

Deploy

OpenAI

YAML

Label

IoT

Bias

Ultimate Artificial Intelligence Glossary 2021

The definitive guide to artificial intelligence terminology.

Signal

AutoML

Cluster

XML

Introduction

As Artificial Intelligence's impact on the world continues to grow, more and more business stakeholders will encounter and grapple with the lexicon.

It can seem daunting at first. While some terms are from our everyday vocabulary, AI has given them entirely new meanings. Others are combinations of ordinary words that now refer to altogether new concepts. And a select few are words and terms you may not have ever used, or even heard, before today.

While AI and its usefulness may still seem a little abstract, there's no denying it's here to stay. This simplified, comprehensive glossary offers a framework to ground you as you explore artificial intelligence and the multiple use cases available to solve your biggest business challenges.

A

Accuracy

Refers to the percentage of correct predictions the classifier made.

Activation

The equation of a neural network cell that transforms data as it passes through the network.

Active Learning

A machine learning term that refers to various methods for actively improving the performance of trained models. Clarifai helps users with advanced workflows to collect prediction data from production environments, auto-annotate high-confidence data, or pipe lower confidence concept predictions into a dataset annotation task queue.

Adversarial Machine Learning

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

Adversarial Example

A very specific transformation of an image, typically featuring very small, deliberate changes to an image that can completely disrupt a previously tuned classifier.

Anchor Box

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

Annotation

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

Annotation Format

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

Annotation Group

Describes what types of objects you are identifying. For example, "chess pieces" or "vehicles". Classes (eg "rook", "pawn") are members of an annotation group.

Application Programming Interface (API)

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

Architecture

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

Artificial General Intelligence (AGI)

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

Artificial Intelligence (or Weak AI)

A computational system that simulates parts of human intelligence but focuses on one narrow task. Also called narrow AI, in contrast to AGI.

Artificial Neural Network

A model for AI and machine learning inspired by the neural network configurations of the human central nervous system, especially the brain.

Augmented Reality

An enhanced version of reality created by the use of technology to overlay digital information on an image of something being viewed through a device (such as a smartphone camera).

Automated Machine Learning

Also known as AutoML. Refers to processes for automating the end-to-end machine learning cycle to help practitioners scale their growing ML efforts efficiently.

Automation Bias

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



B

Baseline

A model used as a reference point for comparing how well another model (typically, a more complex one) is performing. For example, a logistic regression model might serve as a good baseline for a deep model. For a particular problem, the baseline helps model developers quantify the minimal expected performance that a new model must achieve for the new model to be useful.

Batch

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

Batch Inference

Making predictions on many frames at once to take advantage of the GPU's ability to perform parallel operations. This can help improve performance if you are doing offline (as opposed to real-time) prediction.

Batch Size

The number of examples in a batch. For example, the batch size of SGD is 1, while the batch size of a mini-batch is usually between 10 and 1000. Batch size is usually fixed during training and inference

BERT

Stands for Bidirectional Encoder Representations from Transformers, a 2018 research paper published by Google AI Language and created an inflection point for NLP machine learning models.

Bias

1. Stereotyping, prejudice or favoritism towards some things, people, or groups over others. These biases can affect collection and interpretation of data, the design of a system, and how users interact with a system.
2. Systematic error introduced by a sampling or reporting procedure.

Big Data

This is the massive amount of information we now generate about ourselves — our interests and habits — as we move through the digital universe. Some say the term “big data” should be retired, because so much data is collected these days that all data is now part of big data.



Binary Classification

A type of classification task that outputs one of two mutually exclusive classes. For example, a machine learning model that evaluates email messages and outputs either "spam" or "not spam" is a binary classifier.

