



What Is Artificial Intelligence?

Description

How Does Artificial Intelligence Work?

What Is AI?

Less than a decade after helping the Allied forces win World War II by breaking the Nazi encryption machine Enigma, mathematician Alan Turing changed history a second time with a simple question: “Can machines think?”

Turing’s 1950 paper “[Computing Machinery and Intelligence](#)” and its subsequent [Turing Test](#) established the fundamental goal and vision of AI.

At its core, AI is the branch of computer science that aims to answer Turing’s question in the affirmative. It is the endeavor to replicate or simulate human intelligence in machines. The expansive goal of AI has given rise to many questions and debates. So much so that no singular definition of the field is universally accepted.

Can machines think? – Alan Turing, 1950

Defining AI

The major limitation in defining AI as simply “building machines that are intelligent” is that it doesn’t actually explain what AI is and what makes a machine intelligent. AI is an interdisciplinary science with multiple approaches, but advancements in machine learning and deep learning are creating a paradigm shift in virtually every sector of the tech industry.

However, various new tests have been proposed recently that have been largely well received, including a 2019 research paper entitled “On the Measure of Intelligence.” In the paper, veteran deep learning researcher and Google engineer François Chollet argues that intelligence is the “rate at which a learner turns its experience and priors into new skills at valuable tasks that involve uncertainty and

adaptation.” In other words: The most intelligent systems are able to take just a small amount of experience and go on to guess what would be the outcome in many varied situations.

Meanwhile, in their book *Artificial Intelligence: A Modern Approach*, authors Stuart Russell and Peter Norvig approach the concept of AI by unifying their work around the theme of intelligent agents in machines. With this in mind, AI is “the study of agents that receive percepts from the environment and perform actions.”

Check Out the Top AI Companies [View all AI Companies](#)

Norvig and Russell go on to explore four different approaches that have historically defined the field of AI:

Artificial Intelligence Defined: Four Types of Approaches

- Thinking humanly: mimicking thought based on the human mind.
- Thinking rationally: mimicking thought based on logical reasoning.
- Acting humanly: acting in a manner that mimics human behavior.
- Acting rationally: acting in a manner that is meant to achieve a particular goal.

The first two ideas concern thought processes and reasoning, while the others deal with behavior. Norvig and Russell focus particularly on rational agents that act to achieve the best outcome, noting “all the skills needed for the Turing Test also allow an agent to act rationally.”

Former MIT professor of AI and computer science Patrick Winston defined AI as “algorithms enabled by constraints, exposed by representations that support models targeted at loops that tie thinking, perception and action together.”

While these definitions may seem abstract to the average person, they help focus the field as an area of computer science and provide a blueprint for infusing machines and programs with ML and other subsets of [AI](#).

The Future of AI

When one considers the computational costs and the technical data infrastructure running behind artificial intelligence, actually executing on AI is a complex and costly business. Fortunately, there have been massive advancements in computing technology, as indicated by Moore’s Law, which states that the number of transistors on a microchip doubles about every two years while the cost of computers is halved.

Although many experts believe that Moore’s Law will likely come to an end sometime in the 2020s, this has had a major impact on modern AI techniques — without it, deep learning would be out of the question, financially speaking. Recent research found that AI innovation has actually outperformed Moore’s Law, doubling every six months or so as opposed to two years.

By that logic, the advancements artificial intelligence has made across a variety of industries have been major over the last several years. And the potential for an even greater impact over the next several decades seems all but inevitable.

Types of Artificial Intelligence | Artificial Intelligence Explained | What Is AI? | Edureka | Video: edureka!

The Four Types of Artificial Intelligence

AI can be divided into four categories, based on the type and complexity of the tasks a system is able to perform. For example, automated spam filtering falls into the most basic class of AI, while the far-off potential for machines that can perceive people's thoughts and emotions is part of an entirely different AI subset.

What Are the Four Types of Artificial Intelligence?

- Reactive machines: able to perceive and react to the world in front of it as it performs limited tasks.
- Limited memory: able to store past data and predictions to inform predictions of what may come next.
- Theory of mind: able to make decisions based on its perceptions of how others feel and make decisions.
- Self-awareness: able to operate with human-level consciousness and understand its own existence.

