Volume – 3 | Issue – III | May-June 2026 | ISSN: 3049-0278 (Online)

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

AN OVERVIEW OF CONCEPTS AND TERMINOLOGY OF MACHINE LEARNING AND DEEP LEARNING

By
Awinash Singh[1]

Guided By
Dr Manoj Kumar [2]

ABSTRACT
Machine learning and deep learning are changing the world and becoming critical components of it. Articles such as "Why Deep Learning is suddenly changing your life" by Forbes and "Are Computers Already Smarter than Humans?" by Time, demonstrate the impact these new technologies have on our lives today. Because machine learning is not only impacting the world, but is increasingly infiltrating our day-to-day lives, it is important to understand types of machine learning and what they do. Over the last decade of my career, my expertise has been in building scalable systems and data systems (big data & traditional). As I have grown my expertise in artificial intelligence (AI)/machine learning (ML), I have found many introductory resources exist for engineers; however, I've discovered an absence of simple, comprehensive introductory material for a broad audience. This paper presents an overview of machine learning, deep learning and their relationship for a general audience. After the introduction of concepts and terminology, an overview of the process of machine learning is given.

General Terms:- *Machine learning, Artificial intelligence, neural networks*
Keywords :- *Deep learning, artificial intelligence (AI), neural networks, RNN, CNN, machine learning 1.*

ISSN: 3049-0278 (Online) DOI (Crossref) Prefix: 10.63431

AIJITR - Volume- 3, Issue-III, May-June 2026

1. Introduction

Machine learning (ML), a subset of artificial intelligence (AI), is a growing field that is increasingly changing the day-to-day lives of people around the world. This is most visible in products and services that make predictions and act as smart machines, including the following:

- Self driving cars (Tesla Autopilot
- Fraud detection (Paypal
- Product recommendations (Amazon, Netflix
- Image recognition/classification (Google Image Search
- Language translation (Microsoft Translation
- Speech recognition (Apple Siri
- Spam filtering (Gmail

Some of the smartest business leaders in this country agree on AI's impact on the future of the world. The commitment and investment Google has made to AI can be seen in headlines such as, "[Google CEO] Sundar Pichai says the future of google is ai...." . The famous venture capitalist Marc Andreessen is often quoted saying, "software is eating the world," stating how impactful software is on the world . Jensen Huang, CEO of Nvidia, takes this a step further, saying, "Software is eating the world, but AI is going to eat software" . These leaders are not only expressing how impactful software is on the world, but how AI will be just as impactful as software, if not more. Traditionally, algorithms are

1. Research Scholar, Department of Computer Science, Chaudhary Charan Singh University, Meerut, U.P.
2. Assistant Professor, Department of Computer Science, Chaudhary Charan Singh University, Meerut, U.P.

DOI Link (Crossref) Prefix: https://doi.org/10.63431/AIJITR/3.III.2026.31-40
AIJITR, Volume 3, Issue –III, May - June, 2026, PP.31-40
Received on 11th May, 2026 & Accepted on 21st May, 2026,
Published: 17th June, 2026

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

programmed into computers by humans. In machine learning, computers learn to make decisions using algorithms that are self-taught. Ever-increasing amounts of information have enabled ML to learn by applying modern techniques on data. While the most common uses of machine learning include making predictions and acting as smart machines, the field is possible due to a number of approaches, algorithms, mathematics and techniques — which is why it's such a hard field to learn. This paper describes ways to think about ML. Concepts and associated terminology will be introduced throughout this paper. It has been organized into sections, starting with general concepts and terms. Types of machine learning are introduced in section 3, followed by types of machine learning by approach in section 4 and by task in section 5. Neural networks and deep learning are then described in section 6. In section 7, the process of machine learning is briefly covered, and in section 8, a few of the major algorithms used in machine learning are visited. This paper concludes by mentioning a few popular ML algorithms in section 9.

2. Concepts

Several key concepts and terms that are specific to ML are necessary to know in order to understand this paper and will be defined in this section.

Features :- All computer programs take one or more inputs to produce an output or result. In machine learning, individual inputs are named features. A feature is one of many parts of the input to a machine learning algorithm. Features can be thought of as a "...column in the data table for modeling" or as a named input to ML. For example, if AI is trying to predict whether a student will finish college, features could include any one of the following elements: age, SAT score, height, GPA, parents' credit score, number of after school activities, etc… Each element would be a feature of the input to make a prediction. A feature almost always has to be transformed into a numerical form, so gender might be transformed from "male" to "1" and "female" to "0."

Machine Learning Model:- Machine learning model is a term that is colloquially used in multiple ways. The first usage references a working implementation of an ML algorithm that has already been trained.

