helps most employees multiple times a day. These factors lead to tough tradeoffs. But to move at the pace of AI, technology leaders must prioritize accelerated decision-making.

Workforce planning is more difficult than ever

There is still a world of uncertainty to manage. Employers do not know how many AI experts they will need with what type of skills, whether that talent bench even exists, how quickly they can source people, and how they can remain an attractive employer for in-demand hires after they come aboard. On the other hand, they do not know how fast AI may depress demand for other skills and thus require workforce rebalancing and retraining.

Supply chain dependencies can wreak havoc

Fragile supply chains can expose enterprises to disruptions and technical, regulatory, and legal challenges. The AI supply chain is global, with significant R&D concentrated in China, Europe, and North America and with semiconductor and hardware manufacturing concentrated in East Asia and the United States. Today's geopolitics are complex. Furthermore, models and applications are increasingly created in open-source forums spanning many countries.

Demand for greater explainability is a central challenge

Safe AI deployment is increasingly a must-have. Yet most LLMs are often black boxes that do not reveal why or how they came to a certain response, nor what data was used to make it. If AI models cannot provide clear justifications for their responses, recommendations, decisions, or actions—showing the specific factors that led to a credit card application denial, for example—they will not be trusted for critical tasks.

These AI-specific headwinds are formidable but addressable. Companies are pushing ahead. For example, they might use dynamic cost planning or look at procuring NVIDIA clusters to secure the infrastructure they expect to need.¹⁸ Chief HR officers (CHROs) are developing training programs to upskill their current workforces and support some employees in job transitions. But lasting success will take more than that.

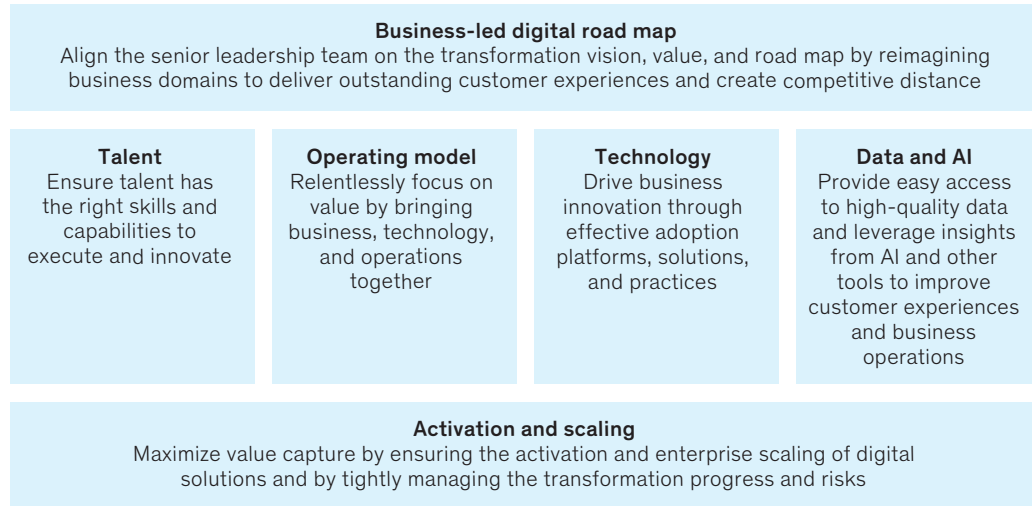
To capture AI value, leaders must rewire their companies

McKinsey's [Rewired](#) framework includes six foundational elements to guide sustained digital transformation: road map, talent, operating model, technology, data, and scaling (Exhibit 18). When companies implement this playbook successfully, they cultivate a culture of autonomy, leverage modern cloud practices, and assemble multidisciplinary agile teams.

¹⁸ "AWS and NVIDIA announce strategic collaboration to offer new supercomputing infrastructure, software and services for generative AI," NVIDIA, press release, November 28, 2023.

Six enduring success factors enable tech-business transformations.

Framework for the coordinated execution of value creation



McKinsey & Company

While these six elements are universally applicable, AI has introduced a few important wrinkles for leaders to address:

- *Adaptability.* AI technology is advancing so rapidly that organizations must adopt new best practices quickly to stay ahead of the competition. Best practices may come in the form of new technologies, talent, business models, or products. For example, a modular approach helps future-proof tech stacks. As natural language becomes a medium for integration, AI systems are becoming more compatible, allowing businesses to swap, upgrade, and integrate models and tools with less friction. This modularity allows enterprises to avoid vendor lock-in and put new AI advancements to use quickly without constantly reinventing their tech stacks.
- *Federated governance models.* Managing data and models can give teams autonomy to develop new AI tools while centrally controlling risk. Leaders can directly oversee high-risk or high-visibility issues, such as setting policies and processes to monitor models and outputs for fairness, safety, and explainability. But they can set direction and delegate other monitoring to business units, including measuring performance-based criteria such as accuracy, speed, and scalability.