Reactive Machines

A reactive machine follows the most basic of AI principles and, as its name implies, is capable of only using its intelligence to perceive and react to the world in front of it. A reactive machine cannot store a memory and, as a result, cannot rely on past experiences to inform decision making in real time.

Perceiving the world directly means that reactive machines are designed to complete only a limited number of specialized duties. Intentionally narrowing a reactive machine's worldview is not any sort of cost-cutting measure, however, and instead means that this type of AI will be more trustworthy and reliable — it will react the same way to the same stimuli every time.

A famous example of a reactive machine is Deep Blue, which was designed by IBM in the 1990s as a chess-playing supercomputer and defeated international grandmaster Gary Kasparov in a game. Deep Blue was only capable of identifying the pieces on a chess board and knowing how each moves based on the rules of chess, acknowledging each piece's present position and determining what the most logical move would be at that moment. The computer was not pursuing future potential moves by its opponent or trying to put its own pieces in better position. Every turn was viewed as its own reality, separate from any other movement that was made beforehand.

Another example of a game-playing reactive machine is Google's AlphaGo. AlphaGo is also incapable

of evaluating future moves but relies on its own neural network to evaluate developments of the present game, giving it an edge over Deep Blue in a more complex game. AlphaGo also bested world-class competitors of the game, defeating champion Go player Lee Sedol in 2016.

Though limited in scope and not easily altered, reactive machine AI can attain a level of complexity, and offers reliability when created to fulfill repeatable tasks.

Limited Memory

Limited memory AI has the ability to store previous data and predictions when gathering information and weighing potential decisions — essentially looking into the past for clues on what may come next. Limited memory AI is more complex and presents greater possibilities than reactive machines.

Limited memory AI is created when a team continuously trains a model in how to analyze and utilize new data or an AI environment is built so models can be automatically trained and renewed.

When utilizing limited memory AI in ML, six steps must be followed: Training data must be created, the ML model must be created, the model must be able to make predictions, the model must be able to receive human or environmental feedback, that feedback must be stored as data, and these steps must be reiterated as a cycle.

There are several ML models that utilize limited memory AI:

- **Reinforcement learning**, which learns to make better predictions through repeated trial and error.
- **Recurrent neural networks (RNN)**, which uses sequential data to take information from prior inputs to influence the current input and output. These are commonly used for ordinal or temporal problems, such as language translation, natural language processing, speech recognition and image captioning. One subset of recurrent neural networks is known as long short term memory (LSTM), which utilizes past data to help predict the next item in a sequence. LSTMs view more recent information as most important when making predictions, and discount data from further in the past while still utilizing it to form conclusions.
- **Evolutionary generative adversarial networks (E-GAN)**, which evolve over time, growing to explore slightly modified paths based off of previous experiences with every new decision. This model is constantly in pursuit of a better path and utilizes simulations and statistics, or chance, to predict outcomes throughout its evolutionary mutation cycle.
- **Transformers**, which are networks of nodes that learn how to do a certain task by training on existing data. Instead of having to group elements together, transformers are able to run processes so that every element in the input data pays attention to every other element. Researchers refer to this as “self-attention,” meaning that as soon as it starts training, a transformer can see traces of the entire data set.

Theory of Mind

Theory of mind is just that — theoretical. We have not yet achieved the technological and scientific capabilities necessary to reach this next level of AI.

The concept is based on the psychological premise of understanding that other living things have

thoughts and emotions that affect the behavior of one's self. In terms of AI machines, this would mean that AI could comprehend how humans, animals and other machines feel and make decisions through self-reflection and determination, and then will utilize that information to make decisions of their own. Essentially, machines would have to be able to grasp and process the concept of "mind," the fluctuations of emotions in decision making and a litany of other psychological concepts in real time, creating a two-way relationship between people and AI.

