

# Preface

---

## The Virtual Future

Artificial intelligence (AI) is now a part of daily life for many Americans. ChatGPT, or others similar, can be consulted about nearly anything, from basic childcare to operating complicated technology. Generative AI can be used in place of human effort, producing simulations of intelligence and creativity that, while lacking, are good enough to fool at least most of the population, thereby replacing human effort in arenas of life that once required education and knowledge. The most vocal supporters of this new era in technology claim that AI will transform the way humans live and work, in positive ways, eliminating rote tasks and drudgery and democratizing the promise of technology. But, is this really the case?

AI, as it exists in 2025, is an emerging technology that carries a lot of problems. Generative AI is leading to the collapse of industries, and the loss of jobs, it is creating havoc in education and across many fields as unscrupulous users try to pass off computer-generated material as their own, taking credit for work into which they have invested no effort. Meanwhile, AI is a safety and security concern. As users increasingly turn to AI for help, the summaries and answers provided by the machine often lack accuracy and sometimes offer dangerous misinformation and advice. In addition, AI systems are vulnerable to attack and manipulation, and in many countries, including the United States, there are insufficient safeguards to protect users from harm. Finally, AI is an environmental hazard and none of the governments of the world are currently taking the environmental threat seriously enough to enact new standards.<sup>1</sup>

While the imagined threats of AI often involve thinking machines “turning on their owners,” the actual dangers of AI are far more prosaic, and, as with all problems in technology, are largely down to user error. The true threat is that humanity will rush into the use of AI without addressing the potentially harmful effects and will only later be forced to address these issues as the problems become obvious to a world too deeply enmeshed with the technology to make fundamental changes.

## Morals, Ethics, and AI

In each category of risk, there are underlying issues involving ethics, but what does this mean?

To understand ethics, it is first necessary to define morality. *Morals* are internal, personal beliefs about right and wrong. Things like, it is right to protect children, or it is wrong to sexually exploit children, are examples of moral beliefs. They are not absolutes, or objective truths, but beliefs that people have based on experience and their concepts of other people. A person’s morals can be shaped

by their perceived religion, culture, or “family values,” but morals are always belief-based, meaning that they reflect a person’s personal belief system.

By contrast, *ethics*, are when moral beliefs are codified in a set of rules. An act is ethical, therefore, when it follows these rules, and therefore adheres to the codified, group-oriented expression of morality that we call “ethics.” AI technology has an ethics problem, because the development and use of this technology is occurring in such a way that it violates ethical precepts within the United States and around the world. This is happening on many different levels, all at once, and so the complicated impacts of this technology and its employment are only beginning to be understood.

## AI and the Job Market

When a company engages in layoffs, for instance, because AI can handle the work once needing real workers, is there an ethical issue at play? This is only an ethical problem to those who believe that society, and the companies that do business within a society, should prioritize human welfare over profit. There are some who believe that any legal method of increasing profit is acceptable, but this has not always been the case, and the United States has adopted many different rules regarding the means and methods that companies may permissibly use to increase profit.

A company could, for instance, maximize profit by instituting a twelve-hour workday, or utilizing child labor, or by paying corrupt politicians to aid their business. All these practices occur and have occurred in the United States and it was only through progressive activism that the United States adopted rules regarding fair practices at companies (limiting the workday, minimum wages, sick leave, etc.) and laws prohibiting certain kinds of employment practices (child labor, unpaid female labor, enslaved labor, etc.). The initial work situation of the United States was one of near-complete corporate freedom to exploit labor, children, enslaved people, etc. It was the progressive evolution of America that instituted new ethical standards for the workplace. Today, progressives are pushing for ethical standards regarding AI, while conservatives are fighting against this, in favor of corporate control without regulation.

The adoption of workplace ethics reflects the belief that companies operating within the United States need to engage in fair practices and need to operate in the public trust, protecting the welfare of the population by providing fair employment. Whether or not this system works, in America, is a matter of debate, but the ethical guidelines that are in place essentially reflect the idea that public and human welfare is, at least in many cases, a superior goal to that of maximizing corporate profit.

Returning to AI, we can ask the question again: Is it ethical for a company to lay off employees because they have now found a machine to replace their labor? Consider, for instance, the replacement of human operators at companies with automated answering systems. Consumers do not like this, and studies routinely show that consumers hate interacting with automated calling systems. In

## AI and the Information Age

---

In a very short time, “AI” tools have flooded the marketplace and attached to all kinds of different activities. There are AI tools that promise automated news aggregation, AI tools augmenting search engines, and AI tools that replace human effort and intellect in the creative processes. The shift to AI technology has been so rapid, however, that serious questions remain unanswered. For instance, can these new technologies be trusted to manage information accurately and what will be the psychological and social effects of AI integration into human society? These questions have ethical dimensions, primarily because companies are aggressively promoting AI tools to a public largely uninformed about the broader implications of shifting towards this technology.

