Artificial Intelligence (Introduction to Machine Learning (ML))

Prof K R Chowdhary

CSE Dept., MBM University

February 06, 2025
Lecture #4

Introduction to ML

- In Machine Learning (ML), we program the computers so that they can "learn" from input to them.
- Input to a learning algorithm is *training data* that represents experience, and output is some expertise, which usually takes the form of another computer program that can perform some task.
- For formal and mathematical understanding of this concept, we should be explicit about following:

$\rightarrow$ What is the training data that our programs uses?

$\rightarrow$ How can the process of learning be automated?

$\rightarrow$ How can we evaluate the success of such a process (i.e., the quality of the output of a learning program)?

Justification for Machine Learning

- In early days of "intelligent" applications, many systems used hand-coded rules of "if" and "else" decisions to process data or adjust to user input.
- In spam filtering, one could make up a *blacklist* of words that would result in an email being marked as spam.
- Manually crafting decision rules is feasible for some applications, in which humans have a *good understanding of the process to model*.
- However, using hand coded rules to make decisions has two major disadvantages:

$\rightarrow$ The logic required to make a decision is specific to a single domain and task. Changing the task even slightly might require a rewrite of the whole system.

$\rightarrow$ Designing rules requires a deep understanding of how a decision should be made by a human expert.

Why hand-coded System are not sufficient?

- One example where this hand-coded approach will fail is in detecting faces in images.
- Today, every smartphone can detect face in an image. However, face detection was an unsolved problem until as recently as 2001.
- The main reason for this: the way in which pixels (which make up an image in a computer) are "perceived" by the computer is very different from how humans perceive a face.
- This difference in representation makes it basically impossible for a human to come up with a good set of rules to describe what constitutes a face in a digital image.
- Using machine learning, however, simply presenting a program with a large collection of images of faces is enough for an algorithm to determine what characteristics are needed to identify a face!!

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.
- In fact, in a functional relation $y = f(x)$, the algorithm learns f, given large numbers of (x, y) pairs.
- Once the learning has taken place by the algorithm, it is able to create an output for an input that it has never seen before. This is done 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 relation." 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 spam filtering, 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 or not (it is desired output).
- Given a future new email, the algorithm will then produce a prediction as to whether the new email is a spam or not (see Fig. 1).

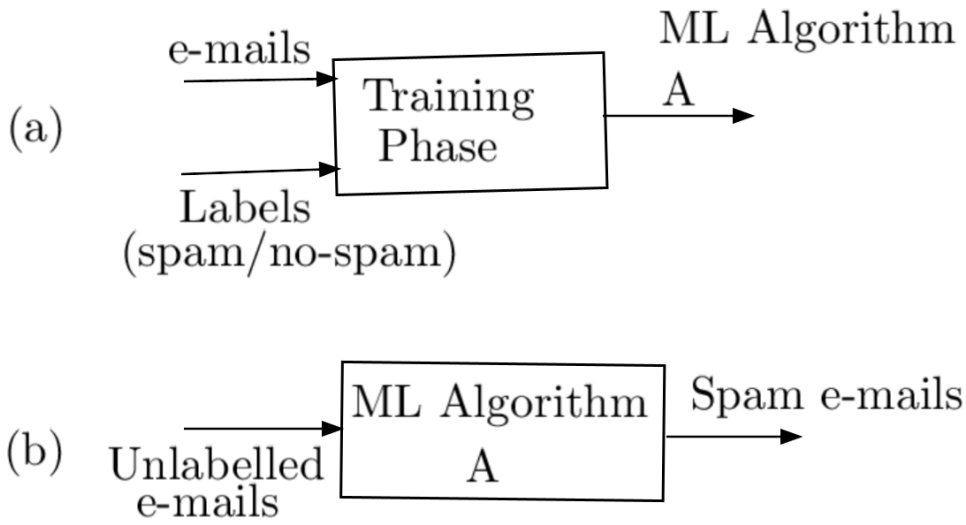

Figure 1: Machine Learning Algorithm: (a) Training phase, (b) Testing Phase

Supervised ML Examples

- *Ex. 1. Identifying pincode from handwritten digits*: Input is scan of handwritten pincode, and desired output is actual digits in the pin-code. To create a dataset for building a machine learning model, you need to collect many postal envelopes, and read the pincodes yourself and store the digits as desired outcomes.
- *Ex.2. Determining whether a tumor is benign*: Input is the medical image, and 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.
- *Ex.3. Detecting fraudulent credit card (CC) transaction*: Input is a record of the credit card transaction, and output is whether it is likely to be fraudulent. CC 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 CCs.

Formal definition of Supervised Learning

• 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.

• 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 $g(\mathbf{x})$. (Fig. 2).

Figure 2: Defining Machine Learning

• The function g comprises all information about relationship between variables **x** and y, excluding 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 this interaction.

• First such distinction to note is, difference between *supervised* and *unsupervised* learning.

• 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 learner receives e-mails on which to manually train the machine. So, 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.

• 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 Learning Examples

• 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*.

References

[1] Chowdhary, K.R. (2020). Machine Learning. In: Fundamentals of Artificial Intelligence. Springer, New Delhi. https://link.springer.com/chapter/10.1007/978-81-322-3972-7_13