

# Lecture 1: What does it mean to learn?

CSE 427: Machine Learning

Md Zahidul Hasan

Lecturer, Computer Science  
BRAC University

Spring 2023

# Learning From Experience

CGPA	GRE	Publications	Admission Offers
3.5	321	3	VT, UC Riverside
4.0	332	4	MIT, UCB, CMU
3.7	321	0	NYU, UMass
3.89	324	1	Purdue, UMass
3.75	329	2	NYU, Yale
3.3	330	0	WSU
3.23	311	1	WSU, FIU

There are various groups and threads on social media that help people with undergraduate/graduate admissions. People post their profiles and the list of universities they got admission offers from. Future candidates can learn from the experiences shared therein and can decide which universities to aim for.

# Applications

- 1 speech recognition, speech synthesis, speaker verification
- 2 document classification, spam detection
- 3 image recognition, object detection, pet and face detection
- 4 machine translation, parsing, parts of speech tagging
- 5 recommendation systems, stock market predictions
- 6 generative art, music synthesis, AI story generator
- 7 computational biology: drug discovery, protein function prediction
- 8 medical diagnosis: skin cancer detection, eye disease detection
- 9 games: chess, go, checker
- 10 self-driving cars, traffic prediction
- 11 online fraud detection, anomaly detection

# Elements of Learning

**Instances:** Examples or instances are the rows of a dataset. For example, in the previous example, a student's profile is an instance.

**Features:** Features or attributes are the columns of a dataset except the last one. For example, a student's profile consists of CGPA, GRE, number of publications, TOEFL, etc.

**Labels:** The last column of the dataset is called the label set. In the given example, the label set contains the list of universities that offered admission. The goal of supervised learning is to learn this column.

**Feature Vector:** A row can be expressed as a single vector excluding the label. In the above scenario there are 7 students, i.e.  $x_1 = (3.5, 321, 3), \dots, x_7 = (3.23, 311, 1)$ . And the labels can be expressed as  $y_1, y_2, \dots, y_7$ .

**Training Data:** Set of instances used to train a learning algorithm. If the training sample has  $m$  instances, it can be expressed as  $S = ((x_1, y_1), (x_2, y_2), \dots, (x_m, y_m))$ .

# Elements of Learning contd.

**Hyperparameters:** Free variables that are independent of the learning algorithm.

**Validation Data:** Examples used to tune the hyperparameters of the learning algorithm.

**Testing Data:** Examples used to evaluate the performance of a learning algorithm. The learning algorithm doesn't have access to the testing data during the learning process.

**Domain:** The domain is a pool of instances from where we build the dataset. For example, CGPA can be any real number from 2 to 4, the GRE score is an integer between 260 and 340, the number of publications can be any non-negative integer. Hence a domain for the student profiles can be  $\mathcal{D} = \mathcal{A} \times \mathcal{B} \times \mathcal{C}$  where  $\mathcal{A}$  is the domain for CGPA,  $\mathcal{B}$  is the domain for GRE,  $\mathcal{C}$  is the domain for number of publications.

**Concept:** A concept is what a learning algorithm is trying to learn. Most of the time it's a function from  $\mathcal{D}$  to another set.

# Elements of Learning contd.

For example, in the case of learning universities,

$$f(3.5, 321, 3) = \{VT, UC\text{Riverside}\}$$

**Hypothesis:** A hypothesis is our assumption about the concept. For example, our learning algorithm thinks that with a 3.5 CGPA, 321 GRE score, and 3 publications, a student can get offers from universities like Virginia Tech, Indiana University, and John Hopkins University. Our hypothesis as well could be a function from  $\mathcal{D}$  to a vector of universities like the following

$$h(3.5, 321, 3) = \{VT, IU, JHU\}.$$

**Loss Function:** During testing, we need to determine the performance of the learning algorithm. If our learning algorithm is a binary classifier that classifies an email into two categories  $\{SPAM, non - SPAM\}$ , then if a testing instance is correctly classified, the loss is 0. Otherwise, it's 1. Let's say we are trying to estimate the rent of a house given a feature vector  $x$ .  $f(x)$  is the actual rent and  $h(x)$  is our estimate. The loss or error could be defined as  $|f(x) - h(x)|$  or  $(f(x) - h(x))^2$ .