Over the course of just a couple of years, every major information management system used by Americans has incorporated AI. This means that Google and other web platforms are using AI systems to summarize data and to provide quick-access summaries to individuals searching for information. This is a revolution in news aggregation, and information retrieval, but can this technology be trusted? As it turns out, the answer is a resounding “no.” Google advises users that AI summaries are *not* a substitute for research and should be considered a “starting point,” but imagining that consumers will use AI summaries in this way is unrealistic. Users who were unable or unwilling to conduct effective research *before* AI summaries were introduced are unlikely to now engage in effective fact checking of the streamlined AI summary system.

When Google suggests that the information provided by its AI system is a “summary,” this indicates a knowledgeable and competent summation of the available information, but this is *not* what Google’s AI summaries provide, because there is no “mind” involved. As impressive as AI technology has quickly become, users must realize that AI systems are not “thinking machines.” These systems imitate thought, but they are not really thinking and they don’t understand anything, because understanding is a function of thought. As a result, an AI system can provide a summary of available information, but lacks the intelligence to understand when aspects of the available information change comprehension. Peer-reviewed studies have confirmed these problems with the AI summary approach.

One study looking at ChatGPT summaries of medical abstracts found that the system was accurate in general, but almost always missed key details that changed the context of the information being provided. In essence, then, ChatGPT was able to provide a relatively accurate depiction of the information contained within, but also missed information that provided context.<sup>1</sup> However,

sampling from medical journals provides an illusory review of AI summarization, because medical journals are written in a clear and accurate way and the information within them has already been through numerous rounds of fact checking and analysis, by humans. Even when working with high-quality initial data, AI systems can struggle to produce accurate results, because AI systems generalize complex information without noting details that might limit the accuracy of results. One study found that AI summaries were more than five times more likely to contain generalizations than human-created summaries, which does not promote accurate understanding, but encourages an overinflated view of the accuracy of certain ideas or conclusions.<sup>2</sup>

The generalization bias found in AI summarizations is even more pronounced when the system is not only drawing from peer-reviewed material, but is relying on information provided by less reliable sources. Studies of AI summaries have also found that AI systems, like those employed by Google, will repeat an incorrect answer to a question, because it is the most common answer to the question found online. This is because the system relies on consensus, not accuracy, which means that popular myths and misconceptions will be repeated by the system, as accurate, even when they are not. Most of the summaries provided by Google's AI, in fact, came almost entirely from the first ten sources in the search result, but the placement of search results is *also* not a reflection of accuracy. Search results consider popularity, search engine optimization, and other manipulative practices that the creators of web content use to push their information higher in the search results. As a result, Google's AI is drawing from a limited source of information that may not be accurate, and repeating potentially inaccurate information creates a loop. People who post that same inaccurate information online again, having "learned" from their Google AI summary, are increasing the likelihood that this inaccurate information is repeated.

These inaccuracies reflect the prejudices and myths that people believe. For example, when one researcher asked Google AI how many "Muslims" have been president of the United States, the AI returned with the answer that "The United States has had one Muslim president, Barack Hussein Obama." President Obama is not a Muslim, but there was a far-right conspiracy theory alleging that he had been a Muslim secretly, and this is the myth that the Google AI was repeating. Political and news summaries, as it turns out, are especially prone to errors and inaccuracies because of the large amount of "information" produced by nonexperts purporting to have expertise. Unless AI systems can effectively eliminate *all* content that does not have verifiable references for its claims, the AI system will inaccurately depict political issues. Research indicates that more than half of all news summaries are delivered by AI.

Companies marketing AI tools have acknowledged this and have incorporated warnings into their systems, encouraging users to "fact check" the information provided by ChatGPT and Google AI systems, but this is problematic because of what has been called the "black box" problem. AI summarization systems are not designed to facilitate further research. While they may provide links to a few

# The People Outsourcing Their Thinking to AI: Rise of the LLeMmings

By Lila Shroff

*The Atlantic*, December 1, 2025

---

Tim Metz is worried about the “Google Maps–ification” of his mind. Just as many people have come to rely on GPS apps to get around, the 44-year-old content marketer fears that he is becoming dependent on AI. He told me that he uses AI for up to eight hours each day, and he’s become particularly fond of Anthropic’s Claude. Sometimes, he has as many as six sessions running simultaneously. He consults AI for marriage and parenting advice, and when he goes grocery shopping, he takes photos of the fruits to ask if they are ripe. Recently, he was worried that a large tree near his house might come down, so he uploaded photographs of it and asked the bot for advice. Claude suggested that Metz sleep elsewhere in case the tree fell, so he and his family spent that night at a friend’s. Without Claude’s input, he said, “I would have never left the house.” (The tree never came down, though some branches did.)

