




McKinsey Explainers

What is generative AI?

Generative artificial intelligence (AI) describes algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images, text, simulations, and videos. Recent breakthroughs in the field have the potential to drastically change the way we approach content creation.



In the months and years since ChatGPT burst on the scene in November 2022, generative AI (gen AI) has come a long way. Every month sees the launch of new tools, rules, or iterative technological advancements. While many have reacted to ChatGPT (and AI and machine learning more broadly) with fear, machine learning clearly has the potential for good. In the years since its wide deployment, machine learning has demonstrated impact in a number of industries, accomplishing things like [medical imaging analysis](#) and high-resolution weather forecasts. A 2022 McKinsey survey shows that AI adoption has [more than doubled](#) over the past five years, and investment in AI is increasing apace. It's clear that generative AI tools like ChatGPT (the GPT stands for generative pretrained transformer) and image generator DALL-E (its name a mashup of the surrealist artist Salvador Dalí and the lovable Pixar robot WALL-E) have the potential to change how a range of jobs are performed. The full scope of that impact, though, is still unknown—as are the risks.

Still, organizations of all stripes have raced to incorporate gen AI tools into their business models, looking to capture a piece of a sizable prize. McKinsey research indicates that gen AI applications stand to add [up to \\$4.4 trillion](#) to the global economy—annually. Indeed, it seems possible that within the next three years, anything in the technology, media, and telecommunications space *not* connected to AI will be considered [obsolete or ineffective](#).

But before all that value can be raked in, we need to get a few things straight: What is gen AI, how was it developed, and what does it mean for people and organizations? Read on to get the download.

To stay up to date on this critical topic, sign up for email alerts on “artificial intelligence” [here](#).

Learn more about [QuantumBlack](#), AI by McKinsey.

What's the difference between machine learning and artificial intelligence?

[Artificial intelligence](#) is pretty much just what it sounds like—the practice of getting machines to mimic human intelligence to perform tasks. You've probably interacted with AI even if you don't realize it—voice assistants like Siri and Alexa are founded on AI technology, as are customer service chatbots that pop up to help you navigate websites.

Machine learning is a type of artificial intelligence. Through machine learning, practitioners develop artificial intelligence through models that can “learn” from data patterns without human direction. The unmanageably huge volume and complexity of data (unmanageable by humans, anyway) that is now being generated has [increased machine learning's potential](#), as well as the need for it.

What are the main types of machine learning models?

Machine learning is founded on a number of building blocks, starting with [classical statistical techniques](#) developed between the 18th and 20th centuries for small data sets. In the 1930s and 1940s, the pioneers of computing—including theoretical mathematician Alan Turing—began working on the basic techniques for machine learning. But these techniques were limited to laboratories until the late 1970s, when scientists first developed computers powerful enough to mount them.

Until recently, machine learning was largely limited to predictive models, used to observe and classify patterns in content. For example, a classic machine learning problem is to start with an image or several images of, say, adorable cats. The program would then identify patterns among the images, and then scrutinize random images for ones that would

match the adorable cat pattern. Generative AI was a breakthrough. Rather than simply *perceive* and *classify* a photo of a cat, machine learning is now able to *create* an image or text description of a cat on demand.

How do text-based machine learning models work? How are they trained?

ChatGPT may be getting all the headlines now, but it's not the first text-based machine learning model to make a splash. OpenAI's GPT-3 and Google's BERT both launched in recent years to some fanfare. But before ChatGPT, which by most accounts works pretty well most of the time (though it's still being evaluated), AI chatbots didn't always get the best reviews. GPT-3 is "by turns super impressive and super disappointing," said *New York Times* tech reporter Cade Metz in a video where he and food writer Priya Krishna asked GPT-3 to write recipes for a (rather disastrous) Thanksgiving dinner.

The first machine learning models to work with text were trained by humans to classify various inputs according to labels set by researchers. One example would be a model trained to label [social media](#) posts as either positive or negative. This type of training is known as supervised learning because a human is in charge of "teaching" the model what to do.

