Machine Learning (Supervised and Unsupervised Learning)

Prof K R Chowdhary

MBM University

September 7, 2024

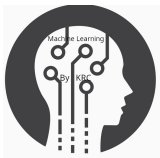

Supervised Learning

Most successful kinds of machine learning algorithms are: those that automate decision-making by generalizing from known examples. In *supervised learning*, the user provides the algorithm with pairs of: *inputs* and *desired* outputs, and the algorithm finds a way to produce the desired output given an input.

The algorithm is able to create an output for an input it has never seen before without any help from a human. Considering that pair of inputs are: $\{(1,2),(2,4),(3,9),(4,16),...,(10,100)\}$, the system learns that the relation is square. Now, this relation (of square) is used to produce output for any input x, and the output will be x^2.

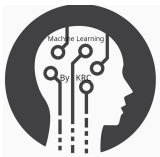

Supervised Learning ...

Considering the example of spam classification, using machine learning, a user provides the algorithm with a large number of emails as input, together with information about whether the email is spam (i.e., the desired output).
Given a future new email, the algorithm will then produce a prediction as to whether the new email is a spam (see Fig. 1).

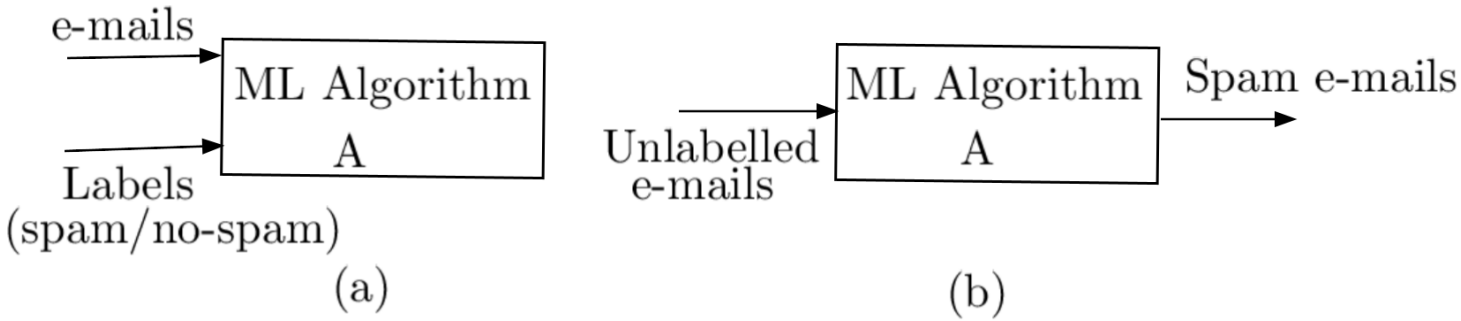

Figure 1: Machine Learning Algorithm: (a) Training phase, (b) Testing Phase

Supervised ML

1. *Identifying zip code from handwritten digits*: Input is a scan of the handwriting, and the desired output is the actual digits in the zip code. To create a dataset for building a machine learning model, you need to collect many postal envelopes, and read the zip codes yourself and store the digits as desired outcomes.

2. *Determining whether a tumor is benign*: Input is the medical image, and the output is whether the tumor is benign. To create a dataset for building a model, you need a database of medical images, and an expert opinion, so a doctor needs to look at all of the images and decide which tumors are benign.

3. *Detecting fraudulent credit card transaction*: Input is a record of the credit card transaction, and the output is whether it is likely to be fraudulent. The credit card issuing company may collect and store all transactions and record it when a user reports a transaction as fraudulent. This result is used to detect frauds in credit cards.

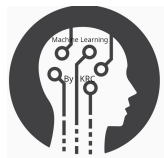

Formal definition of Supervised Learning

Ex1: Given an input or *feature vector* **x**, one of the main goals of machine learning is to predict an output variable y. For example, **x** could be a digitized signature and y, a binary variable that indicates whether the signature is genuine or not.
Ex2: The **x** represents weight and smoking habits of an expecting mother and y the birth weight of the baby. Machine learning prediction here, is encoded in a *prediction function* g, which takes as an input **x** and outputs a guess $g(x)$ for y, denoted by $\hat{y}$, (Fig. 2).

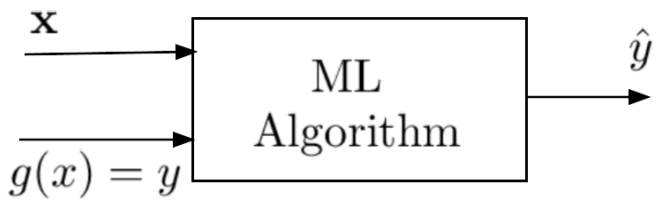

Figure 2: Defining Machine Learning

The g comprises all the information about the relationship between the variables **x** and y, excluding the random and chance cases.

Unsupervised Learning

Since learning involves an interaction between the learner and the environment, it is possible to divide the learning tasks according to the nature of that interaction.

The first such distinction to note is, the difference between *supervised* and *unsupervised* learning. As an example, consider the task of learning to detect spam e-mail versus the task of *anomaly detection*.

For the first, we consider a setting in which the learner receives e-mails to be trained, and each of it is manually labeled as "spam" or "not-spam". On the basis of such training, the learner figures out a rule for labeling a newly arriving e-mail messages as spam or no-spam.

In contrast, for the task of anomaly detection, all that the learner gets as training is a large body of e-mail messages (with no labels) and the learner's task is to detect "unusual" messages.

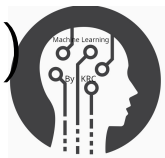

Unsupervised algorithms

In unsupervised learning, only the input data is known, and no known output data or labels are provided. Some examples are:

Identifying topics in a set of blog posts: If you have a large collection of text data, you might want to summarize it and find prevalent themes in it. You might not know beforehand what these topics are, or how many topics there might be. You might cluster similar posts, group-wise.

Segmenting customers into groups: Given a set of customer records, you might want to identify which customers are similar, and whether there are groups of customers with similar preferences. For a shopping site, these might be "parents," "bookworms," or "gamers."

Detecting abnormal access patterns to a website: To identify abuse or bugs, it is often helpful to find access patterns that are different from the norms. Each abnormal pattern might be very different.

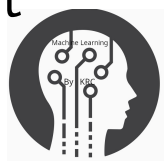

All these processes are *clustering algorithms*.