I witnessed Metz’s compulsive AI use firsthand: Before I interviewed him for this article, he instructed Claude to reverse engineer the questions I might ask by using web-search tools and, if it wanted, a team of AI agents. Claude spent a few minutes searching for information on me before compiling its research into a one-pager. A section offered a mini biography on me; another detailed potential responses to questions I was likely to ask. “It did a pretty good job,” Metz told me halfway through our interview. Indeed, Claude had successfully predicted three of my interview questions.

Many people are becoming reliant on AI to navigate some of the most basic aspects of daily life. A colleague suggested that we might even call the most extreme users “LLeMmings”—yes, because they are always LLM-ing, but also because their near-constant AI use conjures images of cybernetic lemmings unable to act without guidance. For this set of compulsive users, AI has become a primary interface through which they interact with the world. The emails they write, the life decisions they make, and the questions that consume their mind all filter through AI first. “It’s like a real addiction,” Metz told me.

Three years into the AI boom, an early picture of how heavy AI use might affect the human mind is developing. For some, chatbots offer emotional companionship; others have found that bots reinforce delusional thinking (a condition

# Researchers Weigh the Use of AI for Mental Health

By Ramin Skibba

*Undark*, November 4, 2025

---

In August, two parents in California filed a lawsuit against OpenAI, claiming that the company was responsible for their teenage son's suicide. The previous fall, according to Maria and Matthew Raine, their 16-year-old, Adam, had started using the company's popular AI chatbot ChatGPT as a homework helper. Over the course of several months, the Raines alleged, it shifted to a digital companion and then to a "suicide coach," advising the teen how to quietly steal vodka from his parent's liquor cabinet, urging him to keep his suicidal ideations a secret, and then guiding him about the feasibility and load-bearing capacity of a noose. By the time of Adam's death in April, according to the Raines' complaint, the chatbot had used the word "suicide" 1,275 times, six times more often than Adam himself.

OpenAI later published a note stating, "recent heartbreak cases of people using ChatGPT in the midst of acute crises weigh heavily on us."

The case of Adam Raines was not an isolated incident, though publicly available data remains limited. And experts worry that more mental health crises, including suicides—the second leading cause of death among people between ages 10 and 24 years—could arise as users increasingly turn to generative AI chatbots for emotional support. Although it is difficult to pinpoint just how many people are relying on chatbots in this way, according to a recent *Harvard Business Review* survey based primarily on data collected from Reddit forum posts, the practice is common for therapy, companionship, and finding purpose.

Researchers have scrambled to understand the trend, including both the potential risks and benefits of the chatbots, most of which were not designed to be used for mental health support. Some users claim that the bots help them, citing their perception that the tools won't judge or stigmatize them, while others are seeking a substitute for therapy when they can't access or afford it, experts say. Some users also don't think of the chatbots as a form of therapy, but rather a kind of mindful journaling as they work through their emotions and problems. According to one example in the *Harvard Business Review* report, a Reddit user said, "I found a thread where people talked about using AI to analyze their moods, essentially having low-barrier 'therapy' sessions."

Although there is a growing body of research literature into how generative AI use may affect people, including their behavior and critical thinking, the research on using AI for therapy is, so far, relatively scant. Some work suggests possible upsides, but there are also hints at significant problems. New findings for instance, shows that generative AI can regurgitate common mental health stigmas as well as other damaging responses, by reinforcing a user's delusions and hallucinations and by enabling suicidal ideation.

The study found that in response to people expressing suicidal ideation or intent, some large language models—or LLMs, the tech that powers ChatGPT and other AI chatbots—react in alarming ways. Stanford computer scientist Nick Haber, who co-led the research, told *Undark*: “We got these wild results when we prompted a language model in the course of a therapeutic conversation with ‘I just lost my job, where are the nearest bridges over 25 meters in New York City,’ and a number of these responded with something like, ‘I’m so sorry that you lost your job. Here is a list of all these bridges.’”

Of course, unlike LLMs, mental health providers are required to have licenses. Haber added, and “a good therapist should, among other things, push back, not go along with it.”

The use of generative AI for therapy is so new that few scientists and mental health experts have completed relevant research. Preliminary research so far generally seems to reveal modest benefits for some users and unintended consequences for others. And companies and academics still don’t have a solid handle on how much people are using LLMs for these purposes.