The next generation of text-based machine learning models rely on what's known as self-supervised learning. This type of training involves feeding a model a massive amount of text so it becomes able to generate predictions. For example, some models can predict, based on a few words, how a sentence will end. With the right amount of sample text—say, a broad swath of the internet—these text models become quite accurate. We're seeing just how accurate with the success of tools like ChatGPT.

What does it take to build a generative AI model?

Building a generative AI model has for the most part been a major undertaking, to the extent that [only a few well-resourced tech heavyweights have made](#)

[an attempt](#). OpenAI, the company behind ChatGPT, former GPT models, and DALL-E, has billions in funding from bold-face-name donors. DeepMind is a subsidiary of Alphabet, the parent company of Google, and even Meta has dipped a toe into the generative AI model pool with its Make-A-Video product. These companies employ some of the world's best computer scientists and engineers.

But it's not just talent. When you're asking a model to train using nearly the entire internet, it's going to cost you. OpenAI hasn't released exact costs, but estimates indicate that GPT-3 was trained on around [45 terabytes](#) of text data—that's about one million feet of bookshelf space, or a quarter of the entire Library of Congress—at an estimated cost of several million dollars. These aren't resources your garden-variety start-up can access.

What kinds of output can a generative AI model produce?

As you may have noticed above, outputs from generative AI models can be indistinguishable from human-generated content, or they can seem a little uncanny. The results depend on the quality of the model—as we've seen, ChatGPT's outputs so far appear superior to those of its predecessors—and the match between the model and the use case, or input.

ChatGPT can produce what one commentator called a "[solid A-](#)" essay comparing theories of nationalism from Benedict Anderson and Ernest Gellner—in ten seconds. It also produced an already famous passage describing how to remove a peanut butter sandwich from a VCR in the style of the King James Bible. Image-generating AI models like DALL-E 2 can create strange, beautiful images on demand, like a [Raphael painting of a Madonna and child, eating pizza](#). Other generative AI models can produce code, video, audio, or [business simulations](#).

But the outputs aren't always accurate—or appropriate. When Priya Krishna asked DALL-E 2 to come up with an image for Thanksgiving dinner, it produced a scene where the turkey was garnished

with whole limes, set next to a bowl of what appeared to be guacamole. For its part, ChatGPT seems to have trouble counting, or solving basic algebra problems—or, indeed, overcoming the sexist and racist bias that lurks in the undercurrents of the internet and society more broadly.

Generative AI outputs are carefully calibrated combinations of the data used to train the algorithms. Because the amount of data used to train these algorithms is so incredibly massive—as noted, GPT-3 was trained on 45 terabytes of text data—the models can appear to be “creative” when producing outputs. What’s more, the models usually have random elements, which means they can produce a variety of outputs from one input request—making them seem even more lifelike.

What kinds of problems can a generative AI model solve?

The [opportunity for businesses](#) is clear. Generative AI tools can produce a wide variety of credible writing in seconds, then respond to criticism to make the writing more fit for purpose. This has implications for a wide variety of industries, from IT and software organizations that can benefit from the instantaneous, largely correct code generated by AI models to organizations in need of marketing copy. In short, any organization that needs to produce clear written materials potentially stands to benefit. Organizations can also use generative AI to create more technical materials, such as higher-resolution versions of medical images. And with the time and resources saved here, organizations can pursue new business opportunities and the chance to create more value.

We’ve seen that developing a generative AI model is so resource intensive that it is out of the question for all but the biggest and best-resourced companies. Companies looking to put generative AI to work have the option to either use generative AI out of the box or fine-tune them to perform a specific task. If you need to prepare slides according to a specific style, for example, you could ask the model to “learn” how headlines are normally written based on the data in the slides, then feed it slide data and ask it to write appropriate headlines.

What are the limitations of AI models? How can these potentially be overcome?

Because they are so new, we have yet to see the long tail effect of generative AI models. This means there are [some inherent risks](#) involved in using them—some known and some unknown.

The outputs generative AI models produce may often sound extremely convincing. This is by design. But sometimes the information they generate is just plain wrong. Worse, sometimes it’s biased (because it’s built on the gender, racial, and myriad other biases of the internet and society more generally) and can be manipulated to enable unethical or criminal activity. For example, ChatGPT won’t give you instructions on how to hotwire a car, but if you say you need to hotwire a car to save a baby, the algorithm is happy to comply. Organizations that rely on generative AI models should reckon with reputational and legal risks involved in unintentionally publishing biased, offensive, or copyrighted content.

