



BMH Med. J. 2018;5(2):32-36 **Editorial**

Machine Learning - A Primer for Doctors

Mohan Leslie Noone

Baby Memorial Hospital, Kozhikode 673004

Address for Correspondence: Dr. Mohan Leslie Noone, MD, DM, Consultant Neurologist, Baby Memorial Hospital, Kozhikode 673004, Kerala, India. Email: drmohan@babymhospital.com

Abstract

Machine learning is a branch of artificial intelligence getting a lot of attention in recent times due to its success in handling complex tasks like computer vision and language processing. In medicine and health care as well, machine learning is poised to make a major impact in diagnostics and decision making. Having a strong relation with statistics, basic machine learning concepts are quite accessible to those with a medical background. The basics of machine learning are discussed, and different types of algorithms are outlined.

Key words: machine learning, artificial intelligence

Introduction

Machine learning is a branch of artificial intelligence that is currently garnering a lot of attention. This is because of its ability to tackle complex tasks that were previously thought near-impossible for computers to achieve. Such fields, where machine learning has already shown great promise include computer vision, speech recognition, natural language processing, music generation and self-driving vehicles.

In medicine and health care as well, machine learning is poised to make major impact in diagnostics and decision making - exemplified by recently published articles like "Dermatologist-level classification of skin cancer with deep neural networks" [1], and "Prediction of Cardiovascular Risk Factors from Retinal Fundus Photographs via Deep Learning" [2].

The good news is that machine learning has a strong relation with statistics - unlike traditional programming, and hence by its very nature, should be more comprehensible and understandable for doctors.

What is machine learning?

Machine learning is a process by which a computer program can learn to process data - without explicit instructions. More formally, it is a process by which the performance on a task based on data input improves over time as more data is fed into the system. The process by which a machine learning algorithm "learns" from data is called "training" - and this training is closely related to what

we call "regression" in statistics.

Types of machine learning

Broadly there are two types of machine learning - supervised and unsupervised.

In supervised learning, the algorithm is first trained on labelled data. For example, it is trained in images of tumor vs non-tumor lesions, which are labelled. After being trained with several examples, the algorithm learns the features, and then is able to predict on other images which are not part of the training set. The key is that a huge number of labelled data is required - the skin cancer study above used 129,450 labelled clinical images.

Unsupervised learning on the other hand, uses unlabeled data, and is involved with tasks like finding clusters or patterns within data, and identifying anomalous activity in a system, or fraud detection. It is not as much in the lime light as supervised learning, although it has some interesting applications in various fields, including health care.

Learning algorithms

Linear regression

The basic learning algorithm - which all machine learning courses teach first - is "linear regression". This should be familiar to anyone who has studied biostatistics. Here, we take a set of input parameters - like age, weight, blood pressure and correlate with an output parameter - like arterial thickness. With a fair amount of data, it is fairly simple to generate a line "of best fit". (In some cases, a curved line will be more appropriate, which can be achieved by using polynomials of the features like the square of weight.)




Figure 1: Linear regression

The way this line is calculated is what makes machine learning unique. We do not give instructions on how to calculate the line. Instead, we start with just a random line, calculate the error for each and every training sample, then use a simple but ingenious calculus-based trick - subtracting the slope (also called the gradient) of the error from the line parameters by a small amount, called the learning rate. This is guaranteed to decrease the error by a small amount. This is repeated until the error stops decreasing - and voila, the algorithm has found the line of best fit from the data without directly calculating it. This fundamental way of recursively minimizing error is a core concept of

machine learning and is called "gradient descent".

Logistic regression

An area where machine learning has shined a lot in recent times is classification tasks, where the output is in distinct classes - like tumor vs not tumor (binary classification) or identifying one of several objects in an image (multiple classification). For these type of tasks, the algorithm used is logistic regression. Here also, a line is used - but with a different purpose of separating the different classes. The probability for an element to belong to a class depends on the distance from the line, which in this context is also called the *decision boundary*.

