

Executive Summary

Key takeaways from our artificial intelligence coverage



Artificial intelligence is the use of computers to make automated decisions or predictions. Most AI systems perform a single, defined task. Even relatively versatile systems such as ChatGPT are using one relatively narrow task—in ChatGPT’s case, analyzing and generating language and images—to yield impressive results across varied contexts.

How AI works

Modern AI applications fall into two categories:

- **Rule-based systems** follow specific instructions. This requires developers to manually define every potential variable and outcome.
- **Machine learning systems** do not depend entirely on firm rules. Provided with many examples of a task or event, its algorithms learn patterns and apply them to make decisions or predictions.

Though rule-based systems are still common, the most capable AI applications are increasingly reliant on machine learning systems that are structured as *neural networks*, a technique that learns tasks across a series of increasingly complex layers.

Broadly speaking, machine learning systems are trained in one of three ways:

- **Supervised learning**, where all of the training data has been organized, usually with labels, beforehand.
- **Unsupervised learning**, meaning the training data is unlabeled, forcing the system to group and make connections among bits of information.
- **Reinforcement learning**, where the system trains on partially labeled data and learns a given task over time through trial and error.

Benefits

AI can make sense of complex information

By discerning patterns in data too subtle for humans to notice, AI can be useful for analyzing large amounts of information with many complex and interconnected variables.

AI systems can analyze and generate high-quality imagery and speech

Adobe, Google, Microsoft and other tech suppliers are incorporating *generative AI* systems into their products to create detailed text passages and illustrations from simple prompts.

AI can predict some trends or events more accurately than other software can

Instead of relying wholly on historical data to project future events, AI systems can use a range of techniques to make predictions that take into account changing conditions.

AI systems can adapt and improve over time

Many AI applications currently use machine learning algorithms that improve through training with little or no human intervention.

Limitations

AI systems lack common sense, and they don’t perform equally well on different tasks

Although an algorithm might create the illusion of a mind at work, it’s not the fluid intellect of human consciousness. This makes AI systems specialists. Systems adept in one domain, such as detecting fraud or recognizing faces, can’t excel at others.

continued on next page

Most AI systems require immense amounts of data for training

Massive data sets require expensive hardware and lots of energy to gather and train on, and these data sets can lead to or result from privacy violations and copyright infringement.

Problems in training data can lead to poor results, and widespread public distrust of AI

AI systems learn to recreate the patterns they're exposed to in training, including any errors or biases in training data—an issue many experts call the “garbage in, garbage out” problem.

Many sophisticated AI systems are opaque, so that even their developers cannot fully explain them

Referred to as the “black box problem,” this can make it difficult to determine accountability or suggest improvements when systems make mistakes.


Common misconceptions

Soon, AI will become as intelligent as humans

The reality: Researchers are still far from creating a machine that matches or exceeds human intelligence at a wide range of tasks, and experts disagree on how long this will take.

AI is a specific, single technology

The reality: AI is a broad term for many computing systems, and can refer to different techniques in different contexts.



AI is mostly used for high-profile applications where it appears to be human

The reality: AI is often used for mundane, behind-the-scenes business tasks.

AI is objective and free from bias

The reality: Training data can fail to match real-world circumstances, or reflect preexisting imbalances, and so AI systems can impact groups of people differently.

Developments to watch

Organizations will grapple with how to integrate generative AI

Following the launch of ChatGPT and the ensuing hype around generative AI, many organizations are realizing that integrating the technologies into their workflows and products can be difficult.

AI will disrupt jobs, causing instability and unknown long-term outcomes

Advancements in AI have long caused unease about whether the technology will displace workers, making certain jobs obsolete and increasing unemployment. The rise of generative AI is amplifying those concerns.

AI will make robots more adaptable

Recent advancements in AI have enabled roboticists to experiment with new ways to program and train general-purpose machines to perform a wide range of tasks or adapt to new circumstances and actions.

Rising demand for hardware could limit AI development

As demand for computer chips purpose-built for AI outstrips supply researchers, start ups, and small-to medium-size organizations developing their own AI applications could risk falling behind large tech firms.