What if AI became self-aware? | Video: ALLTIME10S

Self-Awareness

Once theory of mind can be established, sometime well into [the future of AI](#), the final step will be for AI to become self-aware. This kind of AI possesses human-level consciousness and understands its own existence in the world, as well as the presence and emotional state of others. It would be able to understand what others may need based on not just what they communicate to them but how they communicate it.

Self-awareness in AI relies both on human researchers understanding the premise of consciousness and then learning how to replicate that so it can be built into machines.

How Is AI Used? Artificial Intelligence Examples

While addressing a crowd at the [Japan AI Experience in 2017](#), [DataRobot](#) CEO Jeremy Achin began his speech by offering the following definition of how AI is used today:

"AI is a computer system able to perform tasks that ordinarily require human intelligence ... Many of these artificial intelligence systems are powered by machine learning, some of them are powered by deep learning and some of them are powered by very boring things like rules."

Other AI Classifications

There are three ways to classify artificial intelligence, based on their capabilities. Rather than types of artificial intelligence, these are stages through which AI can evolve — and only one of them is actually possible right now.

- **Narrow AI:** Sometimes referred to as "weak AI," this kind of AI operates within a limited context and is a simulation of human intelligence. Narrow AI is often focused on performing a single task extremely well and while these machines may seem intelligent, they are operating under far more constraints and limitations than even the most basic human intelligence.
- **Artificial general intelligence (AGI):** AGI, sometimes referred to as "strong AI," is the kind of AI we see in movies — like the robots from *Westworld* or the character Data from *Star Trek: The Next Generation*. AGI is a machine with general intelligence and, much like a human being, it can apply that intelligence to solve any problem.
- **Superintelligence:** This will likely be the pinnacle of AI's evolution. Superintelligent AI will not only be able to replicate the complex emotion and intelligence of human beings, but surpass it in every way. This could mean making judgments and decisions on its own, or even forming its own

ideology.

Narrow AI Examples

Narrow AI, or weak AI as it's often called, is all around us and is easily the most successful realization of AI to date. It has limited functions that are able to help automate specific tasks.

Because of this focus, narrow AI has experienced numerous breakthroughs in the last decade that have had “significant societal benefits and have contributed to the economic vitality of the nation,” according to a [2016 report](#) released by the Obama administration.

Examples of Artificial Intelligence: Narrow AI

- Siri, Alexa and other smart assistants
- Self-driving cars
- Google search
- Conversational bots
- Email spam filters
- Netflix's recommendations

Machine Learning and Deep Learning

Much of narrow AI is powered by breakthroughs in ML and deep learning. Understanding the difference between AI, ML and deep learning can be confusing. Venture capitalist Frank Chen provides a good overview of how to distinguish between them, noting:

“Artificial intelligence is a set of algorithms and intelligence to try to mimic human intelligence. Machine learning is one of them, and deep learning is one of those machine learning techniques.”

Simply put, an ML algorithm is fed data by a computer, and uses statistical techniques to help it “learn” how to get progressively better at a task, without necessarily having been specifically programmed for that task. Instead, ML algorithms use historical data as input to predict new output values. To that end, ML consists of both supervised learning (where the expected output for the input is known thanks to labeled data sets) and unsupervised learning (where the expected outputs are unknown due to the use of unlabeled data sets).

Machine learning is present throughout everyday life. Google Maps uses location data from smartphones, as well as user-reported data on things like construction and car accidents, to monitor the ebb and flow of traffic and assess what the fastest route will be. Personal assistants like Siri, Alexa and Cortana are able to set reminders, search for online information and control the lights in people's homes all with the help of ML algorithms that collect information, learn a user's preferences and improve their experience based on prior interactions with users. Even Snapchat filters use ML algorithms in order to track users' facial activity.

Meanwhile, deep learning is a type of ML that runs inputs through a biologically-inspired neural network architecture. The neural networks contain a number of hidden layers through which the data is processed, allowing the machine to go "deep" in its learning, making connections and weighting input for the best results.