Black Box

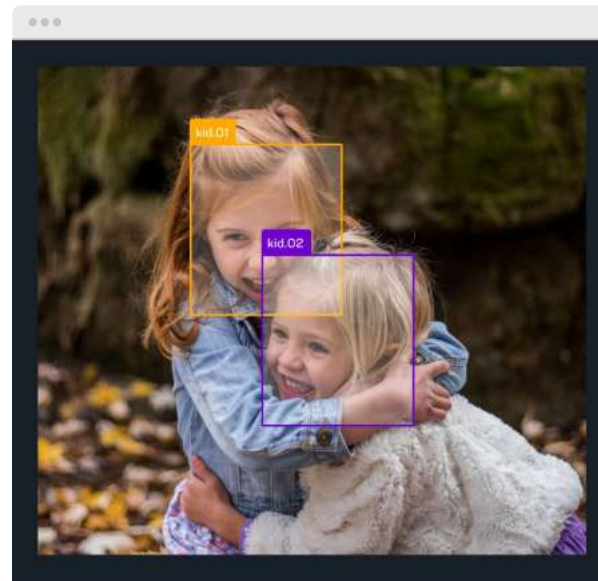
A system that makes it hard to peek behind the curtain to understand what is going on. Neural networks are often described as black boxes because it can be hard to explain "why" they are making a particular prediction. Model explainability is currently a hot topic and field of study.

Boosting

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

Bounding Box

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



Brute Force Search

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

C

Calibration Layer

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

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 (I/O) operations specified by the instructions.

Chatbot

Simulates human conversation, using response workflows or artificial intelligence to interact with people based on verbal and written cues. Chatbots can be the frontline of communication between brands and their users.



Checkpoint

Data that captures the state of the variables of a model at a particular time. Checkpoints enable exporting model weights, as well as performing training across multiple sessions. Checkpoints also enable training to continue past errors (for example, job preemption).

Clarifai Base Workflow/Model

Alters the embeddings over which other images are predicted against. The base workflow/model you choose will optimize custom trained models to use the historical knowledge base from the selected public model.

For example, if you're training a custom model around food, you could choose the Food model to gain optimized results. We recommend choosing the General model if you're not sure which public model would best suit your inputs.

Class

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

Class Balance

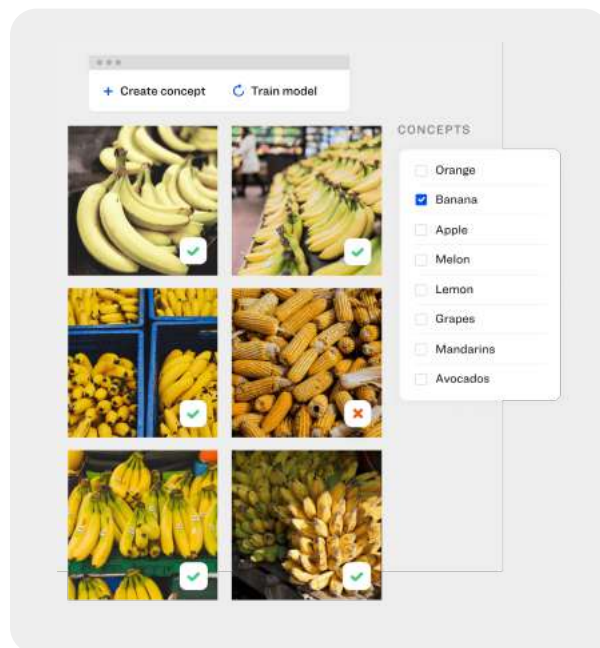
The relative distribution between the number of examples of each class. Models generally perform better if there is a relatively even number of examples for each class. If there are too few of a particular class, that class is “under-represented”. If there are many more instances of a particular class, that class is “over-represented”.

Classifier

An algorithm that implements classification, especially in a concrete implementation. The term often refers to the mathematical function implemented by a classification algorithm that maps input data to a category.

Classification Model

A type of machine learning model for distinguishing among two or more discrete classes. For example, a natural language processing classification model could determine whether an input sentence was in French, Spanish, or Italian. Compare with a regression model.



Cluster

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

Clustering

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

Cohen's Kappa

This is essentially a measure of how well the classifier performed as compared to how well it would have performed simply by chance. In other words, a model will have a high Kappa score if there is a big difference between the accuracy and the null error rate.