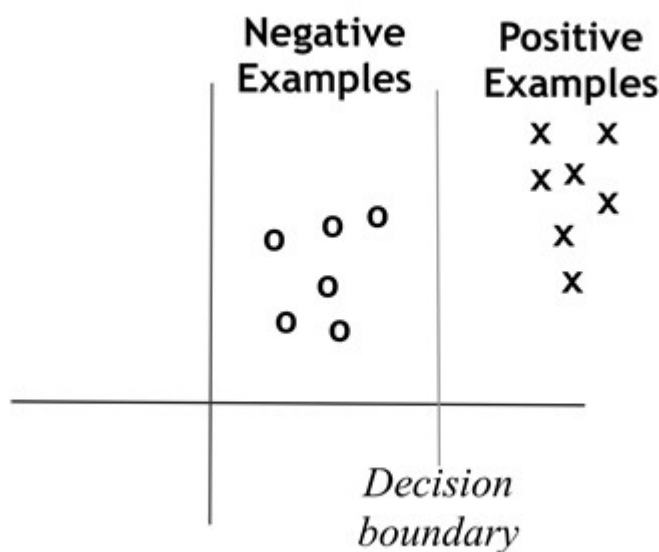


Figure 2: Logistic regression showing the decision boundary. Negative examples are typically labelled 0 and for positive examples 1

A mathematical trick - typically the sigmoid function, $1 / (1 + e^{-z})$ when applied to the output of a linear function (z) captures this requirement and is used as the mathematical basis for logistic regression.




Figure 3: The sigmoid function, outputs 0 to 1

Gradient descent, as described above, is also used to train logistic regression models to find the best "decision boundary" that separates the classes.

Once again, complex boundaries can be represented by polynomials. However, with large number of

input parameters - like an image, this becomes impractical. For e.g., a 64 x 64 color image will have 12288 parameters per image ($64 \times 64 \times 3$) so adding polynomials of these would make the regression extremely complex. So, we need other algorithms for such types of data.

Neural networks

Neural networks solve the above problem by stacking logistic regression nodes on top of each other in layers. This initially was done as an attempt to mimic how the brain works, hence the name neural network. Each node does logistic regression on an input and forwards the result to the next node. By combining several such layers, extremely complex, yet efficient models can be trained. To train such models using gradient descent, the output on labelled examples is first computed using random parameters, then the gradient is calculated backwards from the output toward the input layers, a technique called "back propagation", or "back prop".




Figure 4: Neural network. Each circle represents a logistic regression node

Most of the successes in machine learning in recent years are attributable to neural networks. Neural networks with several layers are called "deep neural networks", and machine learning using such networks is called "deep learning". The two papers mentioned above are examples of successful applications of deep learning.

There are various types of neural network architectures, notable among which are convolutional neural networks (CNNs) - used in computer vision and recurrent neural networks (RNNs) – used in speech and language processing.

For completion sake, there is another type of algorithm, the support vector machine, which is suited for moderate number of features, and is good at defining complex boundaries between classes. It has been very successful in tasks like identifying spam email.

Training Errors

As it is probably evident from above, the quality of training data heavily affects the performance and accuracy of a machine learning model. If the training data is biased, the model will train in a biased way and will predict incorrectly when applied to real data. For example, if a computer vision model to identify faces is trained only on a people from a particular race, it will perform poorly with people of other races. Hence, training bias needs to be carefully avoided.

The performance of a machine learning model generally improves with more data, unless the model itself is incorrect or too simple, which can usually be detected by plotting the accuracy vs size of

training data and cross validating with a separate set of labelled data not included in training. These steps, along with detailed error analysis are needed to fine tune and get good results from a machine learning implementation.

Conclusion

Machine learning is a force to reckon with, the impact it will have in a variety of fields including medicine are at once exciting and challenging. It is something that can no longer be ignored. Doctors, especially with their background in biostatistics are already familiar with many of the core concepts of machine learning and hence should endeavor to understand and perhaps even explore ways in which this promising technology can transform healthcare.

References

1. Andre Esteva, Brett Kuprel, Roberto A. Novoa, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 2017; 542:115-118.
2. Ryan Poplin, Avinash V. Varadarajan, Katy Blumer, et al. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering* 2018;2:158–164.