- *Budget agility.* Given technological advances across models, as well as the opportunity to curate an optimal mix of LLMs, small language models (SLMs), and agents, business leaders should keep their budgets flexible. This helps enterprises optimize their AI deployments simultaneously for costs and performance.
- *AI benchmarks.* These tools can serve as powerful means to quantitatively assess, compare, and improve the performance of different AI models, algorithms, and systems. If technologists come together to adopt standardized public benchmarks—and if more C-level executives start employing benchmarks, including ethical ones—model transparency and accountability will improve and AI adoption will increase, even among more skeptical employees.
- *AI-specific skill gaps.* Notably, 46 percent of leaders identify skill gaps in their workforces as a significant barrier to AI adoption. Leaders will need to attract and hire top-level talent, including AI/ML engineers, data scientists, and AI integration specialists. They will also need to commit to creating an environment that is attractive to technologists. For example, this can mean providing them with plenty of time to experiment, offering access to cutting-edge tools, creating opportunities to engage in open-source communities, and promoting a collaborative engineering culture. Upskilling existing employees is just as critical: [Research](#) from McKinsey's People and Organizational Performance Practice underscores the importance of tailoring training to specific roles, such as offering technical team members bootcamps on library creation while offering prompt engineering classes to specific functional teams.¹⁹
- *Human centrality.* To guarantee both fairness and impartiality, it is important that business leaders incorporate diverse perspectives early and often in the AI development process and maintain transparent communication with their teams. As it stands, less than half of C-suite leaders (48 percent) say they would involve nontechnical employees in the early development stages of AI tools, specifically ideation and requirement gathering. Agile pods and human-centric development practices such as design-thinking and reinforcement learning from human feedback (RLHF) will help leaders and developers create AI solutions that *all* people want to use. In agile pods, technical team members sit alongside employees from business functions such as HR, sales, and product, and from support functions such as legal and compliance. Further, leaders can empathize with employees' uneasiness about AI's impacts on potential job losses by being honest about new skill requirements and head count changes. Forums where employees can provide input on AI applications, voice concerns, and share ideas are valuable for maintaining a transparent, human-first culture.

¹⁹ *People & Organization Blog*, "[Upskilling and reskilling priorities for the gen AI era](#)," blog entry by Sandra Durth, Kiera Jones, Lisa Christensen, and Naveed Rashid, McKinsey, September 30, 2024.

Conclusion: Meeting the AI future

‘Learn from
yesterday, live
for today, hope
for tomorrow.’

– Albert Einstein,
theoretical physicist

The pace at which AI has advanced over the last two years is stunning. Some react to that pace by seeing AI as a challenge to humanity. But what if we take the advice of Reid Hoffman and imagine what could possibly go *right* with AI? Leaders might realize that all the pieces are in place for AI superagency in the workplace.

They might notice that their employees are already using AI and want to use it even more. They may find that millennial managers are powerful change champions ready to encourage their peers. Instead of focusing on the 92 million jobs expected to be displaced by 2030, leaders could plan for the projected 170 million new ones and the new skills those will require.²⁰

This is the moment for leaders to set bold AI commitments and to meet employee needs with on-the-job training and human-centric development. As leaders and employees work together to reimagine their businesses from the bottom up, AI can evolve from a productivity enhancer into a transformative superpower—an effective partner that increases human agency. Leaders who can replace fear of uncertainty with imagination of possibility will discover new applications for AI, not only as a tool to optimize existing workflows but also as a catalyst to solve bigger business and human challenges. Early stages of AI experimentation focused on proving technical feasibility through narrow use cases, such as automating routine tasks. Now the horizon has shifted: AI is poised to unlock unprecedented innovation and drive systemic change that delivers real value.

To meet this more ambitious era, leaders and employees must ask themselves big questions. How should leaders define their strategic priorities and steer their companies effectively amid disruption? How can employees ensure they are ready for the AI transition coming to their workplaces? Questions like the following ones will shape a company's AI future:

²⁰ *Future of jobs report 2025*, World Economic Forum, January 7, 2025.

— For business leaders:

- *Is your strategy ambitious enough?* Do you want to transform your whole business? How can you reimagine traditional cost centers as value-driven functions? How do you gain a competitive advantage by investing in AI?
- *What does successful AI adoption look like for your organization?* What success indicators will you use to evaluate whether your investments are yielding desired ROI?
- *What skills define an AI-native workforce?* How can you create opportunities for employees to develop these skills on the job?

