

STANFORD UNIVERSITY

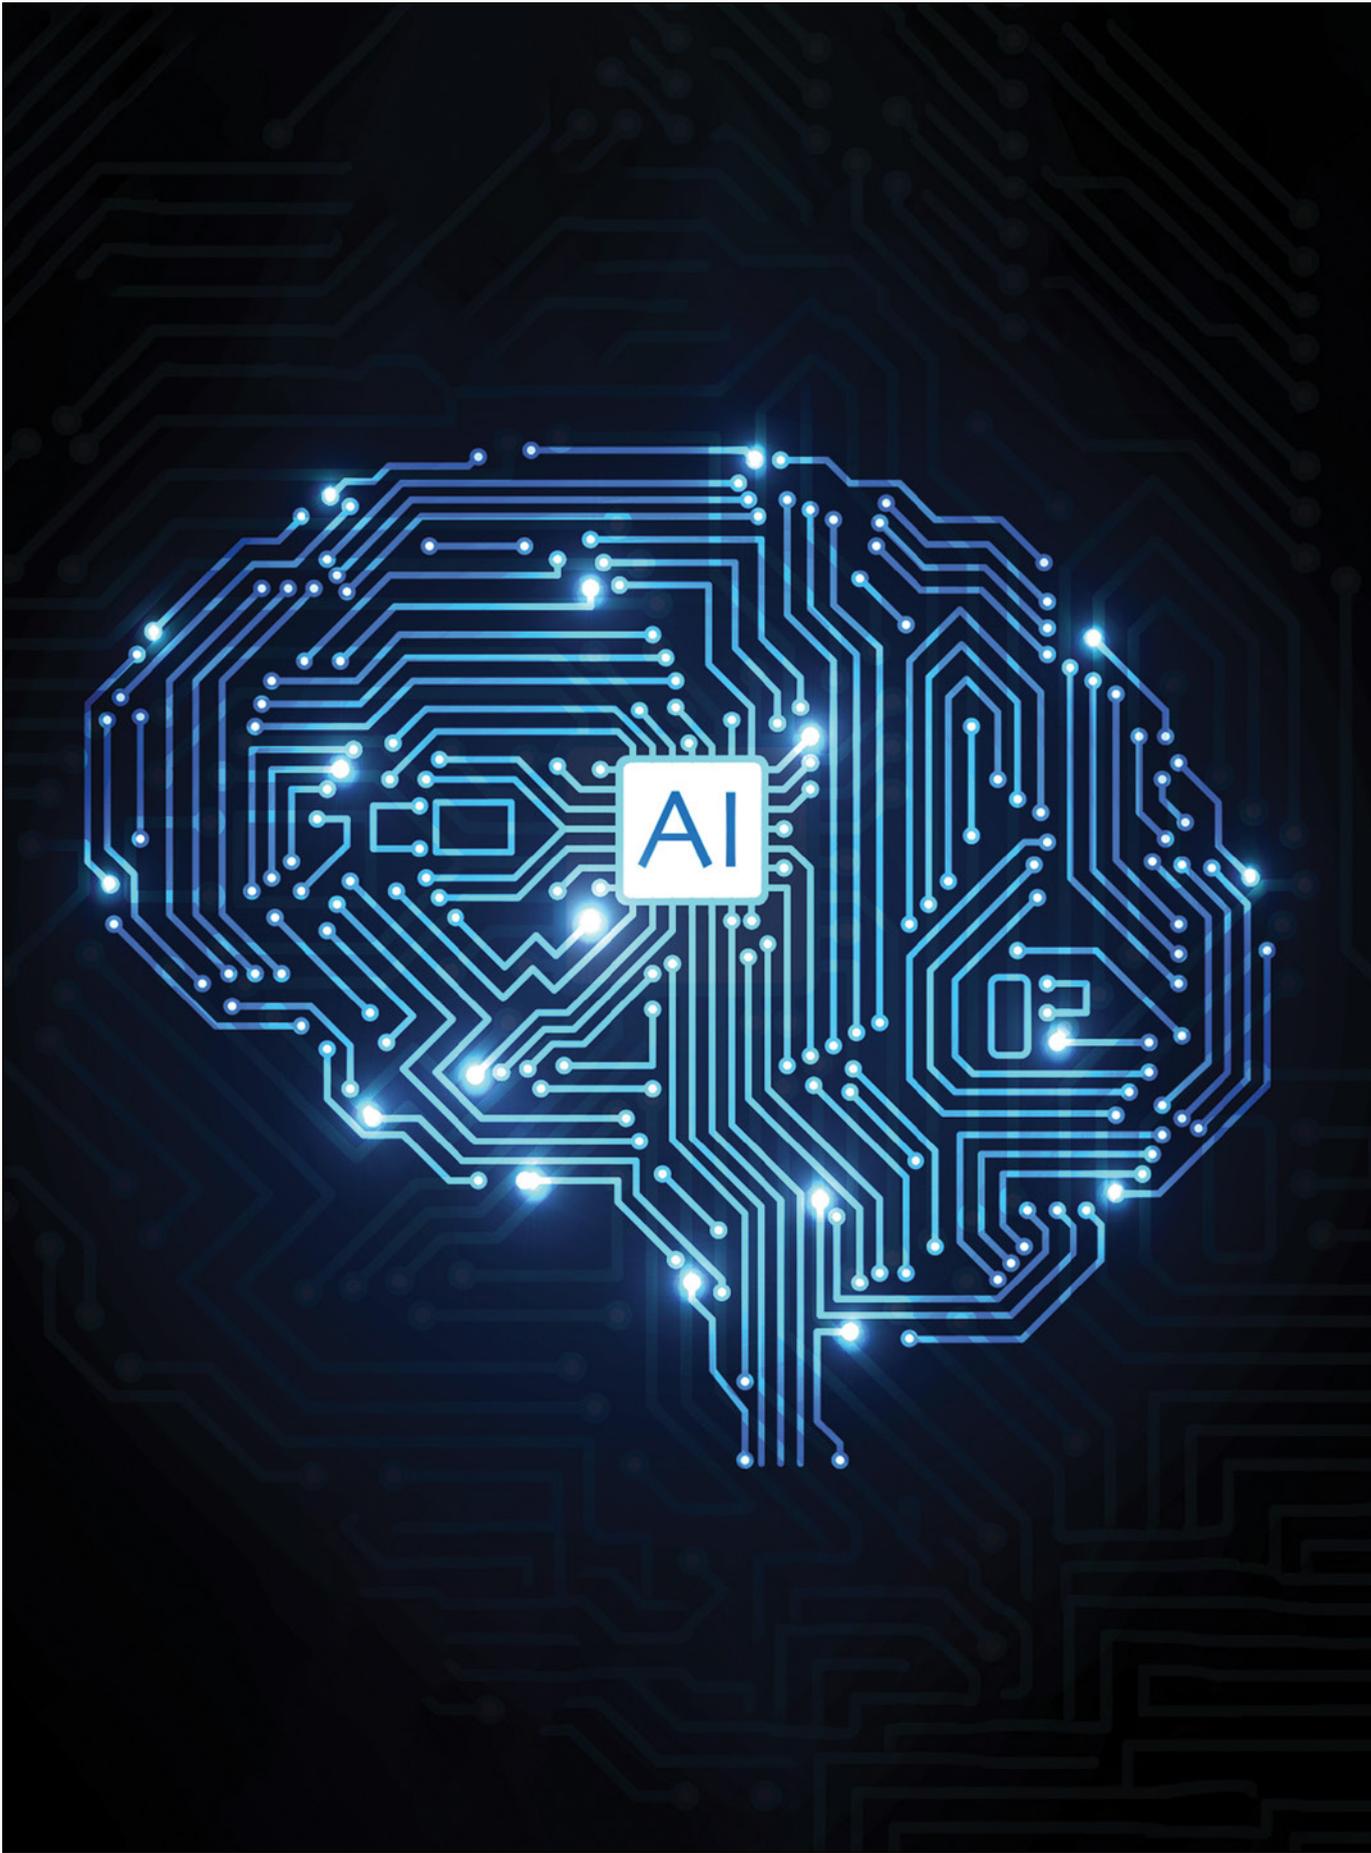
# THE STANFORD EMERGING TECHNOLOGY REVIEW 2026