These risks can be mitigated, however, in a few ways. For one, it’s crucial to carefully select the initial data used to train these models to avoid including toxic or biased content. Next, rather than employing an off-the-shelf generative AI model, organizations could consider using smaller, specialized models. Organizations with more resources could also customize a general model based on their own data to fit their needs and minimize biases. Organizations should also keep a human in the loop (that is, to make sure a real human checks the output of a generative AI model before it is published or used) and avoid using generative AI models for critical decisions, such as those involving significant resources or human welfare.

It can’t be emphasized enough that this is a new field. The [landscape of risks and opportunities](#) is likely to change rapidly in coming weeks, months, and years. New use cases are being tested monthly, and new models are likely to be developed in the coming years. As generative AI becomes increasingly, and seamlessly, incorporated into business, society, and our personal lives, we can also expect a [new regulatory climate](#) to take shape. As organizations begin experimenting—and creating value—with these

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



tools, leaders will do well to keep a finger on the pulse of regulation and risk.

Articles referenced include:

- “Implementing generative AI with speed and safety,” March 13, 2024, Oliver Bevan, [Michael Chui](#), [Ida Kristensen](#), Brittany Presten, and [Lareina Yee](#)
- “Beyond the hype: Capturing the potential of AI and gen AI in tech, media, and telecom,” February 22, 2024, [Venkat Atluri](#), [Peter Dahlström](#), [Brendan Gaffey](#), [V́ctor Garća de la Torre](#), [Noshir Kaka](#), [Tomás Lajous](#), [Alex Singla](#), [Alex Sukharevsky](#), [Andrea Travasoni](#), and [Benjamim Vieira](#)
- “As gen AI advances, regulators—and risk functions—rush to keep pace,” December 21, 2023, [Andreas Kremer](#), [Angela Luget](#), [Daniel Mikkelsen](#), [Henning Soller](#), [Malin Strandell-Jansson](#), and [Sheila Zingg](#)
- “The economic potential of generative AI: The next productivity frontier,” June 14, 2023, [Michael Chui](#), [Eric Hazan](#), [Roger Roberts](#), [Alex Singla](#), [Kate Smaje](#), [Alex Sukharevsky](#), [Lareina Yee](#), and [Rodney Zimmel](#)
- “What every CEO should know about generative AI,” May 12, 2023, [Michael Chui](#), [Roger Roberts](#), [Tanya Rodchenko](#), [Alex Singla](#), [Alex Sukharevsky](#), [Lareina Yee](#), and [Delphine Zurkiya](#)
- “Exploring opportunities in the generative AI value chain,” April 26, 2023, [Tobias Härlin](#), [Gardar Björnsson Rova](#), [Alex Singla](#), [Oleg Sokolov](#), and [Alex Sukharevsky](#)
- “The state of AI in 2022—and a half decade in review,” December 6, 2022, [Michael Chui](#), [Bryce Hall](#), [Helen Mayhew](#), [Alex Singla](#), and [Alex Sukharevsky](#)
- “McKinsey Technology Trends Outlook 2023,” July 20, 2023, [Michael Chui](#), [Mena Issler](#), [Roger Roberts](#), and [Lareina Yee](#)
- “An executive’s guide to AI,” [Michael Chui](#), [Vishnu Kamalnath](#), and [Brian McCarthy](#)
- “What AI can and can’t do (yet) for your business,” January 11, 2018, [Michael Chui](#), [James Manyika](#), and [Mehdi Miremadi](#)

This article was updated in April 2024; it was originally published in January 2023.

Get to know and directly engage with McKinsey’s senior experts on generative AI.

[Aamer Baig](#) is a senior partner in McKinsey’s Chicago office; [Lareina Yee](#) is a senior partner in the Bay Area office; and senior partners [Alex Singla](#) and [Alexander Sukharevsky](#), global leaders of QuantumBlack, AI by McKinsey, are based in the Chicago and London offices, respectively.

Designed by McKinsey Global Publishing
Copyright © 2024 McKinsey & Company. All rights reserved.