

"Machine Learning"

Wikipedia

Machine learning, a branch of artificial intelligence, is a scientific discipline concerned with the design and development of algorithms that allow computers to evolve behaviors based on empirical data, such as from sensor data or databases. A learner can take advantage of examples (data) to capture characteristics of interest of their unknown underlying probability distribution. Data can be seen as examples that illustrate relations between observed variables. A major focus of machine learning research is to automatically learn to recognize complex patterns and make intelligent decisions based on data; the difficulty lies in the fact that the set of all possible behaviors given all possible inputs is too large to be covered by the set of observed examples (training data). Hence the learner must generalize from the given examples, so as to be able to produce a useful output in new cases. Machine learning, like all subjects in artificial intelligence, requires cross-disciplinary proficiency in several areas, such as probability theory, statistics, pattern recognition, cognitive science, data mining, adaptive control, computational neuroscience and theoretical computer science.

Definition

A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .^[1]

Generalization

The core objective of a learner is to generalize from its experience.^[2] The training examples from its experience come from some generally unknown probability distribution and the learner has to extract from them something more general, something about that distribution, that allows it to produce useful answers in new cases.

Human interaction

Some machine learning systems attempt to eliminate the need for human intuition in data analysis, while others adopt a collaborative approach between human and machine. Human intuition cannot, however, be entirely eliminated, since the system's designer must specify how the data is to be represented and what mechanisms will be used to search for a characterization of the data.

Algorithm types

Machine learning algorithms are organized into a taxonomy, based on the desired outcome of the algorithm.

- **Supervised learning** generates a function that maps inputs to desired outputs. For example, in a classification problem, the learner approximates a function mapping a vector into classes by looking at input-output examples of the function.
- **Unsupervised learning** models a set of inputs, like clustering.
- **Semi-supervised learning** combines both labeled and unlabeled examples to generate an appropriate function or classifier.
- **Reinforcement learning** learns how to act given an observation of the world. Every action has some impact in the environment, and the environment provides feedback in the form of rewards that guides the learning algorithm.
- **Transduction** tries to predict new outputs based on training inputs, training outputs, and test inputs.
- **Learning to learn** learns its own inductive bias based on previous experience.

Theory

The computational analysis of machine learning algorithms and their performance is a branch of theoretical computer science known as computational learning theory. Because training sets are finite and the future is uncertain, learning theory usually does not yield absolute guarantees of the performance of algorithms. Instead, probabilistic bounds on the performance are quite common.

In addition to performance bounds, computational learning theorists study the time complexity and feasibility of learning. In computational learning theory, a computation is considered feasible if it can be done in polynomial time. There are two kinds of time complexity results. Positive results show that a certain class of functions can be learned in polynomial time. Negative results show that certain classes cannot be learned in polynomial time.

There are many similarities between machine learning theory and statistics, although they use different terms.

Approaches

Decision tree learning

Decision tree learning uses a decision tree as a predictive model which maps observations about an item to conclusions about the item's target value.

Association rule learning

Association rule learning is a method for discovering interesting relations between variables in large databases.

Artificial neural networks

An artificial neural network (ANN), usually called "neural network" (NN), is a mathematical model or computational model that tries to simulate the structure and/or functional aspects of biological neural networks. It consists of an interconnected group of artificial neurons and processes information using a connectionist approach to computation. Modern neural networks are non-linear statistical data modeling tools. They are usually used to model complex relationships between inputs and outputs or to find patterns in data.

Genetic programming

Genetic programming (GP) is an evolutionary algorithm-based methodology inspired by biological evolution to find computer programs that perform a user-defined task. It is a specialization of genetic algorithms (GA) where each individual is a computer program. It is a machine learning technique used to optimize a population of computer programs according to a fitness landscape determined by a program's ability to perform a given computational task.

Inductive logic programming

Inductive logic programming (ILP) is an approach to rule learning using logic programming as a uniform representation for examples, background knowledge, and hypotheses. Given an encoding of the known background knowledge and a set of examples represented as a logical database of facts, an ILP system will derive a hypothesized logic program which entails all the positive and none of the negative examples.