Computer Vision Platform

A computer vision platform is a (usually cloud-hosted) meta-tool that ties into various other tools to manage all (or part of) your pipeline.

Concept

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

Confidence

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

Confidence Threshold

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

Confusion Matrix

Shows which concepts the model is “confusing” with each other. For example, if a cat/dog/horse classifier thinks many of the images are dogs when they’re actually horses, then we would say that “dog” and “horse” are highly confused, and the entries for dog/horse and horse/dog in the matrix would be high.

Container

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

Convolutional Filter

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

CoreML

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

Convolutional Neural Network

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

CreateML

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

Custom Dataset


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

Custom Model

A small artificial neural network which takes inputs particular to a user, such as images or videos of their products and returns predicted concepts based on what the model is trained to see in the inputs.

Custom Training

The process of teaching a model to make certain predictions.



D

Data

Any collection of information converted into a digital form.

Data Annotation

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

Data Mining

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

Dataset

A collection of data and a ground truth of outputs that you use to train a machine learning model by example. For object detection this would be your set of images (data) and annotations (ground truth) that you would like your model to learn to predict.

Deep Learning

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

Deep Neural Network

An artificial neural network (ANN) with multiple layers between the input and output layers.¹⁵ It uses sophisticated mathematical modeling to process data in complex ways.¹⁶

Deploy

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

Detection


The process of discovering an event or object.

Detection Mode

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

Domain Adaptation

Learning a discriminative classifier or other predictor in the presence of a shift between training and test distributions.



E

Edge Computing

Edge computing is a computing paradigm that moves computation away from data centers to the edge of a network, often on IoT sensor devices or local appliances.

Edge Deployment

Deploying to a device that will make predictions without uploading the data to a central server over the Internet. This could be an iPhone or Android device, a Raspberry Pi, a NVIDIA Jetson, a robot, or even a full computer with a GPU located on-site.

Explorer

A web application that allows you to preview Clarifai models. It is of Clarifai's AI Portal.

F

F Score

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

Facial Recognition

A computer application capable of identifying or verifying a person from a digital image or a video frame from a video source. One of the ways to do this is by comparing selected facial features from the image and a face database.

False Negatives

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

False Positives

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


FastAI

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

Feature Extraction

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

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



Framework

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

G

Generative Adversarial Networks (GANs)

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

GPU Memory

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

Graphics Processing Unit (GPU)

A specialized electronic circuit designed to rapidly manipulate and alter memory to accelerate the creation of images in a frame buffer intended for output to a display device. GPUs are used in embedded systems, mobile phones, personal computers, workstations, game consoles, and (in Clarifai's case) to train deep neural networks.

Grid Search

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

Ground Truth

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

H

Hosted Model

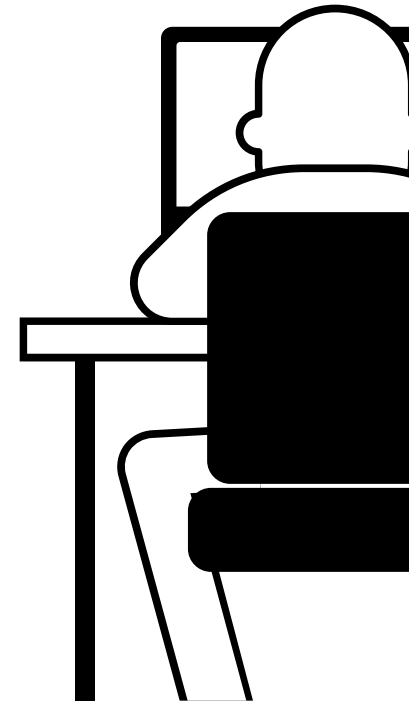

A set of trained weights located in the cloud that you can receive predictions from via an API. (As opposed to an edge-deployed model.)

Human Workforce (“Labelers”)

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

Hyperparameter

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



I

Inference

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

Intelligent Character Recognition (ICR)

Related technology to OCR designed to recognize handwritten characters.

