

The Ultimate Artificial Intelligence Glossary

2024 Edition

Featuring a Dedicated Generative AI Section



What's included in this glossary?

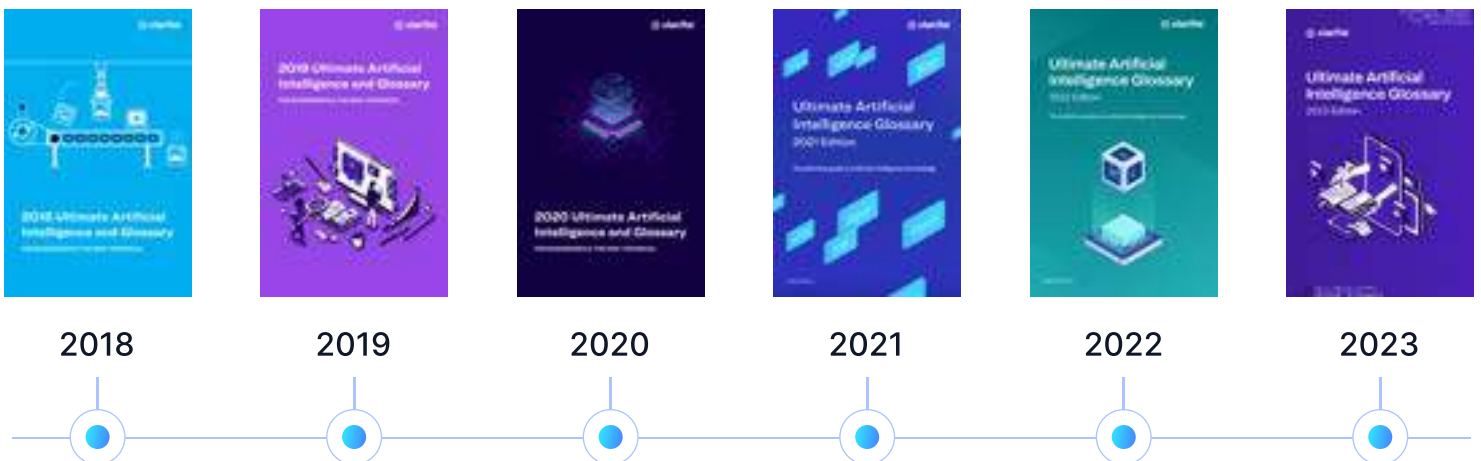
Part 1: General AI Glossary 2-31

Part 2: Dedicated Generative AI Section 32-40

Welcome to your 2024 Ultimate AI Glossary

We've been proud to bring you the annual Ultimate Artificial Intelligence Glossary for the past five years. This year, we've expanded the scope to include many new terms brought into AI conversations across development teams, enterprises, and board rooms thanks to generative AI.

AI has been our core competency for ten years. Our dedication to helping continue to advance the knowledge of AI remains unchanged. As more and more companies explore AI to improve business processes, our glossary enables you to navigate commonly used terms to help accelerate AI readiness. Whether you're just learning about the nuances of AI or want a refresher, this guide is for you. Happy reading!



A

A/B Testing

A statistical way of comparing two (or more) techniques, typically an incumbent against a new rival. It aims to determine which technique performs better, and whether the difference is statistically significant.

Accuracy

Refers to the percentage of correct predictions the classifier made.

Activation Function

In artificial neural networks, the activation function of a node defines the output of that node given an input or set of inputs.

Active Learning

A machine learning term that refers to various methods for actively improving the performance of trained models.



Adversarial Example

Adversarial examples are specialized inputs created with the purpose of confusing a neural network, resulting in the misclassification of a given input. These notorious inputs are indistinguishable to the human eye, but cause the network to fail to identify the contents of the image.

Adversarial Machine Learning

A research field that lies at the intersection of machine learning (ML) and computer security. It enables the safe adoption of ML techniques in adversarial settings like spam filtering, malware detection, etc.

Agents

In the context of AI, agents are software that can independently perform specific tasks without human intervention. They often employ various tools, like calculators or web browsing, to process data and develop solutions.

AI Algorithms

Extended subset of machine learning that tells the computer how to learn to operate on its own through a set of rules or instructions.

AI Ethics in Generative Models

With the advancement of generative AI, the urgency of addressing ethical concerns such as deepfakes, data privacy, and bias within AI has intensified. There are increasing calls for meticulous oversight to guarantee their responsible development and application.

AI-Generated Art and Copyright

The rise of AI in generating art has led to discussions about copyright, ownership, and the definition of creativity.

Anchor Box

The archetypal location, size, and shape for finding bounding boxes in an object detection problem. For example, square anchor boxes are typically used in face detection models.

Annotation

The "answer key" for each image. Annotations are markup placed on an image (bounding boxes for object detection, polygons or a segmentation map for segmentation) to teach the model the ground truth.

Annotation Format

The particular way of encoding an annotation. There are many ways to describe a bounding box's size and position (JSON, XML, TXT, etc) and to delineate which annotation goes with which image.

Annotation Group

Describes what types of objects you are identifying. For example, "chess pieces" or "vehicles".

Application Programming Interface (API)

A set of commands, functions, protocols, and objects that programmers can use to create software or interact with an external system.

Architecture

A specific neural network layout (layers, neurons, blocks, etc). These often come in multiple sizes whose design is similar except for the number of parameters.

Artificial General Intelligence

Computational system that can perform any intellectual task a human can. Also called "Strong AI." At this point, AGI is fictional.

Artificial Intelligence

A computational system that simulates parts of human intelligence but focuses on one narrow task.

Artificial Neural Network

A learning model created to act like a human brain that solves tasks that are too difficult for traditional computer systems to solve.

Artificial Super Intelligence

Artificial Super Intelligence (ASI) refers to a level of AI that surpasses human intelligence across all domains, including creativity, problem-solving, and emotional intelligence.



Audio Speech Recognition (ASR)

A technology that processes human speech into readable text.

Automation Bias

When a human decision maker favors recommendations made by an automated decision-making system over information made without automation, even when the automated decision-making system makes errors.

AutoML

Automates each step of the ML workflow so that it's easier for users with minimal effort and machine learning expertise.

What is **Autonomous AI**?

The most advanced form of AI is autonomous artificial intelligence, in which processes are automated to generate the intelligence that allows machines, bots and systems to act on their own, independent of human intervention. It is often used in autonomous vehicles.

This field of AI is still very new, and researchers are continually refining their algorithms and their approaches to the problem, but it entails multiple layers:

What is **Autonomous AI**? (cont.)

Sensing

Building a model of the constantly shifting world requires a collection of sensors that are usually cameras and often controlled lighting from lasers or other sources. The sensors usually also include position information from GPS or some other independent mechanism.

Fusion

The details from the various sensors must be organized into a single, coherent view of what's happening around the autonomous vehicle. Some images may be occluded. Some may be failing. They may not always be consistent. The sensor fusion algorithms must sort through the details and construct a reliable model that can be used in later stages for planning.

