

SUPERVISED LEARNING

Qiong Liu
FX Palo Alto Laboratory
Palo Alto, California
United States of America
liu@fxpal.com

Ying Wu
EECS Department
Northwestern University
United States of America
yingwu@eecs.northwestern.edu

Synonyms

Supervised Machine Learning; Learning with a Teacher; Learning from Labeled Data; Regression; Classification; Inductive Machine Learning; Active Learning; Semi-supervised Learning;

Definition

Supervised Learning is a machine learning paradigm for acquiring the input-output relationship information of a system based on a given set of paired input-output training samples. As the output is regarded as the label of the input data or the supervision, an input-output training sample is also called labelled training data, or supervised data. Occasionally, it is also referred to as Learning with a Teacher (Haykin 1998), Learning from Labelled Data, or Inductive Machine Learning (Kotsiantis, 2007). The goal of supervised learning is to build an artificial system that can learn the mapping between the input and the output, and can predict the output of the system given new inputs. If the output takes a finite set of discrete values that indicate the class labels of the input, the learned mapping leads to the classification of the input data. If the output takes continuous values, it leads to a regression of the input. The input-output relationship information is frequently represented with learning-model parameters. When these parameters are not directly available from training samples, a learning system needs to go through an estimation process to obtain these parameters. Different from Unsupervised Learning, the training data for Supervised Learning need supervised or labelled information, while the training data for unsupervised learning are unsupervised as they are not labelled (i.e., merely the inputs). If an algorithm uses both supervised and unsupervised training data, it is called a Semi-supervised Learning algorithm. If an algorithm actively queries a user/teacher for labels in the training process, the iterative supervised learning is called Active Learning.

Overview




Figure 1. Block diagram that illustrates the form of Supervised Learning

Figure 1 shows a block diagram that illustrates the form of Supervised Learning. In this diagram, (x_i, y_i) is a supervised training sample, where 'x' represents system input, 'y' represents the system output (i.e., the supervision or labelling of the input x), and 'i' is the index of the training sample. During a Supervised Learning process, a training input x_i is fed to the Learning System, and the Learning System generates an output \tilde{y}_i . The Learning System output \tilde{y}_i is then compared with the ground truth labeling y_i by an arbitrator that computes the difference between them. The difference, termed Error Signal in this diagram, is then sent to the Learning System for adjusting the parameters of the learner. The goal of this learning process is to obtain a set of optimal Learning System parameters that can minimize the differences between \tilde{y}_i and y_i for all i , i.e., minimizing the total error over the entire training data set.

A notable phenomenon is that a minimum training error does not necessarily indicate a good performance in testing. Training is referred to as the learning process that estimates the parameters of the learner based on the ground truth supervised data seen, while testing is to evaluate the predictions of the learner for the data unseen, i.e., the data used in testing have not been included in the training process. Therefore, even if a learner achieves a minimum error on the set of training data, it does not guarantee to perform well to the data unseen. The reason for this is mainly due to the possible overfitting to the training data, i.e., the learner has an unnecessary order of complexity in learning the mapping. This issue is referred to as the generalizability. A good learning algorithm must have a good generalizability. To take into consideration of the generalizability in designing the learner, a learning algorithm needs to balance the objective of minimizing the training error and the complexity of the learner (e.g., the structure and the order of the learner). For example, in the Support Vector Machine, the generalizability of the learner is characterized by the margin of the learned discrimination boundary. The larger the margin, the better the generalization. The support vector machine learns a maximum margin classifier over the training set, and thus it naturally leads to a good generalization performance (Vapnik, 1995).

The Supervised Learning paradigm does not restrict the sources of the input or output data. The input or output may belong to a vector space or a set of discrete values. The learning paradigm does not have special restrictions on the arbitrator either. If y_i is drawn from a continuous space, the error signal is usually computed via $y_i - \tilde{y}_i$. If y_i belongs to a set of discrete values, the arbitrator usually outputs the error signal based on the equality between y_i and \tilde{y}_i . For example, the arbitrator may output 0 for equalled y_i and \tilde{y}_i and output 1 for different y_i and \tilde{y}_i .

There are different approaches to the design of the Learning System in Supervised Learning. Some well-known approaches include the logic-based approach, the multi-layer perceptron approach, the statistical-learning approach, the instance-based learning approach, the Support Vector Machines, and Boosting.

Advantages and Disadvantages

The foremost advantage of Supervised Learning is that all classes or analog outputs manipulated by the algorithm of this paradigm are meaningful to humans. And it can be easily used for discriminative pattern classifi-

cation, and for data regression. But it also has several disadvantages. The first one is caused by the difficulty of collecting supervision or labels. When there is a huge volume of input data, it is prohibitively expensive, if not impossible, to label all of them.. For example, it is not a trivial task to label a huge set of images for image classification. Second, as not everything in the real world has a distinctive label, there are uncertainties and ambiguities in the supervision or labels. For example, the margin for separating the two concepts of “hot” and “cold” is not distinct; and it is difficult to name an object that is a cross between a loveseat and a bed. These difficulties may limit the applications of the Supervised Learning paradigm in some scenarios. To overcome these limitations in practice, other learning paradigms, such as Unsupervised Learning, Semi-supervised Learning, Reinforcement Learning, Active Learning, or some mixed learning approaches can be considered.

Applications

Supervised Learning enables a machine to learn the human behaviour or object behaviour in certain tasks. The learned knowledge can then be used by the machine to perform similar actions on these tasks. Since the computing machinery may perform some input-output mappings much faster and more persistent than the human, machines equipped with a good supervised learner can perform certain tasks much faster and accurate than the human. On the other hand, because of the limitation in hardware, software, and algorithm designs, existing Supervised Learning algorithms still cannot match human’s learning ability on many complicated tasks.

Supervised Learning have been successfully used in areas such as Information Retrieval, Data Mining, Computer Vision, Speech Recognition, Spam Detection, Bioinformatics, Cheminformatics, and Market Analysis (Wikipedia, 2010).

Cross-Reference

Interactive learning, Machine Learning from pairwise relationships, Adaptation and unsupervised learning, Feature selection (unsupervised learning), Classification of learning objects.

References

Haykin, S. 1998 Neural Networks: a Comprehensive Foundation. 2nd. Prentice Hall PTR. ISBN 0-132-73350-1

Vapnik, V. The Nature of Statistical Learning Theory. Springer-Verlag, 1995. ISBN 0-387-98780-0

Kotsiantis, S. Supervised Machine Learning: A Review of Classification Techniques, Informatica Journal 31 (2007) 249-268

Wikipedia, Supervised learning, http://en.wikipedia.org/wiki/Supervised_learning