


AI GLOSSARY



PRODUCED BY

Cyber  Gear

www.cyber-gear.io

A/B TESTING

A controlled, real-life experiment designed to compare two variants of a system or a model, A, and B.

ACCURACY

Accuracy is a scoring system in binary classification (i.e., determining if an answer or output is correct or not) and is calculated as $(\text{True Positives} + \text{True Negatives}) / (\text{True Positives} + \text{True Negatives} + \text{False Positives} + \text{False Negatives})$.

ACTIONABLE INTELLIGENCE

Information you can leverage to support decision making.

www.cyber-gear.io

ACTIVATION FUNCTION

In the context of Artificial Neural Networks, a function that takes in the weighted sum of all of the inputs from the previous layer and generates an output value to ignite the next layer.

ACTIVE LEARNING

(Active Learning Strategy)

A special case of Semi-Supervised Machine Learning in which a learning agent can interactively query an oracle (usually, a human annotator) to obtain labels at new data points.

ADVERSARIAL MACHINE LEARNING

Adversarial machine learning is a technique employed in the field of machine learning that attempts to make models more robust by exposing them to adversarial (and sometimes malicious) input.

ADABOOST

Adaboost is a popular ensemble learning algorithm that combines weak classifiers into a strong classifier by weighting them according to their accuracy.

AGENT

An intelligent agent is an AI system that can independently perceive its environment and act autonomously to reach an objective. The environment may be simulated or physical and doesn't need repeated prompting to achieve a bigger task.

AI ALIGNMENT

AI alignment is a field that aims to guide AI systems to align with humans' intended goals and ethical principles.

AI BIAS

AI bias happens when machine learning algorithms produce prejudiced results. Several factors can result in bias, such as:

When AI systems are either inadvertently or deliberately trained with biased data.

When the algorithm makes incorrect inferences during the learning process.

AI bias is also known as machine learning bias or simply bias.

AI DETECTOR

An AI detector is a tool designed to detect when a piece of text (or sometimes an image or video) was created by generative AI tools (e.g., ChatGPT, DALL-E). These tools aren't 100% reliable, but they can give an indication of the likelihood that a text is AI-generated.

ALPHAGO

AlphaGo is an AI-powered program designed to play the popular board game Go. The game is known for its high level of complexity, meaning that previous computer systems had struggled to play it well.

AUGMENTED INTELLIGENCE

Augmented intelligence is the use of AI to enhance human intelligence and inform decision-making instead of the AI performing tasks and processing data independently to replace humans.

AI ETHICS

A branch of the ethics of technology specific to artificially intelligent systems. Biases are prone to play a significant role in machine learning based on the data that machines are being trained with and range from gender to race to age to economic status and everything in between.

AI SAFETY

An interdisciplinary field that's concerned with the long-term impacts of AI and how it could progress suddenly to a super intelligence that could be hostile to humans.

AI FRAMEWORKS

AI frameworks make the creation of machine learning/deep learning, neural networks, and natural language processing (NLP) applications easier and faster by offering ready solutions. Some of the most popular open-source frameworks include TensorFlow, Theano, PyTorch, Sci-Kit, Keras, Microsoft Cognitive ToolKit, and Apache Mahout.

AI MODEL GOODNESS MEASUREMENT METRICS

The goodness of AI models built for specific purposes such as classification, prediction, and clustering are measured using a set of metrics called AI model goodness measurement metrics. These metrics are called AI model goodness measurement metrics. These include metrics such as accuracy, precision, recall, F-measure, word error rate, sentence error rate, mean absolute error, general language understanding evaluation (GLUE), etc.

AI OPS

Optimizing IT operations using AI. This involves detecting anomalies from IT system logs, metrics, grouping various events or alerts, diagnosing problems, and resolving issues by learning actions from prior incident, tickets, etc. AI ops is also concerned with monitoring and optimizing application performance, and proactively avoiding issues or incidents.

AUTOMATIC SPEECH RECOGNITION (ASR)

ASR is a type of natural language processing that is associated with recognizing human speech such as voice assistants.

AUTOMATION

Processing according to pre-programmed rules.

www.cyber-gear.io

ALGORITHM

An unambiguous specification of a process describing how to solve a class of problems that can perform calculations, process data and automate reasoning.

ALGORITHMIC BIAS

Unfair or discriminatory outcomes in AI systems caused by biased training data or algorithms.

ALIGNMENT

Tweaking an AI to better produce the desired outcome. This can refer to anything from moderating content to maintaining positive interactions toward humans.

ANNOTATION

A metadatum is attached to a piece of data, typically provided by a human annotator.

ANTHROPOMORPHISM

When humans tend to give nonhuman objects humanlike characteristics. In AI, this can include believing a chatbot is more humanlike and aware than it actually is, like believing it's happy, sad or even sentient altogether.

ANAPHORA

In linguistics, an anaphora is a reference to a noun by way of a pronoun. For example, in the sentence, "While John didn't like the appetizers, he enjoyed the entrée," the word "he" is an anaphora.

API

API stands for an application programming interface via which developers can access data and pre-made code solutions.

ARTIFICIAL NARROW INTELLIGENCE (ANI)

Also known as weak AI, ANI is a type of artificial intelligence that can only focus on one task or problem at a given time (e.g. playing a game against a human competitor). This is the current existing form of AI.