Perception

After the model is constructed, the system must begin to identify important areas like any roads or paths or moving objects.

Planning

Finding the best path forward requires studying the model and also importing information from other sources like mapping software, weather forecasts, traffic sensors and more.

Control

After a path is chosen, any device must ensure that the motors and steering work to move along the path without being diverted by bumps or small obstacles.

In general, information flows from the top layer of the sensors down to the control layer as decisions are made. There are feedback loops, though, that bring information from the lower layers back up to the top to improve sensing, planning and perception.

B

Backward Chaining

A method where the model starts with the desired output and works in reverse to find data that might support it.

Base Workflow

One of Clarifai's prebuilt models that can be built upon to create a custom model. It pre-indexes inputs for search and provides a default embedding space.

Baseline

A model used as a reference point for comparing how well another model (typically, a more complex one) is performing. Baseline models help developers quantify the minimal expected performance that a new model must achieve to be useful.

Batch

The set of examples used in one iteration (that is, one gradient update) of model training.

Batch Inference

Asynchronous process that is executing predictions based on existing models and observations, and then stores the output.

Batch Size

The number of training examples utilized in one iteration.

Bayes's Theorem

A famous theorem used by statisticians to describe the probability of an event based on prior knowledge of conditions that might be related to an occurrence.

Bias

When an AI algorithm produces results that are systemically prejudiced due to erroneous assumptions in the machine learning process.

Big Data

Big data refers to data that is so large, fast or complex that it's difficult or impossible to process using traditional methods.

Binary Classification

The task of classifying elements of a set into two groups on the basis of a classification rule i.e. a model that evaluates email messages and outputs either "spam" or "not spam" is a binary classifier.

Black Box AI

An AI system whose inputs and operations are not visible to the user. A black box, in a general sense, is an impenetrable system.

C

Boosting

A machine learning technique that iteratively combines a set of simple and not very accurate classifiers (referred to as "weak" classifiers) into a classifier with high accuracy (a "strong" classifier) by upweighting the examples that the model is currently misclassifying.

Bootstrapping

Bootstrapping is any test or metric that uses random sampling with replacement and falls under the broader class of resampling methods. Bootstrapping assigns measures of accuracy (bias, variance, confidence intervals, prediction error, etc.) to sample estimates.

Bounding Box

In an image, the (x, y) coordinates of a rectangle around an area of interest.



Brute Force Search

A search that isn't limited by clustering/ approximations; it searches across all inputs. Often more time-consuming and expensive, but more thorough.

Calibration Layer

A post-prediction adjustment, typically to account for prediction bias. The adjusted predictions and probabilities should match the distribution of an observed set of labels.

Chatbot

Simulates human conversation, using response workflows or artificial intelligence to interact with people based on verbal and written cues. Chatbots have become increasingly sophisticated in recent years and in the future may be indistinguishable from humans.

Chain-of-Thought Prompting

A technique used in AI language models to produce more reasoned, step-by-step outputs, especially in problem-solving tasks.

Checkpoint

Data that captures the state of the variables of a model at a particular time. Checkpoints enable exporting model weights, performing training across multiple sessions and continuing training past errors.

Class

One of a set of enumerated target values for a label. For example, in a binary classification model that detects spam, the two classes are spam and not spam. In a multi-class classification model that identifies dog breeds, the classes would be poodle, beagle, pug, etc.

Class Balance

The relative distribution between the number of examples of each class used to train a model. A model performs better if there are a relatively even number of examples for each class.

Classification

Process of grouping and categorizing objects and ideas recognized, differentiated, and understood in data.

Classifier

A model that implements classification. It refers to the mathematical function implemented by a classification algorithm that maps input data to a category.

Cluster

A group of observations that show similarities to each other and are organized by similarities.

Clustering

A method of unsupervised learning and common statistical data analysis technique. In this method, observations that show similarities to each other are organized into groups (clusters).

Cognitive Computing

A computerized model that mimics the way the human brain thinks. It involves self learning through the use of data mining, natural language processing, and pattern recognition.

Computer Vision

Field of AI that trains computers to interpret and understand the visual world. Using digital images from cameras and videos and deep learning models, machines can accurately identify and classify objects — and then react to what they “see.”

Concept

Describes an input, similar to a "tag" or "keyword." There are two types: those that you specify to train a model, and those that a model assigns as a prediction.

Confidence

A model is inherently statistical. Along with its prediction, it also outputs a confidence value that quantifies how "sure" it is that its prediction is correct.

Confidence Threshold

We often discard predictions that fall below a certain bar. This bar is the confidence threshold.

Confusion Matrix

A confusion matrix is a table that is often used to describe the performance of a classification model on a set of test data for which the true values are known. The confusion matrix itself is relatively simple to understand, but the related terminology can be... confusing.

Container

A virtualized environment that packages its dependencies together into a portable environment. Docker is one common way to create containers.

Convolutional Filter

A convolution is a type of block that helps a model learn information about relationships between nearby pixels.

Convolutional Neural Network

Convolutional neural networks are deep artificial neural networks that are used primarily to classify images (e.g. name what they see), cluster them by similarity (photo search), and perform object recognition within scenes.

CoreML

A proprietary format used to encode weights for Apple devices that takes advantage of the hardware accelerated neural engine present on iPhone and iPad devices.

CreateML

A no-code training tool created by Apple that will train machine learning models and export to CoreML. It supports classification and object detection along with several types of non computer-vision models (such as sound, activity, and text classification).

Curse of Dimensionality

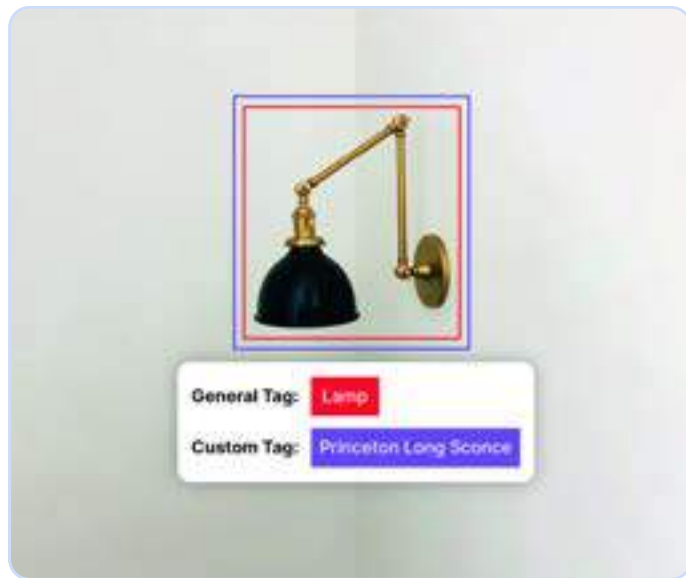
The curse of dimensionality refers to various phenomena that arise when analyzing and organizing data in high-dimensional spaces that do not occur in low-dimensional settings such as the three-dimensional physical space of everyday experience.

