Applications and Limitations of Machine Learning Process

S. P. Meena*
Poornima College of Engineering, Jaipur, Rajasthan, India
***Corresponding Author**
Email Id: meena.sp0035@gmail.com

ABSTRACT

Machine learning is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it to learn for themselves.

***Keywords**: machine learning, algorithm, artificial intelligence, ML.*

INTRODUCTION

Machine learning can help with the diagnosis of diseases. Many physicians use chatbots with speech recognition capabilities to discern patterns in symptoms. Real-world examples for medical diagnosis: Assisting in formulating a diagnosis or recommends a treatment option [1].

Machine learning and deep learning are widely used in many domains to name a few: Medical: For cancer cell detection, brain MRI image restoration, gene printing, etc. Document: Super-resolving historical document images, segmenting text in document images. Banks: Stock prediction, financial decisions [2]. Simply put, machine learning allows the user to feed a computer algorithm an immense amount of data and have the computer analyze and make data-driven recommendations and decisions based on only the input data [3].

APPROACHES OF MACHINE LEARNING

Machine learning approaches are traditionally divided into three broad categories, depending on the nature of the "signal" or "feedback" available to the learning system:

1) Supervised learning: The computer is presented with example inputs and their desired outputs, given by a "teacher", and the goal is to learn a general rule that maps inputs to outputs.
2) Unsupervised learning: No labels are given to the learning algorithm, leaving it on its own to find structure in its input. Unsupervised learning can be a goal in itself (discovering hidden patterns in data) or a means towards an end (feature learning).
3) Reinforcement learning: A computer program interacts with a dynamic environment in which it must perform a certain goal (such as driving a vehicle or playing a game against an opponent). As it navigates its problem space, the program is provided feedback that's analogous to rewards, which it tries to maximize [4].

Supervised Learning

A support-vector machine is a supervised learning model that divides the data into regions separated by a linear boundary. Here, the linear boundary divides the black circles from the white [5].

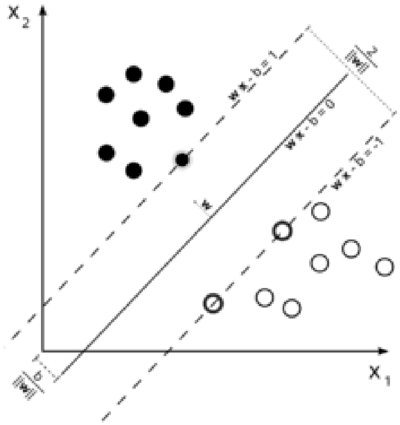

Supervised learning algorithms build a mathematical model of a set of data that contains both the inputs and the desired outputs.[34] The data is known as training data, and consists of a set of training examples. Each training example has one or more inputs and the desired output, also known as a supervisory signal. In the mathematical model, each training example is represented by an array or vector, sometimes called a feature vector, and the training data is represented by a matrix. Through iterative optimization of an objective function, supervised learning algorithms learn a function that can be used to predict the output associated with new inputs.[35] An optimal function will allow the algorithm to correctly determine the output for inputs that were not a part of the training data. An algorithm that improves the accuracy of its outputs or predictions over time is said to have learned to perform that task [6].

Types of supervised learning algorithms include active learning, classification and regression. Classification algorithms are used when the outputs are restricted to a limited set of values, and regression algorithms are used when the outputs may have any numerical value within a range. As an example, for a classification algorithm that filters emails, the input would be an incoming email, and the output would be the name of the folder in which to file the email.

Similarity learning is an area of supervised machine learning closely related to regression and classification, but the goal is to learn from examples using a similarity function that measures how similar or related two objects are. It has applications in ranking, recommendation systems, visual identity tracking, face verification, and speaker verification [7].

Unsupervised Learning

Unsupervised learning algorithms take a set of data that contains only inputs, and find structure in the data, like grouping or clustering of data points. The algorithms, therefore, learn from test data that has not been labeled, classified or categorized. Instead of responding to feedback, unsupervised learning algorithms identify commonalities in the data and react based on the presence or absence of such commonalities in each new piece of data. A central application of unsupervised learning is in the field of density estimation in statistics, such as finding the probability density function.[36] Though unsupervised learning encompasses other domains involving summarizing and explaining data features [8].

Cluster analysis is the assignment of a set of observations into subsets (called *clusters*) so that observations within the same cluster are similar according to one or more predesignated criteria, while observations drawn from different clusters are dissimilar. Different clustering techniques make different assumptions on the structure of the data, often defined by some similarity metric and evaluated, for example, by *internal compactness*, or the similarity between members of the same cluster, and *separation*, the difference between clusters. Other methods are based on estimated density and graph connectivity [8].

Semi-Supervised Learning

Semi-supervised learning falls between unsupervised learning (without any labeled training data) and supervised learning (with completely labeled training data). Some of the training examples are missing training labels, yet many machine-learning researchers have found that unlabeled data, when used in conjunction with a small amount of labeled data, can produce a considerable improvement in learning accuracy. In weakly supervised learning, the training labels are noisy, limited, or imprecise; however, these labels are often cheaper to obtain, resulting in larger effective training sets [9].