AREA UNDER THE CURVE (AUC)

A methodology used in Machine Learning to determine which one of several used models has the highest performance.

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.

ARTIFICIAL GENERAL INTELLIGENCE (AGI)

Also known as strong AI, AGI is a type of artificial intelligence that is considered human-like, and still in its preliminary stages (more of a hypothetical existence in present day)

ARTIFICIAL NEURAL NETWORKS

An architecture composed of successive layers of simply connected units called artificial neurons interweaved with non-linear activation functions, which is vaguely reminiscent of the neurons in an animal brain.

ASSOCIATION RULE LEARNING

A rule-based Machine Learning method for discovering interesting relations between variables in large data sets.

AUTOENCODER

A type of Artificial Neural Network used to produce efficient representations of data in an unsupervised and non-linear manner, typically to reduce dimensionality.

AUTONOMOUS

A machine is described as autonomous if it can perform its task or tasks without needing human intervention.

AUTOML (AUTOMATED MACHINE LEARNING)

Tools and techniques that automate the process of training machine learning models.

AUTOMATED SPEECH RECOGNITION

A subfield of Computational Linguistics interested in methods that enables the recognition and translation of spoken language into text by computers.

BACKPROPAGATION (BACKPROPAGATION THROUGH TIME)

A method used to train Artificial Neural Networks to compute a gradient that is needed in the calculation of the network's weights.

BATCH

The set of examples used in one gradient update of model training.

BAYES'S THEOREM

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

BAYESIAN NETWORKS

Also known as Bayes network, Bayes model, belief network, and decision network, is a graph-based model representing a set of variables and their dependencies.

BAG OF WORDS

The bag of words model is a simple representation of text data used in natural language processing, where the text is represented as a vector of word counts or frequencies.

BATCH SIZE

Number of training samples used in one iteration.

BING CHAT

Bing Chat (also called Bing AI) is a chatbot developed by Microsoft and integrated into their search engine, Bing, from February 7, 2023, onwards. The chatbot was developed in collaboration with OpenAI, based on their GPT-4 technology.

BLACK BOX AI

A black box AI is an AI system that doesn't show how it operates and processes user input. After you provide a black box AI input, it will produce an output. However, you cannot view its code and analyze its logic for producing such results.

BURSTINESS

Burstiness is a measurement of variation in sentence structure in length. AI writing tends to display low levels of burstiness, while human writing tends to have higher burstiness.

BERT

(aka Bidirectional Encoder Representation from Transformers)

Google's technology. A large scale pretrained model that is first trained on very large amounts of unannotated data. The model is then transferred to an NLP task where it is fed another smaller task-specific dataset which is used to fine-tune the final model.

BACKWARD CHAINING

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

BATCH NORMALIZATION

A technique to improve training speed and stability in neural networks.

BIG DATA

Datasets that are too large or complex to be used by traditional data processing applications.

BIAS

(Inductive Bias, Confirmation Bias)

Inductive Bias: the set of assumptions that the learner uses when predicting outputs given inputs that have not been encountered yet.

Confirmation Bias: the tendency to search for, interpret, favor, and recall information in a way that confirms one's own beliefs or hypotheses while giving disproportionately less attention to information that contradicts it.

BIAS-VARIANCE TRADEOFF

A conflict arises when data scientists try to simultaneously minimize bias and variance, which prevents supervised algorithms from generalizing beyond their training set.

BOOSTING

A Machine Learning ensemble meta-algorithm for primarily reducing bias and variance in supervised learning, and a family of Machine Learning algorithms that convert weak learners to strong ones.

BOUNDING BOX

The smallest (rectangular) box fully contains a set of points or an object.

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.

CATAPHORA

In linguistics, a cataphora is a reference placed before any instance of the noun it refers to. For example, in the sentence, "Though he enjoyed the entrée, John didn't like the appetizers," the word "he" is a cataphora.

CATEGORIZATION

Categorization is a natural language processing function that assigns a category to a document.

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.

CHATBOTS

A chat robot that can converse with a human user through text or voice commands. Utilized by e-commerce, education, health, and business industries for ease of communication and to answer user questions.

CHATGPT

The chatbot by OpenAI is the fastest consumer app to ever reach 100m users

www.cyber-gear.io

COMPUTER VISION

When a machine processes visual input from image files (JPEGs) or camera feeds.

www.cyber-gear.io

CHINESE ROOM

The Chinese room is a philosophical thought experiment proposed by John Searle. The reader is asked to imagine an AI system that behaves as if it understands Chinese. It passes the Turing test, convincing a human Chinese speaker that they are speaking to a human being.

CLIP

CLIP, or Contrastive Language-Image Pre-training, is a neural network by OpenAI that learns visual concepts from natural language supervision. Capable of learning from images and text, CLIP was trained on images with text captions from the Internet.

CONSTITUTIONAL AI

Constitutional AI trains AI systems to align with a set of values or principles as defined in a constitution. This approach was developed by AI startup Anthropic.

COPILOT

Copilot is Microsoft 365's AI assistant feature that builds on OpenAI's GPT-4 large language models (LLMs).

CUTOFF DATE

The cutoff date is the date at which the model's information ends. AI models can't recall information past the cutoff date. For example, GPT-3.5's cutoff date is September 2021.

COMPUTATIONAL SEMANTICS

(Semantic Technology)

Computational semantics is the study of how to automate the construction and reasoning of meaning representations of natural language expressions.