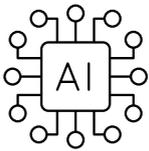
A Report on Ten Key Technologies and Their Policy Implications

**CO-CHAIRS** Condoleezza Rice, Jennifer Widom, and Amy Zegart

**DIRECTOR AND EDITOR IN CHIEF** Herbert S. Lin | **MANAGING EDITOR** Martin Giles







# ARTIFICIAL INTELLIGENCE

## KEY TAKEAWAYS

- Artificial intelligence (AI) is a foundational technology that is supercharging other scientific fields and, like electricity and the internet, has the potential to transform societies, economies, and politics worldwide.
- Despite rapid progress in the past several years, even the most advanced AI models still have many failure modes and vulnerabilities to cyberattacks that are unpredictable, not widely appreciated nor easily fixed, and capable of leading to unintended consequences.
- Nations are competing to shape the global rules and standards for AI, making interoperability, sizeable national compute resources, and international governance frameworks critical levers of geopolitical influence.

## Overview

Artificial intelligence (AI), a term coined by computer scientist and Stanford professor John McCarthy in 1955, was originally defined as “the science and engineering of making intelligent machines.” In turn, intelligence might be defined as the ability to learn and perform suitable techniques to solve problems and achieve goals appropriate to the context in an uncertain, ever-varying world.<sup>1</sup> AI could be said to refer to a computer’s ability to display this type of intelligence.

The emphasis today in AI is on machines that can learn as well as humans can learn, or at least somewhat comparably so. However, because machines are not limited by the constraints of human biology, AI systems may be able to run at much higher speeds and digest larger volumes and types of information than humans are capable of.

---

Today, AI promises to be a fundamental enabler of technological advancement in many fields, arguably of comparable importance to electricity in an earlier era or the internet in more recent years. The science of computing, worldwide availability of networks, and civilization-scale data—everything collectively underlying the AI of today and tomorrow—are poised to have similar impact on technological progress in the future. Moreover, the users of AI will not be limited to those with specialized training; instead, the average person on the street will increasingly interact directly with sophisticated AI applications for a multitude of everyday activities.

The global AI market is projected to be worth \$244.22 billion in 2025, with North America receiving 33.8 percent of total AI revenues.<sup>2</sup> The Stanford Institute for Human-Centered Artificial Intelligence (HAI) *AI Index 2025 Annual Report* found that private investment in all AI start-ups totaled \$150.79 billion in 2024, surpassing the previous record high of over \$120 billion, in 2021, after two consecutive years of decline.<sup>3</sup>

One estimate forecasts that generative AI—which can create novel text, images, video, and audio output and is discussed in more detail later in this chapter—could raise global GDP by \$7 trillion and raise productivity growth by 1.5 percent over a 10-year period if it is adopted widely.<sup>4</sup> Private funding for generative AI start-ups surged to \$33.94 billion in 2024, an 18.7 percent increase from 2023.<sup>5</sup>

The question of what subfields are considered part of AI is a matter of ongoing debate, and the boundaries between these fields are often fluid. Some of the core subfields follow:

- **Machine learning (ML)** Enabling computers to perform tasks without explicit instructions, often by generalizing from patterns in data. This includes deep learning that relies on multilayered artificial neural networks—which process information in a way inspired by the human brain—to model and understand complex relationships within data.

- **Natural language processing** Equipping machines with capabilities to understand, interpret, and produce spoken words and written texts, often using ML techniques.
- **Computer vision** Enabling machines to recognize and understand visual information from the world, convert it into digital data, and make decisions based on these data, often through applying ML.

Much of today's AI is based on ML, though it draws on other subfields as well. ML requires data and computing power—often called compute<sup>6</sup>—and much of today's AI research requires access to these on an enormous scale. The importance of ML is underscored by the award of the 2024 Nobel Prize in Physics (for foundational discoveries enabling ML with artificial neural networks<sup>7</sup>) and the Nobel Prize in Chemistry (for solving the problem of protein structure prediction using AI-based techniques<sup>8</sup>).

In general, a traditional ML model is developed to solve a particular problem, with different problems calling for different models; for problems sufficiently different from each other, entirely new models need to be developed. By contrast, foundation models, a relative newcomer to AI, discussed later in this chapter, may be used across a variety of problems.

ML requires large amounts of data from which it can learn. These data can take various forms, including text, images, videos, sensor readings, and more. Learning from these data is called training the AI model. The quality and quantity of data play a crucial role in determining the performance and capabilities of AI systems. Without sufficient high-quality data, models may generate inaccurate or skewed outcomes. Research continues on how to train systems efficiently: One option is to start from existing models and use a much smaller amount of specially curated data to refine those models' performance for specialized purposes; another option is to compress existing large models into much smaller ones.

For a sense of scale, estimates of the data required to train GPT-4—one of OpenAI’s large language models (LLMs), released in March 2023 and the base on which previous versions of ChatGPT were built—suggest that its training database consisted of the textual equivalent of around 100 million books, or about 10 trillion words, drawn from billions of web pages and scanned books. (LLMs are discussed in more detail below.)

The hardware requirements for computing power are also substantial. For example, reports indicate that the training of GPT-4 took about 25,000 Nvidia A100 GPU deep-learning chips—at a cost of \$10,000 each—running for about 100 days.<sup>9</sup> Including both these chips and other hardware components used, the overall hardware costs for GPT-4 were at least a few hundred million dollars. And the chips underlying this hardware are specialty ones often fabricated offshore.<sup>10</sup> (Chapter 9, on semiconductors, discusses this point at greater length.)

Lastly, AI models consume a lot of energy. Consider, first, the training phase: One estimate of the electricity required to train a foundation model such as GPT-4 puts the figure at about 50 million kilowatt-hours (kWh).<sup>11</sup> The average American household uses about 11,000 kWh per year, meaning the energy needed to train GPT-4 was approximately the same as that used by 4,500 average homes in a year. Paying for this energy adds significant cost and raises environmental concerns, even before a single person actually uses a model.

Once a model is up and running, the cost of energy used to power queries can add up fast. This is known as the inference phase. For ChatGPT, the energy used per query is around 0.002 of a kilowatt-hour, or 2 watt-hours (Wh).<sup>12</sup> (For comparison, a single Google search requires about 0.3 Wh,<sup>13</sup> and an alkaline AAA battery contains about 2 Wh of energy.)