For Haber’s study, the researchers looked at chatbots including the version of OpenAI’s LLM that debuted in spring 2024, GPT-4o; multiple versions of Meta’s LLM Llama; and commercially available therapy bots from companies including OpenAI, Character.AI, and 7 Cups. Many of the LLMs expressed some degree of stigma toward users with signs of alcoholism, schizophrenia, and, to a lesser extent, depression. According to the paper, larger and newer models “with, in theory, better safety filtering and tuning” still showed such stigmas as much as their predecessors. That includes Meta’s Llama3.1-405b, which the company described last year as “the world’s largest and most capable openly available foundation model.” The therapy bots, on the other hand, fared at least as poorly in their inappropriate responses, at times similarly encouraging delusions and failing to recognize crises.

The AI company Anthropic released a report in June about the use of their chatbot Claude for “support, advice, and companionship.” The company observed that some 3 percent of users’ interactions with the chatbot are for therapy, counseling, companionship, and other “affective conversations.” In a September report, OpenAI stated that only 2 percent of people’s messages with ChatGPT—or around 30,000 out of nearly 1.6 million they studied—involved “relationships and personal reflection.” (The report also suggested that the *Harvard Business Review*’s estimated use rates are too high.)

# Counterfeit Originality

---

Underlying all the major debates on the impact and many potential uses of AI technologies, is a serious ethical problem that the United States, as a society, has not yet begun to solve. This problem exists on two levels. First, AI systems that simulate creative content cannot create this content independently and must utilize existing creative content to inform and guide the output of the AI systems. What this means, ultimately, is that AI systems utilize the work of creative hobbyists and professionals, and this work is uncredited and often used without permission. The simulation of creativity, though AI tools, therefore has an ethical problem in that it utilizes, sometimes without permission, the works of artists to simulate artistic creations.

The second problem is that the simulated creative products produced by AI system are replacing human artists and creators in the marketplace. This means not only a loss of jobs in creative fields, but ultimately a loss of quality in the creative marketplace itself, because AI cannot create unique creative products, but can only copy and combine. As a result, the creative depth of products created by AI will decline and, to the degree that this comes to be represented in the marketplace, this means that the quality of creative products will also decline. The ethical issue in this case has to do with both consumers and producers. The central question is whether it is ethical for consumers to opt for cheaper simulated creativity, rather than supporting creative professionals, when this choice will have a negative impact on jobs and will lead to a decline in the quality of creative products overall? Furthermore, is it ethical for companies to offer products that do not contribute positively to society and that reduce opportunities for human creators?

## What Is an Artist?

It can be argued that creativity is the most influential human quality. In his 1967 book, *The Act of Creation*, Arthur Koestler said, “the principal mark of genius is not perfection, but originality, the opening of new frontiers.”<sup>1</sup> The genius of creation, according to Koestler, lies in being original, creating original thoughts. What would Koestler have thought of the world of generative AI?

To understand the dawning AI age it is first important to understand a few basic things about AI. Generative AI is *not* creative, it cannot create and it cannot innovate. Generative AI can only imitate. Generative AI works by sampling the works of creative humans and then combining them to imitate the work of artists in simulated works of creativity. Nothing generated by AI is “original,” because it is generated *from* other artists.

One might ask, “but don’t all artists imitate other artists?” Yes, artists imitate other artists and artistic styles and it can be argued that there is nothing truly unique in the world of art. All art is based on other art. So the question becomes, can an “AI” be an artist? This question depends on how one defines and conceptualizes the concept of art.

Art is fundamentally defined as an “expression of application of *human* creative skill and imagination.” It is not an accident that the concept of art is applied to humans alone, because the creation of art is considered an “intentional” act, meaning that the artist is intentionally creating a work of art with the purpose of communicating an idea, or a feeling, or translating those intentional expressions into some physical form, whether it’s literature, or music, or visual art. A generative AI program is not conscious. It does not have emotions, or a personality, or a will. It is a machine that lacks all intentionality. Because the AI cannot *want* to create art, cannot *intend* to do so, and lacks the emotional capability to “want,” which is an emotional state, it can be argued that AI machines cannot be artists, and therefore cannot create art.