Custom Dataset

A set of images and annotations pertaining to a domain specific problem. In contrast to a research benchmark dataset like coco or Pascal voc.

Custom Training

The process of teaching a model to make certain predictions.



D

Data

In the data science and AI world, any collection of information that is converted into a digital form. It's important to distinguish between structured and unstructured data; structured data is highly specific and is stored in a predefined format such as an spreadsheet table, whereas unstructured data is a conglomeration of many varied types of data that are stored in their native formats, such as images, video, audio, and text. Data is also a plural, with the singular being "datum".

Data Annotation

The process of labeling datasets to be used as inputs for machine learning models.

Data Curation

The process of collecting, organizing, cleaning, labeling, and maintaining data for use in training and testing models.

Data Mining

The process by which patterns are discovered within large sets of data with the goal of extracting useful information from it.

Dataset

A collection of data and a ground truth of outputs that you use to train machine learning models by example.

De-duplication

The removal of identical data, or data that is so similar that for all intents and purposes it can be considered duplicate data. Using visual search, a similarity threshold can be set to decide what should be removed.



Deep Learning

The general term for machine learning using layered (or deep) algorithms to learn patterns in data. It is most often used for supervised learning problems.

Deep Neural Network

An artificial neural network (ANN) with multiple layers between the input and output layers. It uses sophisticated mathematical modeling to process data in complex ways.

E

Deploy

Taking the results of a trained model and using them to do inference on real world data. This could mean hosting a model on a server or installing it to an edge device.

Detection Mode

Also known as object detection. A model that identifies the presence, location and type of objects within images or video frames.

Diversity, Equity and Inclusion (DEI)

Term used to describe policies and programs that promote the representation and participation of different groups of individuals, including people of different ages, races and ethnicities, abilities and disabilities, genders, religions, cultures, and sexual orientations.

Domain Adaptation

A type of transfer learning, domain adaptation is a technique to improve the performance of a model where there is little data in the target domain by using knowledge learned by another model in a related domain. An example could be training a model to recognize taxis using a model that recognizes cars.

Edge AI

Data is processed on the same device that produces it, or at most on a nearby computer. Edge AI means there's no reliance on distant cloud servers or other remote computing nodes, allowing the AI to work faster, and respond more accurately to time-sensitive events.



Edge Computing

A distributed computing framework that brings enterprise applications closer to data sources such as IoT devices or local edge servers.

Embeddings

A categorical feature represented as a continuous-valued feature. Typically, an embedding is a translation of a high-dimensional vector into a low-dimensional space.

Embedding Space

The d-dimensional vector space that features from a higher-dimensional vector space are mapped to. Ideally, the embedding space contains a structure that yields meaningful mathematical results.

Emotional AI

Emotional AI refers to technologies that use affective computing and artificial intelligence techniques to sense, learn about and interact with human emotional life.

Ensemble Models

Ensemble models are a machine learning approach to combine multiple other models in the prediction process. While the individual models may not perform very well, when combined they can be very powerful indeed.

Extensible Markup Language (XML)

A markup language and file format for storing, transmitting, and reconstructing arbitrary data. It defines a set of rules for encoding documents in a format that is both human-readable and machine-readable.

F

F Score

A weighted average of the true positive rate of recall and precision.

Facial Recognition

An application capability of identifying or verifying a person from an image or a video frame by comparing selected facial features from the image and a face database.

False Negatives

An error where a model falsely predicts an input as not having a desired outcome, when one is actually present. (Actual Yes, Predicted No).

False Positives

An error where a model falsely predicts the presence of the desired outcome in an input, when in reality it is not present (Actual No, Predicted Yes).

FastAI

A library built on top of PyTorch for rapid prototyping and experimentation. There is a companion course that teaches the fundamentals of machine learning.

Feature Extraction

When image features at various levels of complexity are extracted from the image data. Typical examples of such features are: Lines, edges, and ridges. Localized interest points such as corners, blobs, or points. More complex features may be related to texture, shape, or motion.

The process by which data that is too large to be processed is transformed into a reduced representation set of features such as texture, shape, lines, and edges.

Folksonomy

User-generated system of classifying and organizing online content into different categories by the use of metadata such as electronic tags.

Framework

Deep learning frameworks implement neural network concepts. Some are designed for training and inference—TensorFlow, PyTorch, FastAI, etc. And others are designed particularly for speedy inference—OpenVino, TensorRT, etc.

G

Generalization

Refers to a model's ability to make correct predictions on new, previously unseen data as opposed to the data used to train the model.

Generative Adversarial Networks (GANs)

A class of artificial intelligence algorithms used in unsupervised machine learning, implemented by a system of two neural networks contesting with each other in a zero-sum game framework. This technique can generate photographs that look at least superficially authentic to human observers, having many realistic characteristics (though in tests people can tell real from generated in many cases).

Generative AI

Models that can be trained using existing content like text, audio files, or images to create new original content.

Graphics Processing Unit (GPU)

A specialized electronic circuit designed to rapidly manipulate and alter memory to accelerate the creation of images in a frame buffer intended for output to a display device. GPUs are used in embedded systems, mobile phones, personal computers, and more.

Green AI

Efforts to make AI more energy-efficient and environmentally friendly are gaining momentum, with a focus on reducing the carbon footprint of training and running AI models.

Grid Search

Grid search is a tuning technique that attempts to compute the optimal values of hyperparameters for training models by performing an exhaustive search through a subset of hyperparameters.

Ground Truth

The "answer key" for your dataset. This is how you judge how well your model is doing and calculate the loss function we use for gradient descent. It's also what we use to calculate our metrics. Having a good ground truth is extremely important. Your model will learn to predict based on the ground truth you give it to replicate.

GPU Memory

The amount of information your GPU can fit on it. A bigger GPU will be able to process more information in parallel which means it can support bigger models (or bigger batch sizes) without running out of memory. If you run out of GPU memory it will crash your program.

H

Hashing

In machine learning, a mechanism for bucketing categorical data, particularly when the number of categories is large, but the number of categories actually appearing in the dataset is comparatively small.

Hidden Layer

A synthetic layer in a neural network between the input layer (that is, the features) and the output layer (the prediction). Hidden layers typically contain an activation function (such as ReLU) for training. A deep neural network contains more than one hidden layer.

Holdout Data

Examples intentionally not used during training. The validation dataset and test dataset are examples of holdout data. It helps evaluate your model's ability to generalize to data other than the data on which it was trained.

Hosted Model

A set of trained weights located in the cloud that you can receive predictions from via an API.

Human Workforce

Workers who can help to complete work on an as-needed basis, which for purposes usually means labeling data (images).



Hyperparameter

The levers by which you can tune your model during training. These include things like learning rate and batch size. You can experiment with changing hyperparameters to see which ones perform best with a given model for your dataset.

I

Inference

Making predictions using the weights you save after training your model.

ImageNet

A large visual database designed for use in visual object recognition software research.

Image Recognition

The ability of software to identify objects, places, people, writing and actions in images.