Given hundreds of millions of queries per day, the operating energy requirement of ChatGPT might be a few hundred thousand kilowatt-hours per day,

at a cost of several tens of thousands of dollars. With the recent focus on what are called reasoning models—foundation models that seemingly “think” through problems step by step before presenting the user with an output—such inference costs have substantially increased in the past year.

AI is expected to automate a wide range of tasks. But it also has particular promise in augmenting human capabilities and further enabling people to do what they are best at doing.<sup>14</sup> AI systems can work alongside humans, complementing and assisting their work rather than replacing them. Some present-day examples follow.

### Healthcare

- **Medical diagnostics** An AI system that can predict and detect the onset of strokes qualified for Medicare reimbursement in 2020.<sup>15</sup>
- **Drug discovery** An AI-enabled search identified a compound that inhibits the growth of a bacterium responsible for many drug-resistant infections, such as pneumonia and meningitis, by sifting through a library of seven thousand potential drug compounds for an appropriate chemical structure.<sup>16</sup>
- **Patient safety** Smart AI sensors and cameras can improve patient safety in intensive care units, in operating rooms, and even at home by improving healthcare providers’ and caregivers’ ability to monitor and react to patient health developments, including falls and injuries.<sup>17</sup>

### Agriculture

- **Production optimization** AI-enabled computer vision helps some salmon farmers pick out fish that are the right size to keep, thus off-loading the labor-intensive task of sorting them.<sup>18</sup>
- **Crop management** Some farmers are using AI to detect and destroy weeds in a targeted

manner, significantly decreasing environmental harm by using herbicides only on undesired vegetation rather than on entire fields. In some cases, this has reduced herbicide use by as much as 90 percent.<sup>19</sup>

## Logistics and Transportation

**Autonomous trucking** From January to August 2024, a consortium of companies autonomously drove trucks carrying tires for over 50,000 long-haul miles.<sup>20</sup> Continued success could automate long-haul drives—the most boring, time-consuming part of a trucker’s job—while keeping human tasks like navigating the first miles from factories and the last miles to customers.

## Law

**Legal review** AI-based systems can reduce the time lawyers spend on contract review by as much as 60 percent. Further, such systems can enable lawyers to search case databases more rapidly than online human searches.<sup>21</sup>

---

# Key Developments

## Foundation Models

Foundation models have dominated the conversation about AI since late 2022. These models are large-scale systems that are trained on vast amounts of diverse data and that can handle a variety of tasks.<sup>22</sup> They often contain billions or trillions of parameters, and this vast scale allows them to capture more complex patterns and relationships. (Parameters are the building blocks of a foundation model, with values set during training. Though not directly interpretable as discrete chunks of knowledge, parameters act like billions of adjustable knobs that collectively guide how the model learns patterns and makes decisions. The most essential part of a parameter is the weights it uses. Weights represent the strength

of connections within the model, encoding learned patterns from training data and determining how much influence each piece of information has on the model’s behavior.)

Trained on large-scale data, foundation models can exhibit broad capabilities<sup>23</sup> and are thus sometimes called general-purpose models. They excel at transfer learning—applying knowledge learned in one context to another—making them more flexible and efficient than traditional task-specific models. A single foundation model can be fine-tuned for various tasks, often reducing the need to train separate models from scratch.

These models are generally classified as closed source, open weight, or open source. A closed-source model is a proprietary one developed and maintained by a specific organization, usually a for-profit company, with its source code, data, and architecture kept confidential. Access to these models is typically restricted through technically enforced usage permissions, such as application programming interfaces, allowing the developers to control the model’s distribution, usage, and updates.

By contrast, an open-source model is one whose code, data, and underlying architecture are publicly accessible, allowing anyone to use, modify, and distribute it freely. Open-weight models fall in between: Their weights are publicly released, but other components, such as training data, are kept confidential.

The most well-known type of foundation model is an LLM—a system trained on very large volumes of textual content. LLMs are an example of generative AI, a type of AI that can produce new material based on how it has been trained and the inputs it is given. Models trained on text can generate new text based on a statistical analysis that makes predictions about what other words are likely to be found immediately after the occurrence of certain words.

These models do not think or feel like humans do, even though their responses may make it seem like

they do. Instead, LLMs use statistical analysis based on training data. For example, because the word sequence “thank you” is far more likely to occur than “thank zebras,” a person’s query to an LLM asking it to draft a thank-you note to a colleague is unlikely to generate the response “thank zebras.”

LLMs generate humanlike language across many subjects, producing potentially useful content such as code, poetry, legal summaries, and medical advice. They outperform median human scores on exams in obstetrics and gynecology;<sup>24</sup> divergent thinking tests;<sup>25</sup> and the LSAT, the GRE, and various Advanced Placement exams.<sup>26</sup> However, they do not necessarily master the underlying skills that these tests assess and still make errors and fail unexpectedly. Developing valid evaluation metrics that accurately capture the true capabilities, limitations, and risks of foundation models remains an open and ongoing research challenge.<sup>27</sup>

Well-known closed-source LLMs include certain of OpenAI’s models, such as those in the GPT series; Anthropic’s Claude Opus 4.1; and Google’s Gemini 2.5 Pro. Well-known open-source or open-weight LLMs include Meta’s Llama 4, Google’s Gemma 2, and Cohere’s Command R.

Specialized foundation models have also been developed in other modalities such as images, audio, and video.

- Foundation models for images are able to generate new images based on a user’s text input. Novel methods for handling images, combined with the use of very large collections of pictures and text for training, have led to models that can turn written descriptions into images that are quickly becoming comparable to—and sometimes indistinguishable from—real-life photographs and artwork created by humans. Examples include OpenAI’s DALL-E 3, the open-source Stable Diffusion, Google’s Imagen, Adobe Firefly, and Meta’s Make-A-Scene.

- An example of a foundation model for audio is UniAudio, which handles all audio types and employs predictive algorithms to generate high-quality speech, sound, and music, surpassing leading methods in tasks such as text to speech, speech enhancement, and voice conversion.

- Foundation models in video such as Meta’s Emu Video and OpenAI’s Sora represent a significant advancement in video generation. Emu first generates an image from a text input and then creates a video based on both the text and the generated image. Sora also enables the user to turn existing images into videos while also editing videos through textual input.

### **Multimodal Models**

AI systems that incorporate multiple modalities—text, images, and sound—within single models are becoming increasingly popular. This multimodal approach, shown in figure 1.1, aims to create more humanlike experiences by leveraging various senses such as sight, speech, and hearing to mirror how humans interact with the world.

Multimodal AI systems have diverse applications across sectors. They can enhance accessibility for people with disabilities through real-time transcription, sign language translation, and detailed image descriptions. They can also eliminate language barriers via cost-effective, near-real-time translation services. In education, multimodal AI can support personalized learning by adapting content to various formats and learner types, improving engagement and comprehension.

When integrated with virtual and augmented reality, AI can create immersive, highly realistic training environments that are particularly valuable in fields like healthcare. The advent of multimodal AI is also set to further transform human–computer interactions, enabling more intuitive communication and expanding the range of tasks that AI systems can handle.