CORPUS

A large dataset of written or spoken material that can be used to train a machine to perform linguistic tasks.

CHATBOT

A computer program or an AI is designed to interact with human users through conversation.

CLASSIFICATION

The task of approximating a mapping function from input variables to discrete output variables, or, by extension, a class of Machine Learning algorithms that determine the classes to which specific instances belong.

CLUSTERING

In Machine Learning, the unsupervised task of grouping a set of objects so that objects within the same group (called a cluster) are more "similar" to each other than they are to those in other groups.

CLOUD COMPUTING

The delivery of computing services, including storage, processing, and networking, over the internet.

www.cyber-gear.io

COLD-START

A potential issue arises from the fact that a system cannot infer anything for users or items for which it has not gathered a sufficient amount of information yet.

COMPOSITE AI

Composite AI refers to a combination of AI techniques for accomplishing the best results.

COLLABORATIVE FILTERING

A method used in the context of recommender systems to make predictions about the interests of a user by collecting preferences from a larger group of users.

COMPUTER VISION

The field of Machine Learning studies how to gain a high-level understanding of images or videos.

www.cyber-gear.io

COGNITIVE SCIENCE

The broader form of study includes AI in addition to philosophy, linguistics, psychology, neuroscience, and anthropology. All of these combine together to learn how the mind functions and when applied to AI, how machines can simulate human thought and action.

CONFIDENCE INTERVAL

A type of interval estimate that is likely to contain the true value of an unknown population parameter. The interval is associated with a confidence level that quantifies the level of confidence of this parameter being in the interval.

CONTRIBUTOR

A human worker providing annotations on the Appen data annotation platform.

www.cyber-gear.io

CONVOLUTIONAL NEURAL NETWORK (CNN)

A class of Deep, Feed-Forward Artificial Neural Networks, often used in Computer Vision.

CENTRAL PROCESSING UNIT (CPU)

The electronic circuitry within a computer carries out the instructions of a computer program by performing the basic arithmetic, logical, control, and input/output operations specified by the instructions.

CROSS-VALIDATION

(k-fold Cross-Validation, Leave-p-out Cross-Validation)

A collection of processes designed to evaluate how the results of a predictive model will generalize to new data sets.

- k-fold Cross-Validation
- Leave-p-out Cross-Validation

DALL-E 2

DALL-E 2 is an impressive text-to-image algorithm that has gained massive attention in the recent past.

www.cyber-gear.io

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.

DEEPPFAKE

A portmanteau of “deep learning” and “fake,” a deepfake is an AI-generated image, audio, or video depicting fake events. It uses powerful machine learning and AI to manipulate or create deceiving content.

DISCRIMINATOR

A discriminator in a generative adversarial network (GAN) is a classifier that tries to identify real data from fake data created by the generator. The goal is for the discriminator to improve its ability to distinguish real from fake, as the generator tries to improve at creating fakes.

DATA ARCHITECT

A data architect is a practitioner of data architecture, a data management discipline concerned with designing, creating, deploying, and managing an organization's data architecture. Data architects often work with data scientists on AI projects.

DATA LAKE

Since data is at the core of every AI use case or solution, aggregating all the data needed to build machine learning and inference models are absolutely critical. The process of assimilating all the data (structured and unstructured) in a consolidated repository is referred to as a data lake. sentiment analysis and natural language processing.

DATA LABELING

The process of annotating data to provide ground truth for supervised learning.

DATA MANAGER

A data manager is an individual concerned with legally acquiring the right kind of data for training AI systems by working with data scientists. A data manager works with data architects to ensure that acquired data is properly versioned and stored for analysis and audit purposes. A data manager is also concerned with the governance of the data per legal and organizational requirements and ensuring that the lifecycle of the data is managed accordingly.

DATA SCIENTIST

A data scientist is an individual, organization, or application that performs statistical analysis, data mining, and retrieval processes on a large amount of data to identify trends, figures, and other relevant information.

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.

DATA

(Structured Data, Unstructured Data, Data augmentation)

The most essential ingredient to all Machine Learning and Artificial Intelligence projects.

www.cyber-gear.io

DATA PREPROCESSING

Cleaning and transforming raw data to prepare it for analysis or modeling.

DIFFUSION

A method of machine learning that takes an existing piece of data, like a photo, and adds random noise. Diffusion models train their networks to re-engineer or recover that photo.

DROPOUT

A regularization technique in neural networks that prevents overfitting by randomly dropping neurons during training.

EDGE COMPUTING

A decentralized computing paradigm where data processing occurs closer to the data source or endpoint devices.

ELIZA

ELIZA was an early chatbot developed in the 1960s by MIT computer scientist Joseph Weizenbaum. ELIZA played the character of a psychotherapist, with the user acting as its patient.

EMBODIED AGENTS

Embodied agents, also referred to as embodied AI, are AI agents with a physical body that perform specific tasks in the physical environment.

EMERGENCE

Emergence describes capabilities that arise in AI systems unpredictably as they become more complex. A system's emergent properties are not observable in its individual parts.

EU AI ACT

The EU AI Act is a regulatory framework for responsible AI deployment in a way that doesn't conflict with data privacy rights.

