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 (True Positives + True Negatives) / (True Positives + True Negatives + False Positives + 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.

ALGORITHM

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

ANNOTATION

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

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.

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 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.

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.

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.

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.

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.

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.

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.

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.

OMPUTER VISION

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

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.

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

DATA MINING

The process of analyzing datasets in order to discover new patterns that might improve the model.

DATA SCIENCE

Drawing from statistics, computer science, and information science, this interdisciplinary field aims to use a variety of scientific methods, processes, and systems to solve problems involving data.

DATASET

A collection of related data points, usually with a uniform order and tags.

DEEP LEARNING

A function of artificial intelligence that imitates the human brain by learning from the way data is structured, rather than from an algorithm that's programmed to do one specific thing.

DATA

(Structured Data, Unstructured Data, Data augmentation)

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

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.

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.

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.

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".

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.

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.

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.

GROUND TRUTH

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

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.

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.

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.

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.

LABEL

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

LINGUISTIC ANNOTATION

Tagging a dataset of sentences with the subject of each sentence, ready for some form of analysis or assessment. Common uses for linguistically annotated data include sentiment analysis and natural language processing.

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.

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.

MODEL

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

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.

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.

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.

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.

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.

PREDICTIVE ANALYTICS

By combining data mining and machine learning, this type of analytics is built to forecast what will happen within a given timeframe based on historical data and trends.

PYTHON

A popular programming language used for general programming.

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.

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.

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.

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.

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.

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.

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.

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.

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.

TEST DATA

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

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.

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.

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.

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.

www.cyber-gear.io