
Machine Learning: A Review of Learning Types

Shagan Sah¹

Abstract

In this paper, various machine learning techniques are discussed. These algorithms are used for many applications which include data classification, prediction, or pattern recognition. The primary goal of machine learning is to automate human assistance by training an algorithm on relevant data. This paper should also serve as a collection of various machine learning terminology for easy reference.

1. Introduction

Machine learning is the study of computer algorithms that provides systems the ability to automatically learn and improve from experience. It is generally seen as a sub-field of artificial intelligence. Machine learning algorithms allow the systems to make decisions autonomously without any external support. Such decisions are made by finding valuable underlying patterns within complex data.

Based on the learning approach, the type of data they input and output, and the type of problem that they solve, there are few primary categories of machine learning algorithms- supervised, unsupervised and reinforcement learning. There are a few hybrid approaches and other common methods that offer natural extrapolation of machine learning problem forms.

In the following sections, all the methods are briefly described. Recommended literature for further reading is also listed. This paper should also serve as a collection of various machine learning terminology for easy reference.

2. Primary Approaches

2.1. Supervised Learning

Supervised learning is applied when the data is in the form of input variables and output target values. The algorithm learns the mapping function from the input to the output.

¹California, USA. Correspondence to: Shagan Sah <sxs4337@rit.edu>.

The availability of large scale labeled data samples makes it an expensive approach for tasks where data is scarce. These approaches can be broadly divided into two main categories-

2.1.1. CLASSIFICATION

The output variable is one of some known number of categories. For example, "cat" or "dog", "positive" or "negative".

2.1.2. REGRESSION

The output variable is a real or a continuous value. For example, "price", "geographical location".

Further reading ([Kotsiantis et al., 2007](#)).




Figure 1. Overview of supervised learning. Input examples are categorized into a known set of classes.

2.2. Unsupervised Learning

Unsupervised learning is applied when the data is available only in the form of an input and there is no corresponding output variable. Such algorithms model the underlying patterns in the data in order to learn more about its characteristics.

One of the main types of unsupervised algorithms is clustering. In this technique, inherent groups in the data are discovered and then used to predict output for unseen inputs. An example of this technique would be to predict customer purchasing behavior.

Further reading ([Bengio et al., 2012](#)).

Machine Learning: A Review of Learning Types




Figure 2. Overview of unsupervised learning. Input samples are grouped into clusters based on the underlying patterns.

2.3. Reinforcement Learning

Reinforcement learning is applied when the task at hand is to make a sequence of decisions towards a final reward. During the learning process, an artificial agent gets either rewards or penalties for the actions it performs. Its goal is to maximize the total reward. Examples include learning agents to play computer games or performing robotics tasks with end goal.

Further reading (François-Lavet et al., 2018).




Figure 3. Overview of reinforcement learning. An agent observes the environment state and performs actions to maximize an overall reward.

3. Hybrid Approaches

3.1. Semi-supervised Learning

As the name suggests, this is an intermediate between supervised and unsupervised learning techniques. These algorithms are trained using a combination of labeled and unlabeled data. In a common setting, there is a small amount of labeled data and a very large amount of unlabeled data. A basic procedure involved is that first similar data is clustered using an unsupervised learning algorithm and then existing labeled data is used to label the rest of the unlabeled data.

Further reading (Chapelle et al., 2009).




Figure 4. Overview of semi-supervised learning. The clusters formed by a large amount of unlabeled data are used to classify a limited amount of labeled data.

3.2. Self-supervised Learning

Self-supervised learning is a form of unsupervised learning where the training data is autonomously (or automatically) labeled. The data is not required to be manually labelled but is labeled by finding and exploiting the relations (or correlations) between different input features. This is done in an unsupervised manner by forcing the network to learn semantic representation about the data. Knowledge is then transferred to the model for the main task. It is sometimes referred to as *pretext learning*.

Further reading (Jing & Tian, 2020).




Figure 5. Overview of self-supervised learning. A model is learned on unlabeled data (data is similar to the labeled data) using a dummy task and then the learned model is used for the main task.