EXPLAINABLE AI

Explainable AI is AI that can explain its decisions and reasoning.

www.cyber-gear.io

EMERGENT BEHAVIOR

When an AI model exhibits unintended abilities.

www.cyber-gear.io

END-TO-END LEARNING, OR E2E

A deep learning process in which a model is instructed to perform a task from start to finish. It's not trained to accomplish a task sequentially but instead learns from the inputs and solves it all at once.

ETHICAL CONSIDERATIONS

An awareness of the ethical implications of AI and issues related to privacy, data usage, fairness, misuse and other safety issues.

ETHICS IN AI

The study of moral and ethical issues related to the design and use of artificial intelligence.

EXPERT SYSTEM

A computer program designed to mimic the decision-making abilities of a human expert in a specific domain.

UNSTRUCTURED DATA

Raw, unprocessed data. Textual data is a perfect example of unstructured data because it is not formatted into specific features.

STRUCTURED DATA

Data processed in a way that it becomes ingestible by a Machine Learning algorithm and, if in the case of Supervised Machine Learning, labeled data; data after it has been processed on the Appen data annotation platform.

DATA AUGMENTATION

The process of adding new information derived from both internal and external sources to a data set, typically through annotation.

DECISION TREE

A category of Supervised Machine Learning algorithms where the data is iteratively split with respect to a given parameter or criteria.

DEEP BLUE

A chess-playing computer developed by IBM, better known for being the first computer chess-playing system to win both a chess game and a chess match against a reigning world champion under regular time controls.

DEEP LEARNING

(Deep Reinforcement Learning)

A broader family of Machine Learning methods based on learning data representations, as opposed to task-specific algorithms. Deep Learning can be supervised, semi-supervised or unsupervised.

DIMENSIONALITY

(Dimensionality Reduction, Curse of Dimensionality)

Dimensionality Reduction: the process of reducing the number of random variables under consideration by obtaining a set of principal variables. Also, see Feature Selection.

Curse of Dimensionality: a phenomenon that arises when analyzing and organizing data in high-dimensional spaces due to the fact that the more the number of dimensions increases, the sparser the amount of available data becomes.

DIGITAL ECOSYSTEM

Several software platforms or cloud services that work in tandem across a network.

www.cyber-gear.io

EMBEDDING

(Word Embedding)

One instance of some mathematical structure contained within another instance, such as a group that is a subgroup.

ENSEMBLE METHODS

In Statistics and Machine Learning, ensemble methods use multiple learning algorithms to obtain better predictive performance that could be obtained from any of the constituent learning algorithms alone. Unlike a statistical ensemble in statistical mechanics, which is usually infinite, a machine learning ensemble consists of only a concrete finite set of alternative models but typically allows for a much more flexible structure to exist among those alternatives.

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.

ENTROPY

The average amount of information conveyed by a stochastic source of data.

EPOCH

In the context of training Deep Learning models, one pass of the full training data set.

ETHICAL AI

Ethical AI refers to AI that refers to ethical guidelines related to fundamental values, like:

Individual rights
Privacy
Non-discrimination
Non-manipulation

EXPERT LEARNING

Expert Learning is the process of using AI to learn from experts in a given field.

FACE RECOGNITION

Face Recognition is the process of identifying a person from an image or video.

www.cyber-gear.io

FRÉCHET INCEPTION DISTANCE (FID)

FID is a metric for evaluating the quality of images created by generative AI.

FEATURE ENGINEERING

The process of selecting and transforming relevant data features to improve model performance.

FEATURE

(Feature Selection, Feature Learning)

A variable is used as an input to a model.

FEATURE LEARNING

An ensemble of techniques meant to automatically discover the representations needed for feature detection or classification from raw data.

FALSE POSITIVE

An error due to the fact a result did reject the null hypothesis when it shouldn't have.

FALSE NEGATIVE

An error is due to the fact a result did not reject the null hypothesis when it should have.

FEED-FORWARD (NEURAL) NETWORKS

An Artificial Neural Network wherein connections between the neurons do not go backward or form a cycle.

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.

FOOM

Also known as fast takeoff or hard takeoff. The concept that if someone builds an AGI that it might already be too late to save humanity.

F-SCORE

A measure of a model's accuracy considering both the precision and the recall to compute the score. More specifically, the F-Score is the harmonic average of the precision and recall, where it reaches its maximal value at 1 (perfect precision and recall) and minimum at 0.

GARBAGE IN, GARBAGE OUT

A principle stating that whenever the input data is flawed, it will lead to misleading results and produces nonsensical output, a.k.a. "garbage".

GENERATIVE AI

Using AI to create text, images, voice, video and computer code.

www.cyber-gear.io

GENERAL DATA PROTECTION REGULATION (GDPR)

A regulation in EU law on data protection and privacy for all individuals within the European Union aims to give control to citizens and residents over their personal data.

GENERATIVE PRE-TRAINED TRANSFORMER (GPT)

Generative pre-trained transformer or GPT is a type of large language model (LLM) that's trained to generate content. First launched by OpenAI in 2018, GPT now has GPT-4 as its latest model as of this writing.

GRAPHICS PROCESSING UNIT (GPU)

A GPU is a type of processor that is suited to powering AI hardware because it can perform more simultaneous computations than a CPU.

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.

GENETIC ALGORITHM

A search heuristic inspired by the Theory of Evolution reflects the process of natural selection where the fittest individuals are selected to produce offspring of the following generation.

GENERATIVE ADVERSARIAL NETWORKS (GANS)

A class of Artificial Intelligence algorithms used in Unsupervised Machine Learning implemented as the combination of two Neural Networks competing with each other in a zero-sum game framework.