Image Recognition

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

Image Segmentation

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

Put simply, segmentation is to put a bounding box around the desired object in an image and do a pixel-by-pixel outline of that object, removing the background.

ImageNet

A large visual database designed for use in visual object recognition software research. Over 14 million URLs of images have been hand-annotated by ImageNet to indicate what objects are pictured; in at least one million of the images, bounding boxes are also provided.

ImageNet Challenge

A competition where research teams evaluate their algorithms on the given data set and compete to achieve higher accuracy on visual recognition tasks.

IoT

The Internet of things is a network of interrelated devices.

Input

Any form of data - text, audio, code, that can be encoded digitally.

J

Jetson

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

JSON

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


Jupyter Notebook

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

L

Label

The class of a specific object in your dataset. In classification, this is the entirety of the prediction. In object detection, it is the non-spatial component of the bounding box.



Labeler

AI powered labeling tool for human workforces, backed by Clarifai's.

Labeling

The process of annotating datasets to train machine learning models.

M

Machine Learning (ML)

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

Metadata

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

Misclassification Rate


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

MLOps

Also known as Machine Learning Operations. Best practices for organizations to operationalize machine learning. Often involves collaboration between data scientists and devops professionals to help manage production ML, but platforms like Clarifai remove the need for devops.

Model

A processing block that takes inputs, such as images or videos, and returns predicted concepts.



Model Size

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

N

Natural Language Processing (NLP)

A branch of artificial intelligence that helps computers understand, interpret, and manipulate human language. This field of study focuses on helping machines to better understand human language in order to improve human-computer interfaces with use cases like moderation, information extraction, summarization, etc.

Neural Architecture Search

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

Neural Network

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

Noise

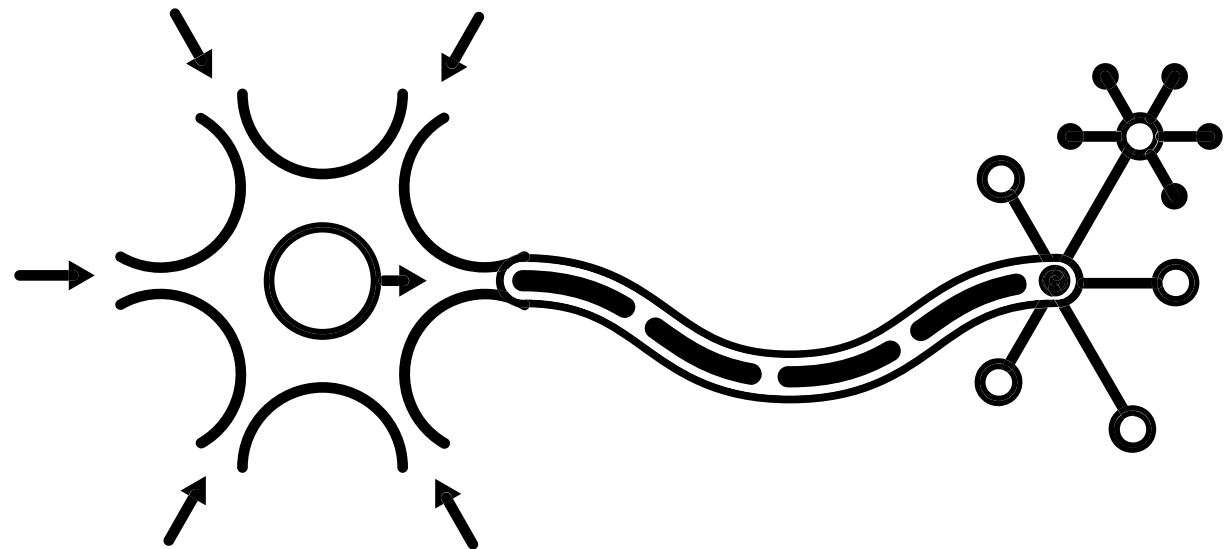
Signals with no causal relation to the target function.

Not Suitable For Work (NSFW)

Shorthand tag used to mark certain content as being profane, offensive, and/or otherwise potentially disturbing, which a platform may not wish to have posted on their site or may want to mark as mature.