## Limitations and Risks of Current AI Systems

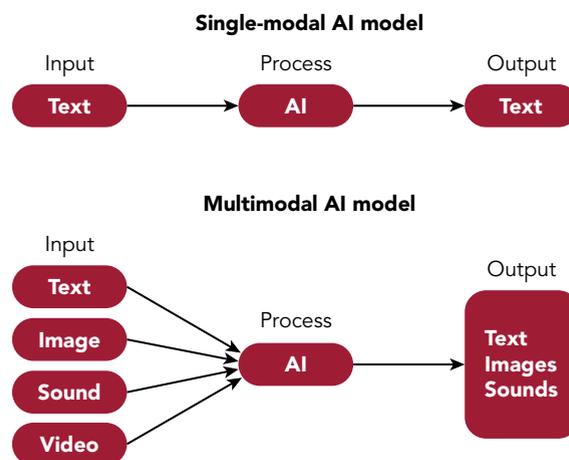
Potential positive impacts of AI will likely come from societal applications. But no technology is an unalloyed good. Negative impacts are expected both from current AI limitations and future advances. Key issues include the following:

- **Explainability** AI systems generally cannot explain their reasoning or data sources. While explanations aren't always needed, in critical domains, like medical decision making, they are essential for user confidence and trust.
- **Bias and fairness** Models trained on biased datasets reproduce those biases. For example, a facial recognition system trained mainly on one ethnic group may perform poorly on others, likely leading to disproportionate harms.<sup>28</sup> Because data reflects historical inequities, models inevitably embed them, too.
- **Vulnerability to attacks** Small changes to data or inputs can trick AI into false conclusions. For example, small changes, invisible to the naked

eye, made to the individual pixels of a stop sign image can cause an AI to classify it as a yield sign.<sup>29</sup> (Chapter 3, on cryptography and computer security, discusses this point in greater detail.) This could prove particularly dangerous for systems used in medicine or the military. Newer models (e.g., multimodal models and agents, covered later in this chapter) expand possible attack vectors.

- **Deepfakes** AI provides the capability for generating highly realistic but entirely inauthentic audio and video imagery. This has obvious implications for evidence presented in courtrooms and for efforts to manipulate political contests. However, despite widespread concerns expressed in 2024 about the potential effects of deepfakes on elections, fake audio, images, and videos did not play as transformative or disruptive a role as feared in the 2024 US elections.<sup>30</sup> Of particular interest is the observation that traditional “cheap fakes”—relatively crude attempts made by manipulating videos and other content—were more prevalent than AI-generated deepfakes. Nevertheless, concerns remain about the potential impact on future democratic processes as the sophistication and usage of AI-generated deepfakes increase.

**FIGURE 1.1** Multimodal AI systems can transform one type of input into a different type of output



# AI agents ideally operate by executing tasks with minimal human input and oversight. . . . Yet, from a technical standpoint, present-day [agents] face major limitations.

- **Privacy** Many LLMs are trained on vast amounts of internet data, often without careful filtering, and such data can include individuals' personal information. Once incorporated into training datasets, this information may be reproduced or disclosed by the model. Additionally, as AI handles sensitive tasks like mental health support, privacy concerns will grow.
- **Overtrust and overreliance** Familiarity increases user trust, but people may become too complacent when employing AI tools. For example, a recent study showed that developers who used AI coding assistants wrote less secure code—yet they believed that what they were producing was more secure.<sup>31</sup>
- **Hallucinations** Hallucinations occur whenever models generate plausible but false outputs, leaving users unaware that the outputs are fabricated. In September 2024, a Stanford professor asked an AI to list ten of her publications. It returned five real publications and five invented ones, complete with convincing titles and summaries. When she flagged the errors as hallucinations, the model simply produced two new fabricated results.

Researchers are aware of these problems and are working on fixes, but solutions often don't generalize beyond specific cases.

## Over the Horizon

### AI Agents

AI agents ideally operate by executing tasks with minimal human input and oversight, such as setting people's daily agendas and coordinating software tools. Such agents are gaining traction in activities within enterprises such as customer service and invoice processing. OpenAI's strategy for ChatGPT involves it becoming a "super assistant" that handles everyday tasks such as drafting emails, offering medical advice, and managing finances.<sup>32</sup> Yet, from a technical standpoint, present-day AI agents face major limitations.

- **Memory** The basic unit that an LLM reads and generates to process text—whether a whole word, part of a word, or punctuation—is known as a token. For AI agents to be effective, they need to remember things, such as when a meeting was scheduled in a previous session. An agent's memory is limited by context length, which is the maximum number of input and output tokens the system can handle at any one time. Although the context length of top systems has expanded dramatically in recent years, it is still not enough to remember all the details needed to execute many multistep tasks, especially across different sessions. Efforts to increase cross-session memory and enhance long-term storage are still nascent.

- **Reliability** Even with adequate memory, agents can suffer from goal drift, infinite loops, and resource exhaustion, which undermine their usefulness in real-world settings.
  - Goal drift occurs when an AI agent fails to consistently pursue its original objective during task execution and instead begins to focus on less relevant objectives, which can result in undesired behavior or outcomes.
  - Infinite loops occur when an agent “gets stuck,” repeatedly performing the same actions or reasoning steps without making progress toward its goal.
  - Resource exhaustion happens when an agent consumes excessive computational or memory resources due to complex processing, repeated retries, or ineffective algorithms, resulting in degraded performance or system failure.
- **Interoperability** Most agents cannot seamlessly communicate with other agents or external systems. However, in November 2024, the Model Context Protocol (MCP), an open standard for secure, efficient agent-to-system integration, was introduced by Anthropic; since then, it has been adopted by OpenAI, Google DeepMind, Microsoft, and others. The MCP provides a universal interface for an agent to read files, execute functions, handle contextual prompts, and connect with external tools, data sources, and applications, enabling AI agents to access real-time data and perform actions beyond their training data.

### **Efficiency, Specialization, and Synthetic Data**

Progress in AI is shifting from building ever more resource-intensive models to using resources more efficiently. For example, because of limited supplies of real data, synthetic data is increasingly used. (In this context, synthetic data is artificially generated data that are designed to mimic the statistical

properties and patterns of real-world data. However, such data do not themselves reveal anything real about the world.)

In addition, because large models consume significant amounts of energy, AI engineers are focusing on models that can be constructed with fewer computational resources. For example, one approach is to do individual calculations at a lower level of precision (that is, using fewer bits to represent a number), thus reducing the volume of data processed and memory used. Individual calculations are therefore less accurate, but the model is trained using a huge number of calculations on different though similar data; averaging over many such calculations can compensate for the loss of accuracy in individual ones. A variety of other approaches also rely on eliminating parts of model training that may be unnecessary or redundant for a given purpose.

Future AI gains will increasingly depend not just on large compute capacity and large amounts of data but also on domain-specific data and efficiency-focused innovations. Quantum computing (discussed in chapter 7, on quantum technologies) may lower the requirements of ML for energy and compute, although it is unclear whether such enhancements are within reach of current technologies.<sup>33</sup>

### **Embodied AI**

Embodied AI means AI that is able to sense and act in the real world (e.g., through its integration into robots or other physical devices). This has the potential to enhance robotic capabilities and expand the range of interactions robots have with the physical world. Systems combining robots and AI could potentially address knowledge tasks, physical tasks, or combinations of both. (This topic is explored further in chapter 8, on robotics.) As research progresses in AI autonomy and reasoning, embodied AI systems may be able to handle increasingly complex tasks with greater independence. This could lead to applications in various fields such as logistics and domestic assistance.

