

	<b>Artificial Intelligence Terms</b>
Artificial intelligence	This refers to the general concept of machines acting in a way that simulates or mimics human intelligence. AI can have a variety of features, such as human-like communication or decision making.
Autonomous	A machine is described as autonomous if it can perform its task or tasks without needing human intervention.
Backward chaining	A method where the model starts with the desired output and works in reverse to find data that might support it.
Bias	Assumptions made by a model that simplify the process of learning to do its assigned task. Most supervised machine learning models perform better with low bias, as these assumptions can negatively affect results.
Big data	Datasets that are too large or complex to be used by traditional data processing applications.
Bounding box	Commonly used in image or video tagging, this is an imaginary box drawn on visual information. The contents of the box are labeled to help a model recognize it as a distinct type of object.
Chatbot	A chatbot is program that is designed to communicate with people through text or voice commands in a way that mimics human-to-human conversation.
Cognitive computing	This is effectively another way to say artificial intelligence. It's used by marketing teams at some companies to avoid the science fiction aura that sometimes surrounds AI.
Computational learning theory	A field within artificial intelligence that is primarily concerned with creating and analyzing machine learning algorithms.
Corpus	A large dataset of written or spoken material that can be used to train a machine to perform linguistic tasks.
Data mining	The process of analyzing datasets in order to discover new patterns that might improve the model.
Data science	Drawing from statistics, computer science and information science, this interdisciplinary field aims to use a variety of scientific methods, processes and systems to solve problems involving data.
Dataset	A collection of related data points, usually with a uniform order and tags.
Deep learning	A function of artificial intelligence that imitates the human brain by learning from the way data is structured, rather than from an algorithm that's programmed to do one specific thing.
Entity annotation	The process of labeling unstructured sentences with information so that a machine can read them. This could involve labeling all people, organizations and locations in a document, for example.
Entity extraction	An umbrella term referring to the process of adding structure to data so that a machine can read it. Entity extraction may be done by humans or by a machine learning model.
Forward chaining	A method in which a machine must work from a problem to find a potential solution. By analyzing a range of hypotheses, the AI must determine those that are relevant to the problem.

General AI	AI that could successfully do any intellectual task that can be done by any human being. This is sometimes referred to as strong AI, although they aren't entirely equivalent terms.
Hyperparameter	Occasionally used interchangeably with parameter, although the terms have some subtle differences. Hyperparameters are values that affect the way your model learns. They are usually set manually outside the model.
Intent	Commonly used in training data for chatbots and other natural language processing tasks, this is a type of label that defines the purpose or goal of what is said. For example, the intent for the phrase "turn the volume down" could be "decrease volume".
Label	A part of training data that identifies the desired output for that particular piece of data.
Linguistic annotation	Tagging a dataset of sentences with the subject of each sentence, ready for some form of analysis or assessment. Common uses for linguistically annotated data include sentiment analysis and natural language processing.
Machine intelligence	An umbrella term for various types of learning algorithms, including machine learning and deep learning.
Machine learning	This subset of AI is particularly focused on developing algorithms that will help machines to learn and change in response to new data, without the help of a human being.
Machine translation	The translation of text by an algorithm, independent of any human involvement.
Model	A broad term referring to the product of AI training, created by running a machine learning algorithm on training data.
Natural language generation (NLG)	The umbrella term for any machine's ability to perform conversational tasks, such as recognizing what is said to it, understanding the intended meaning and responding intelligibly.
Natural language processing (NLP)	As a subset of natural language processing, natural language understanding deals with helping machines to recognize the intended meaning of language — taking into account its subtle nuances and any grammatical errors.
Natural language understanding (NLU)	Also called a neural net, a neural network is a computer system designed to function like the human brain. Although researchers are still working on creating a machine model of the human brain, existing neural networks can perform many tasks involving speech, vision and board game strategy.
Neural network	This refers to the process by which a machine turns structured data into text or speech that humans can understand. Essentially, NLG is concerned with what a machine writes or says as the end part of the communication process.
Overfitting	An important AI term, overfitting is a symptom of machine learning training in which an algorithm is only able to work on or identify specific examples present in the training data. A working model should be able to use the general trends behind the data to work on new examples.
Parameter	A variable inside the model that helps it to make predictions. A parameter's value can be estimated using data and they are usually not set by the person running the model.
Pattern recognition	The distinction between pattern recognition and machine learning is often blurry, but this field is basically concerned with finding trends and patterns in data.

Predictive analytics	By combining data mining and machine learning, this type of analytics is built to forecast what will happen within a given timeframe based on historical data and trends.
Python	A popular programming language used for general programming.
Reinforcement learning	A method of teaching AI that sets a goal without specific metrics, encouraging the model to test different scenarios rather than find a single answer. Based on human feedback, the model can then manipulate the next scenario to get better results.
Semantic annotation	Tagging different search queries or products with the goal of improving the relevance of a search engine.
Sentiment analysis	The process of identifying and categorizing opinions in a piece of text, often with the goal of determining the writer's attitude towards something.
Strong AI	This field of research is focused on developing AI that is equal to the human mind when it comes to ability. General AI is a similar term often used interchangeably.
Supervised learning	This is a type of machine learning where structured datasets, with inputs and labels, are used to train and develop an algorithm.
Test data	The unlabeled data used to check that a machine learning model is able to perform its assigned task.
Training data	This refers to all of the data used during the process of training a machine learning algorithm, as well as the specific dataset used for training rather than testing.
Transfer learning	This method of learning involves spending time teaching a machine to do a related task, then allowing it to return to its original work with improved accuracy. One potential example of this is taking a model that analyzes sentiment in product reviews and asking it to analyze tweets for a week.
Turing test	Named after Alan Turing, famed mathematician, computer scientist and logician, this tests a machine's ability to pass for a human, particularly in the fields of language and behavior. After being graded by a human, the machine passes if its output is indistinguishable from that of human participant's.
Unsupervised learning	This is a form of training where the algorithm is asked to make inferences from datasets that don't contain labels. These inferences are what help it to learn.
Validation data	Structured like training data with an input and labels, this data is used to test a recently trained model against new data and to analyze performance, with a particular focus on checking for overfitting.
Variance	The amount that the intended function of a machine learning model changes while it's being trained. Despite being flexible, models with high variance are prone to overfitting and low predictive accuracy because they are reliant on their training data.
Variation	Also called queries or utterances, these work in tandem with intents for natural language processing. The variation is what a person might say to achieve a certain purpose or goal. For example, if the intent is "pay by credit card," the variation might be "I'd like to pay by card, please."

Weak AI

Also called narrow AI, this is a model that has a set range of skills and focuses on one particular set of tasks. Most AI currently in use is weak AI, unable to learn or perform tasks outside of its specialist skill set.