Someone saying, "Let's get a prediction from the ML model," is an example of this usage. The second use of the term machine learning model often generically references which algorithm is used to solve a machine learning problem. DataRobot , a company providing AI capabilities to enterprises, states on their website that they'll help customers pick an ML model, as an example of the second usage. The last usage refers to an artifact created during the training of an ML algorithm. As Amazon explains,

"The process of training an ML model involves providing an ML algorithm (that is, the learning algorithm) with training data to learn from. The term ML model refers to the model artifact that is created by the training process" . For example, if someone says they want to export or import an ML model, it's a reference to an artifact—a collection of settings and mathematical components that configure the ML algorithm to be usable. An ML algorithm with an ML model artifact enables ML to be applied as in the first usage, "Let's get a prediction from the ML model."

Parameters and Hyperparameters

Parameters are configuration information necessary for an ML algorithm to run. A parameter must be learned through the training of an ML algorithm. Clarifying examples of a parameter are presented [69] by Jason Brownlee, PhD in AI: A model parameter is a configuration variable that is internal to the model and whose value can be estimated from data.

- They are required by the model when making predictions.
- The values define the skill of the model on your problem.
- They are estimated or learned from data.
- They are often not set manually by the practitioner.
- They are often saved as part of the learned model.

Parameters are one of the three major inputs needed to run an ML algorithm and derive its results. The three major inputs are parameters, hyperparameters and input data. Parameters and hyperparameters together are the ML model artifact referenced in the third usage of machine learning model above. A type of parameter is the hyperparameter, that , "... is a parameter whose value is set before the learning process begins. By contrast, the values of other parameters are derived via training". It is important to understand that

"Hyperparameters and parameters are often used interchangeably but there is a difference between them. You call something a 'hyperparameter' if it cannot be learned within the estimator directly. However, 'parameters' is more general term. When you say 'passing the parameters to the model', it generally means a combination of hyperparameters along with some other parameters that are not directly related to your estimator but are required for your model" .

An example of a hyperparameter is the number that tells an ML model how many times to execute training for the model or how many examples/inputs to put in a batch for a single round, also called an epoch, of training.

Training:- Training is the process of running through a training dataset to produce an ML model. This is considered the way that a machine learning algorithm learns. The term training frequently refers to the process of training a supervised

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

learning algorithm (defined in section 4) which is one of the most commonly used types of machine learning algorithms.
Label:- Label is the name of the expected result of an output. Labeling a photo as having a cat in it is an example of a label. These initial concepts are necessary to learn about machine learning.

3. Types of Machine Learning

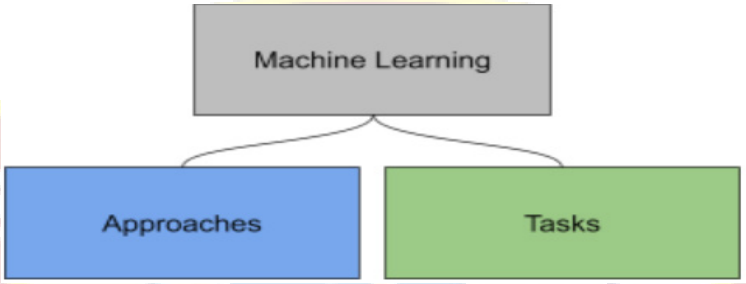

The subject of machine learning can be difficult to understand for people new to the field. There are many new and abstract concepts to learn which, if not presented clearly, makes them difficult to conceptualize. To make matters worse, introductions to the field often mix the concepts of approaches to machine learning with the problems they solve. The next sections will separate the approaches to machine learning from the tasks they perform and the problems they solve.

4. Machine Learning by Approach

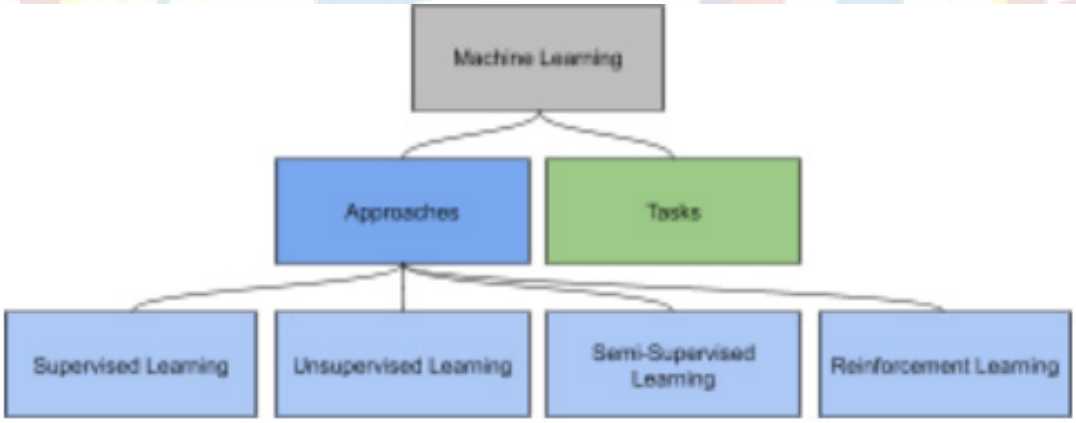

There are four main approaches to machine learning: supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning. Approaches to machine learning will be defined with examples in this section.