3.3. Self-taught Learning

Self-taught learning is applicable in solving a supervised learning task given both labeled and unlabeled data, where the unlabeled data does not share the class labels or the generative distribution of the labeled data. In simple words, it applies transfer learning from unlabeled data. Once the representation has been learned in the first stage, it can then be applied repeatedly to different classification tasks.

Further reading (Raina et al., 2007).

Machine Learning: A Review of Learning Types




Figure 6. Overview of self-taught learning. A model is learned on unlabeled data (data may be from a dissimilar domain as the data used in main task) and then trained with small amounts of labeled data.

4. Other Common Approaches

4.1. Multi-task Learning

Multi-task learning refers to a training paradigm where multiple tasks are learned at the same time by a single model. This allows leveraging of useful relationships contained in related tasks. They improve generalization across all the tasks and hence improve prediction accuracy for specific tasks compared to models trained individually.

Further reading (Zhang & Yang, 2017).




Figure 7. Overview of multi-task learning. Model learning is achieved through multiple tasks that represent properties of the data.

4.2. Active Learning

This algorithm proactively selects a subset of data samples that it wants to learn from. The samples are selected from a large pool of unlabeled samples and are then labeled. This allows the algorithm to perform better than traditional methods with substantially less labeled data for training. Such methods are highly useful where unlabeled data may be abundant but labels are difficult, time-consuming, or expensive to obtain.

Further reading (Settles, 2009).




Figure 8. Overview of active learning. From a large pool of unlabeled data, a model selects the samples that it can learn most from for a required task. The selected data is labeled and then used to train the model.

4.3. Online Learning

Online learning involves training using data that becomes available in a sequential order. This technique contrasts with batch sampling based learning where the complete training data is always available. It is useful in scenarios where algorithms are required to dynamically adapt to novel data patterns from all incoming data.

4.3.1. INCREMENTAL LEARNING

Incremental learning strategy is very similar to (or at times same as) online learning. The main difference is that in online learning a training sample is used only once from an incoming data stream. In incremental learning, samples are usually picked from a finite dataset and the same samples can be processed multiple times.

4.3.2. SEQUENTIAL LEARNING

Sequential learning is a term widely used for learning with data that has a temporal ordering to it. Under certain conditions, it can be also interpreted as a type of online learning.

Further reading (Losing et al., 2018).




Figure 9. Overview of online learning. The model learns from a continuous incoming stream of data.

4.4. Transfer Learning

Transfer learning refers to training (or fine-tuning) a developed algorithm on a different yet related task. The main idea is about transferring knowledge from one supervised learning task to another and hence it generally requires further

Machine Learning: A Review of Learning Types

labeled data from a different but related task. One limitation of this approach is the requirement of additional labeled data, rather than unlabeled data, for the new supervised learning tasks.

Further reading (Weiss et al., 2016).




Figure 10. Overview of transfer learning. A model is learned on a dataset and the knowledge is transferred to a new task by retaining a part of the model.

4.5. Federated Learning

Federated learning enables training in a distributed manner using a large corpus of data residing on independent devices. It de-centralizes model training without sharing data samples among individual entities. This addresses the fundamental problems of privacy, ownership, and locality of data.

Further reading (Bonawitz et al., 2019).




Figure 11. Overview of federated learning. The data resides with individual entities which provides model updates to a centralized server without sharing their data.

4.6. Ensemble Learning

Ensemble learning is a machine learning paradigm where multiple learners are trained to solve the same problem. They obtain better predictive performance than could be obtained from any one of the constituent learning algorithms. An ensemble contains a number of learners which are usually called base learners. The generalization ability of an ensemble is usually much stronger than that of the base learners.

Further reading (Polikar, 2012).




Figure 12. Overview of ensemble learning. Multiple models are learned for the same task and their individual predictions are used to obtain the best result.

4.7. Adversarial Learning

In adversarial machine learning, a model is explicitly trained on a lot of adversarial data such that it is not fooled by those examples. When a standard machine learning model is deployed in the real world, it is susceptible to failures due to presence of intelligent and adaptive adversaries. This is because common machine learning techniques are designed for stationary environments where the training and test data are assumed to be generated from the same statistical distribution. Adversarial learning enhances the model capability against a malicious adversary by surreptitiously manipulating the input data.