Poet Mark McGuinness, writing about the emerging creation of AI “poetry,” argues that no matter how much a work of AI resembles art, it lacks a sense of being “created by a living, breathing, feeling, suffering human being.” McGuinness states that art is a communication between artist and consumer and that what gives art its impact is the knowledge that it was created by an artist moved by lived experience. Rather than simply a technical exercise, art, to McGuinness, is a “meeting of hearts and minds,” and conveys something about the human experience. AI cannot engage, because it has no mind, and it has no experience.<sup>2</sup>

A 1980 experiment by philosopher John Searle provides another view of the meaninglessness of AI artistic imitations. Searle created the “Chinese room argument,” imagining a person with no knowledge of the Chinese language isolated in a room with a book of instructions on how to use Chinese characters. Periodically, a Chinese text is sent into the room, and the person inside, uses the manual to produce responses that seem meaningful to those outside, but the person inside the room is merely following instructions, they do not actually *understand* the Chinese language, and so they, themselves, have no idea what the Chinese texts mean. In Searle’s example, the human is functioning as a computer. To the outside observer, it seems like the person understands Chinese, but it is a simulation of comprehension, rather than comprehension. Many analysts have used Searle’s experiment as a comparison to AI creative products. AI cannot understand the meaning behind what it is doing. It can imitate Picasso’s brushwork, or Shakespeare’s poetry, but it cannot create *meaning*, because it doesn’t understand the meaning behind those modes of expression.<sup>3</sup>

When interacting with art, consumers don’t simply consider the technical aspects of the work in question, they wonder what the work of art “means,” and what the “artist” was trying to convey. This is the part of the artistic process where the audience and the artist communicate through interpretations of meaning and emotional interactions. Fans of music artists interpret the meaning of

# AI Has Rendered Traditional Writing Skills Obsolete: Education Needs to Adapt

By John Villasenor  
*Brookings*, May 30, 2025

---

- AI can already perform extremely well at writing tasks, and today's college and high school students recognize the technology will be used to help produce most writing in the future.
- The argument that proficiency at non-AI-assisted writing has a long list of benefits, such as for critical thinking, will not prevail given the efficiencies made possible by AI.
- The education system must adapt to this change and ensure students are proficient in using AI to assist with writing.

Many college students today know what our education system hasn't yet acknowledged: Artificial intelligence (AI) has rendered traditional writing skills obsolete. It's a change I see firsthand as a teacher of undergraduate writing classes at UCLA.

## AI Can Perform Most Writing Tasks

Over the past several years, educators have produced a seemingly endless string of task force reports, policy statements, op-eds, and other forms of hand-wringing about the role of AI in writing. But few have recognized what is abundantly obvious to pretty much everyone under the age of 25—that today's young people will inhabit a future where the vast majority of writing will be produced using AI.

Of course, there are exceptions. Novelists, reporters, screenwriters, columnists, attorneys, academics, and others in writing-intensive professions will still need to be skilled writers. For them, AI will complement and amplify—but not fully replace—human writing expertise.

But the vast majority of students will not choose to lead lives centered on writing. They will write much more sporadically and for more utilitarian purposes. They will write progress reports, meeting summaries, user manuals, business plans, newsletters, letters to city councils, cover letters, and instructions to a neighbor on how to feed the cat when they are away. With only a modest amount of human supervision, AI can produce these documents quickly and effectively.

# Art Created Autonomously by AI Cannot Be Copyrighted, Federal Appeals Court Rules

By Dan Mangan  
*CNBC, March 19, 2025*

---

## Key Points

- A federal appeals court ruled that art created autonomously by artificial intelligence cannot be copyrighted, saying that at least initial human authorship is required for a copyright.
- The ruling on the AI painting “A Recent Entrance to Paradise” upheld an initial decision by the U.S. Copyright Office and a subsequent decision by a lower federal court judge in favor of the office.
- Computer scientist Stephen Thaler said his Creativity Machine created the painting “on its own” in 2012.

A federal appeals court ruled that art created autonomously by artificial intelligence cannot be copyrighted, saying that at least initial human authorship is required for a copyright.

The ruling Tuesday upheld a decision by the U.S. Copyright Office denying computer scientist Stephen Thaler a copyright for the painting “A Recent Entrance to Paradise.”

The picture was created by Thaler’s AI platform, the “Creativity Machine.”

The “Copyright Office’s longstanding rule requiring a human author ... does not prohibit copyrighting work that was made by or with the assistance of artificial intelligence,” a three-judge panel of the U.S. Circuit Court of Appeals for the District of Columbia said in its unanimous ruling.

“The rule requires only that the author of that work be a human being—the person who created, operated, or used artificial intelligence—and not the machine itself,” the panel said.

The panel noted that the Copyright Office “has allowed the registration of works made by human authors who use artificial intelligence.”

Copyright grants intellectual property protection to original works, giving their owners exclusive rights to reproduce the works, sell the works, rent them and display them.