Supervised Learning

Supervised learning is an approach of machine learning where an algorithm is trained using a series of examples. Those examples as a whole are known as a training data set, which is made up of pairs of inputs and outputs. In essence, the algorithm is provided with both the input and labeled answer so it can learn to predict the results of future inputs. "Therefore, the goal of supervised learning is to learn a function that, given a sample of data and desired outputs, best approximates the relationship between input and output observable in the data" . Supervised learning could, for example, predict if a photo contains a picture of a cat. This ML model would be created by a training data set that is made up of individual photos labeled as "containing a cat" or "not containing a cat."

Unsupervised Learning

Unsupervised learning takes the opposite approach to supervised learning where it learns from data with only inputs (unlabeled data). In essence, unsupervised learning is about applying techniques that analyze a dataset to find patterns (structure, relationships, features) within and between data elements. There are a number of tasks that unsupervised learning solves. Unsupervised learning tasks are compression, generalization, anomaly detection, and dimensionality reduction, which involve reducing the number of features or inputs for machine learning. The most common example problem is how to group related items, which is solved by clustering. Clustering is often used as an example because the concept of finding groups (or clusters) of items is easy to explain. An example of clustering is to find groups of property for real estate modeling. With a list of property sales and data about those properties, clustering can determine groups of properties with similar characteristics. An example group could be: three-bedroom brick detached homes in Queens, New York are trending along with two-bedroom condos in Manhattan. Pricing and user demographic analysis can then be performed on groupings like this. Clustering can also be used to identify other "customers like you" to make product recommendations on Netflix or Amazon.

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

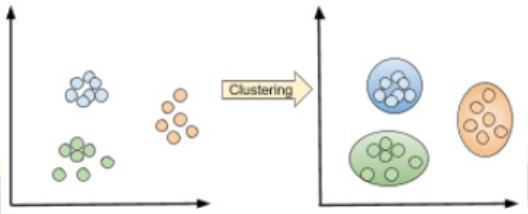

Semi-supervised Learning

Semi-supervised learning is an approach that falls somewhere between supervised and unsupervised learning. Semi-supervised learning is trained with both labeled and unlabeled data. This data is made up of two types of examples, ones with pairs of inputs and outputs (or labeled) and others with just inputs and no outputs (or unlabeled). There are different ways to approach semi-supervised learning. One approach is to use clustering (unsupervised learning), in combination with supervised learning, to group related examples of labeled and unlabeled data together before training. Another approach is to use the labeled data to predict labels for the unlabeled data. This newly generated training data only consists of labeled data which will be used to retrain the ML model. This second approach is used to augment datasets where there is a lot of data but the extra manual labeling is too expensive. Determining if a photo contains a picture of a cat is one example. The dataset contains a set of 1,000 photos that are each labeled as "cat" or "no cat." There are 9,000 more photos that are unlabeled for a total of 10,000 photos. The set of 1,000 labeled photos are used to train an algorithm to predict if a photo it has never seen contains a cat. The trained ML model is used to predict if each of the 9,000 remaining photos contains a cat and to label them. At this point, the ML algorithm is trained using the 10,000 photos that are now all labeled.

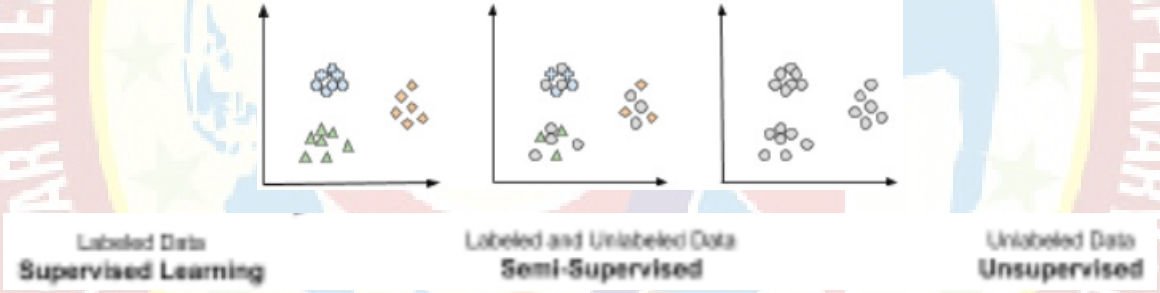

Reinforcement Learning

Reinforcement learning is defined as taking actions in the environment, then learning from the results of those actions in the form of feedback or changes to the environment. Shweta Bhatt at KDnuggets, a site about AI, said, "Reinforcement Learning(RL) is a type of machine learning technique that enables an agent to learn in an interactive environment by trial and error using feedback from its own actions and experiences". To learn more about reinforcement learning .

