

APPENDIX IV. AI/ML GLOSSARY*

➤ **A/B Testing**

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

➤ **Accuracy**

Percentage of correct predictions made by the model.

➤ **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 is able to 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. In ML, a method, function, or series of instructions used to generate a machine learning model. Examples include linear regression, decision trees, support vector machines, and neural networks.

➤ **Annotation**

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

➤ **Area Under the Curve (AUC)**

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

➤ **Artificial Intelligence**

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

➤ **Artificial Neural Network**

An architecture composed of successive layers of simple 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.

➤ **Attribute**

A quality describing an observation (e.g., color, size, weight). In Excel terms, these are column headers.

➤ **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.

➤ **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' 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 (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 metric**

What is the average difference between your predictions and the correct value for that observation?

- Low bias could mean every prediction is correct. It could also mean half of your predictions are above their actual values and half are below, in equal proportion, resulting in low average difference.
- High bias (with low variance) suggests your model may be underfitting and you're using the wrong architecture for the job.

➤ **Bias term**

Allows models to represent patterns that do not pass through the origin. For example, if all my features were 0, would my output also be zero? Is it possible there is some base value upon which my features have an effect? Bias terms typically accompany weights and are attached to neurons or filters.

➤ **Bias-Variance Tradeoff**

A conflict arising when data scientists try to simultaneously minimize bias and variance, that 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 containing a set of points or an object.

➤ **Categorical Variables**

Variables with a discrete set of possible values. Can be ordinal (order matters) or nominal (order doesn't matter).

➤ **Chatbot**

A computer program or an AI 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. Classification

➤ **Predicting a categorical output.**

- Binary classification predicts one of two possible outcomes (e.g., is the email spam or not spam?)
- Multi-class classification predicts one of multiple possible outcomes (e.g., is this a photo of a cat, dog, horse or human?)

➤ **Classification Threshold**

The lowest probability value at which we're comfortable asserting a positive classification. For example, if the predicted probability of being diabetic is > 50%, return True, otherwise return False.

➤ **Clustering**

In Machine Learning, the unsupervised grouping of data into buckets 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 arising 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.

➤ **Computer Vision**

The field of Machine Learning that studies how to gain high-level understanding from 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.

➤ **Confusion Matrix**

Table that describes the performance of a classification model by grouping predictions into 4 categories.

- True Positives: we correctly predicted they do have diabetes
- True Negatives: we correctly predicted they don't have diabetes
- False Positives: we incorrectly predicted they do have diabetes (Type I error)
- False Negatives: we incorrectly predicted they don't have diabetes (Type II error)

➤ **Continuous Variables**

- Variables with a range of possible values defined by a number scale (e.g., sales, lifespan).

➤ **Contributor**

A human worker providing annotations on a data annotation platform.

➤ **Convergence**

A state reached during the training of a model when the loss changes very little between each iteration.

➤ **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 that 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.

➤ **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 in respect to a given parameter or criteria.

➤ **Deduction**

A top-down approach to answering questions or solving problems. A logic technique that starts with a theory and tests that theory with observations to derive a conclusion, e.g., We suspect X, but we need to test our hypothesis before coming to any conclusions.

➤ **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)**

Deep Learning is derived from one machine learning algorithm called perceptron or multi-layer perceptron that gain more and more attention nowadays because of its success in different fields like, computer vision to signal processing and medical diagnosis to self-driving cars. As all other AI algorithms deep learning is from decades, but now today we have more and more data and cheap computing power that make this algorithm really powerful to achieve state of the art accuracy. In modern world this algorithm knowns as artificial neural network. deep learning is much more than traditional artificial neural network. But it was highly influenced by machine learning's neural network and perceptron network. It is 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.

➤ **Dimension**

Dimension for machine learning and data scientists differs from physics, here Dimension of data means how much feature is in a data ocean (data-set). e.g., in case of object detection application, flatten image size and color channel (e.g., 28*28*3) is a feature of the input set. In case of house price prediction (maybe) house size is the data-set so we call it one-dimensional data.

➤ **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: phenomena that arise 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.