Image Segmentation

The process of dividing a digital image into multiple segments with the goal of simplifying the representation of an image into something that is easier to analyze. Segmentation divides whole images into pixel groupings, which can then be labeled and classified.

Implicit Bias

Automatically making an association or assumption based on one's mental models and memories. Implicit bias can affect how data is collected and classified, and how machine learning systems are designed and developed.

Information Retrieval

The area of Computer Science studying the process of searching for information in a document, searching for documents themselves, and also searching for metadata that describes data and for databases of texts, images or sounds.

Input

Any information or data sent to a computer for processing is considered input.

Input Layer

The first layer (the one that receives the input data) in a neural network.

Intelligent Character Recognition (ICR)

Related technology to OCR designed to recognize handwritten characters.

IoT

The interconnection via the internet of computing devices embedded in everyday objects, enabling them to send and receive data.

J

Jetson

An edge computing device created by NVIDIA that includes an onboard GPU.

JSON

A freeform data serialization format originally created as part of JavaScript but now used much more broadly. Many annotation formats use JSON to encode their bounding boxes.

Jupyter Notebook

A common data science tool that enables you to execute Python code visually. Each "cell" in the notebook is a block of code that you can execute by hitting "Ctrl+Enter". The results of the execution are displayed below the cell.

K

Knowledge Graph

A knowledge graph is not the kind of graph you may be used to showing the relationship between two variables. Rather, it's a collection of nodes and edges where the nodes represent concepts, entities, relationships, and events, and the edges represent the connections between them.

L

Label

Assigning a class or category to a specific object in your dataset.

Labeler

AI-automated tool using end-to-end workflows to label images and video at scale to create high-quality training data.



Labeling

Also known as data labeling; the process of annotating datasets to train machine learning models.

Labeling Criteria

An essential guide to the labeling requirements for a project. It should include instructions for the labeling process itself as well as written definitions and a multitude of visual examples for each concept. These examples help the labelers anticipate the visual differences and edge cases that can occur within one single concept.

M

Machine Intelligence

An umbrella term that encompasses machine learning, deep learning and classical learning algorithms.

Machine Learning

A general term for algorithms that can learn patterns from existing data and use these patterns to make predictions or decisions with new data.

Masked Language Model

A language model that predicts the probability of candidate tokens to fill in blanks in a sequence.

Metadata

Information about an analog or digital object, a component of an object, or a coherent collection of objects. Metadata describing digital content is often structured (e.g., with tagging or markup).

Misclassification Rate

Rate used to gauge how often a model's predictions are wrong.

MLOps

Also known as Machine Learning Operations. Best practices for organizations to operationalize machine learning. Often involves collaboration between data scientists and devops professionals to manage production ML.

Modality

A high-level data category. For example, numbers, text, images, video and audio are five different modalities.

Model

The representation of what a machine learning system has learned from the training data.

ModelOps

As defined by Gartner, ModelOps is focused primarily on the governance and life cycle management of a wide range of operationalized artificial intelligence (AI) and decision models, including machine learning, knowledge graphs, rules, optimization, linguistic and agent-based models.

Model Size

The number of parameters (or neurons) a model has and can be measured in terms of the size of the weights file on disk.

Types of Models

It's important to understand the different types of models for different data types.

Audio Transcription Model

These models take audio containing speech and convert it into text. These can be extremely useful as they allow audio to be searched for key terms, or AI models to transmit text instead of audio over networks, which is much smaller and faster.

Classification Model

Reads an input such as text, image, audio, or video data and generates an output that classifies it into a category. For example, a language classification model might read a sentence and determine whether it's in French, Spanish, or Italian.

Detection Model

Detection comprises two tasks; listing "what" things appear in an image, and "where" they appear. Results are returned as bounding boxes along with the names of the detected items.

Domain Model

Focuses on understanding a single domain, such as travel, weddings, food, not-safe-for-work (NSFW), etc.

Embedding Model

Computers and models can't understand images and text like humans do. Embedding models take unstructured input like images, audio, text, and video and transform them into a series of numbers called vectors which can then be input into the prediction models.

Masked Language Model

A language model that predicts the probability of which words make the most sense to fill in blanks in a sequence. A simple example could be "Good ____, how are you?" where probable candidate words could be "morning", "day", or "evening".

Types of Models (cont.)

Multimodal Model

A model whose inputs and/or outputs include more than one modality. For example, consider a model that takes both an image and a text caption (two modalities) as features, and outputs a score indicating how appropriate the text caption is for the image.



Named Entity Recognition (NER) Model

A method that is used for recognizing entities such as people, dates, organizations, and locations that are present in a text document.

Pre-trained Model

A model, or the component of a model, that has been preliminary trained, generally using another data set. (for example, finding lines, corners, and patterns of colors). Pre-training on a large dataset like the huge Common Objects in Context (COCO), which has 330,000 images with 1.5 million objects to detect, can reduce the number of custom images you need to obtain satisfactory results.

Predictive Model

A model that uses observations measured in a sample to gauge the probability that a different sample or remainder of the population will exhibit the same behavior or have the same outcome.

Segmentation Model

Instead of bounding boxes returned for each concept, this model indicates via a heat map and trace (think coloring book) of regions for each concept.

Model Training

The process of determining the best model.

Monte Carlo Simulation

A Monte Carlo simulation is used to model the probability of different outcomes in a process that cannot easily be predicted due to the intervention of random variables. It is a technique used to understand the impact of risk and uncertainty. It was developed while working on nuclear weapons in the 1940s, and was given the code name “Monte Carlo” in reference to the Monte Carlo Casino in Monaco, where one of the inventor’s uncles would borrow money from relatives to gamble.

Multi-class Classification

Classification problems that distinguish between more than two classes. For example, there are approximately 53 species of maple trees, so a model that categorized maple tree species would be multi-class.

N

Named Entity Recognition (NER)

A sub-task of information extraction that seeks to identify and classify named entities in text into predetermined categories such as the names, locations, parts-of-speech, etc.

Natural Language Processing (NLP)

A branch of AI that helps computers understand, interpret, and manipulate human language. This field of study focuses on helping machines understand human language in order to improve human-computer interfaces.

Natural Language Understanding

Determining a user’s intentions based on what the user typed or said. For example, a search engine uses natural language understanding to determine what the user is searching for based on what the user typed or said.

Neural Architecture Search

Automatically trying many variations of model layouts and hyperparameters to find the optimal configuration.

Neural Network

Series of algorithms that endeavors to recognize underlying relationships in a set of data through a process that mimics the way the human brain operates.

Neuro-Symbolic

The combining of neural and symbolic AI architectures to address complementary strengths and weaknesses of each, providing a robust AI capable of reasoning, learning, and cognitive modeling.

Neuron

A unit in an Artificial Neural Network processing multiple input values to generate a single output value.

Noise

Signals with no causal relation to the target function.

Normalization

The process of converting an actual range of values into a standard range of values, typically -1 to +1 or 0 to 1.