---

## Policy Issues

### *The Future of Work*

Within five to ten years, more workers will have AI integrated into their workflows or will have their jobs replaced entirely by AI systems, a potential disruption to the job market.<sup>34</sup> LLMs have already shown they are sometimes useful in fields like law, customer support, coding, and journalism. These developments raise concerns about AI's significant impact on many knowledge-based jobs and employment overall. However, uncertainty abounds. What and how many present-day jobs will disappear? Which tasks could best be handled by AI? And what new jobs might be created by the technology today and in the future?

Some broad outlines and trends are clear.

- Individuals whose jobs entail routine white-collar work may be more affected than those whose jobs require physical labor; some will experience painful shifts in the short term.<sup>35</sup>
- AI is helping some workers to increase their productivity and job satisfaction.<sup>36</sup> At the same time, other workers are already losing their jobs as AI demonstrates some competence for business operations—despite potentially underperforming the humans it replaces.<sup>37</sup> In some cases, companies are deciding that the cost savings of eliminating human workers outweigh the drawbacks of mediocre AI performance.
- Training displaced workers to be more competitive in an AI-enabled economy does not solve the problem if new jobs are not available. The nature and extent of new roles resulting from widespread AI deployment are not clear at this point. However, historically, the introduction of new technologies has not resulted in a long-term net loss of jobs.<sup>38</sup>

### *Governance and Regulation of AI*

Governments around the world have been increasingly focused on establishing regulations and guidelines for AI.

In the United States, the Trump administration revoked the Biden administration's Executive Order (EO) 14110, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*. It subsequently issued a new EO titled *Removing Barriers to American Leadership in Artificial Intelligence*, aiming to promote AI innovation and leadership by eliminating the restrictions and requirements on AI contained in EO 14110. In August 2025, the administration launched America's AI Action Plan, which outlines a policy road map to "accelerate innovation, build American AI infrastructure, and lead in international diplomacy and security."<sup>39</sup> This plan faces challenges, especially regarding alignment with concurrent proposals to scale back broader scientific research funding.

In addition, friction between state and federal approaches to AI governance is growing. States are experimenting with their own AI legislation, often proposing requirements that go well beyond federal guidance. Advocates view this "policy laboratory" approach as essential for innovation in governance; critics warn it creates a fragmented compliance landscape that hampers interstate commerce. Recent state action, as described in table 1.1, underscores the emerging rule-making patchwork.

In the European Union, the most ambitious attempts to regulate AI came into force in August 2024 with the EU Artificial Intelligence Act. The act forbids certain applications of AI, such as individual predictive policing based solely on a person's data profile or tracking of their emotional state in the workplace and educational institutions, unless for medical or safety reasons.<sup>40</sup> Additionally, it imposes a number of requirements on high-risk AI systems and foundation models, addressing transparency and explainability, human oversight, cybersecurity, and robustness.

**TABLE 1.1** Selected states' actions regarding AI legislation

State	Legislation	Key provisions
Colorado	Senate Bill 24-205 (Colorado AI Act)	Mandates duties on developers and deployers of "high-risk" AI systems to prevent algorithmic discrimination
Texas	Texas Responsible Artificial Intelligence Governance Act	Prohibits AI systems used for behavioral manipulation, discrimination, or deployment of deepfakes
California	Multiple AI bills (15+), including Assembly Bill 2013	Targets generative AI systems and requires training data disclosure for AI systems used by Californians

To help operationalize these rules, the European AI Office oversaw one of the earliest and most formalized multi-stakeholder consultations in AI policy to date: the General-Purpose AI (GPAI) Code of Practice. Developed by thirteen independent experts with input from nearly one thousand participants across member states, academia, civil society, and industry, the code supplements the EU AI Act by offering detailed provisions on transparency, copyright, and safety and security—giving foundation model developers a recognized pathway to meeting certain requirements of the EU AI Act with greater legal clarity and reduced enforcement risk.

Other important international AI governance developments include the AI Summit series. The first AI Safety Summit, held in November 2023 at Bletchley Park in the United Kingdom,<sup>41</sup> issued the Bletchley Declaration. In it, the European Union and twenty-eight nations collectively endorsed international cooperation to manage risks associated with highly capable general-purpose AI models. The summit also led to the establishment of the United Kingdom's AI Safety Institute (now the UK AI Security Institute) and the US AI Safety Institute (now the Center for AI Standards and Innovation), located within the National Institute of Standards and Technology. Similar institutions have since been established in a number of other countries, including Japan, Singapore, South Korea, Canada, France, Kenya, Australia, and the European Union.

The Seoul Declaration, from the AI Seoul Summit 2024, built on the Bletchley Declaration to acknowledge the importance of interoperability between national AI governance frameworks to maximize benefits and minimize risks from advanced AI systems. In contrast, the 2025 Artificial Intelligence Action Summit, in France, marked a notable shift from the safety-oriented tone of earlier gatherings, placing greater emphasis on accelerating innovation and industrial adoption of AI.

### **National Security and Geopolitics**

The intensifying technological race between the United States and China regarding AI is entering a new phase. While the United States continues to push the technical frontier with increasingly capable models, China is aggressively diffusing existing AI capabilities across every sector—from education to manufacturing to governance—aiming to lock in large-scale network advantages at home and abroad. China's open-source model releases, such as DeepSeek, further challenge America's frontier status, accelerate global adoption, and undermine US containment efforts.

Infrastructure—especially compute capacity—has become a critical, strategic global resource that is vital to remaining competitive in economic, military, cyber, and intelligence domains. For example, countries like Canada and the United Kingdom have

announced major compute infrastructure projects.<sup>42</sup> In the United States, the privately funded Stargate AI infrastructure initiative was launched in January 2025. First reports suggested this would invest as much as \$500 billion over the next few years, although a more recent report suggests a scaling back of initial objectives.<sup>43</sup>

The United States also began piloting the National Artificial Intelligence Research Resource (NAIRR) in January 2024. Supported by the Trump administration, NAIRR is a federally backed initiative to give academics and civil-society researchers shared access to advanced compute, high-quality data, and AI tools. However, it is noteworthy that investments for AI from high-tech companies exceeded \$27 billion in 2023 alone. This figure is far larger than the \$2.6 billion authorized for appropriation over six years for NAIRR under the Creating Resources for Every American To Experiment with Artificial Intelligence (CREATE AI) Act of 2025.