Self-driving cars are a recognizable example of deep learning, since they use deep neural networks to detect objects around them, determine their distance from other cars, identify traffic signals and much more. The wearable sensors and devices used in the healthcare industry also apply deep learning to assess the health condition of the patient, including their blood sugar levels, blood pressure and heart rate. They can also derive patterns from a patient's prior medical data and use that to anticipate any future health conditions.

Artificial General Intelligence

The creation of a machine with human-level intelligence that can be applied to any task is the Holy Grail for many AI researchers, but the quest for artificial general intelligence has been fraught with difficulty.

The search for a "universal algorithm for learning and acting in any environment," as Russel and Norvig put it, isn't new. In contrast to weak AI, strong AI represents a machine with a full set of cognitive abilities, but time hasn't eased the difficulty of achieving such a feat.

AGI has long been the muse of dystopian science fiction, in which super-intelligent robots overrun humanity, but experts agree it's not something we need to worry about anytime soon.

Although, for now, AGI is still a fantasy, there are some remarkably sophisticated systems out there now that are approaching the AGI benchmark. One of them is GPT-3, an autoregressive language model designed by OpenAI that uses deep learning to produce human-like text. GPT-3 is not intelligent, but it has been used to create some extraordinary things, including a chatbot that lets you talk to historical figures and a question-based search engine. MuZero, a computer program created by DeepMind, is another promising frontrunner in the quest to achieve true AGI. It has managed to master games it has not even been taught to play, including chess and an entire suite of Atari games, through brute force, playing games millions of times.

Superintelligence

Besides narrow AI and AGI, some consider there to be a third category known as superintelligence. For now, this is a completely hypothetical situation in which machines are completely self-aware, even

surpassing the likes of human intelligence in practically every field, from science to social skills. In theory, this could be achieved through a single computer, a network of computers or something completely different, as long as it is conscious and has subjective experiences.

Nick Bostrom, a founding professor and leader of Oxford's Future of Humanity Institute, appears to have coined the term back in 1998, and predicted that we will have achieved superhuman artificial intelligence within the first third of the 21st century. He went on to say that the likelihood of this happening will likely depend on how quickly neuroscience can better understand and replicate the human brain. Creating superintelligence by imitating the human brain, he added, will require not only sufficiently powerful hardware, but also an "adequate initial architecture" and a "rich flux of sensory input."

EXAMPLES OF ARTIFICIAL INTELLIGENCE

artificial intelligence examples

Image not found or type unknown

AI HAS MANY USES. EXAMPLES INCLUDE EVERYTHING FROM AMAZON ALEXA TO SELF-DRIVING CARS. | IMAGE; SHUTTERSTOCK

Why Is Artificial Intelligence Important?

AI has many uses — from [boosting vaccine development](#) to [automating detection of potential fraud](#).

AI private market activity saw a record-setting year in 2021, [according to CB Insights](#), with global funding up 108 percent compared to 2020. Because of its fast-paced adoption, AI is making waves in a variety of industries.

Business Insider Intelligence's [2022 report on AI in banking](#) found more than half of financial services companies already use AI solutions for risk management and revenue generation. The application of AI in banking could lead to upwards of \$400 billion in savings.

As for medicine, a [2021 World Health Organization](#) report noted that while integrating AI into the healthcare field comes with challenges, the technology “holds great promise,” as it could lead to benefits like more informed health policy and improvements in the accuracy of diagnosing patients.

AI has also made its mark on entertainment. The global market for AI in media and entertainment is estimated to reach \$99.48 billion by 2030, growing from a value of \$10.87 billion in 2021, according to [Grand View Research](#). That expansion includes [AI uses](#) like recognizing plagiarism and developing high-definition graphics.

Artificial Intelligence Pros and Cons

While AI is certainly viewed as an important and quickly evolving asset, this emerging field comes with its share of [downsides](#).

The Pew Research Center [surveyed 10,260 Americans](#) in 2021 on their attitudes toward AI. The results found 45 percent of respondents are equally excited and concerned, and 37 percent are more concerned than excited. Additionally, more than 40 percent of respondents said they considered driverless cars to be bad for society. Yet the idea of using AI to identify the spread of false information on social media was more well received, with close to 40 percent of those surveyed labeling it a good idea.