Examples of reinforcement learning have been used for tasks like self-driving cars with Tesla's Autopilot and game play with games like Go with Google's Alpha Go . The above approaches to machine learning all use different mechanisms to approach learning. Some approaches to machine learning can be used to solve the same problems using the right supervised or semi-supervised learning algorithms, while other approaches to ML can be complimentary like unsupervised learning is with semi-supervised learning. Each approach is a high-level classification of groups of algorithms that, at times, also contain sub-families of algorithms. For example, clustering is a form of unsupervised learning, but within the classification of clustering, there are sub-families like centroid-based clustering, density-based clustering and hierarchical clustering. While this section gives a basic overview of the major approaches to ML, the next section will review tasks that can be solved by ML.

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Volume – 3 | Issue – III | May-June 2026 | ISSN: 3049-0278 (Online)

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

5. Machine Learning by Task

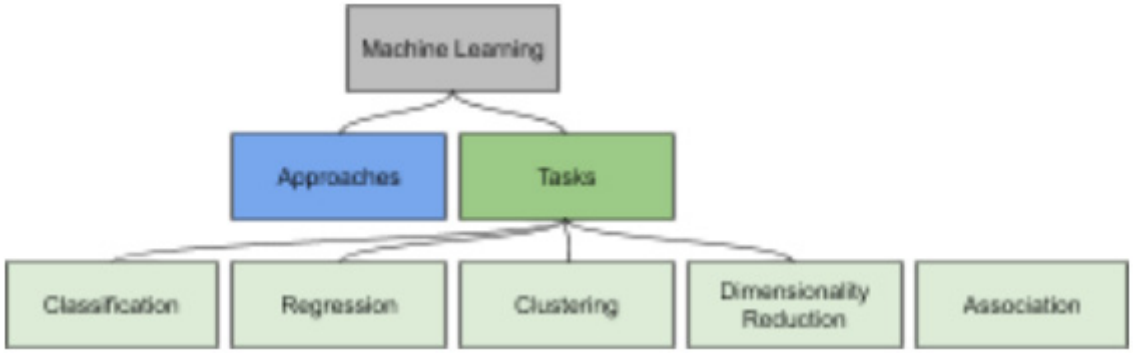

The last section discussed approaches to machine learning and the next section discusses tasks/problems that ML addresses. There are five main tasks of machine learning: classification, regression, clustering, dimensionality reduction and association. Each approach will be defined with examples in this section.

Classification

Classification is the "...process of putting people or things into groups based on ways that they are alike". It's the use of machine learning to determine if an input is in one group or another and is probably the simplest of the five main tasks to understand. Classification is the process of identifying an individual object and labeling the category it belongs to, which can be applied to text, audio or image data. Answering the question, "Does an image contain a cat?" is an example of classification. Categorization of multiple objects in a photo is also possible. A photo that is labelled as containing a person, car, hat, cat and dog is an example of multi-item classification.

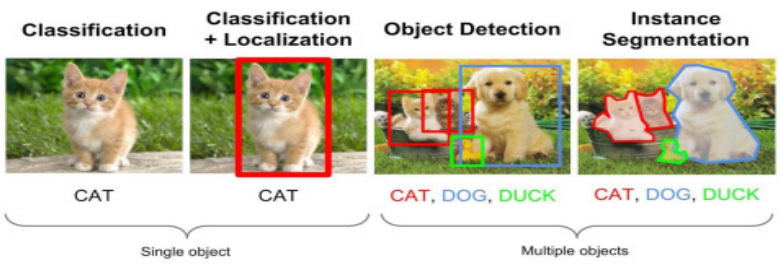

Image Credit: Arthur Ouaknine's DL post on Medium.

Regression

Regression is used to make predictions. Each prediction is the estimation of a resultant value associated with one or more input values, also known as features. Training regression results in the creation of a function which, given an input(s), will produce an output numeric value. The function is a mathematical formula that describes the relationship between the input and output. The prediction that regression makes is based on the mathematical relationship between the features and result. If there is a weak relationship, or no relationship, a prediction will not be able to be made. A common form of regression is linear regression. "Linear regression is a linear model, e.g. a model that assumes a linear relationship between the input variables (x) and the single output variable (y)" . Model in this quote refers to the function described earlier, and the linear relationship explains that the function will be the equation of a straight line. The function, also known as the equation of the line, is the regression function. Finding the equation of the line that represents the average is "training" the linear regression which is a form of learning. Examples of linear regression include using the length of a person's pants to predict a person's height, or using someone's credit score to predict the chance someone will default on a loan payment. The function for this linear regression is depicted by the red line drawn in the graph. The x-axis represents the inputs and the y-axis represents the results, also known as the prediction. The grey circles on the graph depict examples used for training. The distance a grey dot is from the line represents how inaccurate the prediction is, also known as the error.