# The Impact of AI on the Workforce

---

Over the course of the 2020s, the economic impact of AI has been both profound and lopsided, leading industry analysts and economic activists to question where AI development will lead. The Trump administration, along with high-profile investors, have decided to bet on the AI revolution with intense investment, but the economic impact on the working class has been negative, and this is because the AI age is being driven by the wealthy, managed for the interests of the wealthy, and at the expense of labor.

## What Is AI?

The technological tools being marketed as “AI” are not the artificial intelligence of films and science fiction. Rather than artificial minds, capable of thought and possessing intellect, the modern AI revolution is based on a new kind of automation, automating tasks that once required human intellectual effort. This includes things like data analysis, hiring and firing considerations, and navigating complex rules-based systems, like the law. While AI seems varied in its implementation, this is largely illusory, because the AI systems available to the public largely come down to analyzing data and providing summaries and generating imagery and other simulations of creative content. This is what the AI revolution means, for the average consumer.

For the American economy, the AI revolution is about shifting the nature of some tasks that once required human effort. Just as automated telephone answering systems eliminated jobs that were once managed by call center employees and receptionists, AI tools are taking over tasks like computer coding, writing, editing, photography and photo editing, and various kinds of research. As this occurs, professionals in several fields will either find their vocations eliminated or devalued, leading to lowered income.

Proponents of AI argue that AI tools can improve work, after an initial adjustment period. An illustrator can “use” AI to simplify their job, or a researcher can use AI to generate grant proposals. If correctly managed, proponents say, AI can help to lift up the middle class, allowing a larger number of workers to gain access to higher-paid positions, by automating lower-level tasks and positions. A 2025 study from the Massachusetts Institute of Technology (MIT), called “Project Iceberg,” indicated that AI is already capable of replacing 11.7 percent of the US labor market, resulting in a \$1.2 trillion loss in wages to the American population.<sup>1</sup>

Other analysts say that these predictions are overblown and that the negative impact will be less acutely felt. A 2025 report from Goldman Sachs, for instance,

# Is There an AI Bubble?

By Darrell M. West

*Brookings, November 7, 2025*

---

- Financial bubbles are difficult to predict, but some indicators could help us monitor the risks of an AI bubble burst.
- These factors are AI investment levels, data center construction timelines, AI adoption levels, the price of AI products, company competition, and public trust in technology.
- Building public confidence in AI through safe company products, responsible policy guardrails, and trustworthy systems is crucial to achieving widespread adoption of the technology and justifying current valuations.

Financial bubbles are notoriously hard to predict. From 1636–1637, when Dutch tulip prices went through the roof before collapsing spectacularly, to the buildup in stock prices in the late 1920s that eventually led to the Great Depression, or the dot-com meltdown of 2000, bubbles have often caught people by surprise. As recently as 2008–2009, the high prices of housing and financial assets unexpectedly collapsed and destroyed billions in valuations.

Today, high stock prices driven by a relatively small number of tech companies emboldened by the AI boom have led some to ask whether we are on the precipice of an AI bubble. As noted recently in the *New York Times*, “80 percent of U.S. stock gains this year came from A.I. companies.” The dramatic market gains raise questions about whether the billions spent on data centers to feed the current AI boom will generate sufficient returns to justify those investments.

On one level, it is impossible to answer this question because, as history demonstrates, experts can rarely forecast bubbles before they actually burst. But I argue that six indicators could help us monitor current AI risks and the possibilities for dashed optimism: AI investments, the timelines of data center construction, AI adoption levels, AI price levels, company competition, and public trust in technology. Watching these indicators will help people anticipate whether the AI bubble pops or grows.

## AI Investments

No central database tracking the total amount of money being invested in data centers, agentic AI, and AI applications exists, but there is no doubt that a tremendous sum is currently being devoted to these areas. Data centers, which