Support vector machines

Support vector machines (SVMs) are a set of related supervised learning methods used for classification and regression. Given a set of training examples, each marked as belonging to one of two categories, an SVM training algorithm builds a model that predicts whether a new example falls into one category or the other.

Clustering

Cluster analysis or clustering is the assignment of a set of observations into subsets (called *clusters*) so that observations in the same cluster are similar in some sense. Clustering is a method of unsupervised learning, and a common technique for statistical data analysis.

Bayesian networks

A Bayesian network, belief network or directed acyclic graphical model is a probabilistic graphical model that represents a set of random variables and their conditional independencies via a directed acyclic graph (DAG). For example, a Bayesian network could represent the probabilistic relationships between diseases and symptoms. Given symptoms, the network can be used to compute the probabilities of the presence of various diseases. Efficient algorithms exist that perform inference and learning.

Reinforcement learning

Reinforcement learning is concerned with how an *agent* ought to take *actions* in an *environment* so as to maximize some notion of long-term *reward*. Reinforcement learning algorithms attempt to find a *policy* that maps *states* of the world to the actions the agent ought to take in those states. Reinforcement learning differs from the supervised learning problem in that correct input/output pairs are never presented, nor sub-optimal actions explicitly corrected.

Applications

Applications for machine learning include machine perception, computer vision, natural language processing, syntactic pattern recognition, search engines, medical diagnosis, bioinformatics, brain-machine interfaces and cheminformatics, detecting credit card fraud, stock market analysis, classifying DNA sequences, speech and handwriting recognition, object recognition in computer vision, game playing, software engineering, adaptive websites, robot locomotion, and structural health monitoring.

Machine learning techniques helped win a major software competition: in 2006, the online movie company Netflix held the first "Netflix Prize" competition to find a program to better predict user preferences and beat its existing Netflix movie recommendation system by at least 10%. The AT&T Research Team BellKor won over several other teams with their machine learning program called Pragmatic Chaos. After winning several minor prizes, it won the 2009 grand prize competition for \$1 million.^[3]

Software

RapidMiner, KNIME, Weka, ODM, Shogun toolbox, Orange and Apache Mahout are software suites containing a variety of machine learning algorithms.

Journals and conferences

- *Machine Learning* (journal)
- *Journal of Machine Learning Research*
- *Neural Computation* (journal)
- International Conference on Machine Learning (ICML) (conference)
- Neural Information Processing Systems (NIPS) (conference)
- List of upcoming conferences in Machine Learning and Artificial Intelligence ^[4] (conference)

References

- [1] Tom M. Mitchell (1997) *Machine Learning* p.2
- [2] Christopher M. Bishop (2006) *Pattern Recognition and Machine Learning*, Springer ISBN 0-387-31073-8.
- [3] "BelKor Home Page" (<http://www2.research.att.com/~volinsky/netflix/>) research.att.com
- [4] <http://sites.google.com/site/fawadsyed/upcoming-conferences>

