

Glossary

Common terms in artificial intelligence



Algorithm

A set of rules or instructions that a human or a computer follows to perform a task.

Artificial general intelligence (AGI)

A theoretical AI system that could process, interpret, reason, and make judgments in a wide variety of situations. Also called strong AI. In contrast, weak or narrow AI systems can only complete a limited task or set of tasks. All current AI systems are narrow; no AGI yet exists.

Artificial intelligence (AI)

A computer system trained to extrapolate from data in order to make automated decisions or predictions.

Automation

Reproducing manual processes with robotics and sometimes artificial intelligence in a way that does not require human intervention.

Backpropagation

A type of algorithm used in the training process for artificial neural networks. When combined with a function called gradient descent, backpropagation allows developers to adjust the weights of individual neurons and improve the network's performance.

Black box algorithm

An algorithm that produces results for which the process and parameters the algorithm used to create them cannot be easily deciphered by humans.

Computer vision

A field of AI focused on extracting information from visual data, including photos and videos.

Convolutional neural networks

A class of neural networks that make use of multiple convolution layers, comprising sets of filters that work on portions of an input and then pool their results to produce an overall interpretation. The technique has proved effective in image processing applications, such as feature extraction.

Deep learning

A subfield of *machine learning* that involves using a multilayered neural network to analyze data or automate complex tasks such as feature extraction or natural language processing.

Feature extraction

The process of reducing raw data into more manageable groups by identifying features, or common attributes.

Foundation model

A general-purpose AI program that can be adapted for many uses. Coined by researchers at Stanford University, the term describes a category of large machine learning models—including the models that power ChatGPT and other generative AI systems—that carry out a broad range of tasks.

Generative adversarial networks (GANs)

Machine learning systems that generate data on their own. One part of the system (the generative network) produces data, such as a set of colored pixels on a grid, and another (the discriminative network) evaluates whether the product meets certain criteria. The process is repeated until the first system is consistently producing data that satisfies the criteria of the second. GANs are commonly used to alter or generate realistic images.

continued on next page

Generative AI

A subfield of artificial intelligence focused on generating content, such as text, images, and audio, using *machine learning* models. Generative models are trained on databases of human-created media, and some are among the largest AI systems to date.

GPU cluster

A grouping of *graphics processing units* (GPUs) that work in tandem to train AI algorithms. Ongoing increases in the scale and complexity of AI systems are partially attributable to the use of warehouse-size GPU clusters.

Hallucination

A false or nonsensical error in the content produced by a *generative AI* model. Generative models sometimes produce text that can be misconstrued as fact, such as a false citation of an article or paper that doesn't exist. The cause of a hallucination is often hard to determine, as most generative models based their content on vast amounts of training data, as opposed to a single source.

Large language model (LLM)

A type of *generative AI* model that produces text, usually in response to written prompts from users. LLMs are trained on massive quantities of text, such as books, webpages, and transcripts of audio recordings. AI applications that rely on LLMs, such as ChatGPT, are able to generate complex and nuanced text, such as a cover letter for a given role or explanation for why a line of computer code isn't working as intended.

Machine learning

A subfield of artificial intelligence involving systems that can be trained to interpret or extrapolate from data without depending on explicit, preprogrammed rules.

Natural language processing

The use of algorithms that can interpret, analyze, or generate written or spoken human language.

Neural network

A set of algorithms that can analyze and learn from data using a process inspired by the interaction of neurons in the brain. Neural networks are often used in *machine learning* systems, and particularly as components of algorithms trained with deep learning techniques.

Neural processing unit (NPU)

A computer hardware component that is specifically engineered to run a *neural network* faster and more efficiently than with traditional processors.

Parameters

In the context of *machine learning*, parameters are the parts of a model that change during the training process and that determine how well the model can turn data inputs into outputs, such as predictions. Experts sometimes present the relative size and complexity of different models by their total parameters, which can exceed 1 trillion.

Pretraining

In a step that precedes the standard *machine learning* training process, engineers feed an AI system large amounts of data, enabling it to learn general patterns and relationships among a general class of examples, such as imagery or human language. Pretraining can give an AI model a head start during subsequent training to learn more specific tasks. See *training*.

Reinforcement learning

A method of training algorithms that, like unsupervised learning, uses unlabeled training sets but also provides predefined incentives, or rewards, when the system (programmed to maximize those rewards) behaves in the desired way.

continued on next page

Rule-based systems

A class of AI systems that process inputs according to preprogrammed conditions: a preset knowledge base and a series of if-then-else statements. Every action such a system takes can be explained.

Speech recognition

The use of algorithms to identify and interpret spoken language and transform it into computer functions or readable text.

Supervised learning

Training a *machine learning* model using a training set that has been given labels by a human.

Text-to-image generator

A type of *generative AI* model that, after training on large numbers of existing images, can transform text prompts from users—such as “A tree in a field with a house in the distance”—into new images. Whether abstract or photorealistic, the resulting images can be difficult to distinguish from images captured or created by humans. .

Training

In the context of AI, *training* refers to the process of exposing a *machine learning* model to large numbers of examples that are relevant to a task, such as images of handwritten numbers or current and past home values in a given area. Training often requires massive data sets of thousands or even trillions of examples and groups of computer processors running for days at a time.

Training set

A data set that is used to train a *machine learning* model, or adjust its algorithms in order to improve performance.

Turing test

The best-known test of computer intelligence. A theoretical or practical exercise in which humans interact with a computer, for instance a chatbot, and the computer is judged on how closely its interactions resemble what the human would expect from another person.

Unsupervised learning

Training a *machine learning* model using a training set that has not been given labels.

Validation set

A data set that is isolated from the training set and is used to evaluate the performance of a *machine learning* model.