Null Error Rate

How often one would be wrong if one always predicted the majority prediction.³² (e.g. if you make 100 predictions, 60 “yes” and 40 “no”, the null error rate would be $40/100=0.40$ because if you always predicted yes, you would only be wrong for the 40 “no” cases).



O

Object Detection

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

Object Recognition

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

Object Tracking

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

On-premise Software

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

One Shot Classification

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

Open Neural Network Exchange (ONNX)

ONNX is an open format to represent machine learning models.

OpenAI

Mission is to ensure that artificial general intelligence benefits all of humanity.

Optical Character Recognition (OCR)

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

Output

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

Outsourced Labeling

Paying people to annotate and/or label your images. Its effectiveness can depend on domain expertise of annotators, but companies provide custom education/training sessions prior to starting high-volume annotation.

Overfitting

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

P

Parameter


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

Pipeline

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

Polygon

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



Portal

Clarifai's user interface for its AI solutions and its end-to-end AI tools.

Precision (Recognition)

A rate that measures how often a model is correct when it predicts 'yes.'

Prediction

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

Predictive Model

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

Pre-trained Model

A model that has already been trained on another dataset. Many things it learns will be broadly applicable to images in other datasets (for example, finding lines, corners, and patterns of colors). Pre-training on a large dataset like COCO can reduce the number of custom images you need to obtain satisfactory results.

Positive Predictive Value (PPV)

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

Prevalence

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

Production

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

PyTorch

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

R

Recall (Sensitivity)

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

Receiver Operating Characteristic (ROC) Curve

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

Recurrent Neural Network

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

Regression

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

Reinforcement Learning

A type of machine learning in which machines are “taught” to achieve their target function through a process of experimentation and reward receiving positive reinforcement when its processes produce the desired result and negative reinforcement when they do not. This is differentiated from supervised learning, which would require an annotation for every individual action the algorithm would take.

S

Search Query

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

Segmentation Model

Models help identify objects or boundaries by segmenting images into distinct regions based on pixel characteristics.

Selective Filtering

When a model ignores “noise” to focus on valuable information.

Siamese Networks

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

Signal

Inputs, information, data.

Software Development Kit (SDK)


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

Specificity

The rate of how often a model predicts “no,” when it’s actually “no.”

Standard Classification

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



Structured Data

Data that resides in a fixed field within a file or record. Structured data is typically stored in a relational database (RDBMS). It can consist of numbers and text, and sourcing can happen automatically or manually, as long as it’s within an RDBMS structure.

Supervised Learning

1) A type of machine learning in which human input and supervision are an integral part of the machine learning process on an ongoing basis. In supervised learning, there is a clear outcome to the machine’s data mining and its target function is to achieve this outcome, nothing more.

2) A class of machine learning algorithms that learn patterns from outcome data. Supervised learning algorithms make predictions based on a set of examples.

Synthetic Data

Images that are created rather than collected. There are several strategies for creating more training data including using 3D models, GAN synthesis, and context augmentation.

T

Target Function

The end goal of an algorithm.

Taxonomy

The formal structure of all the types of objects within a particular domain. They can follow either a flat or hierarchical format and provide names for each object in relation to the other objects, often capturing the membership properties of each. There are usually specific, complete, consistent, and definitive rules for classifying all objects in the domain. This ensures any newly discovered object fits into one and only one category of the structure.

TensorFlow

An open-source software library also used for machine learning applications such as neural networks. It is used for both research and production at Google and was released under the Apache 2.0 open source license in 2015.

Tensor Processing Unit (TPU)

Google's hardware accelerator for performing operations on tensors. It is much faster than a GPU for some workloads. Most often they are run on Google Cloud or Google Colab but there is also an edge-TPU that can be deployed in the field.

Test Data Set

In machine learning, the test data set is the data given to the machine after the training and validation phases have been completed. This data set is used to check the performance characteristics of the algorithms produced after the completion of the first two phases when presented with unknown data. This will give a good indication of the accuracy, sensitivity, and specificity of the algorithm's predictive powers.