Further reading (Lowd & Meek, 2005).




Figure 13. Overview of adversarial learning. The model is trained to discriminate between real and synthetic data samples.

4.8. Meta Learning

In a meta-learning paradigm, the machine learning model gains experience over multiple learning episodes that often cover a distribution of related tasks and then uses this experience to improve any future learning performance. The goal is to solve new tasks with only a small number of training samples. In contrast to conventional machine learning approaches where a given task is learned from scratch using a fixed learning algorithm, meta-learning aims to improve the learning algorithm itself, given the experience of multiple learning episodes, hence is also referred to as learn the learning process. Examples include *few-shot learning* and *metric learning*.

Machine Learning: A Review of Learning Types

Further reading (Hospedales et al., 2020).

4.8.1. METRIC LEARNING

Metric learning is a form of machine learning that utilizes distances between data samples. It learns from similarity or dis-similarity among the examples. It is often used for dimensionality reduction, recommendation systems, identity verification etc.

Further reading (Suárez et al., 2018).




Figure 14. Overview of meta learning. The model gains experience by learning over multiple learning episodes on related tasks before using the knowledge on the main task.

4.9. Targeted Learning

Targeted learning methods build machine learning based models that estimate the features of the probability distribution of the data. In simple words, they target the learning towards a certain parameter of interest. These methods are also used to obtain influence statistics about the model parameters. They are popular since the estimated parameter selection allows the subjective choices made by machines to mimic human behavior.

Further reading (van der Laan & Petersen, 2012).

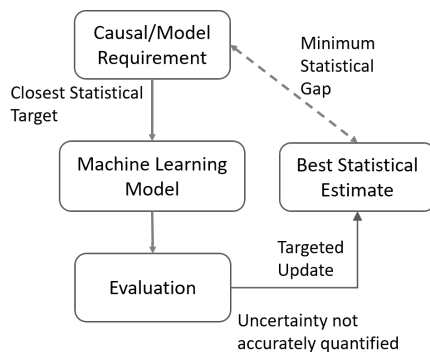


Figure 15. Overview of targeted learning. The model makes targeted updates to specific parameters that minimize the statistical goal.

4.10. Concept Learning

This approach involves learning from concepts to identify whether a sample belongs to a specific category or not. This

is done by processing the training data to find a hypothesis (or a function) that best fits the training examples. The goal is to classify a data point as either belonging to or not belonging to a particular concept or idea. In this context, a concept can be viewed as a boolean-valued function defined over a large data set. A common approach is using the *Find-S Algorithm*.

Further reading (Mitchell et al., 1997).




Figure 16. Overview of concept learning. The model finds the best hypothesis that satisfies all the boolean concepts in the data.

4.11. Bayesian Learning

Bayesian learning uses *Bayes'* theorem to determine the conditional probability of a hypotheses given some evidence or observations. In contrast to maximum likelihood learning, Bayesian learning explicitly models uncertainty over both the input data and the model parameters. The initial or prior knowledge is incorporated through a distribution over the parameters.

Further reading (Bernardo & Smith, 2009).




Figure 17. Overview of Bayesian learning. The model used initial knowledge (prior) and the data observations to determine the conditional probability of for data using *Bayes'* theorem.

4.12. Analytical Learning

The goal is to use logical reasoning to identify features that can distinguish among different input examples. It is a non-statistical learning approach that allows a learner to process

Machine Learning: A Review of Learning Types

information, break it into component parts (features), and generate hypotheses by using critical and logical thinking skills. These approaches analyze each problem instance individually, rather than a set of problem instances. Such approaches do not require large amounts of training data to work well.

4.12.1. INDUCTIVE LEARNING

The goal is to use statistical reasoning to identify features that empirically distinguish different input examples. The performance is highly dependent on the number of training samples.

Further reading (Kawaguchi et al., 2019; Ruiz, 2012).




Figure 18. Overview of analytical and inductive learning. This terminology is used to distinguish between models learning using logical or statistical reasoning.