AI is a boon for improving productivity and efficiency while at the same time reducing the potential for human error. But there are also some disadvantages, like development costs and the possibility for automated machines to [replace human jobs](#). It's worth noting, however, that the artificial intelligence industry stands to create jobs, too — some of which have not even been invented yet.

AI TIMELINE: HISTORY OF ARTIFICIAL INTELLIGENCE

history of artificial intelligence

Image not found or type unknown

THE HISTORY OF ARTIFICIAL INTELLIGENCE IS LONG AND ROBUST, GOING BACK TO THE 1940S. | IMAGE: SHUTTERSTOCK

A Brief History of Artificial Intelligence

[Intelligent robots](#) and artificial beings [first appeared](#) in ancient Greek myths. And Aristotle's development of syllogism and its use of deductive reasoning was a key moment in humanity's quest to understand its own intelligence. While the roots are long and deep, the history of AI as we think of it today spans less than a century. The following is a quick look at some of the most important events in AI.

1940s

- (1943) Warren McCulloch and Walter Pitts publish the paper "[A Logical Calculus of Ideas Immanent in Nervous Activity](#)," which proposes the first mathematical model for building a neural network.
- (1949) In his book [The Organization of Behavior: A Neuropsychological Theory](#), Donald Hebb proposes the theory that neural pathways are created from experiences and that connections between neurons become stronger the more frequently they're used. Hebbian learning continues to be an important model in AI.

1950s

- (1942) Isaac Asimov publishes the [Three Laws of Robotics](#), an idea commonly found in science fiction media about how artificial intelligence should not bring harm to humans.
- (1950) Alan Turing publishes the paper "Computing Machinery and Intelligence," proposing what is now known as the Turing Test, a method for determining if a machine is intelligent.
- (1950) Harvard undergraduates Marvin Minsky and Dean Edmonds build [SNARC](#), the first neural network computer.
- (1950) Claude Shannon publishes the paper "[Programming a Computer for Playing Chess](#)."
- (1952) Arthur Samuel develops a self-learning program to play checkers.
- (1954) The Georgetown-IBM machine translation experiment automatically translates 60 carefully selected Russian sentences into English.
- (1956) The phrase "artificial intelligence" is coined at the Dartmouth Summer Research Project on Artificial Intelligence. Led by John McCarthy, the conference is widely considered to be the birthplace of AI.
- (1956) Allen Newell and Herbert Simon demonstrate Logic Theorist (LT), the first reasoning program.
- (1958) John McCarthy develops the AI programming language Lisp and publishes "[Programs with Common Sense](#)," a paper proposing the hypothetical Advice Taker, a complete AI system with the ability to learn from experience as effectively as humans.
- (1959) Allen Newell, Herbert Simon and J.C. Shaw develop the General Problem Solver (GPS), a program designed to imitate human problem-solving.
- (1959) Herbert Gelernter develops the Geometry Theorem Prover program.
- (1959) Arthur Samuel coins the term "machine learning" while at IBM.
- (1959) John McCarthy and Marvin Minsky found the MIT Artificial Intelligence Project.

1960s

- (1963) John McCarthy starts the AI Lab at Stanford.
- (1966) The Automatic Language Processing Advisory Committee (ALPAC) report by the U.S. government details the lack of progress in machine translations research, a major Cold War initiative with the promise of automatic and instantaneous translation of Russian. The ALPAC report leads to the cancellation of all government-funded MT projects.
- (1969) The first successful expert systems are developed in DENDRAL, a XX program, and MYCIN, designed to diagnose blood infections, are created at Stanford.

1970s

- (1972) The logic programming language PROLOG is created.
- (1973) The Lighthill Report, detailing the disappointments in AI research, is released by the British government and leads to severe cuts in funding for AI projects.
- (1974-1980) Frustration with the progress of AI development leads to major DARPA cutbacks in academic grants. Combined with the earlier ALPAC report and the previous year's Lighthill Report, AI funding dries up and research stalls. This period is known as the "[First AI Winter](#)."