GOOGLE BARD

An AI chatbot by Google that functions similarly to ChatGPT but pulls information from the current web, whereas ChatGPT is limited to data until 2021 and isn't connected to the internet.

GUARDRAILS

Policies and restrictions placed on AI models to ensure data is handled responsibly and that the model doesn't create disturbing content.

GRAPHIC PROCESSING UNIT [GPU]

A specialized electronic circuit designed to rapidly manipulate and alter memory to accelerate the rendering of images thanks to its parallel processing architecture, which allows it to perform multiple calculations simultaneously.

GRADIENT DESCENT

An optimization algorithm used to train machine learning models by minimizing a loss function.

GROUND TRUTH

A piece of information obtained through direct observation as opposed to inference.

HALLUCINATION

An incorrect response from AI. Can include generative AI producing answers that are incorrect but stated with confidence as if correct. The reasons for this aren't entirely known.

HEURISTIC

A computer science technique designed for quick, optimal, solution-based problem solving.

www.cyber-gear.io

HUMAN-IN-THE-LOOP

Human-in-the-loop (HITL) is a branch of artificial intelligence that leverages both human and machine intelligence to create machine learning models. In a traditional human-in-the-loop approach, people are involved in a virtuous circle where they train, tune, and test a particular algorithm.

HUMAN FEEDBACK

Human Feedback is being used to improve the results of the machine learning models.

HYPERPARAMETER

(Hyperparameter Tuning)

A configuration, external to the model and whose value cannot be estimated from data, that data scientists continuously tweak during the process of training a model.

- The process of manually determining the optimal configuration to train a specific model.

HYBRID AI

Hybrid AI is a combination of human insight and AI, such as machine learning and deep learning.

www.cyber-gear.io

HYPERPARAMETERS

Parameters that aren't learned from the data but set prior to training.

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".

IMAGENET

A large visual dataset made of 14 million URLs of hand-annotated images organized in twenty-thousand (20,000) different categories, designed for use in visual

IMAGE RECOGNITION

The problem in Computer Vision of determining whether an image contains some specific object, feature, or activity.

INFERENCE

The process of making predictions by applying a trained model to new, unlabeled instances.

INFORMATION RETRIEVAL

The area of Computer Science studies 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.

JASPER

Jasper is an AI platform that enables businesses to build, deploy, and manage AI applications.

www.cyber-gear.io

JUPYTER NOTEBOOK

An open-source web application that allows the creation and sharing of documents containing live code, equations, visualizations, and narrative text.

www.cyber-gear.io

K-MEANS CLUSTERING

A popular clustering algorithm that partitions data into clusters based on similarity.

K-NEAREST NEIGHBORS

K-Nearest Neighbors is a supervised machine learning algorithm for classification and regression.

KERAS

A high-level neural networks API running on top of TensorFlow, Theano, or CNTK.

KNOWLEDGE ENGINEERING

Knowledge engineering is the field of AI that aims to emulate a human expert's knowledge in a certain field.

LABELED DATA

Data that has been labeled with a specific category or class.

LAYER

(Hidden Layer)

A series of neurons in an Artificial Neural Network process a set of input features, or, by extension, the output of those neurons. Hidden Layer: a layer of neurons whose outputs are connected to the inputs of other neurons, therefore not directly visible as a network output.

LARGE LANGUAGE MODEL

AI models like GPT that can understand and generate text.

LARGE LANGUAGE MODEL META AI (LLAMA)

LLaMA is a family of open-source large language models (LLMs) launched by Meta AI in 2023.

LABEL

A part of training data that identifies the desired output for that particular piece of data.

LINEAR ALGEBRA

Linear algebra is the key branch in mathematics when it comes to artificial intelligence and machine learning algorithms.

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.

LIMITED MEMORY

Systems with short-term memory limited to a given timeframe.

LEARNING-TO-LEARN

A new direction within the field of Machine Learning investigating how algorithms can change the way they generalize by analyzing their own learning process and improving on it.

LEARNING-TO-RANK

The application of Machine Learning to the construction of ranking models for Information Retrieval systems.

LEARNING RATE

A scalar value is used by the gradient descent algorithm at each iteration of the training phase of an Artificial Neural Network to multiply with the gradient.

LOGIT FUNCTION

The inverse of the sigmoidal "logistic" function is used in mathematics, especially in statistics.

LONG SHORT-TERM MEMORY NETWORKS

A variation of the Recurrent Neural Network was proposed as a solution to the vanishing gradient problem.

LOSS FUNCTION

A mathematical function that measures the difference between predicted and actual values, used in training models.

MACHINE LEARNING

The subfield of Artificial Intelligence often uses statistical techniques to give computers the ability to "learn", i.e., progressively improve performance on a specific task, with data, without being explicitly programmed.

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 a text by an algorithm, independent of any human involvement.

MIDJOURNEY

MidJourney is a new AI art generator that turns text into an image. And not just any image, but a realistic, creative, or abstract masterpiece unlike we've never seen before!

www.cyber-gear.io

MODEL

A broad term referring to the product of AI training, created by running a machine learning algorithm on training data.

MOATS

Moats are mechanisms that prevent competitors from copying a proprietary LLM. An LLM's moats are training data, model weights and the cost of training.