— For employees:

- *What does achieving AI mastery mean for you?* Does it extend to confidently using AI for personal productivity tasks such as research, planning, and brainstorming?
- *How do you plan to expand your understanding of AI?* Which news sources, podcasts, and video channels can you follow to remain informed about the rapid evolution of AI?
- *How can you rethink your own work?* Some of the most innovative ideas often emerge from within teams, rather than being handed down from leadership. How would you redesign your work to drive bottom-up innovation?

These questions have no easy answers, but a consensus is emerging on how to best address them. For example, some companies deploy both bottom-up and top-down approaches to drive AI adoption. Bottom-up actions help employees experiment with AI tools through initiatives such as hackathons and learning sessions. Top-down techniques bring executives together to radically rethink how AI could improve major processes such as fraud management, customer experience, and product testing.

These kinds of actions are critical as companies seek to move from AI pilots to AI maturity. Today only 1 percent of business leaders report that their companies have reached maturity. Over the next three years, as investments in the technology grow, leaders must drive that percentage way up. They should make the most of their employees' readiness to increase the pace of AI implementation while ensuring trust, safety, and transparency. The goal is simple: capture the enormous potential of gen AI to drive innovation and create real business value.

Acknowledgments

Hannah Mayer is a partner in McKinsey's Bay Area office, where [Lareina Yee](#) is a senior partner, [Michael Chui](#) is a knowledge developer and senior fellow, and [Roger Roberts](#) is a partner.

The authors wish to thank [Alex Panas](#), a senior partner in the Boston office; [Eric Kutcher](#), a senior partner in the Bay Area office; [Kate Smaje](#), a senior partner in the London office; [Noshir Kaka](#), a senior partner in the Mumbai office; [Robert Levin](#), a senior partner in the Boston office; and [Rodney Zimmel](#), a senior partner in the New York office, for their contributions to this report.

The authors were inspired by the impact delivered by our QuantumBlack, AI by McKinsey, colleagues, led by [Alex Singla](#) and [Alex Sukharevsky](#), and our gen AI lab leaders, especially [Carlo Giovine](#) and Stephen Xu.

The research was led by consultants Akshat Gokhale, Amita Mahajan, Begum Ortaoglu, Estee Chen, Hailey Bobsein, Katharina Giebel, Mallika Jhamb, Noah Furlonge-Walker, and Sabrina Shin. This report was edited by executive editors Kristi Essick and Rick Tetzeli.

Thank you to Reid Hoffman, along with his chief of staff Aria Finger and representatives at [Superagency: What Could Possibly Go Right with Our AI Future](#) publisher Authors Equity, especially Deron Triff, for their ongoing collaboration. Working together with Hoffman, who brings the distinctive perspective of being both an investor in and a mentor to the creators of AI, we looked at a central question: How can businesses win with AI in the medium and long terms? We benefited from working sessions with Hoffman, CEOs, and AI industry thought leaders.

We would like to thank the members of the Stanford Institute for Human-Centered Artificial Intelligence (HAI) who challenged our thinking and provided valuable feedback.

This report contributes to McKinsey's ongoing research on AI and aims to help business leaders understand the forces transforming ways of working, identify strategic impact areas, and prepare for the next wave of growth. As with all McKinsey research, this work is independent and has not been commissioned or sponsored in any way by any business, government, or other institution. The report and views expressed here are ours alone. We welcome your comments on this research at SuperagencyReport@McKinsey.com. Learn more about [our gen AI insights](#) and sign up for our [newsletter](#).

Methodology

To create our report, we surveyed 3,613 employees (managers and independent contributors) and 238 C-level executives in October and November 2024. Of these, 81 percent came from the United States, and the rest came from five other countries: Australia, India, New Zealand, Singapore, and the United Kingdom. The employees spanned many roles, including business development, finance, marketing, product management, sales, and technology.

All the survey findings discussed in the report, aside from two sidebars presenting international nuances, pertain solely to US workplaces. The findings are organized in this way because the responses from US employees and C-suite executives provide statistically significant conclusions about the US workplace. Analyzing global findings separately allows a comparison of differences between US responses and those from other regions.

Due to rounding, percentages may not always sum to 100 percent.

Three-quarters of survey respondents in the United States work at organizations generating at least \$100 million in annual revenue, and half work at companies with annual revenues exceeding \$1 billion. All US C-suite leader respondents work at organizations with annual revenues of at least \$1 billion. Looking at workforce size, 20 percent of US respondents work at companies with fewer than 10,000 employees, 49 percent work at companies with between 10,000 and 50,000, and 31 percent work at companies with more than 50,000.