➤ **Entropy**

The average amount of information conveyed by a stochastic (having a random probability distribution or pattern that may be analyzed statistically but may not be predicted precisely) source of data.

➤ **Epoch**

In the context of training Deep Learning models, one pass of the full training data set, i.e., the number of times the algorithm sees the entire data set.

➤ **Extrapolation**

Making predictions outside the range of a dataset, e.g., My dog barks, so all dogs must bark. In machine learning trouble often arises from extrapolation outside the range of training data.

➤ **False Positive**

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

➤ **False Positive Rate**

Defined as:

$FPR = 1 - \text{Specificity} = \frac{\text{False Positives}}{\text{False Positives} + \text{True Negatives}}$
The False Positive Rate forms the x-axis of the ROC curve.

➤ **False Negative**

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

➤ **Feature (Feature Selection, Feature Learning)**

Feature selection is the process of selecting relevant features from a data-set for creating a Machine Learning model. A variable is used as an input to the model. With respect to a dataset, a feature represents an attribute and value combination. Color is an attribute. "Color is blue" is a feature. In Excel terms, features are similar to cells. The term feature has other definitions in different contexts.

➤ **Feature Learning**

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

Feature Selection

Feature selection is the process of selecting relevant features from a data-set for creating a Machine Learning model.

➤ **Feature Vector**

A list of features describing an observation with multiple attributes. In Excel we call this a row.

➤ **Feed-Forward (Neural) Networks**

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

➤ **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 aiming to give control to citizens and residents over their personal data.

➤ **Genetic Algorithm**

A search heuristic inspired by the Theory of Evolution that 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.

➤ **Gradient Accumulation**

A mechanism to split the batch of samples—used for training a neural network—into several mini-batches of samples that will be run sequentially. This is used to enable using large batch sizes that require more GPU memory than available.

➤ **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)**

Hyperparameters are higher-level properties of a model such as how fast it can learn (learning rate) or complexity of a model. The depth of trees in a Decision Tree or number of hidden layers in a Neural Networks are examples of hyperparameters. They are configurations, external to models 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.

➤ **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 object recognition research.

➤ **Image Recognition**

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

➤ **Induction**

A bottoms-up approach to answering questions or solving problems. A logic technique that goes from observations to theory, e.g., ‘We keep observing X, so we infer that Y must be True.’

➤ **Inference**

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

➤ **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.

➤ **Instance**

A data point, row, or sample in a dataset. Another term for observation

Label, it is the ‘answer’ portion of an observation in supervised learning. For example, in a dataset used to classify flowers into different species, the features might include the petal length and petal width, while the label would be the flower’s species.

➤ **Layer (Hidden Layer)**

A series of neurons in an Artificial Neural Network that 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.

➤ **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, with the size of the update steps to take during optimization loops like Gradient Descent. A high learning rate enables more ground to be covered at each step, but at risk of overshooting the lowest point since the slope of the hill is constantly changing. With a very low learning rate, we can confidently move in the direction of the negative gradient since we are recalculating it so frequently. A low learning rate is more precise, but calculating the gradient is time-consuming, so it will take us a very long time to get to the bottom.

➤ **Learning Rate**

A scalar value 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 used in mathematics, especially in statistics.

➤ **Long Short-Term Memory Networks**

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

➤ **Loss**

Loss = true value (from data-set) - predicted value (from ML-model) The lower the loss, the better a model (unless the model has over-fitted to the training data). The loss is calculated on training and validation and its interpretation is how well the model is doing for these two sets. Unlike accuracy, loss is not a percentage. It is a summation of the errors made for each example in training or validation sets.

➤ **Machine Learning**

The subfield of Artificial Intelligence that 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. Mitchell (1997) provides a succinct definition: “A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.” In simple language machine learning is a field in which human made algorithms have an ability learn by itself or predict future for unseen data.

Machine Learning Lifecycle Management

DevOps for Machine Learning systems.

➤ **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 abstracted representation of what a Machine Learning system has learned from the training data during the training process and is a data structure that stores a representation of a dataset (weights and biases). Models are created/learned when you train an algorithm on a dataset.