O

Object Detection

A computer technology related to computer vision and image processing that deals with detecting instances of semantic objects of a certain class (such as humans, buildings, or cars) in digital images and videos. This technique also involves localizing the object in question, which differentiates it from classification, which only tells the type of object.

Object Recognition

Also known as object classification. A computer vision technique for identifying objects in images or videos.

Object Tracking

The process of following a specific object of interest, or multiple objects, in a given scene. It traditionally has applications in video and real-world interactions where observations are made following an initial object detection.

On-premise Software

Software that is installed and runs on computers located on the premises of the organization using that software versus at a remote facility such as a server farm or on the cloud.

One Shot Classification

A model that only requires that you have one training example of each class you want to predict on. The model is still trained on several instances, but they only have to be in a similar domain as your training example.

Open Neural Network Exchange (ONNX)

ONNX is an open format to represent machine learning models.

OpenAI

U.S. AI research organization whose mission is to ensure that artificial general intelligence benefits all humanity.

Optical Character Recognition (OCR)

A computer system that takes images of typed, handwritten, or printed text and converts them into machine-readable text.

P

Optimization

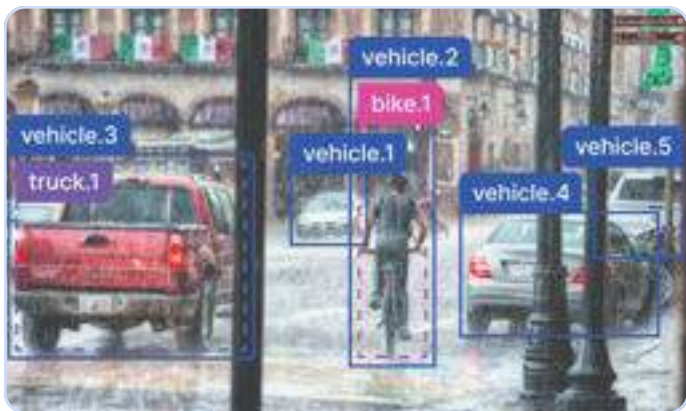
The selection of the best element (with regard to some criterion) from some set of available alternatives.

Output

Predictions made after the input uploaded to or fed into a model are processed by the model.

Outsourced Labeling

Paying people to annotate, or label, your data. Its effectiveness can depend on the domain expertise of annotators. Providing a comprehensive labeling criteria is crucial for training annotators before beginning a project.



Overfitting

A machine learning problem where an algorithm is unable to discern information that is relevant to its assigned task from information which is irrelevant within training data. Overfitting inhibits the algorithm's predictive performance when dealing with new data.

Parameter

Any characteristic that can be used to help define or classify a system. In AI, they are used to clarify exactly what an algorithm should be seeking to identify as important data when performing its target function.

Pattern Recognition

A branch of machine learning that focuses on the recognition of patterns and regularities in data, although it is in some cases considered to be nearly synonymous with machine learning.

Pipeline

The process of going from raw images to prediction. Usually this encompasses collecting images, annotation, data inspection and quality assurance, transformation, preprocessing and augmentation, training, evaluation, deployment, inference (and then repeating the cycle to improve the predictions).

Polygon

A (usually non-rectangular) region defining an object with more detail than a rectangular bounding box. Polygon annotations can be used to train segmentation models or to enhance performance of object-detection models by enabling a more accurate bounding box to be maintained after augmentation.



Positive Predictive Value (PPV)

Very similar to precision, except that it takes prevalence into account. In the case where the classes are perfectly balanced (meaning the prevalence is 50%), the positive predictive value is equivalent to precision.

Precision

Indicator of a machine learning model's performance – the quality of a positive prediction made by the model. Refers to the number of true positives divided by the total number of positive predictions.

Prediction

An attempt by a model to replicate the ground truth. A prediction usually contains a confidence value for each class.

Prevalence

The rate of how often the “yes” condition actually occurs in a sample.

Production

The deployment environment where the model will run in the wild on real-world images (as opposed to the testing environment where the model is developed).

Prompt Engineering

The practice of crafting effective prompts to elicit better outputs from language models has become an area of interest, especially with models like ChatGPT and DALL-E.

Proximal Policy Optimization (PPO)

PPO is a simplified and efficient reinforcement learning algorithm that refines policy functions to optimize performance in a given environment. It builds on TRPO's approach to ensuring stable policy updates but simplifies the process, making it more practical for implementation.

Pruning

The use of a search algorithm to cut off undesirable solutions to a problem in an AI system. It reduces the number of decisions that can be made by the AI system.

PyTorch

A popular open source deep learning framework developed by Facebook. It focuses on accelerating the path from research prototyping to production deployment.

Q

Quick Training

A process that accelerates the training of machine learning models by utilizing advanced techniques and optimized architectures. It often involves transfer learning, where a pre-trained model is fine-tuned with a specific, smaller dataset, making the training more efficient and less resource-intensive.

R

RAG (Retrieval Augmented Generation)

This combines a retrieval system with a generator to improve the factual accuracy of generated text. It retrieves relevant documents and uses them to inform the generation process.

RAG-Related Terms

Augmented Language Models with Memory Networks

Incorporating memory networks into language models to provide a form of external knowledge storage and retrieval.

Contextualized Embeddings for Retrieval

The use of embeddings that capture the context within which the query is made, improving the relevance of retrieved information in RAG systems.

Cross-Modal Retrieval

Expanding the retrieval component to include different types of data, such as images or audio, alongside text.

Dense Passage Retrieval (DPR)

This is a technique for retrieving relevant document passages using dense vector representations. DPR is often used in RAG to retrieve information from a knowledge base.

End-to-End Trainable Retrieval Models

RAG and similar systems often involve training the retrieval and generation components jointly, leading to more cohesive and effective models.

Few-Shot Learning in Information Retrieval

The application of few-shot learning techniques to improve the retrieval component in RAG, enabling it to adapt to new domains or tasks with minimal training data.

RAG-Related Terms (cont.)

Hybrid Retrieval-Generation Models

These models combine retrieval and generation capabilities to enhance the output. RAG is a prime example, but there are other hybrid models that follow a similar approach.

Knowledge-Intensive Language Tasks (KILTs)

These are tasks that require external knowledge beyond the training data. RAG is particularly relevant for these tasks as it augments generation with retrieved information.

Latent Knowledge Models

These models infer knowledge that's not explicitly stated in the text. RAG can be seen as a method to externalize latent knowledge through retrieval.

Neural Information Retrieval (NeuIR)

This involves using neural network-based approaches for information retrieval. NeuIR is a key component in RAG systems, where the retrieval model is often powered by neural networks.

Question Answering with External Knowledge Bases

Using RAG for QA tasks where the answers are not contained within the training data but require external information sources.

Recall (Sensitivity)

The fraction of relevant instances that have been retrieved over the total amount of relevant instances.

Receiver Operating Characteristic (ROC) Curve

This is a commonly used graph that summarizes the performance of a classifier over all possible thresholds. It is generated by plotting the True Positive Rate (y-axis) against the False Positive Rate (x-axis) as you vary the threshold for assigning observations to a given class.