Further reading

- Sergios Theodoridis, Konstantinos Koutroumbas (2009) "Pattern Recognition", 4th Edition, Academic Press, ISBN 978-1-59749-272-0.
- Ethem Alpaydm (2004) *Introduction to Machine Learning (Adaptive Computation and Machine Learning)*, MIT Press, ISBN 0262012111
- Bing Liu (2007), *Web Data Mining: Exploring Hyperlinks, Contents and Usage Data* (<http://www.cs.uic.edu/~liub/WebMiningBook.html>). Springer, ISBN 3540378812
- Toby Segaran, *Programming Collective Intelligence*, O'Reilly ISBN 0-596-52932-5
- Ray Solomonoff, " An Inductive Inference Machine (<http://world.std.com/~rjs/indinf56.pdf>)" A privately circulated report from the 1956 Dartmouth Summer Research Conference on AI.
- Ray Solomonoff, *An Inductive Inference Machine*, IRE Convention Record, Section on Information Theory, Part 2, pp., 56-62, 1957.
- Ryszard S. Michalski, Jaime G. Carbonell, Tom M. Mitchell (1983), *Machine Learning: An Artificial Intelligence Approach*, Tioga Publishing Company, ISBN 0-935382-05-4.
- Ryszard S. Michalski, Jaime G. Carbonell, Tom M. Mitchell (1986), *Machine Learning: An Artificial Intelligence Approach, Volume II*, Morgan Kaufmann, ISBN 0-934613-00-1.
- Yves Kodratoff, Ryszard S. Michalski (1990), *Machine Learning: An Artificial Intelligence Approach, Volume III*, Morgan Kaufmann, ISBN 1-55860-119-8.
- Ryszard S. Michalski, George Tecuci (1994), *Machine Learning: A Multistrategy Approach, Volume IV*, Morgan Kaufmann, ISBN 1-55860-251-8.
- Bhagat, P.M. (2005). *Pattern Recognition in Industry*, Elsevier. ISBN 0-08-044538-1.
- Bishop, C.M. (1995). *Neural Networks for Pattern Recognition*, Oxford University Press. ISBN 0-19-853864-2.
- Richard O. Duda, Peter E. Hart, David G. Stork (2001) *Pattern classification* (2nd edition), Wiley, New York, ISBN 0-471-05669-3.
- Huang T.-M., Kecman V., Kopriva I. (2006), *Kernel Based Algorithms for Mining Huge Data Sets, Supervised, Semi-supervised, and Unsupervised Learning* (<http://learning-from-data.com>), Springer-Verlag, Berlin, Heidelberg, 260 pp. 96 illus., Hardcover, ISBN 3-540-31681-7.

- KECEMAN Vojislav (2001), Learning and Soft Computing, Support Vector Machines, Neural Networks and Fuzzy Logic Models (<http://support-vector.ws>), The MIT Press, Cambridge, MA, 608 pp., 268 illus., ISBN 0-262-11255-8.
- MacKay, D.J.C. (2003). *Information Theory, Inference, and Learning Algorithms* (<http://www.inference.phy.cam.ac.uk/mackay/itila/>), Cambridge University Press. ISBN 0-521-64298-1.
- Mitchell, T. (1997). *Machine Learning*, McGraw Hill. ISBN 0-07-042807-7.
- Ian H. Witten and Eibe Frank *Data Mining: Practical machine learning tools and techniques* Morgan Kaufmann ISBN 0-12-088407-0.
- Sholom Weiss and Casimir Kulikowski (1991). *Computer Systems That Learn*, Morgan Kaufmann. ISBN 1-55860-065-5.
- Mierswa, Ingo and Wurst, Michael and Klinkenberg, Ralf and Scholz, Martin and Euler, Timm: *YALE: Rapid Prototyping for Complex Data Mining Tasks*, in Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-06), 2006.
- Trevor Hastie, Robert Tibshirani and Jerome Friedman (2001). *The Elements of Statistical Learning* (<http://www-stat.stanford.edu/~tibs/ElemStatLearn/>), Springer. ISBN 0387952845.
- Vladimir Vapnik (1998). *Statistical Learning Theory*. Wiley-Interscience, ISBN 0471030031.

External links

- Ruby implementations of several machine learning algorithms (<http://ai4r.rubyforge.org>)
- Andrew Ng's Stanford lectures and course materials (<http://see.stanford.edu/see/courseinfo.aspx?coll=348ca38a-3a6d-4052-937d-cb017338d7b1>)
- The Encyclopedia of Computational Intelligence (http://scholarpedia.org/article/Encyclopedia_of_Computational_Intelligence)
- International Machine Learning Society (<http://machinelearning.org/>)
- Kmining List of machine learning, data mining and KDD scientific conferences (http://kmining.com/info_conferences.html)
- Machine Learning Open Source Software (<http://mloss.org/about/>)
- Machine Learning Video Lectures (http://videlectures.net/Top/Computer_Science/Machine_Learning/)
- Open Source Artificial Learning Software (<http://salproject.codeplex.com>)
- The Computational Intelligence and Machine Learning Virtual Community (<http://cimlcommunity.org/>)
- R Machine Learning Task View (<http://cran.r-project.org/web/views/MachineLearning.html>)
- Machine Learning Links and Resources (<http://www.reddit.com/r/MachineLearning/>)
- Idiap Research Institute - Links and Resources (<http://www.idiap.ch/scientific-research/themes/machine-learning/>)
- OpenCV implementations of several machine learning algorithms (http://opencv.willowgarage.com/documentation/cpp/ml__machine_learning.html) part of OpenCV (open source computer vision library) (<http://opencv.willowgarage.com/wiki/>)