### **Export Controls**

Export controls are often used to slow the advances of rivals, but open-source diffusion and breakthroughs in compute-efficient training may erode their impact. This leaves policymakers with a strategic choice: focus on restraining competitors through export controls, or accelerate domestic innovation and the global adoption of domestic products—or attempt both simultaneously. Under the Biden administration, and initially under the Trump administration, the United States has taken the export

control route. However, in August 2025 reports suggested a shift, with President Trump and Chinese President Xi Jinping considering an arrangement allowing US chipmakers Nvidia and AMD (Advanced Micro Devices) to sell certain chips to China in exchange for a 15 percent revenue share to the US government.<sup>44</sup>

### **Use in the Military**

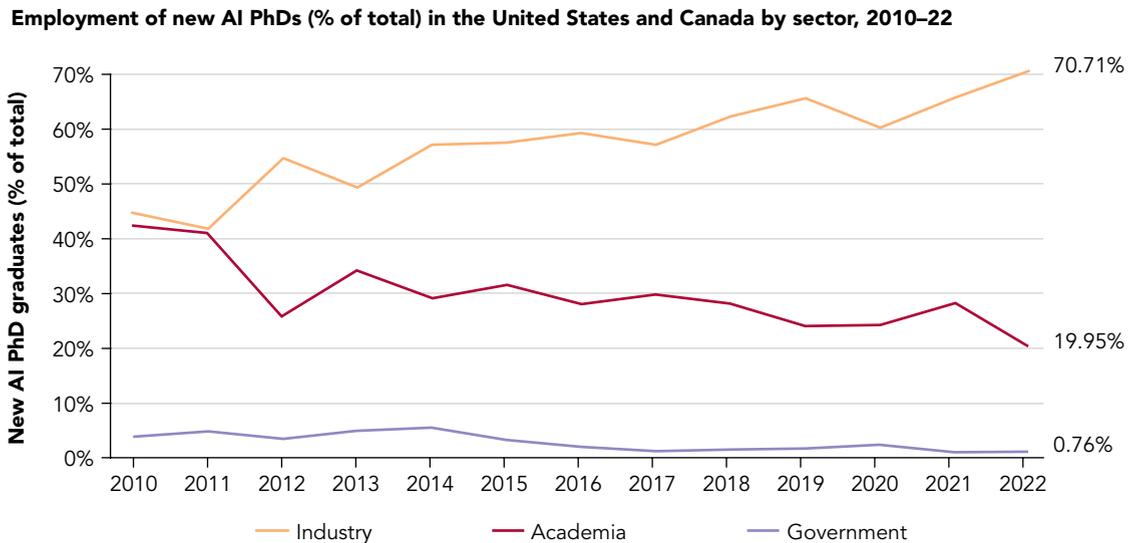
AI is expected to have a profound impact on militaries worldwide.<sup>45</sup> Weapons systems, command and control, logistics, acquisition, and training will all seek to leverage AI to operate more effectively and efficiently, at lower cost and with less risk to friendly forces. Trying to overcome decades of institutional inertia, the US Department of Defense is dedicating billions of dollars to institutional reforms and research advances aimed at integrating AI into its warfighting and war preparation strategies. Senior military officials are concerned that failure to adapt to the emerging opportunities and challenges presented by AI would pose significant national security risks, particularly considering that both Russia and China are investing heavily in AI capabilities.

### **Talent**

The United States is eating its seed corn with respect to the AI talent pool. Faculty at Stanford and other universities report that the number of students studying in AI who are joining industry, particularly startups, is increasing at the expense of those pursuing academic careers and contributing to foundational AI

**The resources needed to train GPT-4 far exceed those available through grants or any other sources to any reasonably sized group of the top US research universities.**

**FIGURE 1.2** Most new AI PhDs hired in North America are flocking to industry



Source: Adapted from Nestor Maslej, Loredana Fattorini, Raymond Perrault, et al., *The AI Index 2024 Annual Report*, AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April 2024. Data from CRA Taulbee Survey, 2023

research. The United States is thus experiencing an AI “brain drain” that does not favor the US research enterprise or its innovation capacity (figure 1.2).

Many factors are contributing to this trend. One is that industry careers offer compensation far exceeding academic packages. Academic researchers must also secure money for equipment, compute, and staff, often relying on government funding that is typically small compared to what large companies might be willing to invest in their own researchers. Consider, for example, that the resources needed to build and train GPT-4 far exceed those available through grants or any other sources to any reasonably sized group of the top US research universities, let alone any single university. This gap is exacerbated by recent cuts in federal research funding.

Industry often makes decisions more rapidly than government grant makers and imposes fewer regulations on the conduct of research. Large companies

are also at an advantage because they have research-supporting infrastructure in place, such as compute facilities and data warehouses.

Finally, other nations are actively recruiting talent that in the past tended to favor employment in the United States. China’s recruitment of top scientific talent, especially ethnic Chinese in the United States, is being driven through offers of benefits such as high salaries and generous research funding. Countries like Canada and the United Kingdom actively recruit US-based researchers and offer AI-focused visas. The brain drain has been exacerbated by changes in US immigration policy that have caused top researchers to leave and deterred talented international students from studying in the United States.

### Copyright

Many foundation models have been trained on vast amounts of data found on the internet. These data

have generally been used without the consent or permission of their owners, and a number of lawsuits were filed over this issue in 2023–24. They include *Getty Images v. Stability AI* over alleged infringement on copyrights of photographs;<sup>46</sup> *The New York Times v. OpenAI* and *Microsoft* over alleged use of millions of articles published by the *Times*;<sup>47</sup> and *Sony Music, Universal Music Group, and Warner Records v. AI start-ups Suno and Udio* over alleged use of protected content to train their music-generation systems.<sup>48</sup> Such cases have raised important questions about appropriately compensating and acknowledging data creators whose data is used to train AI models.

At the time of this writing (September 2025), one major copyright lawsuit against an AI company has reached a settlement. Anthropic, an AI developer, agreed to pay \$1.5 billion to settle a class-action suit brought by authors and publishers;<sup>49</sup> the settlement was approved on a preliminary basis by the judge overseeing the case as this publication went to press.<sup>50</sup> Even if it is approved, it will not create a legal precedent because of the out-of-court nature of the settlement. Whether the use of publicly available web data to train AI models is legally permissible remains, and will remain, unsettled across jurisdictions until more court decisions are rendered.

---

## Glossary

**Compute:** The processing power, typically measured in number of specialized chips and scale of energy use, required to train and run AI models.

**Distillation:** A technique for compressing a large, complex AI model into a smaller one that is faster and more efficient while retaining most of its capabilities.

**Generative AI/Foundation models:** Large, general-purpose AI systems trained on vast datasets that can

generate text, images, code, or other outputs and be adapted to many downstream applications.

**Inference:** The stage when a trained AI model is used to generate predictions, outputs, or decisions in response to new inputs.

**Model training:** The process of teaching an AI system by exposing it to large datasets and adjusting its parameters until it can perform a given task (or a set of tasks) effectively.

**Multimodal AI:** AI systems designed to process and integrate multiple kinds of data such as text, images, audio, or video into a single model.

**Open source/Open weight/Closed source:** Terms describing how freely AI models are shared; open-source makes training data, code, and weights public; open-weight shares only the trained parameters; and closed-source restricts access entirely.

**Scaling laws:** Observed patterns showing that, as AI models are trained with more data, parameters, and compute, their performance improves in predictable ways but also with potentially diminishing returns and real-world limits.

**Synthetic data:** Artificially generated data such as simulated text, images, or code used to supplement or replace real-world data for training and testing AI systems.

## NOTES

1. Christopher Manning, “Artificial Intelligence Definitions,” Stanford Institute for Human-Centered AI, September 2020, <https://hai.stanford.edu/sites/default/files/2020-09/AI-Definitions-HAI.pdf>.