Recurrent Neural Network

A type of artificial network with loops in them, allowing recorded information, like data and outcomes, to persist by being passed from one step of the network to the next. They can be thought of as multiple copies of the same network with each passing information to its successor.

Regression

A statistical measure used to determine the strength of the relationships between dependent and independent variables.

Reinforcement Learning

A type of machine learning in which machines are "taught" to achieve their target function through a process of experimentation and reward receiving positive reinforcement when its processes produce the desired result and negative reinforcement when they do not.

This is differentiated from supervised learning, which would require an annotation for every individual action the algorithm would take.



ReLU

In the context of artificial neural networks, the ReLU (rectified linear unit) activation function is an activation function which outputs the same as its input if the input is positive, and zero if the input is negative. A related function is the leaky rectified linear unit (leaky rectified linear unit) which assigns a small positive slope for $x < 0$. (🔗)

Responsible AI

Umbrella term for aspects of making appropriate business and ethical choices when adopting AI, including business and societal value, risk, trust, transparency, fairness, bias mitigation, explainability, accountability, safety, privacy and regulatory compliance.

S

Search Query

A query that a user feeds into a search engine to satisfy his or her information needs. If the query itself is a piece of visual content then that is what is known as a "visual search query."

Selective Filtering

When a model ignores "noise" to focus on valuable information.

Siamese Networks

A different way of classifying image where instead of training one model to learn to classify image inputs it trains two neural networks that learn simultaneously to find similarity between images.

Signal

Inputs, information, data.

Software Development Kit (SDK)

A set of software development tools that allows for the creation of applications on a specific platform.

Specificity

The rate of how often a model predicts "no," when it's actually "no."

Standard Classification

The process by which an input is assigned to one of a fixed set of categories. In machine learning, this is often achieved by learning a function that maps an input to a score for each potential category.

Strong AI

A theoretical form of AI that replicates human functions, such as reasoning, planning, and problem-solving.

Structured Data

Data that resides in a fixed field within a file or record. Structured data is typically stored in a relational database.

It can consist of numbers and text, and sourcing can happen automatically or manually, as long as it's within an RDBMS structure.

Supervised Learning

A machine learning approach that's defined by its use of labeled datasets. These datasets are designed to train or "supervise" algorithms into classifying data or predicting outcomes accurately. Using labeled inputs and outputs, the model can measure its accuracy and learn over time.

Symbiotic Intelligence

A combination of human and artificial intelligence. Instead of relying on memory, or having to open a book, or visit a website, an enhanced human could have access to all of the information that is stored on the internet, and an advanced AI could feed the relevant data points to the human brain, enabling the human to be fully in control.

Synthetic Data

Training data that is artificially generated rather than produced by real-world events.

Synthetic Intelligence

Synthetic intelligence (SI) is an alternative term for artificial intelligence emphasizing that the intelligence of machines need not be an imitation or in any way artificial; it can be a genuine form of intelligence. An analogy can be made with simulated diamonds (such as cubic zirconia) versus synthetic diamonds (real diamonds made of carbon created by humans).

T

Target Function

The end goal of an algorithm.

Taxonomy

In essence, a taxonomy is a model's worldview, or the framework for how your model sees its training data. In practice, it's a list of visually-distinct model concepts and the definitions of those concepts.

Temporal Data

Data recorded at different points in time.

TensorFlow

An open-source software library also used for machine learning applications such as neural networks.

Test Dataset

The sample of data used to provide an unbiased evaluation of a final model fit on the training dataset.

Torch

A scientific computing framework with wide support for machine learning algorithms, written in C and lua.

Train

Training a model is the process of iteratively adjusting its parameters to converge on the weights that optimally mimic the training data.

Training Dataset

Training data is the initial dataset used to train machine learning algorithms. Models create and refine their rules using this data. It's a set of data samples used to fit the parameters of a machine learning model to training it by example.

Transfer Learning

Transferring information from one machine learning task to another. It might involve transferring knowledge from the solution of a simpler task to a more complex one, or involve transferring knowledge from a task where there is more data to one where there is less data.

Transformer

A neural network that transforms a sequence of elements (like words in a sentence) into another sequence to solve sequence-to-sequence tasks.

U

True Negatives

Actual negatives that are correctly identified as such (Actual No, Predicted No).

True Positives

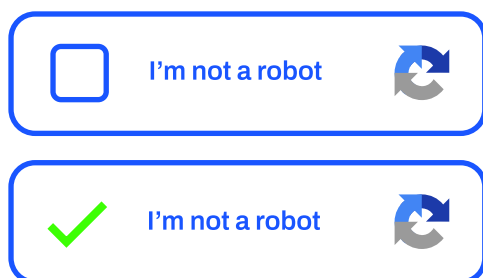
Actual positives that are correctly identified as such (Actual Yes, Predicted Yes).

Trust Region Policy Optimization (TRPO)

TRPO is a reinforcement learning technique that iteratively updates policy functions to enhance performance. It guarantees that each new policy is at least as good as the previous one, avoiding large updates that could degrade performance.

Turing Test

A test developed by Alan Turing in 1950, used to identify true artificial intelligence. It tested a machine's ability to exhibit intelligent behavior equivalent to, or indistinguishable from, that of a human.



Unstructured Data

Unstructured data is information that either does not have a pre-defined data model or is not organized in a pre-defined manner. Unstructured data may include documents, images, video and audio.

Unsupervised Learning

Uses machine learning algorithms to analyze and cluster unlabeled datasets. These algorithms discover hidden patterns or data groupings without the need for human intervention. Its ability to discover similarities and differences in information make it the ideal solution for exploratory data analysis, cross-selling strategies, customer segmentation, and image recognition.

V

Validate

The model is given new, previously unseen data, and then metrics are collected on how well it performs predictions on them. This is analogous to a human learning math problems using one set of questions, then tested to see if they learned properly with a different set of questions.

Validation Data Set

The sample of data used to provide an unbiased evaluation of a model fit on the training dataset while tuning model hyperparameters. The evaluation becomes more biased as skill on the validation dataset is incorporated into the model configuration.

Variance

The error due to sensitivity to fluctuations in the training set computed as the expectation of the squared deviation of a random variable from its mean.

Verify/Verification

The process of verifying that labeled data has been labeled correctly in adherence to the ground truth.

Video Frame Interpolation

The process of synthesising several frames of video, between two originally adjacent frames of video.

Visual Dictionary

A document that defines every model concept with a written definition and also a wide array of visual examples. This helps establish ground truth by providing confirmation that each involved party understands the model's worldview, or taxonomy.

Visual Match

Instead of doing a search which returns the items in the database in sorted order, a visual match couple be considered returning a yes/no answer of whether the query is close enough to any item in the database to be considered a "match".

Visual Recognition

The ability of software to identify objects, places, people, writing, and actions in images and videos.

Visual Search

The ability of software to find visually similar content based on an image or video query.



