Principles of Machine Learning

Evan Halloran

January 24, 2026

Contents

1 Introduction

1.1 What is Machine Learning?

Perhaps the most important concept in the whole of computer science is that of the *algorithm*. An algorithm, simply put, is a series of finite, unambiguous steps that transforms an input into an output. Much of the standard undergraduate curriculum nowadays is devoted to the implementation and analysis of algorithms. However, there are a multitude of tasks that humans can do fairly easily that cannot be reduced to a simple step-by-step procedure. Our abilities to drive vehicles, move through crowded rooms without hitting objects, and speak foreign languages are all examples of non-algorithmic tasks. Machine learning is the process by which a computer is trained on a plethora of data so that it too may learn to do these tasks for which we have no devised algorithm. The basic idea is to start with a general model equipped with a set list of parameters corresponding to the different features we'd like to capture. The model is then trained on a data set during which the parameters are tuned to better accomplish the task. After training, the model becomes specialized in the task it set out to learn, and the finalized values of the parameters coupled with the general model itself become our algorithm for the task. Machine learning is the process by which machines are taught to do something humans can do easily without an explicit step-by-step procedure.

As an example, take the transaction data acquired by a large supermarket chain. Each consumer transaction is stored as a log consisting of the customer's ID, items bought, amount spent, time of day, etc. The millions of transactions daily add up to be an impossibly large amount of data to parse by any given team of researchers. That is what makes tasks like these perfect for machine learning—if we can train a model to observe this data and make predictions about future spending habits, then the supermarket's profits will increase.

Indeed, data mining is an application of machine learning, specifically the training of models on ginormous data sets in the hopes of gleaning valuable insights and patterns. In the same way, machine learning is a subfield of artificial intelligence. All ML models must be intelligent in some part, for they must be able to learn from different types of training sets. If the environment the machine is observing changes, we would hope it changes its learning accordingly.

Machine learning does not aim to fully articulate the workings of a complex, dynamic system. Such a task is likely impossible. Rather, ML seeks to independently learn the underlying patterns and symmetries present in the data. Training for ML models is usually separated into two different purposes. A *predictive* ML model aims to predict the future given data, while a *descriptive* ML model aims to gain knowledge from data.

Mathematics and computer science play an important role in the development of machine learning. Firstly, models are evaluated by how faithfully they can predict/describe the data it is trying to learn. This performance criterion is a mathematical function of the given parameters of the model. The function is then optimized using multivariate calculus to find which parameters correspond to the most successful model. The model is built from these parameters by using the theory of statistics. After such a model is created, its performance is further tested by evaluating its time/space complexity. It is very much possible that a slow and spatially complex model is not as desirable as a faster but less-accurate one.

1.2 Examples of Machine Learning Applications

1.2.1 Classification

Suppose we are a bank tasked with determining the risk associated with loaning credit to a customer. There are several factors that play into risk-classification of customers such as age, savings, income, profession, and past financial history. Based on these factors, we want to separate all customers into one of two mutually exclusive classes: low-risk and high-risk. The bank will be able to save itself financial losses by only extending credit to low-risk customers. As an example, suppose our machine trains on data and comes up with the following classification rule:

$$\text{IF income } > \theta_1 \wedge \text{ savings } > \theta_2 \text{ THEN low-risk ELSE high-risk}$$

Such a rule is called a discriminant, which is the name for a mathematical function that separates its arguments into different cases. For instance, the prototypical discriminant in math is the function $\Delta(a, b, c) = b^2 - 4ac$ which, of course, is used to determine the number of real roots of a general quadratic $y = ax^2 + bx + c$.

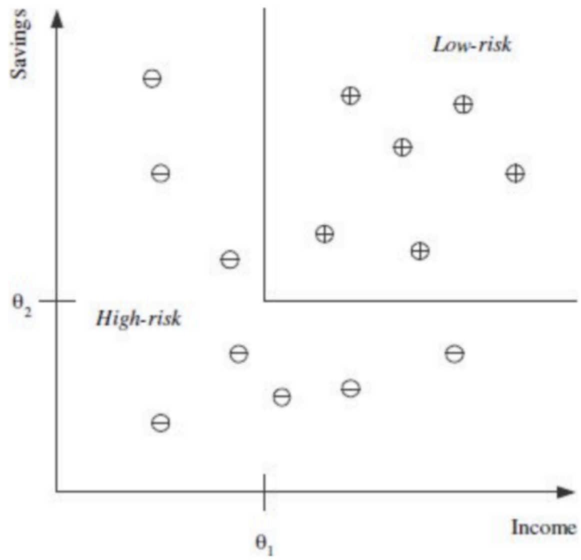

Figure 1: Training set of low-risk (-) and high-risk (+) potential credit loan customers. The bold line is the learned discriminant

2 Supervised Learning

3 Bayesian Decision Theory

4 Parametric Methods

5 Multivariate Methods

6 Dimensionality Reduction

7 Clustering

8 Nonparametric Methods

9 Decision Trees

10 Linear Discrimination

11 Multilayer Perception

12 Design and Analysis of Machine Learning Experi-