

AI GLOSSARY

Artificial intelligence, in so many words

Just what do people mean by artificial intelligence (AI)? The term has never had clear boundaries. When it was introduced at a seminal 1956 workshop at Dartmouth College, it was taken broadly to mean making a machine behave in ways that would be called intelligent if seen in a human. An important recent advance in AI has been machine learning, which shows up in technologies from spellcheck to self-driving cars and is often carried out by computer systems called neural networks. Any discussion of AI is likely to include other terms as well.

ALGORITHM A set of step-by-step instructions. Computer algorithms can be simple (if it's 3 p.m., send a reminder) or complex (identify pedestrians).

BACKPROPAGATION The way many neural nets learn. They find the difference between their output and the desired output, then adjust the calculations in reverse order of execution.

BLACK BOX A description of some deep learning systems. They take an input and provide an output, but the calculations that occur in between are not easy for humans to interpret.

DEEP LEARNING How a neural network with multiple layers becomes sensitive to progressively more abstract patterns. In parsing a photo, layers might respond first to edges, then paws, then dogs.

EXPERT SYSTEM A form of AI that attempts to replicate a human's expertise in an area, such as medical diagnosis. It combines a knowledge base with a set of hand-coded rules for applying that knowledge. Machine-learning techniques are increasingly replacing hand coding.

GENERATIVE ADVERSARIAL NETWORKS A pair of jointly trained neural networks that generates realistic new data and improves through competition. One net creates new examples (fake Picassos, say) as the other tries to detect the fakes.

MACHINE LEARNING The use of algorithms that find patterns in data without explicit instruction. A system might learn how to associate features of inputs such as images with outputs such as labels.

NATURAL LANGUAGE PROCESSING A computer's attempt to "understand" spoken or written language. It must parse vocabulary, grammar, and intent, and allow for variation in language use. The process often involves machine learning.

NEURAL NETWORK A highly abstracted and simplified model of the human brain used in machine learning. A set of units receives pieces of an input (pixels in a photo, say), performs simple computa-

tions on them, and passes them on to the next layer of units. The final layer represents the answer.

NEUROMORPHIC CHIP A computer chip designed to act as a neural network. It can be analog, digital, or a combination.

PERCEPTRON An early type of neural network, developed in the 1950s. It received great hype but was then shown to have limitations, suppressing interest in neural nets for years.

REINFORCEMENT LEARNING A type of machine learning in which the algorithm learns by acting toward an abstract goal, such as "earn a high video game score" or "manage a factory efficiently." During training, each effort is evaluated based on its contribution toward the goal.

STRONG AI AI that is as smart and well-rounded as a human. Some say it's impossible. Current AI is weak, or narrow. It can play chess or drive but not both, and lacks common sense.

SUPERVISED LEARNING A type of machine learning in which the algorithm compares its outputs with the correct outputs during training. In unsupervised learning, the algorithm merely looks for patterns in a set of data.

TENSORFLOW A collection of software tools developed by Google for use in deep learning. It is open source, meaning anyone can use or improve it. Similar projects include Torch and Theano.

TRANSFER LEARNING A technique in machine learning in which an algorithm learns to perform one task, such as recognizing cars, and builds on that knowledge when learning a different but related task, such as recognizing cats.

TURING TEST A test of AI's ability to pass as human. In Alan Turing's original conception, an AI would be judged by its ability to converse through written text.

Matthew Hutson is a science journalist in New York City.

At biotech startup Zymergen, robotic fingers are poised to pick microbe colonies in an AI-controlled quest for strains that crank out more chemicals.

never touches it. Instead, 500 times per second, a pulse of sound waves causes the liquid itself to ripple and launch a droplet a thousand times smaller than one a human can transfer.

Yet none of that is the really futuristic part. Big bio labs have used robots and barcodes for years. Even the liquid-moving technology—called acoustic droplet ejection—has existed for decades. The real giveaway comes when I ask Dean what experiment this robot is working on right now. "I have no idea," he says. He could easily find out, but he didn't design the experiment. Instead, it was the output of a computer program.

"I want to be very clear," says Zymergen CEO Joshua Hoffman, heading off a persistent misunderstanding. "There is a human scientist in the loop, looking at the results and reality checking them." But for interpreting data, generating hypotheses, and planning experiments, he says, the ultimate goal is "to get rid of human intuition."

Zymergen is one of several companies

AI Glossary: Artificial intelligence, in so many words

Matthew Hutson

Science **357** (6346), 19.

DOI: 10.1126/science.357.6346.19

ARTICLE TOOLS

<http://science.sciencemag.org/content/357/6346/19>

PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science (print ISSN 0036-8075; online ISSN 1095-9203) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. 2017 © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. The title *Science* is a registered trademark of AAAS.