MODEL COLLAPSE

Model collapse is when low-quality, AI-generated content contaminates the training set for future models.

MACHINE LEARNING LIFECYCLE MANAGEMENT

DevOps for Machine Learning systems.

www.cyber-gear.io

MACHINE TRANSLATION

A subfield of computational linguistics that studies the use of software to translate text or speech from one language to another.

MICROSOFT BING

A search engine by Microsoft that can now use the technology powering ChatGPT to give AI-powered search results. It's similar to Google Bard in being connected to the internet.

ML OPS

ML ops or machine learning operations is the process of taking an experimental machine learning model into a production web system. Machine learning models are tested and developed in isolated experimental systems. When an algorithm is ready to be launched, ML ops is practiced between data scientists, DevOps, and machine learning engineers to transition the algorithm to production systems.

MODEL

A model is an abstract representation of what a Machine Learning system has learned from the training data during the training process.

MONTE CARLO

An approximate methodology that uses repeated random sampling to generate synthetic simulated data.

MULTI-MODAL LEARNING

A subfield of Machine Learning aiming to interpret multimodal signals together and build models that can process and relate information from multiple types of data.

MULTIMODAL AI

A type of AI that can process multiple types of inputs, including text, images, videos and speech.

MULTI-TASK LEARNING

A subfield of Machine Learning that exploits similarities and differences across tasks to solve multiple tasks at the same time.

NAIVE BAYES

A family of simple probabilistic classifiers based on applying Bayes' theorem with strong independence assumptions between the features.

NAMED ENTITY RECOGNITION

A subtask 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)

The area of Artificial Intelligence studies the interactions between computers and human languages, in particular how to process and analyze large amounts of natural language data.

NEURAL NETWORK

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 STYLE TRANSFER

A technique that combines the content of one image with the style of another.

NEURAL STYLE TRANSFER

A technique that combines the content of one image with the style of another.

NATURAL LANGUAGE GENERATION (NLG)

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.

NATURAL LANGUAGE UNDERSTANDING (NLU)

As a subset of natural language processing, natural language understanding deals with helping machines recognize the intended meaning of language – taking into account its subtle nuances and any grammatical errors.

NEURON

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

OBJECT DETECTION

Object Detection is the process of identifying and locating objects in an image or video.

OPENAI

The start-up behind ChatGPT and Dall-E is already valued at \$29b.

www.cyber-gear.io

OPTICAL CHARACTER RECOGNITION

The conversion of images of printed, handwritten, or typed text into a machine-friendly textual format.

Abandonment: When the collection gets enough demand, the owner removes the website, the Discord group, and the media content. Once holders find out, they'll try to sell the NFTs as quickly as possible. The creator profits from the sale rush.

OPTIMIZATION

The selection of the best element (concerning some criterion) from some set of available alternatives.

OVERFITTING

The fact that a model unknowingly identified patterns in the noise and assumed those represented the underlying structure; is the production of a model that corresponds too closely to a particular set of data, and therefore fails to generalize well to unseen observations.

PATTERN RECOGNITION

An area of Machine Learning focusing on the (supervised or unsupervised) recognition of patterns in the data.

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.

PARAMETERS

Numerical values that give LLMs structure and behavior, enabling it to make predictions.

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.

PROMPT CHAINING

An ability of AI to use information from previous interactions to color future responses.

PROMPT ENGINEER

A person that specializes in interacting with Generative AI models in order to generate the best output.

www.cyber-gear.io

PYTHON

A popular programming language used for general programming.

www.cyber-gear.io

PYTORCH

An open-source machine learning framework developed by
Facebook.

www.cyber-gear.io

PATHWAYS LANGUAGE MODE (PALM)

PaLM is Google's transformer-based large language model.
It can perform various tasks, including:

- Commonsense and arithmetic reasoning
- Explaining jokes
- Generating codes

PATHWAYS LANGUAGE MODEL (PALM)

PaLM is Google's transformer-based LLM, based on similar technology to GPT-3 and GPT-4. The Google Bard chatbot runs on PaLM.

PERPLEXITY

Perplexity is a measurement of how unpredictable (perplexing) a text is. A text with high perplexity is more likely to read unnaturally or be nonsensical than a text with low perplexity. AI writing tools tend to produce text with relatively low perplexity, as this gives them a higher chance of making sense.

PREDICTIVE AI

Predictive AI is a method of analyzing data using statistical algorithms to predict upcoming outcomes.

PROMPT

A prompt is a human user's input to an AI system, which will generate an output or result.

PROMPT ENGINEERING

Prompt engineering is the process of crafting and refining prompts for a generative AI model. AI users utilize prompt engineering to improve the output from the AI model.

POOLING

(Max Pooling)

The process of reducing a matrix generated by a convolutional layer to a smaller matrix.

PERSONALLY IDENTIFIABLE INFORMATION

Any piece of information that can be used on its own or in combination with some other information to identify a particular individual.

PRECISION

The number of correct positive results is divided by the number of all positive results returned by a classifier.

PREDICTION

The inferred output of a trained model is provided with an input instance.

PREPROCESSING

The process of transforming raw data into a more understandable format.

PRE-TRAINED MODEL

A model, or the component of a model, that has been preliminary trained, generally using another data set. See also: Transfer Learning.