➤ **Monte Carlo**

An approximate methodology that uses repeated random sampling in order 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 in order to solve multiple tasks are 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 that studies the interactions between computers and human languages, in particular how to process and analyze large amounts of natural language data.

➤ **Neural Networks**

See Artificial Neural Networks. Neural Networks are mathematical algorithms modeled after the brain’s architecture, designed to recognize patterns and relationships in data.

➤ **Neuron**

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

➤ **Node**

See Neuron

➤ **Noise**

Any irrelevant information or randomness in a dataset which obscures the underlying pattern.

➤ **Normalization**

Restriction of the values of weights in regression to avoid overfitting and improving computation speed.

➤ **Null Accuracy**

Baseline accuracy that can be achieved by always predicting the most frequent class (“B has the highest frequency, so let’s guess B every time”).

➤ **Observation**

A data point, row, or sample in a dataset. Another term for instance.

➤ **Optical Character Recognition**

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

➤ **Optimization**

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

➤ **Outlier**

An observation that deviates significantly from other observations in the dataset.

➤ **Overfitting**

Overfitting occurs when your model learns the training data too well and incorporates details and noise specific to your dataset. You can tell a model is overfitting when it performs great on your training/validation set, but poorly on your test set (or new real-world data). The fact that a model unknowingly identified patterns in the noise and assumed those represented the underlying structure; the production of a model that corresponds too closely to a particular set of data, and therefore fails to generalize well to unseen observations.

➤ **Parameters**

Parameters are properties of training data learned by training a machine learning model or classifier. They are adjusted using optimization algorithms and unique to each experiment.

Examples of parameters include weights in an artificial neural network; support vectors in a support vector machine; and coefficients in a linear or logistic regression.

➤ **Pattern Recognition**

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

➤ **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 in order to identify a particular individual.

➤ **Precision**

The number of correct positive results divided by the number of all positive results returned by a classifier. In the context of binary classification (Yes/No), precision measures the model’s performance at classifying positive observations (i.e., ‘Yes’). In other words, when a positive value is predicted, how often is the prediction correct? We could game this metric by only returning positive for the single observation we are most confident in.

$$P = \frac{\text{TruePositives}}{\text{TruePositives} + \text{FalsePositives}}$$

➤ **Prediction**

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

➤ **Preprocessing**

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

➤ **Pre-trained Model**

A model or the components of a model that have 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 that 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 individual trees.

➤ **Recall**

The fraction of all relevant samples that are correctly classified as positive. Also called sensitivity. In the context of binary classification (Yes/No), recall measures how “sensitive” the classifier is at detecting positive instances. In other words, for all the true observations in our sample, how many did we “catch.” We could game this metric by always classifying observations as positive.

➤ **$R = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$**

➤ **Recall vs. Precision**

Say we are analyzing Brain scans and trying to predict whether a person has a tumor (True) or not (False). We feed it into our model and our model starts guessing. Precision is the % of True guesses that were actually correct! If we guess 1 image is True out of 100 images and that image is actually True, then our precision is 100%! Our results aren't helpful however because we missed 10 brain tumors! We were super precise when we tried, but we didn't try hard enough. Recall, or Sensitivity, provides another lens which with to view how good our model is. Again, let's say there are 100 images, 10 with brain tumors, and we correctly guessed 1 had a brain tumor. Precision is 100%, but recall is 10%. Perfect recall requires that we catch all 10 tumors!

➤ **Rectified Linear Unit**

A unit employing the rectifier function as an activation function.

➤ **Recurrent Neural Networks**

A class of Artificial Neural Network 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 their internal state (memory) to process sequential signals.

➤ **Regression (Linear Regression, Logistic Regression)**

A set of statistical processes for estimating the relationships among variables, e.g., Predicting a continuous output (e.g., price, sales). 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, an explanatory variable used as an input to a model.

➤ **Regularization**

Regularization is a technique utilized to combat the overfitting problem. This is achieved by adding a complexity term to the loss function that gives a bigger loss for more complex models. It is thus the process of introducing additional information in order to prevent overfitting.

➤ **Reinforcement Learning**

The subfield of Machine Learning inspired by human behavior studying how an agent should act 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.