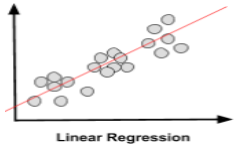

Clustering

Clustering is the process of taking data and then assigning that data to a number of groups. Wikipedia says, "clustering

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

is the task of grouping a set of objects in such a way that objects in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups (clusters)". A familiar example of clustering is used in product recommendation systems for Amazon and Netflix. It's also used in marketing to create segments of users that have similar buying habits. Another example of clustering would be looking at data of cancer patients to see if there are hot spots neighborhoods for cancer or lack of cancer.

Dimensionality reduction

Dimensionality reduction is used to decide the smallest number of inputs/features necessary to make an accurate decision in ML. If an estimate of a person's height is needed and there is information available on pant length, dress length, jacket size, favorite color, and age, then a prediction could be attempted using the available information. Based on the set of available features, the minimum number of features necessary to make a prediction will be chosen by dimensionality reduction. In the above example, applying dimensionality reduction might indicate that only two features, pant length and jacket size, are necessary to determine a person's height. Dress length is not needed because it is the same as knowing a person's pant length and jacket size. The analysis based on the previous result shows that a person's favorite color and age are not correlated to the prediction (common sense also indicates this). Dimensionality reduction is important because using less information reduces calculations and storage requirements. When there are hundreds or thousands of features, the resulting cost savings will be much greater.

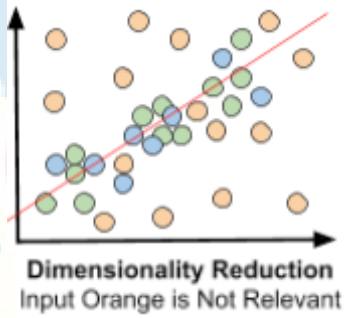

The above image shows that the feature/input that's represented in orange doesn't correlate with the red line. The blue and green dots do correlate with the red line, which means they help produce the prediction. Dimensionality reduction in this illustration would indicate that the blue and green features should be kept for use, while the orange should be "reduced" (i.e. removed).

Association rule learning

Association rule learning is the process of learning associations between data in large datasets by analyzing them. Association rule learning finds antecedent (preceding) data and consequences (outcomes) from data . These associations are considered rules. New rules are learned when new data is available. Rules are supported by three major concepts:

- Confidence is the probability of X with the co-occurrence of Y (X -> Y). E.g. probability of a person who buys milk (X) also buys cereal (Y).
- Support is the fraction of times (X -> Y) occurs out of all data entries.
- Lift is the "ratio of confidence to baseline probability of occurrence of {Y}"

The most popular example of association rule learning is market basket analysis, a technique used to uncover associations between purchased items. If you run association rule learning on shopping carts at a supermarket a rule might identify that if a person buys bread, eggs, and cereal, they will have a high probability of buying milk. {Bread, Eggs, Cereal} -> {Milk}. Catalogue or store layout design relies on analysis like this to optimize the layout of catalogues/stores to increase consumer buying behavior as an example. Association rule learning isn't about learning the habits of specific shoppers but global associations across all shoppers. This is not the best approach for designing personalized recommendations for online shopping carts since algorithms like collaborative filtering work better for this use-case.

6. Neural Networks and Deep Learning

The hottest AI topic during 2019 was deep learning and neural networks. Neural network "Neural networks are a set of algorithms, modeled loosely after the human brain, that are designed to recognize patterns". Neural networks are a general type of architecture that conceptually mimics a set of interconnected neurons in the brain, each of which is called a node. Each individual node (neuron) is connected to a number of other nodes. Like neurons in the brain, the nodes are interconnected, and each node has a number of inputs and a number of outputs. The source of the inputs and destination for the output of each individual node can be configured arbitrarily into a number of architectures. An architecture describes the sources of inputs to each node and the destinations of the outputs of each node. Each node calculates its output using a simple mathematical formula based on the inputs it receives.

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Volume – 3 | Issue – III | May-June 2026 | ISSN: 3049-0278 (Online)

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

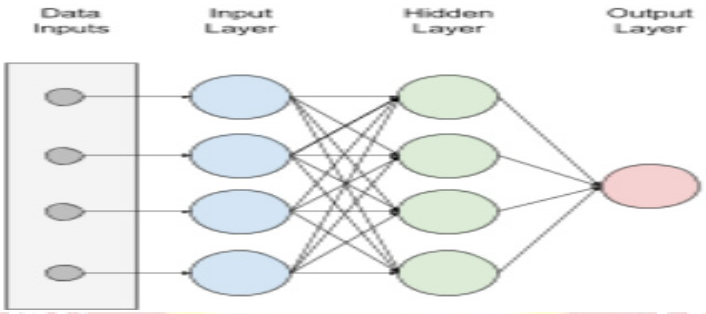

"Most of today's neural nets are organized into layers of nodes, and they're 'feed-forward,' meaning that data moves through them in only one direction. An individual node might be connected to several nodes in the layer beneath it, from which it receives data, and several nodes in the layer above it, to which it sends data." Neural networks are very flexible and can be used with all major approaches and most of the tasks defined above. They have been applied to computer vision, speech recognition, image processing, medical diagnosis and many other areas. A popular use of deep learning is to detect objects in photos (like a cat).