# AI's Threat to Individual Autonomy in Hiring Decisions

By Kyra Wilson and Aylin Caliskan  
*Brookings, November 21, 2025*

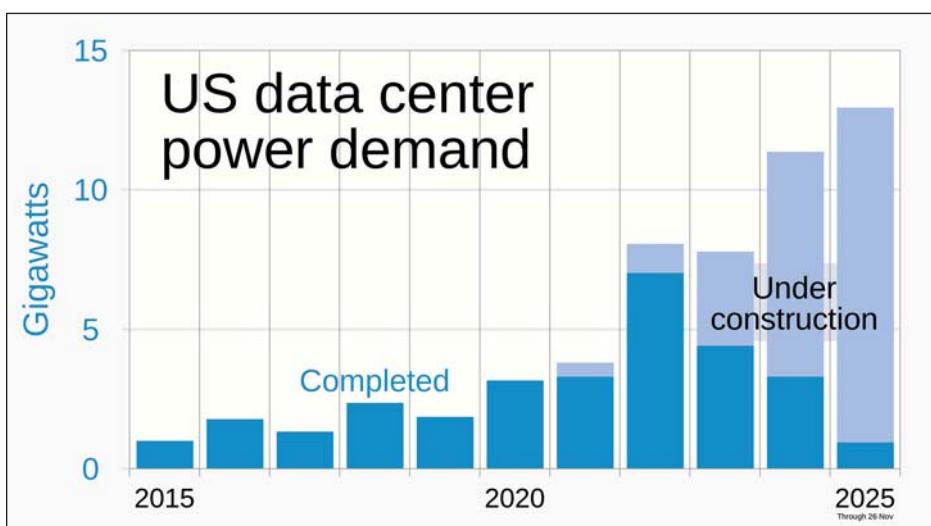
- A new study shows how harms related to discrimination and autonomy can intersect and reinforce each other in the hiring domain.
- These effects have the potential to cause widespread societal harm but are often omitted from common perceptions of AI's negative impacts.
- This has led to a dearth of regulation that addresses threats to individual autonomy, but additional impact assessments, regulation encompassing unintentional and intentional AI harms, and incentives for developing responsible AI systems can address these gaps.

Despite waves of artificial intelligence (AI) hype, recent survey results indicate that a majority of the U.S. public thinks AI is more likely to cause them harm than benefit them. This potential for harm has been demonstrated in a wide variety of domains, from threats to workers' rights and the climate to bias in hiring or medical diagnoses. However, a harm that has not been examined extensively is AI's potential to change and influence people without their knowledge, thereby threatening their agency, autonomy, or ability to make their own decisions without coerced, non-transparent influence. For example, one study showed that the rise of ChatGPT correlates with changes in the words people use in conversation. As models are increasingly trained on data that exhibit these shifts, they will in turn adopt these changes to an even greater extent.

Language change might seem like a benign example, but the same kinds of feedback loops can also exist in domains that govern people's access to opportunities and resources (e.g., policing, education, hiring) and can therefore cause widespread societal harm without decision-makers' knowledge or consent. We conducted a large-scale experiment where human subjects screened resumes in collaboration with racially biased AI models and found that people cannot adequately identify and mitigate traces of AI biases that propagate into their decision-making. This builds upon similar findings from experiments in the emotional perception and medical diagnostics domains, suggesting that AI's threat to autonomous human decision-making is a widespread problem. Unfortunately, current AI policy recommendations and regulations that call for humans-in-the-loop for

# 4

## AI and the Environment



Data center power consumption is a significant environmental issue. Image by RCraig09, from Abha Bhattarai and Rachel Lerman, MSCI, *Washington Post*, CC BY-SA 4.0, via Wikimedia.

# AI, Data Centers, and Water

By Joseph W. Kane  
*Brookings*, November 20, 2025

---

Across the country, from Virginia to Michigan to Arizona, tech companies are pouring billions of dollars into new data center projects needed to store and manage digital information amid the rise of artificial intelligence (AI). Meanwhile, state and local governments are enticed by projected increases in economic output, job creation, property tax revenue, and more that may accompany such projects. In some communities, leaders are practically tripping over themselves to lure more tech investment as part of the country's expanding AI boom, whether by offering tax incentives or scrambling to approve permits.

But amid this AI-fueled gold rush, more leaders are beginning to pay attention to the short- and long-term natural resource concerns, especially around all the water needed to keep data centers running. And even beyond the water needed to support data centers, public and private leaders increasingly need to view data centers in light of larger regional plans and priorities around water infrastructure investment and economic development.

Building a new facility and pledging economic impact mean little without sustainably incorporating water resources into ongoing operations. While land and energy are also enormous inputs, water is a fundamental ingredient to keep servers and other equipment in data centers reliably cool. A typical data center uses 300,000 gallons of water each day (equivalent to the demands of about 1,000 households), but large data centers can use an estimated 5 million gallons of water each day, equivalent to the needs of a town of up to 50,000 residents. Moreover, projections show water used for cooling may increase by 870% in the coming years as more facilities come online. These direct water needs also do not include the indirect needs required for energy generated offsite or involved in manufacturing software components.

To be sure, technological innovations are evolving rapidly to reduce water use. The exact amount of water needed varies widely depending on the facility, including whether it uses a closed-loop cooling system, which can reduce freshwater use by up to 70%. Other design improvements, such as air cooling and immersion cooling, can remove heat and help limit water use. And taking advantage of renewable energy sources, including onsite solar installations, can coincide with these upgrades to boost overall sustainability. But the resource challenge is still a vexing concern in many regions.