W

WeakAI

Also known as narrow AI, weak AI refers to a non-sentient computer system that operates within a predetermined range of skills and usually focuses on a singular task or small set of tasks. Most AI in use today is weak AI.

Weight

A coefficient for a feature in a linear model, or an edge in a deep network. The goal of training a linear model is to determine the ideal weight for each feature. If a weight is 0, then its corresponding feature does not contribute to the model.

Width

The number of neurons in a particular layer of a neural network.

Y

Workflows

Workflows enable users to make predictions on a graph that combines one or more pre-trained, custom models and fixed function model operators using a single API call.

YAML

A markup language originally invented by Yahoo that is now commonly used as a format for configuration files.



Generative Artificial Intelligence Glossary



A

Adversarial Autoencoder (AAE)

A type of autoencoder which combines the principles of adversarial loss, integral to GANs, and the architecture of an autoencoder. This combination empowers the model to learn complex distributions of data effectively.

Audio Synthesis

This involves using AI to create new, artificial sounds or voice outputs. Such sounds can be as simple as a specific tone or as complex as a mimicked form of speech.

Auto-regressive Models

These are generative models that produce data by conditioning each element's probability on previous elements in a sequence. For example, WaveNet and PixelCNN are autoregressive models for creating music and images, respectively.

Autoencoder

An autoencoder is an artificial neural network utilized for learning efficient encodings of input data. It has two crucial components: an encoder that compresses the input data and a decoder that reconstructs the data from its reduced form.

Autoregressive Generative Models

These models predict the distribution of subsequent sequence elements using prior sequence elements to implicitly establish a distribution across sequences using Conditional Probability's Chain Rule. The main architectures for autoregressive models are causal convolutional networks and recurrent neural networks.

B

BERT (Bidirectional Encoder Representations from Transformers)

BERT, developed by Google, is a pre-trained transformer-based language model. It stands out for its bidirectional training approach, which allows it to understand the context of a word based on all of its surroundings (left and right of the word).

Bloom

Developed by The BLOOM project, Bloom is a large-scale language model that can execute a vast array of natural language understanding and generation tasks accurately.

C

ChatGPT

Developed by OpenAI, ChatGPT is a specialized large-scale language model that generates human-like text. It's a popular choice for developing AI-powered chatbots due to its convincing conversation-generation capabilities.

CLIP (Contrastive Language—Imagen Pretraining)

This involves using AI to create new, artificial sounds or voice outputs. Such sounds can be as simple as a specific tone or as complex as a mimicked form of speech.

D

Conditional GANs (cGANs)

These are a type of GAN where a conditional variable is introduced to the input layer, allowing the model to generate data conditioned on certain factors. This augmentation provides the model the capability to generate data with desired characteristics.



Crossmodal

Cross-modal learning refers to using information from one modality to understand or make predictions in another modality. This could involve translating or transforming the data in some way. For example, a cross-modal learning system might be designed to accept text input and output a related image or vice versa.

CycleGAN

A type of GAN that can translate an image from a source domain to a target domain without paired examples. It's particularly useful in tasks like photo enhancement, image colorization, and style transfer for unpaired photo-to-photo translation.

DALL-E 2

This is an updated version of DALL-E, an AI model developed by OpenAI to generate images from textual descriptions. It's an excellent example of a multi-modal AI system.

Data Distribution

In machine learning, data distribution refers to the overall layout or spread of data points within a dataset. In the case of generative models such as GANs, the generator seeks to mimic the actual data distribution.

Deepfake

Synthetic media in which a person in an existing image or video is replaced with someone else's likeness using machine learning techniques. While they could serve interactive entertainment purposes, deepfakes may mislead viewers, often with unintended consequences.

Diffusion

In AI, 'diffusion' refers to a technique used for generating new data by starting with a portion of actual data, then gradually adding random noise. This process is generally reversed, with a neural network trained to predict the reverse process of noise addition to the data.

Discriminator

In a GAN, the discriminator is the component that tries to differentiate real data instances from the fictitious ones fabricated by the generator. It helps refine the generator's ability to create realistic data.

E

Embedding

An embedding represents data in a new form, often a vector space, facilitating comparisons and calculations with other data points. Similar items should have similar embeddings, making it an essential feature for many AI tasks, like recommendation systems and natural language processing.

Emergence/Emergent Behavior

("sharp left turns," intelligence explosions). In artificial intelligence, emergence refers to complex phenomena that arise from simple rules or processes. Radical concepts like "sharp left turns" and "intelligence explosions" denote sudden, dramatic developments in AI, often related to AGI's emergence.

F

Few-Shot Learning

A machine learning method where the model learns to perform a task from a few examples per class. For instance, it can correctly categorize new data after being shown only a few samples from each category.



Fine-tuning

A form of transfer learning wherein a pre-trained model is slightly modified or adjusted to perform a new task. This process allows for more efficient use of the pre-trained models by adjusting them to solve tasks similar to ones they were originally trained on.

Foundation Model

In AI, foundation models are large-scale AI models trained on diverse and extensive data meant to be fine-tuned or adapted for more specific tasks. These are called foundation models, as they offer a robust and broad foundation that can be built upon for various AI tasks.

G

Generative Models for Images

These are generative models like GANs, VAEs, and DALL-E, trained on image data and capable of generating new images that reflect the patterns found in the training data.

Generative Pre-trained Transformer (GPT)

GPT is a family of neural network models trained to generate content. These models are pre-trained on vast amounts of text data, allowing them to generate coherent and relevant text based on user prompts. GPT models can automate content creation and analyze customer feedback for insights, fostering personalized interactions.

Generator

In a Generative Adversarial Network, the generator is the component that creates new instances of data by learning to mimic the real data distribution.

GPT-1, GPT-2, GPT-3, and GPT-4

Progressive versions of the generative pre-trained transformers developed by OpenAI. Each model sees improvements and expansions on its predecessors, offering advanced text generation capabilities and greater application versatility. GPT-3, for instance, is an extremely sophisticated model known for its wide-ranging applicability, including translation, question-answering, and text completion tasks.

GPT-J

GPT-J is an open-source large language model developed by EleutherAI in 2021. It is a generative pre-trained transformer model with 6 billion parameters, similar to GPT-3, but with some architectural differences. GPT-J was trained on a large-scale dataset called The Pile, a mixture of sources from different domains.

GPT-Neo

GPT-Neo is a family of transformer-based language models from EleutherAI based on the GPT architecture. It is an open-source alternative to GPT-3 that can generate natural language texts using deep learning. GPT-Neo has two pre-trained models with 1.3 billion and 2.7 billion parameters and was also trained using The Pile dataset.

Grounding

Is the process of linking a model's output to factual and verifiable information sources. This technique enhances the accuracy and reliability of the model, especially in applications where factual correctness is critical. Grounding reduces the risk of the model generating unfounded or incorrect content.

H

Hallucination

In AI, a hallucination occurs when a model makes erroneous conclusions and generates content that doesn't correspond to reality. These erroneous outputs indicate problems in the workings of the AI model. Team vigilance is necessary to maintain the accuracy and reliability of AI systems in identifying and mitigating hallucinations.