Article Sources and Contributors

Machine learning *Source:* <http://en.wikipedia.org/w/index.php?oldid=415826400> *Contributors:* APH, AXRL, Aaron Kauppi, Aaronbrick, Aceituno, Adiel, Adoniscik, Ahoerstermeier, Ahyeek, Aiwing, AnAj, André P Ricardo, Arcenciel, Arvindn, Ataulf, BMF81, Baguasquirrel, Beetstra, BenKovitz, BertSeghers, Biochaos, BlaiseFEgan, Blaz.zupan, Bonadea, Boxplot, Bumbulski, Buridan, Businessman332211, Calltech, Casia wyq, Celendin, Centrx, Cfallin, ChangChienFu, ChaoticLogic, CharlesGillingham, Chriblo, Chris the speller, Chrisonaall, Clemwang, Clickey, Cmbishop, ComodiCast, Ctacmo, CultureDrone, Cvdwalt, Damienfrancois, Dana2020, Dancter, Darnelr, Dave Runger, DaveWF, DavidCBryant, Debejyo, Defza, Delirium, Denoir, Devantheryvv, Dicklyon, Dsilver, Dzkd, Edouard.darchimbaud, Essjay, Evansad, Examtester, Fabiform, FidesLT, Fram, Funandtrvl, Furrykef, Gareth Jones, Gene s, Genius002, Giftlite, GordonRoss, Grafen, Graytay, Gtfjbl, Haham hanuka, Helwr, Hike395, Innohead, Intgr, IradBG, J04n, James Kidd, Jbmurray, Jcautilli, Jdizzle123, Jim15936, JimmyShelter, Jmartinezot, Joerg Kurt Wegner, JonHarder, Jrennie, Jrjrl, Jroudh, Jwojt, Jyoshimi, KYN, Keefaas, KellyCoinGuy, Khalid hassani, Kinimod, Kku, KnightRider, Kumioko, Kyhui, L Kensington, Lars Washington, Levin, LokiClock, Lordvolton, Lovok Sovok, MTJM, Masatran, Mdd, Mereda, Michael Hardy, Misterwindupbird, Mneser, Moorejh, Movado73, MrOllie, Mxn, Nesbit, Netalarm, Nk, NotARuski, Nowozin, Pebkac, Penguinbroker, Peterjones, Pgr94, Philpraxis, Piano non troppo, Pintaio, Pmbhagat, Pranjic973, Predictor, Proffviktork, Quebec99, QuickUkie, Quintopia, RJASE1, Rajah, Ralf Klinkenberg, Redgecko, RexSurvey, Rjwilmsi, Ronz, Ruud Koot, Ryszard Michalski, Scigrex14, Scorpion451, Seabcan, Seaphoto, Sebastjanmm, Shirik, Shizhao, Silvonen, Sina2, Soultaco, Spiral5800, Srinivasasha, StaticGull, Stephen Turner, Superbacana, Swordsmankirby, Tedickey, Topbanana, Trondtr, Utcursch, VKokielov, Vilapi, Vivohobson, WMod-NS, Webidiap, WhatWasDone, Wht43, Why Not A Duck, Wikinacious, Winnerdy, WinterSpw, Wjbean, Wrdieter, Yoshua.Bengio, YrPolishUncle, Yworo, ZeroOne, Zosoin, Иъ Лю Ха, 287 anonymous edits

License

Creative Commons Attribution-Share Alike 3.0 Unported
<http://creativecommons.org/licenses/by-sa/3.0/>