The analysis extended far beyond surveys. The researchers also conducted interviews with dozens of C-level executives and industry experts to understand their perspectives on AI's transformative potential and the steps they are taking to lead their organizations through this transition. The report was further enriched by discussions with experts at Stanford HAI, the Digital Economy Lab at HAI, and McKinsey's leading AI experts. Our survey and research primarily focus on gen AI; however, it is important to note that participants in the survey may not have consistently differentiated between gen AI and other forms of AI.

Additionally, we developed a comprehensive database featuring more than 250 potential AI use cases, building on the 63 gen AI use cases identified by McKinsey's Digital Practice. This database also incorporates proprietary McKinsey research on personal productivity as well as industry reports, along with secondary research from the US government's Federal AI Use Case Inventories, NASA, press articles, and public interviews with technology leaders.

Glossary

The following terms in this report are defined specifically for its context.

Adoption and deployment: Deployment typically refers to the extent to which an organization rolls out a technology product (whether developed in-house or purchased off the shelf), and adoption reflects how extensively these products are used to generate measurable business value. Given the emerging nature of AI, many companies are simultaneously deploying and adopting, iterating as they go. Therefore, this report often uses adoption and deployment interchangeably to refer to the overall uptake of AI tools.

Agentic AI: Systems with autonomy and goal-directed behavior capable of making independent decisions, planning, and adapting to achieve specific objectives without direct, ongoing human input.

Application programming interface (API): Intermediary software components that allow two applications to talk to each other; a structured way for AI systems to programmatically access (usually external) models, data sets, or other pieces of software.

Artificial Intelligence (AI): The ability of software to perform tasks that traditionally require human intelligence, mirroring some cognitive functions usually associated with human minds.

Deep learning: A subset of machine learning that uses deep neural networks, which are layers of connected “neurons” whose connections have parameters or weights that can be trained. Deep learning is especially effective at learning from unstructured data such as images, text, and audio.

Digital workforce: A collaborative ecosystem where humans and automated agents work together, leveraging digital platforms, AI, and cloud computing to enhance productivity, efficiency, and scalability across various industries.

Employee: A worker in a corporate setting, either a manager or independent contributor. Examples of the type of employees represented in this report include people working in product management, marketing, technology, business development, sales, and finance.

Foundation models: Deep learning models trained on vast quantities of unstructured, unlabeled data that can be used for a wide range of tasks out of the box or adapted to specific tasks through fine-tuning. Examples of these models are DALL-E 2, GPT-4, PaLM, and Stable Diffusion.

Generative AI (gen AI): AI that is typically built using foundation models and has capabilities that earlier forms lacked, such as the ability to generate content. Foundation models can also be used for nongenerative purposes (for example, classifying user sentiment as negative or positive based on call transcripts).

Graphics processing units (GPUs): Computer chips originally developed for producing computer graphics, such as for video games, that are also useful for deep learning applications. In contrast, traditional machine learning usually runs on central processing units (CPUs), normally referred to as a computer’s “processor.”

Hallucination: A scenario where an AI system generates outputs that lack grounding in reality or a provided context. For instance, an AI chatbot may fabricate information or present a false narrative.

Large language models (LLMs): A class of foundation models that can process massive amounts of unstructured text and learn the relationships between words or portions of words, known as tokens. This enables LLMs to generate natural-language text, performing tasks such as summarization or knowledge extraction. GPT-4 (which underlies ChatGPT) and the Llama family of models from Meta are examples of LLMs.

Modality: A high-level data category such as numbers, text, images, video, and audio.

Multimodal capabilities: The ability of an AI system to process and generate various types of data (text, images, audio, video) simultaneously, enabling complex tasks and rich outputs.

Productivity (from labor): The ratio of GDP to total hours worked in the economy. Labor productivity growth generally comes from increases in the amount of capital available to each worker, the education and experience of the workforce, and improvements in technology.

Prompt engineering: The process of designing, refining, and optimizing input prompts to guide a gen AI model toward producing desired and accurate outputs.

Reasoning AI: AI systems that perform logical thinking, step-by-step planning, problem solving, and decision making using structured or unstructured data, going beyond pattern recognition to draw conclusions and solve complex problems.

Superagency: A state where individuals, empowered by AI, amplify their creativity, productivity, and positive impact. Even those not directly engaging with AI can benefit from its broader effects on knowledge, efficiency, and innovation.

Unstructured data: Data that lack a consistent format or structure (for example, text, images, video, and audio files) and typically require more advanced techniques to extract insights.



January 2025

Copyright © McKinsey & Company

Designed by McKinsey Global Publishing

www.mckinsey.com