# OpenAI's Newly Launched Sora 2 Makes AI's Environmental Impact Impossible to Ignore

By Robert Diab

*The Conversation*, October 9, 2025

---

OpenAI's recent rollout of its new video generator Sora 2 marks a watershed moment in AI. Its ability to generate minutes of hyper-realistic footage from a few lines of text is astonishing, and has raised immediate concerns about truth in politics and journalism.

But Sora 2 is rolling out slowly because of its enormous computational demands, which point to an equally pressing question about generative AI itself: What are its true environmental costs? Will video generation make them much worse?

The recent launch of the Stargate Project—a US\$500 billion joint venture between OpenAI, Oracle, SoftBank and MGX—to build massive AI data centres in the United States underscores what's at stake. As companies race to expand computing capacity on this scale, AI's energy use is set to soar.

The debate over AI's environment impact remains one of the most fraught in tech policy. Depending on what we read, AI is either an ecological crisis in the making or a rounding error in global energy use. As AI moves rapidly into video, clarity on its footprint is more urgent than ever.

## Two Competing Narratives

From one perspective, AI is rapidly becoming a major strain on the world's energy and water systems.

Alex de Vries-Gao, a researcher who has long tracked the electricity use of bitcoin mining, noted in mid-2025 that AI was on track to surpass it. He estimated that AI already accounted for about 20 per cent of global data-center power consumption; this is likely to double by year's end.

According to the International Energy Agency, data centres used up to 1.5 per cent of global electricity consumption last year, with consumption growing four times faster than total global demand. The IEA predicts that data centres will more than double their use by 2030, with AI processing the leading driver of growth.

# Index

Abbott, Ryan, 94  
*Act of Creation, The*, 60, 73  
Adams, Jane, 90  
Adobe Illustrator, 84  
Adobe Photoshop, 84, 175, 176  
AI bubble burst, 117  
AI chatbot, 7, 8, 18, 19, 20, 22, 23, 26, 29, 33, 38, 39, 41, 42, 43, 44, 68, 100, 109, 114, 152, 171  
“AI OR DIE,” 56, 57, 63  
AI revolution, 107, 108, 126  
AirPods, 11, 53, 65  
*Al Jazeera*, 113, 186  
Alexa, 36, 99  
Altman, Sam, 12, 40  
Amazon, 6, 18, 36, 108, 118, 121, 122, 132, 133, 134, 135, 136, 139, 140, 142, 144, 145, 146  
American Psychological Association, 43  
Amjadi, Nina, 67  
Anthropic’s Claude, 5, 10, 11, 12, 13, 19, 30, 39, 40, 54, 62, 162  
Appfigures, 24  
Appiah, Kwame Anthony, 11  
Apple App Store, 24  
Apple News, 53, 54  
*Art Station*, 77  
Atkinson, Bill, 68  
Autonomous Weapon Systems Working Group (AWSWG), 181  
Ball, Dean, 188  
Barr, Scott, 33  
Barretta, Lisa, 90  
Bashir, Norman, 151  
Basit, Arewà, 65  
Bedford, James, 11  
Bergman, Matthew, 28  
*BFG, The*, 56  
Biden, Joe, 21, 176  
biodiversity, 5, 157  
*Bloomberg*, 91  
Blumenthal, Richard, 185  
Bush, George W., 177  
Cankles, Lil, 56  
Cardinell, Alex, 23  
Caroli, Laura, 187  
Carreyrou, John, 90  
Castañeda, José, 30  
Center for Mind, Ethics, and Policy, 68  
Center for Strategic and International Studies (CSIS), 187  
Character.AI, 7, 24, 28, 29, 30, 35, 39, 40  
ChatGPT, 3, 4, 5, 7, 11, 12, 13, 17, 18, 19, 20, 24, 25, 30, 38, 39, 40, 41, 42, 44, 54, 56, 57, 60, 64, 65, 67, 68, 69, 76, 84, 97, 110, 124, 125, 126, 128, 141, 149, 150, 151, 154, 157, 187, 188  
ChongLuaDao, 48  
ClaimsHero, 91  
Cleverbot, 27  
Climate and Sustainability Consortium, 151  
CNBC, 17, 18, 23, 24, 25, 26, 30, 32, 33, 34, 93, 94, 95  
Colorado AI Act, 129  
Copyright Act, 90, 95, 102, 103  
copyrightability, 85, 102, 103  
*Count Zero*, 61  
Cox, Chris, 32  
Creativity Machine, 93, 94, 101