Reinforcement learning

Reinforcement learning is an area of machine learning concerned with how software agents ought to take actions in an environment so as to maximize some notion of cumulative reward. Due to its generality, the field is studied in many other disciplines, such as game theory, control theory, operations research, information theory, simulation-based optimization, multi-agent systems, swarm intelligence, statistics and genetic algorithms. In machine learning, the environment is typically represented as a Markov decision process (MDP). Many reinforcement learning algorithms use dynamic programming techniques. Reinforcement learning algorithms do not assume knowledge of an exact mathematical model of the MDP, and are used when exact models are infeasible. Reinforcement learning algorithms are used in autonomous vehicles or in learning to play a game against a human opponent [9].

APPLICATIONS OF MACHINE LEARNING PROCESS [10]

There are many applications for machine learning, including:

1) Agriculture
2) Anatomy
3) Adaptive website
4) Affective computing
5) Astronomy
6) Banking
7) Bioinformatics
8) Brain–machine interfaces
9) Cheminformatics
10) Citizen science
11) Computer networks
12) Computer vision
13) Credit-card fraud detection
14) Data quality
15) DNA sequence classification
16) Economics
17) Financial market analysis
18) General game playing
19) Handwriting recognition
20) Information retrieval

21) Insurance
22) Internet fraud detection
23) Knowledge graph embedding
24) Linguistics
25) Machine learning control
26) Machine perception
27) Machine translation
28) Marketing
29) Medical diagnosis
30) Natural language processing
31) Natural language understanding
32) Online advertising
33) Optimization
34) Recommender systems
35) Robot locomotion
36) Search engines
37) Sentiment analysis
38) Sequence mining
39) Software engineering
40) Speech recognition
41) Structural health monitoring
42) Syntactic pattern recognition
43) Telecommunication
44) Theorem proving
45) Time-series forecasting
46) User behavior analytics
47) Behaviorism

In 2006, the media-services provider Netflix held the first "Netflix Prize" competition to find a program to better predict user preferences and improve the accuracy of its existing Cinematch movie recommendation algorithm by at least 10%. A joint team made up of researchers from AT&T Labs-Research in collaboration with the teams Big Chaos and Pragmatic Theory built an ensemble model to win the Grand Prize in 2009 for $1 million. Shortly after the prize was awarded, Netflix realized that viewers' ratings were not the best indicators of their viewing patterns ("everything is a recommendation") and they changed their recommendation engine accordingly [10].

In 2010 The Wall Street Journal wrote about the firm Rebellion Research and their use of machine learning to predict the financial crisis. In 2012, co-founder of Sun Microsystems, Vinod Khosla, predicted that 80% of medical doctors' jobs would be lost in the next two decades to automated machine learning medical diagnostic software. In 2014, it was reported that a machine learning algorithm had been applied in the field of art history to study fine art paintings and that it may have revealed previously unrecognized influences among artists. In 2019 Springer Nature published the first research book created using machine learning. In 2020, machine learning technology was used to help make diagnoses and aid researchers in developing a cure for COVID-19. Machine learning is recently applied to predict the green behavior of human-being. Recently, machine learning technology is also applied to optimise smartphone's performance and thermal behaviour based on the user's interaction with the phone [11].

ADVANTAGES OF MACHINE LEARNING [12]

1) Automation of Everything. Machine Learning is responsible for cutting the workload and time.
2) Wide Range of Applications.
3) Scope of Improvement.
4) Efficient Handling of Data. ...
5) Best for Education and Online Shopping.
6) Possibility of High Error.
7) Algorithm Selection.
8) Data Acquisition.

LIMITATIONS OF MACHINE LEARNING PROCESS

Although machine learning has been transformative in some fields, machine-learning programs often fail to deliver expected results. Reasons for this are numerous: lack of (suitable) data, lack of access to the data, data bias, privacy problems, badly chosen tasks and algorithms, wrong tools and people, lack of resources, and evaluation problems. In 2018, a self-driving car from Uber failed to detect a pedestrian, who was killed after a collision. Attempts to use machine learning in healthcare with the IBM Watson system failed to deliver even after years of time and billions of dollars invested. Machine learning has been used as a strategy to update the evidence related to systematic review and increased reviewer burden related to the growth of biomedical literature. While it has improved with training sets, it has not yet developed sufficiently to reduce the workload burden without limiting the necessary sensitivity for the findings research themselves [13].

Bias

Machine learning approaches in particular can suffer from different data biases. A machine learning system trained specifically on current customers may not be able to predict the needs of new customer groups that are not represented in the training data. When trained on man-made data, machine learning is likely to pick up the constitutional and unconscious biases already present in society. Language models learned from data have been shown to contain human-like biases. Machine learning systems used for criminal risk assessment have been found to be biased against black people. In 2015, Google photos would often tag black people as gorillas, and in 2018 this still was not well resolved, but Google reportedly was still using the workaround to remove all gorillas from the training data, and thus was not able to recognize real gorillas at all. Similar issues with recognizing non-white people have been found in many other systems. In 2016, Microsoft tested a chatbot that learned from Twitter, and it quickly picked up racist and sexist language. Because of such challenges, the effective use of machine learning may take longer to be adopted in other domains. Concern for fairness in machine learning, that is, reducing bias in machine learning and propelling its use for human good is increasingly expressed by artificial intelligence scientists, including Fei-Fei Li, who reminds engineers that "There’s nothing artificial about AI...It’s inspired by people, it’s created by people, and—most importantly—it impacts people. It is a powerful tool we are only just beginning to understand, and that is a profound responsibility” [14].