2. “Artificial Intelligence—Worldwide,” Statista, accessed August 15, 2025, <https://www.statista.com/outlook/tmo/artificial-intelligence/worldwide>.

3. Nestor Maslej, Loredana Fattorini, Raymond Perrault, et al., *The AI Index 2024 Annual Report*, AI Index Steering Committee, Stanford Institute for Human-Centered AI, Stanford University, May 2024, [https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI\\_AI-Index-Report-2024.pdf](https://aiindex.stanford.edu/wp-content/uploads/2024/05/HAI_AI-Index-Report-2024.pdf).

4. "Generative AI Could Raise Global GDP by 7%," Goldman Sachs, April 5, 2023, <https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html>.
5. Maslej, Fattorini, Perrault, et al., *The AI Index 2024 Report*.
6. Jafar Alzubi, Anand Nayyar, and Akshi Kumar, "Machine Learning from Theory to Algorithms: An Overview," *Journal of Physics: Conference Series* 1142, Second National Conference on Computational Intelligence (December 2018), <https://doi.org/10.1088/1742-6596/1142/1/012012>.
7. "Summary," the Nobel Prize in Physics 2024, the Nobel Prize, October 12, 2024, <https://www.nobelprize.org/prizes/physics/2024/summary/>.
8. "Summary," the Nobel Prize in Chemistry 2024, the Nobel Prize, October 12, 2024, <https://www.nobelprize.org/prizes/chemistry/2024/summary/>.
9. Kif Leswing, "Meet the \$10,000 Nvidia Chip Powering the Race for A.I.," CNBC, February 23, 2023, <https://www.cnbc.com/2023/02/23/nvidias-a100-is-the-10000-chip-powering-the-race-for-ai-.html>; Kasper Groes Albin Ludvigsen, "The Carbon Footprint of GPT-4," Medium, July 18, 2023, <https://towardsdatascience.com/the-carbon-footprint-of-gpt-4-d6c676eb21ae>.
10. Darian Woods and Adrian Ma, hosts, *The Indicator from Planet Money*, podcast, "The Semiconductor Founding Father," NPR, December 21, 2021, 10:14, <https://www.npr.org/transcripts/1066548023>.
11. Kasper Groes Albin Ludvigsen, "The Carbon Footprint of ChatGPT," Medium, December 21, 2022, <https://medium.com/data-science/the-carbon-footprint-of-chatgpt-66932314627d>.
12. Kasper Groes Albin Ludvigsen, "ChatGPT's Electricity Consumption," Medium, July 12, 2023, <https://towardsdatascience.com/chatgpts-electricity-consumption-7873483feac4>. Different sources provide somewhat different numbers for the energy cost per query, but they all are in the range of a few watt-hours.
13. "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption," Technology Innovation, EPRI, May 28, 2024, <https://www.epri.com/research/products/3002028905>.
14. Hope Reese, "A Human-Centered Approach to the AI Revolution," Stanford University Institute for Human-Centered AI, October 17, 2022, <https://hai.stanford.edu/news/human-centered-approach-ai-revolution>.
15. "Viz.ai Receives New Technology Add-on Payment (NTAP) Renewal for Stroke AI Software from CMS," news release, Viz.ai, August 4, 2021, <https://www.viz.ai/news/ntap-renewal-for-stroke-software>.
16. Gary Liu, Denise B. Catacutan, Khushi Rathod, et al., "Deep Learning-Guided Discovery of an Antibiotic Targeting *Acinetobacter baumannii*," *Nature Chemical Biology* 19 (2023): 1342–50, <https://doi.org/10.1038/s41589-023-01349-8>.
17. Albert Haque, Arnold Milstein, and Fei-Fei Li, "Illuminating the Dark Spaces of Healthcare with Ambient Intelligence," *Nature* 585 (2020): 193–202, <https://doi.org/10.1038/s41586-020-2669-y>.
18. "Innovasea Launches AI-Powered Biomass Camera for Salmon," *The Fish Site*, August 17, 2023, <https://thefishsite.com/articles/innovasea-launches-ai-powered-biomass-camera-for-salmon>.
19. "Machine Learning in Agriculture: Use Cases and Applications," *ltransition.com*, February 1, 2023, <https://www.itransition.com/machine-learning/agriculture>.
20. "J.B. Hunt, Bridgestone and Kodiak Surpass 50,000 Autonomous Long-Haul Trucking Miles in Delivery Collaboration," Kodiak.ai, August 7, 2024, <https://kodiak.ai/news/jb-hunt-and-kodiak-collaborate>.
21. Steve Lohr, "A.I. Is Doing Legal Work. But It Won't Replace Lawyers, Yet," *New York Times*, March 19, 2017, <https://www.nytimes.com/2017/03/19/technology/lawyers-artificial-intelligence.html>.
22. Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, et al., "On the Opportunities and Risks of Foundation Models," preprint, arXiv, Stanford University, July 12, 2022, <https://doi.org/10.48550/arXiv.2108.07258>.
23. Bommasani, Hudson, Adeli, et al., "On the Opportunities and Risks."
24. Sarah W. Li, Matthew W. Kemp, Susan J. S. Logan, Sebastian E. Illanes, and Mahesh A. Choolani, "ChatGPT Outscored Human Candidates in a Virtual Objective Structured Clinical Examination in Obstetrics and Gynecology," *American Journal of Obstetrics & Gynecology* 229, no. 2 (August 2023): 172.E1–12, <https://doi.org/10.1016/j.ajog.2023.04.020>.
25. Kent F. Hubert, Kim N. Awa, and Darya L. Zabelina, "The Current State of Artificial Intelligence Generative Language Models Is More Creative Than Humans on Divergent Thinking Tasks," *Scientific Reports* 14, no. 3440 (February 2024), <https://doi.org/10.1038/s41598-024-53303-w>.
26. Josh Achiam, Steven Adler, Sandhini Agarwal, et al., "GPT-4 Technical Report," preprint, arXiv, March 4, 2024, <https://doi.org/10.48550/arXiv.2303.08774>.
27. Olawale Salaudeen, Anka Reuel, Ahmed Ahmed, et al., "Measurement to Meaning: A Validity-Centered Framework for AI Evaluation," preprint, arXiv, 2025, <https://arxiv.org/pdf/2505.10573>.
28. Joy Buolamwini and Timnit Gebru, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification," *Proceedings of Machine Learning Research* 81, Conference on Fairness, Accountability, and Transparency (February 2018): 1–15, <https://www.media.mit.edu/publications/gender-shades-intersectional-accuracy-disparities-in-commercial-gender-classification>.
29. Dan Boneh, "Cryptography," chap. 3 in *The Stanford Emerging Technology Review 2025*, Hoover Institution and Stanford School of Engineering (Stanford, CA: Board of Trustees of the Leland Stanford Junior University, 2025), <https://doi.org/10.64576/0103>.
30. Sayash Kapoor and Arvind Narayanan, "We Looked at 78 Election Deepfakes. Political Misinformation Is Not an AI Problem," Knight First Amendment Institute, December 13, 2024, <http://knightcolumbia.org/blog/we-looked-at-78-election-deepfakes-political-misinformation-is-not-an-ai-problem>.
31. Neil Perry, Megha Srivastava, Deepak Kumar, and Dan Boneh, "Do Users Write More Insecure Code with AI Assistants?," preprint, arXiv, Cornell University, December 18, 2023, <https://doi.org/10.48550/arXiv.2211.03622>.
32. Alex Heath, "OpenAI Wants ChatGPT to Be a 'Super Assistant' for Every Part of Your Life," *The Verge*, May 30, 2025, <https://www.theverge.com/command-line-newsletter/677705/openai-chatgpt-super-assistant>.
33. Zhenghao Yin, Iris Agresti, Giovanni de Felice, et al., "Experimental Quantum-Enhanced Kernel-Based Machine Learning on a Photonic Processor," *Nature Photonics* 19, no. 9 (2025): 1020–27, <https://doi.org/10.1038/s41566-025-01682-5>.