4.13. Multi-modal Learning

These are types of algorithms that learn features over multiple modalities. Examples of modalities include visual, auditory, kinesthetic among other sensory data. By combining such modes, learners are able to combine information from different sources and hence yield better feature extraction and predictions at a large scale.

Further reading (Baltrušaitis et al., 2018).




Figure 19. Overview of multi-modal learning. The model is learned using data from multiple modalities to exploit their relationships.

4.14. Deep Learning

Deep learning is a technique to implement various machine learning algorithms using multi-layers neural networks. These multiple processing layers learn representations of data with multiple levels of abstraction for understanding the input data.

Further reading (LeCun et al., 2015).




Figure 20. Overview of deep learning. A term used for a multi-layered neural network that learns feature extraction and classification (or other discrimination task) in an end-to-end manner.

4.15. Curriculum Learning

In the curriculum learning paradigm, the training data is organized in a meaningful order which gradually illustrates more complex concepts. The idea is analogous to human learning in an organized education system that introduces different concepts at different times. This technique allows exploitation of previously learned concepts to ease the learning of new abstractions.

Further reading (Bengio et al., 2009).




Figure 21. Overview of curriculum learning. The model is learned in stages where data is organized in an meaningful order such that the complexity gradually increases.

References

- Baltrušaitis, T., Ahuja, C., and Morency, L.-P. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41 (2):423–443, 2018.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.

Machine Learning: A Review of Learning Types

- Bengio, Y., Courville, A. C., and Vincent, P. Unsupervised feature learning and deep learning: A review and new perspectives. *CoRR, abs/1206.5538*, 1:2012, 2012.
- Bernardo, J. M. and Smith, A. F. *Bayesian theory*, volume 405. John Wiley & Sons, 2009.
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konecny, J., Mazzocchi, S., McMahan, H. B., et al. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046*, 2019.
- Chapelle, O., Scholkopf, B., and Zien, A. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., and Pineau, J. An introduction to deep reinforcement learning. *arXiv preprint arXiv:1811.12560*, 2018.
- Hospedales, T., Antoniou, A., Micaelli, P., and Storkey, A. Meta-learning in neural networks: A survey. *arXiv preprint arXiv:2004.05439*, 2020.
- Jing, L. and Tian, Y. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Kawaguchi, K., Bengio, Y., Verma, V., and Kaelbling, L. P. Generalization in machine learning via analytical learning theory. *stat*, 1050:6, 2019.
- Kotsiantis, S. B., Zaharakis, I., and Pintelas, P. Supervised machine learning: A review of classification techniques. *Emerging artificial intelligence applications in computer engineering*, 160:3–24, 2007.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *nature*, 521(7553):436–444, 2015.
- Losing, V., Hammer, B., and Wersing, H. Incremental online learning: A review and comparison of state of the art algorithms. *Neurocomputing*, 275:1261–1274, 2018.
- Lowd, D. and Meek, C. Adversarial learning. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pp. 641–647, 2005.
- Mitchell, T. M. et al. Machine learning. 1997. *Burr Ridge, IL: McGraw Hill*, 45(37):870–877, 1997.
- Polikar, R. Ensemble learning. In *Ensemble machine learning*, pp. 1–34. Springer, 2012.
- Raina, R., Battle, A., Lee, H., Packer, B., and Ng, A. Y. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th international conference on Machine learning*, pp. 759–766, 2007.
- Ruiz, C. Using domain knowledge in machine learning inductive vs analytical learning, 2012. URL http://web.cs.wpi.edu/~cs539/f12/Slides/ilp_ebl.pdf.
- Settles, B. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- Suárez, J. L., García, S., and Herrera, F. A tutorial on distance metric learning: mathematical foundations, algorithms and software. *arXiv preprint arXiv:1812.05944*, 2018.
- van der Laan, M. J. and Petersen, M. L. Targeted learning. In *Ensemble Machine Learning*, pp. 117–156. Springer, 2012.
- Weiss, K., Khoshgoftaar, T. M., and Wang, D. A survey of transfer learning. *Journal of Big data*, 3(1):9, 2016.
- Zhang, Y. and Yang, Q. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*, 2017.