Neural network formula

Different formulas may be used to calculate the output of a neuron. y=w*X+b is the most commonly used example to calculate the output of each neural network node.

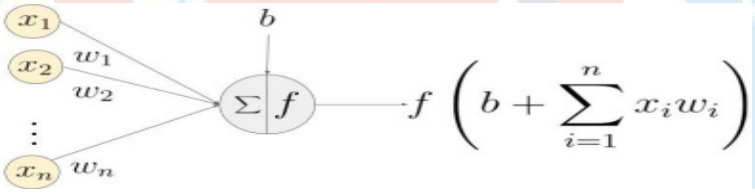

An example of a neuron showing the input (x_1 - x_n), their corresponding weights (w_1 - w_n), a bias (b) and the activation function f applied to the weighted sum of the inputs.

Image Credit: [100] The formula of a neuron

y=w X+b is the formula for a line, where w is a weight and b is a parameter of the model (called a bias). w and b are values that are "learned" when training the neural network, and initial values are assigned at random. The number of parameters can vary depending on the individual mathematical formula used for the neural network algorithm. The output of this formula is supplied as an input to another function called the activation function. Each neuron might use a different activation function, the output of which is sent as an input to nodes in the next layer. The process of calculating each neuron's value is called forward propagation, and its output makes a prediction about the world. The next step in the process is backward propagation, which is how the algorithm gets "trained" (i.e. the process of updating weights and biases). Deep learning What is deep learning, and how does it relate to neural networks? In essence deep learning is a new name for neural networks. A number of major advances since the early 2000s have led to neural network architectures that have grown enormously in both the number of nodes and layers, hence the name "deep learning."

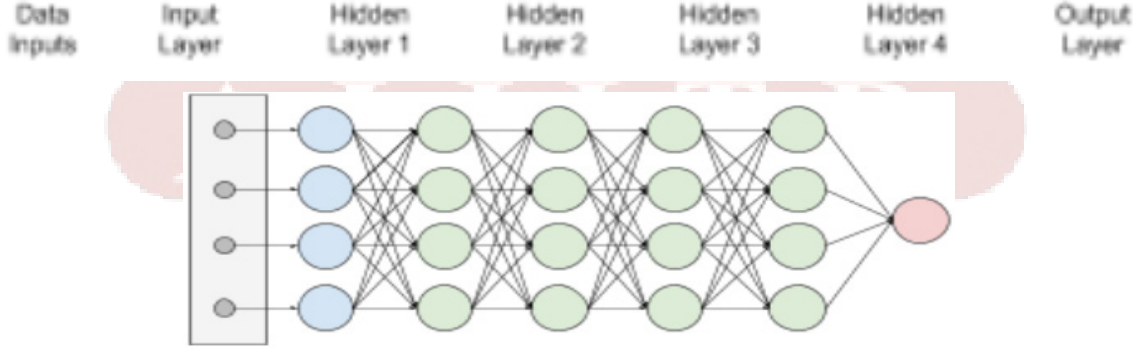

The "deep" neural network architecture was made possible by advances in both how the mathematical formulas were computed and improvements in the hardware (GPU processors) used to run these formulas. The transformation of

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Volume – 3 | Issue – III | May-June 2026 | ISSN: 3049-0278 (Online)

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

neural network mathematical formulas into equivalent matrix calculations has enabled the use of GPU processes for performance gains. The speed increases in running these algorithms/formulas are many orders of magnitude faster than earlier versions. These increases have continued over the years and are made possible by research developments in both hardware and mathematics. Major companies like Google and Nvidia are among hundreds of companies and universities that are contributing to these advances.

7. Process of Machine Learning

The high level process of training and using machine learning is relatively simple. The process can be highly interactive and can loop back on itself in a number of steps. This next section will present a simplified model of the process of machine learning.

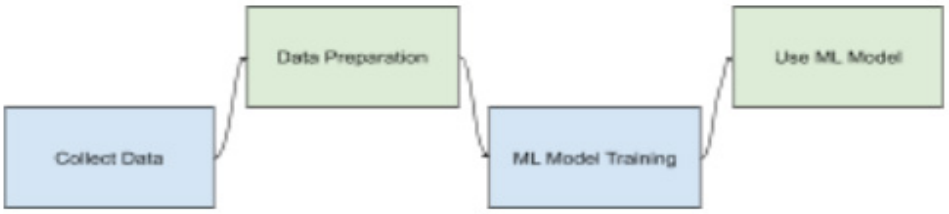

The machine learning process is made up of four steps. Those four steps most accurately capture the process of supervised machine learning. The generalized process is applicable to multiple approaches but some, like semi-supervised learning, differ more than others.