Overfitting

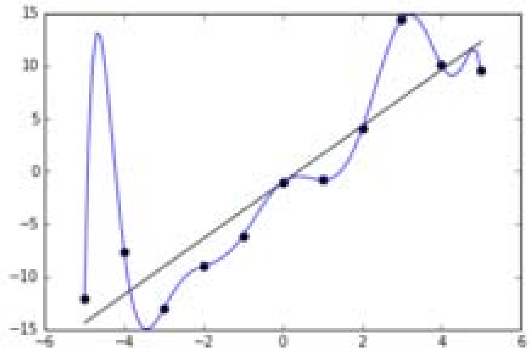

The blue line could be an example of overfitting a linear function due to random noise. Settling on a bad, overly complex theory gerrymandered to fit all the past training data is known as overfitting. Many systems attempt to reduce overfitting by rewarding a theory in accordance with how well it fits the data, but penalizing the theory in accordance with how complex the theory is [15].

Other Limitations

Learners can also disappoint by "learning the wrong lesson". A toy example is that an image classifier trained only on pictures of brown horses and black cats might conclude that all brown patches are likely to be horses. A real-world example is that, unlike humans, current image classifiers often don't primarily make judgments from the spatial relationship between components of the picture, and they learn relationships between pixels that humans are oblivious to, but that still correlate with images of certain types of real objects. Modifying these patterns on a legitimate image can result in "adversarial" images that the system misclassifies [16].

Adversarial vulnerabilities can also result in nonlinear systems, or from non-pattern perturbations. Some systems are so brittle that changing a single adversarial pixel predictably induces misclassification [17].

REFERENCES

1) Nilsson N. Learning Machines, McGraw Hill, 1965.
2) Duda, R., Hart P. Pattern Recognition and Scene Analysis, Wiley Interscience, 1973
3) S. Bozinovski "Teaching space: A representation concept for adaptive pattern classification" COINS Technical Report No. 81-28, Computer and Information Science Department, University of Massachusetts at Amherst, MA, 1981. https://web.cs.umass.edu/publication/docs/1981/UM-CS-1981-028.pdf
4) Mitchell, T. (1997). Machine Learning. McGraw Hill. p. 2. ISBN 978-0-07-042807-2.
5) Harnad, Stevan (2008), "The Annotation Game: On Turing (1950) on Computing, Machinery, and Intelligence", in Epstein, Robert; Peters, Grace (eds.), The Turing Test Sourcebook: Philosophical and Methodological Issues in the Quest for the Thinking Computer, Kluwer, pp. 23–66, ISBN 9781402067082
6) "Introduction to AI Part 1". Edzion. 2020-12-08. Retrieved 2020-12-09.
7) "AN EMPIRICAL SCIENCE RESEARCH ON BIOINFORMATICS IN MACHINE LEARNING – Journal". Retrieved 28 October 2020.

8) "rasbt/stat453-deep-learning-ss20" (PDF). GitHub.
9) Sarle, Warren (1994). "Neural Networks and statistical models". CiteSeerX 10.1.1.27.699.
10) Russell, Stuart; Norvig, Peter (2003) [1995]. Artificial Intelligence: A Modern Approach (2nd ed.). Prentice Hall. ISBN 978-0137903955.
11) Langley, Pat (2011). "The changing science of machine learning". Machine Learning. **82** (3): 275–279. doi:10.1007/s10994-011-5242-y.
12) Alpaydin, Ethem (2010). Introduction to Machine Learning. MIT Press. p. 9. ISBN 978-0-262-01243-0.
13) Cornell University Library (August 2001). "Breiman: Statistical Modeling: The Two Cultures (with comments and a rejoinder by the author)". Statistical Science. **16** (3). doi:10.1214/ss/1009213726. Retrieved 8 August 2015.
14) Gareth James; Daniela Witten; Trevor Hastie; Robert Tibshirani (2013). An Introduction to Statistical Learning. Springer. p. vii.
15) Mohri, Mehryar; Rostamizadeh, Afshin; Talwalkar, Ameet (2012). Foundations of Machine Learning. USA, Massachusetts: MIT Press. ISBN 9780262018258.
16) Alpaydin, Ethem (2010). Introduction to Machine Learning. London: The MIT Press. ISBN 978-0-262-01243-0. Retrieved 4 February 2017.
17) Le Roux, Nicolas; Bengio, Yoshua; Fitzgibbon, Andrew (2012). In Sra, Suvrit; Nowozin, Sebastian; Wright, Stephen J. (eds.). Optimization for Machine Learning. MIT Press. p. 404. ISBN 9780262016469.