I

Image Translation

A task in computer vision where the goal is to map or translate one image into another, often using a model known as GANs. For example, translating a daytime scene into a nighttime scene.

Inpainting

A generative task where the AI is meant to fill in missing or corrupted parts of an image. Typical applications include photo restoration and the completion of unfinished art.

Latent Space

In generative models, latent space refers to a compressed input data representation. It is the transition medium between the noise injected into the GAN's generator and its output.

L

Langchain

Langchain is a concept in AI and machine learning that affects reasoning capability. When prompting an LLM, the "chain-of-thought" technique improves the model's reasoning by breaking tasks into smaller, discrete steps. A more complex approach, "tree-of-thought," allows logical steps to branch and backtrack.

Large Language Models (LLMs)

Large-scale AI models trained on extensive text data, such as GPT-3 and BERT. They can respond to prompts, generate text, answer questions, create poetry, and even generate code. This ability can enable personalized and authentic customer interactions and assist in automating customer-facing content.



M

Machine Learning Bias

Bias in machine learning can occur from intentional or unintentional biased data or algorithms making incorrect assumptions, leading to skewed decisions. Understanding and addressing this bias ensures fair and accurate treatment for all customers.

MidJourney

Midjourney is a text-to-image AI service developed by an independent research lab. It allows users to generate images based on textual descriptions, creating a wide range of art forms, from realistic to abstract styles and is especially known for its high-quality, well-structured, and detailed images.

Mixture of Experts

Is a machine learning method where specialized models, or "experts", handle different parts of data distribution. The final prediction is a blend of these expert outputs, adjusted by a "gating" system that determines each expert's relevance. This leverages individual strengths to form a more robust model.

N

Modalities

Refer to the various types of data that a model can process and interpret. These include text, images, audio, video, and other forms of sensory data. Each modality represents a unique form of information, offering distinct insights and characteristics that can be utilized in AI applications.

Mode Collapse

This phrase refers to a situation when the Generator in a Generative Adversarial Network begins to produce the same output (or a narrow set of outputs) repetitively rather than generating diverse outputs. It destabilizes the learning process and poses a challenge in GAN training.

Multi-modal AI

This type of AI has the capability to process and understand inputs from different data types, like text, speech, images, and videos. Thus, these AI models can deal with diverse data inputs, enhancing their applicability in various contexts.

Multimodal

A multi-modal learning model makes predictions by accepting and analyzing various types of input, such as audio and video data, improving its understanding of scenarios like movie scenes.

NeRF (Neural Radiance Fields)

A method for creating a three-dimensional scene from two-dimensional images using a neural network. NeRF can create a photorealistic rendering, synthesize views, and offer more capabilities in understanding and reconstructing scenes from 2D images.

O

Outpainting

A generative task where the AI is asked to extend the existing content of an image. It fills the areas beyond the image boundaries with plausible content that seamlessly connects with the original image context.

P

Parameters

Are the fundamental elements that define the behavior and output of a model. They are akin to settings or dials that can be adjusted to control various aspects of the model's performance, such as its responsiveness, creativity, and accuracy. Parameters play a crucial role in fine-tuning the model to achieve desired results, whether it's generating text, images, or other forms of content. They are essential for optimizing the model to suit specific tasks or applications.

S

Plugins / tools

AI agents such as LLMs may have the ability to use 'tools' via APIs that give them new capabilities. For example, LLMs equipped with web search capabilities can access data not present in their training dataset, which can significantly reduce the risk of hallucinations.

Prompt

A prompt is the initial input or direction given to an AI model to execute a task or answer a query. It sets the starting context for the model's generation process.

Prompt Tokens / Sampled Tokens / Completion

These terms relate to how the AI uses tokens or units of data as input or output. A prompt token starts the model's data generation process, the model chooses sampled tokens during this process, and completion signifies the model's output following the prompt.

R

RLHF (Reinforcement Learning from Human Feedback)

This technique incorporates human feedback into the learning process of an AI model. Evaluators provide feedback on the model's outputs, which can help improve the model's performance over time.

Self-Supervised Learning

This is a type of machine learning that trains an algorithm to analyze and infer from test data that hasn't been classified, labeled or categorized, as opposed to supervised learning that usually involves training data which has been labeled, classified and categorized.

Sequence Generation

This is a task in natural language processing where the model generates a sequence of words or symbols, such as in text generation. It's one of the capabilities of auto-regressive language models, including GPT and BERT.



Style Transfer

This generative model application involves capturing the artistic style of one image (the style source) and transferring it onto another image (the content source).

T

StyleGAN

Developed by NVIDIA, StyleGAN is a GAN-based model known for its high-quality and consistent outputs. Particularly, it gained attention for its capability to generate hyper-realistic images of human faces.

Super Resolution

An application of generative models that involves increasing the resolution of an image. Using these models, lower-quality images can be enhanced successfully.

Symbolic Artificial Intelligence

A type of AI that leverages symbolic reasoning mechanisms to solve problems and represent knowledge. It typically involves the usage of symbols to represent concepts and the relationships between them.

System Prompt

This refers to the predefined instructions that set the general behavior of an AI system, like a chatbot. It begins every interaction and influences how the AI responds to user inputs.

Temperature

Is a parameter that controls the randomness and creativity of a model's output. A higher temperature setting results in more varied and unpredictable responses, fostering creativity. Conversely, a lower temperature yields more deterministic and predictable outputs, enhancing coherence and reliability. This parameter is essential for fine-tuning the balance between novelty and accuracy in generated content.

Token

In the context of neural networks, tokenization is the process of encoding text into numerical values. Tokens may represent letters, groups of letters, or whole words.

V

Variational Autoencoder (VAE)

Unlike traditional autoencoders, which can learn any function to reconstruct data, a VAE places additional constraints on encoded representations, so they learn parameters of a probability distribution representing the data.

W

Weakly Supervised Learning

This is an approach to supervised learning in which the training data is noisy, limited, or, imprecise; however, these weakly labeled samples are often easier and cheaper to obtain, resulting in larger effective training sets.

Z

Zero-Shot Learning

A type of machine learning where the model can make predictions about data it has never encountered during its training. It leverages similarities between what it has seen and the novel data to make predictions.

About Clarifai

Clarifai provides a full stack AI platform for developers and teams to quickly get vision, language, and audio AI into production. Clarifai enables enterprises and public sector organizations to extract insights rapidly from their unstructured data and, in doing so, increase the value of their data and human capital, reduce AI expenses, and do so with enterprise-friendly AI governance. By significantly reducing the necessity for specialist data science knowledge, we democratize AI technology, lowering entry barriers and encouraging adoption across different sectors. Founded in 2013 by Matt Zeiler, Ph.D., Clarifai has been recognized by leading industry analysts, Forrester, Gartner, and IDC, and has been a market leader in AI since winning the top five places in image classification at the 2013 ImageNet Challenge. For more information, please visit:

www.clarifai.com.



clarifai
The World's AI™

www.clarifai.com.