Manual: 7.2. Machine Learning

Machine learning, in its basic definition, is a way of representing empirical data in an efficient form, so that it can either be reproduced, generalized, recognized or some pattern in it can be discovered. The goal of machine learning is, then, to summarize data in the form of (a set of) equations.

A very simple example of machine learning would be measuring a series of points that lie on a straight line. The straight line is a model that represents the data we have measured. We use a simple equation to represent the straight line: y, the vertical axis, is equal to m, the slope of that straight line, multiplied by x, the horizontal axis, plus a constant b that is the intercept of the line with the vertical axis. So $y = mx + b$ is a summary of the data and it includes two parameters: m and b. As the measurements for both x and y are imprecise, the model parameters m and b also have uncertainties. This allows the model to cope with the imprecise nature of real life data sets.

Finding the values of m and b from a particular set of empirical data is, as we said, a simple example of machine learning. Thousands of data points have been reduced to two parameters of an equation.

The method of how to compute the parameters of a model from data is called an algorithm. Beyond the algorithm needed to compute the parameters of a model from data, machine learning often needs a second algorithm to update the parameters whenever additional data becomes available. This is not always possible, depending on the model, and represents a serious advantage as initial learning is usually time consuming.

The way a machine approaches learning is not the same as a human, but considering an analogy can be quite useful in the case of neural networks.

The brain is essentially a network of neurons that are connected by synapses. Each neuron and each synapse are individually quite simple objects but in a network they are able to carry out some astonishingly complex actions. Consider putting a name to the face of a person, for example, and connecting that to memories with that person. All this happens unconsciously, within fractions of a second. That's a neural network at work.

As a person is not born knowing all the people they will meet in their life, the brain's neural network gets trained as it experiences things, learns them, and can then draw on this knowledge quickly.

That is similar to how an artificial neural network learns. You start with a prototypical neural network and you start feeding it experiences. The more experiences you show the neural network the more it becomes capable of correctly representing those experiences and recalling them in the future.

The learning of a straight line we discussed is a simple example of machine learning. Neural networks, on the other hand, are capable of representing extremely complex data sets.

© 2005 — 2023 algorithmica technologies Inc., all rights reserved