34. Maja S. Svanberg, Wensu Li, Martin Fleming, Brian C. Goehring, and Neil C. Thompson, "Beyond AI Exposure: Which Tasks Are Cost-Effective to Automate with Computer Vision?," *Future Tech*, Working Paper, January 18, 2024, [https://futuretech-site.s3.us-east-2.amazonaws.com/2024-01-18+Beyond\\_AI\\_Exposure.pdf](https://futuretech-site.s3.us-east-2.amazonaws.com/2024-01-18+Beyond_AI_Exposure.pdf).
35. Claire Cain Miller and Courtney Cox, "In Reversal Because of A.I., Office Jobs Are Now More at Risk," *New York Times*, August 24, 2023, <https://www.nytimes.com/2023/08/24/upshot/artificial-intelligence-jobs.html>.
36. Martin Neil Baily, Erik Brynjolfsson, and Anton Korinek, "Machines of Mind: The Case for an AI-Powered Productivity Boom," Brookings Institution, May 10, 2023, <https://www.brookings.edu/articles/machines-of-mind-the-case-for-an-ai-powered-productivity-boom>.
37. Pranshu Verma and Gerrit De Vynck, "ChatGPT Took Their Jobs: Now They Walk Dogs and Fix Air Conditioners," *Washington Post*, June 5, 2023, <https://www.washingtonpost.com/technology/2023/06/02/ai-taking-jobs/>; Challenger, Gray & Christmas, Inc., "Challenger Report," May 2023, <https://omscgcinc.wpenginepowered.com/wp-content/uploads/2023/06/The-Challenger-Report-May23.pdf>.
38. David Autor, Caroline Chin, Anna M. Salomons, and Bryan Seegmiller, "New Frontiers: The Origins and Content of New Work, 1940–2018," National Bureau of Economic Research, Working Paper 30389, August 2022, <https://doi.org/10.3386/w30389>.
39. "White House Unveils America's AI Action Plan," The White House, July 23, 2025, <https://www.whitehouse.gov/articles/2025/07/white-house-unveils-americas-ai-action-plan/>.
40. Parliament and Council Regulation 2024/1689 Artificial Intelligence Act, art. 6–7, 2024 O.J. L 2024/1689, <https://artificialintelligenceact.eu/section/3-1/>.
41. "The Bletchley Declaration by Countries Attending the AI Safety Summit, 1–2 November 2023," Department for Science, Innovation, and Technology; Foreign, Commonwealth, and Development Office, Prime Minister's Office, November 1, 2023, <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>.
42. "AI Sovereign Compute Infrastructure Program," Government of Canada, May 5, 2025, <https://ised-isde.canada.ca/site/ised/en/ai-sovereign-compute-infrastructure-program>; "UK Compute Roadmap," Department for Science, Technology, and Innovation, July 17, 2025, <https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap>.
43. Eliot Brown and Jin Berber, "SoftBank and OpenAI's \$500 Billion AI Project Struggles to Get Off Ground," *Wall Street Journal*, July 21, 2025, <https://www.wsj.com/tech/ai/softbank-openai-a3dc57b4>.
44. Michael Schuman, "Trump Wants a China Deal That Benefits Him, Not the U.S.," *The Atlantic*, August 14, 2025, <https://www.theatlantic.com/ideas/archive/2025/08/china-us-trade-deal-ai-chips/683855/>.
45. Final Report, National Security Commission on Artificial Intelligence, March 19, 2021, <https://apps.dtic.mil/sti/pdfs/AD1124333.pdf>.
46. Charlotte Hill, Charlotte Allen, Tom Perkins, and Harriet Campbell, "Generative AI in the Courts: Getty Images v Stability AI," Penningtons Manches Cooper, February 16, 2024, <https://www.penningtonslaw.com/news-publications/latest-news/2024/generative-ai-in-the-courts-getty-images-v-stability-ai>.

[www.penningtonslaw.com/news-publications/latest-news/2024/generative-ai-in-the-courts-getty-images-v-stability-ai](https://www.penningtonslaw.com/news-publications/latest-news/2024/generative-ai-in-the-courts-getty-images-v-stability-ai).

47. Michael M. Grynbaum and Ryan Mac, "The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work," *New York Times*, December 27, 2023, <https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html>.

48. Blake Brittain, "Music Labels Sue AI Companies Suno, Udio for US Copyright Infringement," Reuters, June 24, 2024, <https://www.reuters.com/technology/artificial-intelligence/music-labels-sue-ai-companies-suno-udio-us-copyright-infringement-2024-06-24/>.

49. Cade Metz, "Anthropic Agrees to Pay \$1.5 Billion to Settle Lawsuit with Book Authors," *New York Times*, September 5, 2025, <https://www.nytimes.com/2025/09/05/technology/anthropic-settlement-copyright-ai.html>.

50. Blake Brittain, "US Judge Preliminarily Approves \$1.5 Billion Anthropic Copyright Settlement," Reuters, September 25, 2025, <https://www.reuters.com/sustainability/boards-policy-regulation/us-judge-approves-15-billion-anthropic-copyright-settlement-with-authors-2025-09-25/>.

---

## STANFORD EXPERT CONTRIBUTORS

### Dr. Fei-Fei Li

SETR Faculty Council, Sequoia Professor in the Computer Science Department, and Professor, by courtesy, of Operations, Information, and Technology at the Graduate School of Business

### Dr. Christopher Manning

Thomas M. Siebel Professor of Machine Learning, and Professor of Linguistics and of Computer Science

### Anka Reuel

SETR Fellow and PhD Candidate in Computer Science

---

Copyright © 2026 by the Board of Trustees of the Leland Stanford Junior University

This publication reflects updates through December 2025

32 31 30 29 28 27 26      7 6 5 4 3 2 1

Designer: Howie Severson

Typesetter: Maureen Forys

Image credits: Linda A. Cicero/Stanford News and iStock.com/PTC-KICKCAT92 (cover); iStock.com/mofuku (p. 22); iStock.com/wacomka (p. 38); iStock.com/FeelPic (p. 56); iStock.com/JONGHO SHIN (p. 70); iStock.com/Chartchai San-saneeyashewin (p. 88); iStock.com/ArtemisDiana (p. 102); iStock.com/PhonlamaiPhoto (p. 116); iStock.com /imaginima (p. 142); iStock.com/Floriana (p. 156); iStock.com/dima\_zel (p. 170); Tim Griffith (p. 225)