Torch


A scientific computing framework with wide support for machine learning algorithms, written in C and lua. The main author is Ronan Collobert, and it is used at Facebook AI Research and Twitter.

Train

The process iteratively of adjusting your model's parameters to converge on the weights that optimally mimic the training data.

Training Data Set

In machine learning, the training data set is the data given to the machine during the initial "learning" or "training" phase. From this data set the machine is meant to gain some insight into options for the efficient completion of its assigned task through identifying relationships between the data.



Transfer Learning

A machine learning process that transfers features (knowledge) learned from training large models on general datasets, into new models without having to start from nothing.

Transformer

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

True Negatives

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

True Positives

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

Turing Test

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

U

Unstructured Data

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

Unsupervised Learning

A class of machine learning algorithms that learns patterns in data without knowing outcomes. Here, the machine is presented with totally unlabeled data, then asked to find the intrinsic patterns.

V

Validate

During the training process of a neural network, the validation set is used to assess how well the model is generalizing. These examples are not used to calculate the gradient; they are the ones used to calculate your metrics and see how well they are improving over time.

Validation Data Set

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

Video Frame Interpolation

Is to synthesize several frames in the middle of two adjacent frames of the original video. Video Frame Interpolation can be applied to generate slow motion video, increase video frame rate, and frame recovery in video streaming.



Vision Processing Unit (VPU)

As of 2016, it is an emerging class of microprocessor and a specific type of AI accelerator, designed to accelerate machine vision tasks.

Visual Recognition

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

Visual Search

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

W

Web Crawler (Spider)

An internet bot that systematically browses the World Wide Web, typically for the purpose of Web indexing, copying pages for processing by a search engine which indexes the downloaded pages, allowing users to search more efficiently.

Web Scraper

The automated processes implemented using a bot or web crawler. It is a form of copying, in which specific data is gathered and copied from the web, typically into a central local database or spreadsheet, for later retrieval or analysis.

Workflows

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

X

XML

A hierarchical data format (HTML, the markup language defining the layout and content of the page you are currently reading, is a subset of XML). In computer vision XML is most commonly used with the Pascal VOC XML annotation format.

Y

YAML

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

About Clarifai

Clarifai is the leading independent provider of deep learning AI for unstructured data. Headquartered in NYC, the company was founded in 2013 by Matt Zeiler, Ph.D. after winning the top 5 places at ImageNet with the goal of delivering state-of-the-art AI technologies to enterprises and organizations around the world. Clarifai offers the most powerful platform for the end-to-end AI lifecycle, UIs that unleash deep learning for every skill set, and the best solutions to important use cases. The company has raised \$40M from Union Square Ventures, Menlo Ventures, Lux Capital, NVIDIA, Google Ventures, and Qualcomm.

Clarifai continues to grow with 100+ employees at its headquarters in New York City, and offices in San Francisco, Washington, DC, and Tallinn, Estonia.



Terms

A

- 3 Accuracy
- 3 Activation
- 3 Active Learning
- 3 Adversarial Machine Learning
- 3 Adversarial Example
- 3 Anchor Box
- 3 Annotation
- 3 Annotation Format
- 3 Annotation Group
- 4 Application Programming Interface (API)
- 4 Architecture
- 4 Artificial General Intelligence (AGI)
- 4 Artificial Intelligence (or Weak AI)
- 4 Artificial Neural Network
- 4 Augmented Reality
- 4 Automated Machine Learning
- 4 Automation Bias

B

- 5 Baseline
- 5 Batch
- 5 Batch Interface
- 5 Batch Size
- 5 BERT
- 6 Bias
- 6 Big Data
- 6 Binary Classification