1980s

- (1980) Digital Equipment Corporations develops R1 (also known as XCON), the first successful commercial expert system. Designed to configure orders for new computer systems, R1 kicks off an investment boom in expert systems that will last for much of the decade, effectively ending the first AI Winter.
- (1982) Japan's Ministry of International Trade and Industry launches the ambitious Fifth Generation Computer Systems project. The goal of FGCS is to develop supercomputer-like performance and a platform for AI development.
- (1983) In response to Japan's FGCS, the U.S. government launches the Strategic Computing Initiative to provide DARPA funded research in advanced computing and AI.
- (1985) Companies are spending more than a billion dollars a year on expert systems and an entire industry known as the Lisp machine market springs up to support them. Companies like Symbolics and Lisp Machines Inc. build specialized computers to run on the AI programming language Lisp.
- (1987-1993) As computing technology improved, cheaper alternatives emerged and the Lisp machine market collapsed in 1987, ushering in the "[Second AI Winter](#)." During this period, expert systems proved too expensive to maintain and update, eventually falling out of favor.

1990s

- (1991) U.S. forces deploy DART, an automated logistics planning and scheduling tool, during the Gulf War.
- (1992) Japan terminates the FGCS project in 1992, citing failure in meeting the ambitious goals outlined a decade earlier.
- (1993) DARPA ends the Strategic Computing Initiative in 1993 after spending nearly \$1 billion and falling far short of expectations.
- (1997) IBM's Deep Blue beats world chess champion Gary Kasparov.

2000s

- (2005) [STANLEY](#), a self-driving car, wins the DARPA Grand Challenge.
- (2005) The U.S. military begins investing in autonomous robots like Boston Dynamics' "Big Dog" and iRobot's "PackBot."
- (2008) Google makes breakthroughs in speech recognition and introduces the feature in its iPhone app.

2010s

- (2011) [IBM's Watson](#) handily defeats the competition on *Jeopardy!*.
- (2011) Apple releases Siri, an AI-powered virtual assistant through its iOS operating system.
- (2012) Andrew Ng, founder of the Google Brain Deep Learning project, feeds a neural network using deep learning algorithms 10 million YouTube videos as a training set. The neural network learned to recognize a cat without being told what a cat is, ushering in the breakthrough era for neural networks and deep learning funding.
- (2014) Google makes the first [self-driving car](#) to pass a state driving test.
- (2014) Amazon's Alexa, a virtual home [smart device](#), is released.
- (2016) Google DeepMind's [AlphaGo](#) defeats world champion Go player Lee Sedol. The complexity of the ancient Chinese game was seen as a major hurdle to clear in AI.
- (2016) The first "robot citizen," a humanoid robot named Sophia, is created by Hanson Robotics and is capable of facial recognition, verbal communication and facial expression.
- (2018) Google releases natural language processing engine [BERT](#), reducing barriers in translation and understanding by ML applications.
- (2018) [Waymo](#) launches its Waymo One service, allowing users throughout the Phoenix metropolitan area to request a pick-up from one of the company's self-driving vehicles.

2020s

- (2020) Baidu releases its LinearFold AI algorithm to scientific and medical teams working to develop a vaccine during the early stages of the SARS-CoV-2 pandemic. The algorithm is able to predict the RNA sequence of the virus in just 27 seconds, 120 times faster than other methods.
- (2020) [OpenAI](#) releases natural language processing model [GPT-3](#), which is able to produce text modeled after the way people speak and write.
- (2021) OpenAI builds on GPT-3 to develop [DALL-E](#), which is able to create images from text prompts.
- (2022) The [National Institute of Standards and Technology](#) releases the first draft of its [AI Risk Management Framework](#), voluntary U.S. guidance "to better manage risks to individuals, organizations, and society associated with artificial intelligence."
- (2022) DeepMind unveils Gato, an AI system trained to perform hundreds of tasks, including playing Atari, captioning images and using a robotic arm to stack blocks.

Written by Alyssa Schroer

Category

1. Main
2. Science-Tech-AI-Medical & Gen. Research

Date Created

03/02/2023