- Data collection
- Data preparation
- ML model training
- ML Model use

Data collection

Data collection is the first step of the process of machine learning. In many scenarios, data has not been collected using strict and definable criteria. In some cases, the data collection process may need to be evaluated to see if the data is usable for its intended purpose. Depending on the use-case and algorithm, you might collect anywhere from a few hundred to a few million examples. Labels for collected data, if not applied consistently during collection, will produce poor results. If labels don't exist then part of data collection might be the process of manually or automatically adding labels to the collected data. An example of this is an application that accepts text to be translated and has a human marketplace of people that translate the submitted text manually. In such a case, the data that's being collected manually is used to train a supervised learning model.

Data preparation

The next step is data preparation. Data typically needs to be checked for quality, normalized and transformed into a representation that works better for machine learning training. Parts of this process are typically called feature engineering. "Feature engineering is the process of using domain knowledge to extract features from raw data via data mining techniques" or analysis, and is part of data preparation. Cleaning the data is typically done as part of data preparation - for example, removing incomplete data, data that is clearly wrong and data that might not pass automated or manual quality checks. Parts of feature engineering are an art while others are scientific. Feature engineering can use algorithms to evaluate which features to include in the training of an ML model. Part of feature engineering includes converting features into numerical representations. Those features may or may not be normalized, brought to a common scale (e.g. between 0 and 1), and converted from many features to a single feature or visa-versa.

ML model training

ML model training is a machine learning algorithm's process of learning. This process differs by type of learning, but in the case of supervised learning, typically it is training the algorithm using batches of data. Each batch is called an epoch. Neural networks are trained using the following steps for each epoch:

- Calculate current loss (forward propagation) by applying the algorithm.
- Calculate current gradient/slope (backward propagation)
- Update parameters (gradient descent)
- Repeat

The loss is the difference between the expected result and predicted result for a node of the neural network. The gradient calculation calculates the slope of every parameter of the mathematical formula for each node in the neural network architecture. The gradient (calculated slope) for each parameter is then used to update that parameter. After a fixed number of epochs or a measured error of the results, the training ends.

Published by: Amitrakshar Publishers – Amitrakshar International Journal of Interdisciplinary and Transdisciplinary Research (AIJITR)
https://www.amitrakshar.co.in/journal/ : editor@amitrakshar.co.in DOI (Crossref) Prefix: 10.63431

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

ML model use

Using the ML model is the final and simplest step. Using the deployed machine learning model is just as simple as giving an input and waiting for output (aka the result.

8. Popular Machine Learning Algorithms

So how do algorithms relate to machine learning approaches and tasks? Approaches and tasks are just two ways to bucket or group types of machine learning algorithms. Specific implementations of a machine learning algorithm might have one or more approaches or tasks that they fall into. As people might want to learn and research more below is a list of popular machine algorithms providing the approaches and tasks they correspond to. This should make it easier to find starting points for further research

- k-means - unsupervised learning, clustering
- svm - supervised learning, classification or regression
- transductive svm - semi-supervised learning, classification or regression
- random forests - supervised learning, classification or regression
- q-learning - reinforcement learning
- linear regression - supervised learning, regression
- principal component analysis - unsupervised, dimensionality reduction
- apriori - unsupervised learning, association

9. Conclusion

The goal of this paper was to provide a quick and easy-to-understand overview of machine learning. If you're a new learner in the field, please provide feedback where things aren't clear. If you're an industry expert, please provide feedback where you feel anything might not be accurate. If your next step is to become a deep learning practitioner or just learn a lot more about the field I would start with the following resources:

- Deep Learning Specialization from Coursera and Deeplearning.ai
- Practical Deep Learning for Coders by Fast.ai