- 6 Black Box
- 6 Boosting
- 6 Bounding Box
- 6 Brute Force Search

C

- 6 Calibration Layer
- 6 Central Processing Unit (CPU)
- 7 Chatbot
- 7 Checkpoint
- 7 Clarifai Base Workflow/Model
- 7 Class
- 7 Class Balance
- 7 Classifier
- 8 Classification Model
- 8 Cluster
- 8 Clustering
- 8 Cohen's Kappa
- 8 Computer Vision Platform
- 8 Concept
- 8 Confidence
- 8 Confidence Threshold
- 8 Confusion Matrix
- 9 Container
- 9 Convolutional Filter
- 9 CoreML
- 9 Convolutional Neural Network
- 9 CreateML
- 9 Custom Dataset
- 9 Custom Model
- 9 Custom Training

D

- 9 Data
- 9 Data Annotation
- 9 Data Mining
- 10 Dataset
- 10 Deep Learning
- 10 Deep Neural Network
- 10 Deploy
- 10 Detection
- 10 Detection Mode
- 10 Domain Adaptation

E

- 10 Edge Computing
- 10 Edge Deployment
- 10 Explorer

F

- 11 F Score
- 11 Facial Recognition
- 11 False Negatives
- 11 False Positives
- 11 FastAI
- 11 Feature Extraction
- 11 Framework

G

- 11 Generative Adversarial Networks (GANs)
- 12 GPU Memory
- 12 Graphics Processing Unit (GPU)
- 12 Grid Search
- 12 Ground Truth

H

- 12 Hosted Model
- 12 Human Workforce (“Labelers”)
- 12 Hyperparameter

I

- 13 Inference
- 13 Intelligent Character Recognition (ICR)
- 13 Image Recognition
- 13 Image Segmentation
- 13 ImageNet
- 13 ImageNet Challenge
- 13 IoT
- 13 Input

J

- 13 Jetson
- 13 JSON
- 13 Jupyter Notebook

L

- 10 Label
- 14 Labeler
- 14 Labeling

M

- 14 Machine Learning (ML)
- 14 Metadata
- 14 Misclassification Rate
- 14 MLOps
- 14 Model
- 14 Model Size

N

- 15 Natural Language Processing (NLP)
- 15 Neural Architecture Search
- 15 Neural Network
- 15 Noise
- 15 Not Suitable For Work (NSFW)
- 15 Null Error rate

O

- 16 Object Detection
- 16 Object Recognition
- 16 Object Tracking
- 16 On-premise Software
- 16 One Shot Classification
- 16 Open Neural Network Exchange (ONNX)
- 16 OpenAI
- 16 Optical Character Recognition (OCR)
- 16 Output
- 16 Outsourced Labeling
- 16 Overfitting

P

- 17 Parameter
- 17 Pipeline
- 17 Polygon
- 17 Portal
- 17 Precision (Recognition)
- 17 Prediction
- 17 Predictive Model
- 18 Pre-trained Model
- 18 Positive Predictive Value (PPV)
- 18 Prevalence
- 18 Production
- 18 PyTorch

R

- 18 Recall (Sensitivity)
- 18 Receiver Operation Characteristic (ROC) Curve
- 18 Recurrent Neural Network
- 18 Regression
- 18 Reinforcement Learning

S

- 19 Search Query
- 19 Segmentation Model
- 19 Selective Filtering
- 19 Siamese Networks
- 19 Signal
- 19 Software Development Kit (SDK)
- 19 Specificity
- 19 Standard Classification
- 19 Structured Data
- 19 Supervised Learning
- 19 Synthetic Data

T

- 20 Target Function
- 20 Taxonomy
- 20 TensorFlow
- 20 Tensor Processing Unit (TPU)
- 20 Test Data Set
- 20 Torch
- 20 Train
- 20 Training Data Set
- 21 Transfer Learning
- 21 Transformer
- 21 True Negatives
- 21 True Positives
- 21 Turing Test

U

- 21 Unstructured Data
- 21 Unsupervised Learning

V

- 21 Validate
- 21 Validation Data Set
- 22 Video Frame Interpolation
- 22 Vision Processing Unit (VPU)
- 22 Visual Recognition
- 22 Visual Search

W

- 22 Web Crawler (Spider)
- 22 Web Scraper
- 22 Workflows

X

- 22 XML

Y

- 22 YAML