PRESCRIPTIVE ANALYTICS

Prescriptive analytics is a type of data analytics, the use of technology to help businesses make better decisions through the analysis of raw data. Specifically, prescriptive analytics factors information about possible situations or scenarios, available resources, past performance, and current performance, and suggests a course of action or strategy. It can be used to make decisions on any time horizon, from immediate to long term.

PRINCIPAL COMPONENT ANALYSIS

A process that uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of linearly uncorrelated variables called principal components.

PRIOR

The probability distribution would represent the preexisting beliefs about a specific quantity before new evidence is considered.

QUANTUM COMPUTING

Quantum computing has the potential to dramatically accelerate the evolution of AI.

www.cyber-gear.io

Q-LEARNING

Q-learning is a type of reinforcement learning that enables AI models to learn and improve iteratively over time.

QUILLBOT

QuillBot is a company known for its powerful paraphrasing tool. It also offers a variety of other AI writing tools (e.g., grammar checker, summarizer) and a plagiarism checker.

www.cyber-gear.io

RECOMMENDATION ENGINE

A recommendation engine is an AI algorithm that is used to serve users content based on their preferences. Social sites, such as TikTok, and streaming platforms, such as Spotify and YouTube, use recommendation engines to personalize user feeds.

REINFORCEMENT LEARNING FROM HUMAN FEEDBACK (RLHF)

Reinforcement learning from human feedback (RLHF) is a tactic that trains AI models with direct human feedback. Instead of getting a reward or punishment, the AI models get feedback from humans, typically in the form of ranking the models' behaviors.

ROBOT

A robot is a machine that is capable of carrying out actions automatically (sometimes autonomously). Robots usually contain computer systems that are programmed to allow them to carry out their tasks. The study and design of robots is called robotics.

RANDOM FOREST

An ensemble learning method that operates by constructing a multitude of decision trees at training time and outputting a combined version (such as the mean or the mode) of the results of each tree.

REACTIVE MACHINES

Reactive machines can analyze, perceive, and make predictions about experiences, but do not store data; they react to situations and act based on the given moment.

RECALL

The fraction of all relevant samples that are correctly classified as positive.

RECTIFIED LINEAR UNIT

A unit employing the rectifier function as an activation function.

RECURRENT NEURAL NETWORKS

A class of Artificial Neural Networks where connections between neurons form a directed graph along a sequence, allowing it to exhibit dynamic temporal behavior for a time sequence and to use its internal state (memory) to process sequential signals.

REGRESSION

(Linear Regression, Logistic Regression)

A set of statistical processes for estimating the relationships among variables.

Linear Regression: a simple type of regression taking a linear combination of features as an input, and outputting a continuous value.

Logistic Regression: a type of regression generating a probability for each possible discrete label value in a classification problem by applying a sigmoid function to a linear prediction.

REGRESSOR

A feature, is an explanatory variable used as an input to a model.

REGULARIZATION

The process of introducing additional information to prevent overfitting.

REINFORCEMENT LEARNING

The subfield of Machine Learning is inspired by human behaviour studying how an agent should take action in a given environment to maximize some notion of cumulative reward.

REINFORCEMENT LEARNING AGENT

An entity that interacts with an environment and learns to make decisions to maximize a reward.

REPRODUCIBILITY (CRISIS OF)

A methodological crisis in science in which scholars have found that the results of many scientific studies are difficult or impossible to replicate or reproduce on subsequent investigation, either by independent researchers or by the original researchers themselves.

RESTRICTED BOLTZMANN MACHINES

A restricted Boltzmann machine (RBM) is a generative stochastic artificial neural network that can learn a probability distribution over its set of inputs.

RECOMMENDATION SYSTEM

AI systems that provide personalized recommendations, often used in e-commerce and content platforms.

ROBOTICS

Focused on the design and manufacturing of robots that exhibit and/or replicate human intelligence and actions.

www.cyber-gear.io

ROBOTIC PROCESS AUTOMATION [RPA]

Uses software with artificial intelligence and machine learning capabilities to perform repetitive tasks once completed by humans.

SCIKIT-LEARN

An open-source machine learning library for Python.

www.cyber-gear.io

SEMI-SUPERVISED LEARNING

A class of supervised learning techniques that also leverages available unlabeled data for training, typically using a small number of labeled instances in combination with a larger amount of unlabeled rows. See also Supervised Learning and Unsupervised Learning.

SEMANTIC ANALYSIS

A more sophisticated form of Natural Language Processing, this concept is focused on the process of stringing words together as well as the way that language is understood through cultural context. This could function as a help to create eBooks and blog posts, with the potential to replace human writers or content marketers (!). Wait, what?

SENTIMENT ANALYSIS

The use of natural language processing, text analysis, computational linguistics, and biometrics to systematically identify, extract, quantify, and study affected states and subjective information.

SEMANTIC ANNOTATION

Tagging different search queries or products with the goal of improving the relevance of a search engine.

SINGULARITY

In the field of AI, singularity refers to the event where the AI becomes self-aware and starts to evolve on its own out of control.

SPEECH RECOGNITION

The technology that enables machines to transcribe and understand spoken language.

www.cyber-gear.io

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.

STATISTICAL DISTRIBUTION

In statistics, an empirical distribution function is the distribution function associated with the empirical measure of a sample. This cumulative distribution function is a step function that jumps up by $1/n$ at each of the n data points. Its value at any specified value of the measured variable is the fraction of observations of the measured variable that are less than or equal to the specified value.