REFERENCES

[1] https://towardsdatascience.com/supervised-vs-unsupervised-l earning-14f68e32ea8d
[2] https://www.tesla.com/autopilot
[3] https://www.forbes.com/sites/bernardmarr/2016/09/30/whatare-the-top-10-use-cases-for- machine-learning-and-ai/54f19 00694c9
[4] https://www.mckinsey.com/industries/retail/our-insights/how -retailers-can-keep-up-with-consumers
[5] https://ai.googleblog.com/2014/09/building-deeper-understan ding-of-images.html#uds-search-results
[6] https://www.microsoft.com/en-us/translator/business/machin e-translation/
[7] https://machinelearning.apple.com/2018/08/09/regionally-spe cific-language-models.html
[8] https://techcrunch.com/2017/05/31/google-says-its-machine-l earning-tech-now-blocks-99-9-of-gmail-spam-and-phishingmessages/
[9] http://fortune.com/ai-artificial-intelligence-deep-machine-lea rning/
[10] http://time.com/4960778/computers-smarter-than-humans/
[11] https://a16z.com/2011/08/20/why-software-is-eating-the-wor ld/
[12]https://www.technologyreview.com/s/607831/nvidia-ceo-soft ware-is-eating-the-world-but-ai-is-going-to-eat-software/
[13]https://www.theverge.com/2017/10/4/16405174/ceo-sundarpichai-interview-google-ai-artificial-intelligence-interface
[14]https://towardsdatascience.com/simple-explanation-of-semi-supervised-learning-and-pseudo -labeling-c2218e8c769b
[15] https://www.datascience.com/blog/what-is-semi-supervised-l earning
[16] https://spinningup.openai.com/en/latest/spinningup/rl_intro2. html#citations-below
[17] https://www.kdnuggets.com/2018/03/5-things-reinforcementlearning.html
[18] https://deepmind.com/blog/alphago-zero-learning-scratch/
[19] https://www.technologyreview.com/s/608155/teslas-new-ai-g uru-could-help-its-cars-teach-themselves/
[20] https://medium.com/app-affairs/9-applications-of-machine-le arning-from-day-to-day-life-112a47a429d0
[21] https://medium.com/machine-learning-for-humans/unsupervi sed-learning-f45587588294 [22] https://medium.com/deep-math-machine-learning-ai/different -types-of-machine-learning-and-their-types-34760b9128a2
[23] https://en.wikipedia.org/wiki/Reinforcement_learning

Amitrakshar International Journal
of Interdisciplinary and Transdisciplinary Research (AIJITR)
(A Social Science, Science and Indian Knowledge Systems Perspective)
Open-Access, Peer-Reviewed, Refereed, Bi-Monthly, International E-Journal

[24] https://en.wikipedia.org/wiki/Machine_learning#Processes_and_techniques
[25] https://www.edureka.co/blog/what-is-machine-learning/
[26] https://towardsdatascience.com/types-of-machine-learning-al gorithms-you-should-know-953a08248861
[27] https://towardsdatascience.com/kmeans-clustering-for-classif ication-74b992405d0a
[28]https://medium.freecodecamp.org/deep-reinforcement-learni ng-where-to-start-291fb0058c01
[29]https://towardsdatascience.com/a-tour-of-the-top-10-algorith ms-for-machine-learning-newbies-dde4edffae11
[30] https://en.wikipedia.org/wiki/Training,_validation,_and_test_ sets
[31] https://en.wikipedia.org/wiki/Supervised_learning
[32]https://machine learning mastery.com/supervised-and-unsupervised - machine-learning-algorithms/
[33] https://www.techopedia.com/definition/30389/supervised-learning
[34] http://dataaspirant.com/2014/09/19/supervised-and-unsupervi sed-learning/
[35] https://blog.algorithmia.com/introduction-to-unsupervised-learning/
[36] https://medium.com/machine-learning-for-humans/unsupervi sed-learning-f45587588294
[37] https://towardsdatascience.com/unsupervised-learning-with-p ython-173c51dc7f03
[38] https://dzone.com/articles/10-interesting-use-cases-for-the-kmeans-algorithm
[39] https://www.kdnuggets.com/2017/08/recommendation-syste m-algorithms-overview.html [40] https://en.wikipedia.org/wiki/Unsupervised_learning
[41] https://en.wikipedia.org/wiki/Semi-supervised_learning
[42]https://medium.freecodecamp.org/an-introduction-to-reinforc ement-learning-4339519de419
[43] https://skymind.ai/wiki/deep-reinforcement-learning
[44] https://blog.openai.com/spinning-up-in-deep-rl/
[45]https://www.theverge.com/2018/8/28/17787610/openai-dota2-bots-ai-lost-international-reinforcement-learning
[46] https://deepmind.com/research/publications/playing-atari-dee p-reinforcement-learning/
[47] http://karpathy.github.io/2016/05/31/rl/
[48] https://en.wikipedia.org/wiki/Reinforcement_learning
[49] https://en.wikipedia.org/wiki/Statistical_classification
[50] https://www.merriam-webster.com/dictionary/classification
[51] https://en.wikipedia.org/wiki/Regression_analysis
[52]https://medium.com/zylapp/review-of-deep-learning-algorith ms-for-object-detection-c1f3d437b852
[53] https://machinelearningmastery.com/linear-regression-for-ma chine-learning/
[54] https://www.geeksforgeeks.org/regression-classification-supe rvised-machine-learning/
[55]https://towardsdatascience.com/selecting-the-best-machine-learning-algorithm-for-your-regression-problem-20c330bad4ef