➤ **ROC (Receiver Operating Characteristic) Curve**

A plot of the true positive rate against the false positive rate at all classification thresholds. This is used to evaluate the performance of a classification model at different classification thresholds. The area under the ROC curve can be interpreted as the probability that the model correctly distinguishes between a randomly chosen positive observation (e.g., 'spam') and a randomly chosen negative observation (e.g., 'not spam').

➤ **Segmentation**

It is the process of partitioning a data set into multiple distinct sets. This separation is done such that the members of the same set are similar to each other and different from the members of other sets.

➤ **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 number 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.

➤ **Specificity**

In the context of binary classification (Yes/No), specificity measures the model's performance at classifying negative observations (i.e., 'No'). In other words, when the correct label is negative, how often is the prediction correct? We could game this metric if we predict everything as negative.

$$S = \frac{\text{TrueNegatives}}{\text{TrueNegatives} + \text{FalsePositives}}$$

➤ **Speech Recognition**

See Automated Speech Recognition

➤ **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**

Training a model using a labeled dataset. 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 which categorizes new examples.

➤ **Synthetic Data**

Data 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 Set**

A set of observations used at the end of model training and validation to assess the predictive power of your model. How generalizable is your model to unseen data?

➤ **Time Series (Time Series Data)**

A sequence of data points 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. Testing Data: The subset of available data that a data scientist selected for the testing phase of the development of a model.

➤ **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. Training Data: The subset of available data that a data scientist selected for the training phase of the development of a model. Also, see What is Training Data?

➤ **Training Set**

A set of observations used to generate machine learning models.

➤ **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, and where a model developed for a task is reused as the starting point for a model on a second task. In transfer learning, we take the pre-trained weights of an already trained model (one that has been trained on millions of images belonging to 1000's of classes, on several high power GPUs for several days) and use these already learned features to predict new classes.

➤ **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.

➤ **True Positive Rate**

Another term for recall, i.e.

$$TPR = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

The True Positive Rate forms the y-axis of the ROC curve.

➤ **Type 1 Error**

False Positives. Consider a company optimizing hiring practices to reduce false positives in job offers. A type 1 error occurs when candidate seems good and they hire her, but she turns out to be bad.

➤ **Type 2 Error**

False Negatives. The candidate was great, but the company passed on her.

➤ **Uncertainty**

A range of values likely to enclose the true value.

➤ **Underfitting**

Occurs when a Machine Learning algorithm fails to capture the underlying structure of the data properly, typically because the model is either insufficiently sophisticated, or not appropriate for the task at hand. Underfitting is when a model over-generalizes and fails to incorporate relevant variations in data that would give the model more predictive power. A model is underfitting when it performs poorly on both training and test sets. This is the opposite of Overfitting.

➤ **Universal Approximation Theorem**

A neural network with one hidden layer can approximate any continuous function but only for inputs in a specific range. If you train a network on inputs between -2 and 2, then it will work well for inputs in the same range, but you can't expect it to generalize to other inputs without retraining the model or adding more hidden neurons.

➤ **Unsupervised Learning**

The area of Machine Learning that consists in inferring a function that describes the structure of unlabeled data, thus training a model to find patterns in an unlabeled dataset (e.g., clustering).

➤ **Validation**

The process of using hold-out data in order 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 Set**

A set of observations used during model training to provide feedback on how well the current parameters generalize beyond the training set. If training error decreases but validation error increases, a model is likely overfitting and training should thus be paused.

➤ **Vanishing/Exploding Gradients**

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

➤ **Variance**

An error due to sensitivity to small fluctuations in the training set computed as the expectation of the squared deviation of a random variable from its mean. Low variance suggests a model is internally consistent, with predictions varying little from each other after every iteration. High variance (with low bias) suggests your model may be overfitting and reading too deeply into the noise found in every training set.

* This glossary is a compendium derived with acknowledgements from the following sources:

<http://robotics.stanford.edu/~ronnyk/glossary.html> <https://developers.google.com/machine-learning/glossarys>

<https://appen.com/ai-glossary/>