STOCHASTIC PARROT

An analogy of LLMs that illustrates that the software doesn't have a larger understanding of meaning behind language or the world around it, regardless of how convincing the output sounds. The phrase refers to how a parrot can mimic human words without understanding the meaning behind them.

STYLE TRANSFER

The ability to adapt the style of one image to the content of another, allowing an AI to interpret the visual attributes of one image and use it on another. For example, taking the self-portrait of Rembrandt and re-creating it in the style of Picasso.

SUPERVISED LEARNING

The Machine Learning task of learning a function mapping an input to an output based on example input-output pairs.

SUPPORT VECTOR MACHINES (SVM)

A class of discriminative classifiers formally defined by a separating hyperplane, where for each provided labeled training data point, the algorithm outputs an optimal hyperplane that categorizes new examples.

SYNTHETIC DATA

Data is generated artificially when real data cannot be collected in sufficient amounts, or when original data doesn't meet certain requirements.

TENSORFLOW

An open-source library, popular among the Machine Learning community, for data flow programming across a range of tasks. It is a symbolic math library and is also used for machine learning applications such as neural networks.

www.cyber-gear.io

TEST DATA

The unlabeled data is used to check that a machine learning model is able to perform its assigned task.

TEMPERATURE

Parameters set to control how random a language model's output is. A higher temperature means the model takes more risks.

TEXT-TO-IMAGE GENERATION

Creating images based on textual descriptions.

www.cyber-gear.io

TECHNOLOGICAL SINGULARITY

The singularity describes a point in the future where advanced AI becomes more intelligent than humans and technological growth becomes uncontrollable.

TEST SET

An error due to the fact a result did reject the null hypothesis when it shouldn't have. Subset of data used to evaluate the final model's performance.

TOKEN

A token is the basic unit of text that an LLM uses to understand and generate language. It may be a word or parts of a word. Paid LLMs, such as GPT-4's API, charge users by token.

TRAINING SET

Subset of data used to train a model.

VALIDATION SET

Subset of data used to tune hyperparameters and model choices.

VARIATIONAL AUTOENCODER

Variational autoencoders are a generative AI model architecture commonly used for signal analysis and finding efficient coding of input data. They are comparable to GANs in that they pit two neural networks against each other -- an encoder and a decoder.

TIME SERIES

(Time Series Data)

A sequence of data points is recorded at specific times and indexed accordingly to their order of occurrence.

TESTING

(Testing Data)

In the context of Supervised Machine Learning, the process of assessing the final performance of a model using hold-out data.

TOPIC MODELING

A category of Unsupervised Machine Learning algorithms that uses clustering to find hidden structures in textual data, and interpret them as topics.

TRAINING DATA

In the context of Supervised Machine Learning, the construction of algorithms that can learn from and make predictions from data.

TRANSFORMER

A type of neural network architecture used for natural language processing. It's the T in GPT.

TRANSFER LEARNING

An area of Machine Learning that focuses on using knowledge gained to solve a specific problem and apply this knowledge to a different but related problem.

TRANSFORMER MODEL

A neural network architecture and deep learning model that learns context by tracking relationships in data, like in sentences or parts of images. So, instead of analyzing a sentence one word at a time, it can look at the whole sentence and understand the context.

TURING TEST

A test developed by Alan Turing to evaluate a machine's ability to exhibit intelligent behavior equivalent to that of a human. The test consists in having the machine chat with a human. If a human evaluator witnessing the conversation from outside the room where the test takes place can't reliably tell the machine from the human apart, the machine is said to have passed the Turing test.

UNCERTAINTY

A range of values is likely to enclose the true value.

UNDERFITTING

The fact that a Machine Learning algorithm fails to capture the underlying structure of the data properly, typically because the model is either not sophisticated enough, or not appropriate for the task at hand; the opposite of Overfitting.

UNSUPERVISED LEARNING

The area of Machine Learning consists in inferring a function that describes the structure of unlabeled data.

VALIDATION

The process of using hold-out data to evaluate the performance of a trained model; by opposition to the testing phase which is used for the final assessment of the model's performance, the validation phase is used to determine if any iterative modification needs to be made to the model.

VALIDATION DATA

Structured like training data with 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.

VARIATION

Also called queries or utterances, these work in tandem with intents for natural language processing. 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."

VANISHING/EXPLODING GRADIENTS

A dreaded difficulty and a major obstacle to recurrent net performance that data scientists face when training Artificial Neural Networks with gradient-based learning methods and back propagation, due to the neural network's weights receiving an update proportional to the partial derivative of the error function concerning the current weight in each iteration of training.

VARIANCE

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

VOICE RECOGNITION

The ability of a computer to recognize and respond to spoken commands.

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.

WEAK SUPERVISION

Weak Supervision is a form of AI training that uses noisy data.

WEAK AI, AKA NARROW AI

AI that's focused on a particular task and can't learn beyond its skill set. Most of today's AI is weak AI.

WHISPER

In AI speak, Whisper is a multilingual automatic speech recognition tool developed and launched by OpenAI in 2022. It can recognize speech in different languages, identify a language, and translate speech into another language.


YIELD OPTIMIZATION

Maximizing the output of a process while minimizing the resources used.

www.cyber-gear.io

ZERO SHOT LEARNING

Zero Shot Learning is a type of AI that can learn without any labeled data.

Cyber  Gear

www.cyber-gear.io

Last updated on 30 October 2023