**Empirical Error:** Given a concept  $f$ , and a training sample  $S = (x_1, x_2, \dots, x_m)$  of size  $m$ , then the empirical error of a hypothesis  $h$  is defined by:

$$R_S(h) = \frac{\sum_{i=1}^m 1_{f(x_i) \neq h(x_i)}}{m}$$

**Generalization Error:** The training samples are only a partial view of the data. The true error or generalization error is the error of our learning algorithm if we could test it against all the data instances in the domain. Any instance  $x$  is sampled from the domain  $\mathcal{D}$ . It's defined by:

$$R(h) = \mathbb{E}_{x \sim \mathcal{D}}[1_{f(x) \neq h(x)}]$$

# Machine Learning Tasks

Some of the most common machine learning tasks are the following:

**Classification:** This involves assigning a label to each data instance. If we are classifying documents based on their category, then we need to assign each document a label such as {sports, politics, entertainment, science,  $\dots$ }. If we are trying to classify a document based on whether it gives off a positive vibe or not, then the labels should be {POSITIVE, NEGATIVE}.

**Classification with missing inputs:** If we want to diagnose a patient's disease based on their history, bodily features and test results, it could be possible that some people can't afford some expensive medical tests. In that case, we end up with features that have no value for a particular patient. Even in those cases, we need to be able to diagnose their illnesses. In the usual learning scenario, we learn one function to classify the testing samples. In this case, we might need to learn  $2^n$  different functions for the classification where  $n$  is the total number of features.



# Machine Learning Tasks contd.

**Regression:** Unlike classification, we are not assigning a label to each instance here. We want to estimate a value which is commonly a real number. For example, we might be trying to estimate the rent of an apartment based on the total square feet, number of rooms, proximity to the city, utility facilities etc.

**Transcription:** Transcription converts unstructured data into textual forms. For example, optical character recognition, speech to text etc.

**Ranking:** Given a search prompt, a search engine generates a list of webpages which need to be sorted for users' convenience.

**Clustering:** Clustering is applied to identify groups of similar entities such as community detection in social media or streaming services.

**Manifold learning:** Manifold learning is the task of reducing the feature space intelligently so that learning can occur more efficiently.

# Machine Learning Tasks contd.

**Anomaly detection:** We can use machine learning to detect unusual behaviors from a credit card owner to identify theft.

**Synthesis and sampling:** With synthesis and sampling, we can generate new music or new artwork by learning from the past works of an artist.

**Imputation of missing values:** If the value of some features are lacking, we can use our wisdom gained from experience to impute the missing values.

**Denoising:** Suppose, a clean example  $x$  is observed to be a corrupted  $\bar{x}$  due to some noisy process. From experience, given  $\bar{x}$ , we can predict the original clean  $x$ .

**Probability mass/density estimation:** In this scenario, we are just trying to learn  $p(x)$  for an instance  $x$ .

# Learning Scenarios

**Supervised learning:** In this case, we are always given either a set of labels or a value or some ranking. Classification, ranking, regression, etc are supervised tasks.

**Unsupervised learning:** In unsupervised learning, there is no column for labels. Clustering and dimensionality reduction are some of the unsupervised learning techniques.

**Semi-supervised learning:** Semi-supervised learning is ideal when it's expensive to get labeled data. In this case, we can learn from the labels that are given, and utilize similarity measures as in unsupervised learning to make predictions.

**Transductive inference:** Just like the semi-supervised scenario, we get some labeled data and some unlabeled data. But in the case of induction, we learn from the training set to make general predictions. We make no assumption about the test data. But in this case, the test data are the unlabeled instances only. The goal is to minimize error in the test samples only.

# Learning Scenarios contd.

**Offline learning:** In the offline scenario, all the learning happens before testing. It learns from the training set before it is tested.

**Online learning:** In online learning, the training and testing phases are intertwined. Here, the algorithm makes a prediction when it encounters a data instance. Then it learns from the mistake.

**Reinforcement learning:** In reinforcement learning, the environment doesn't provide an explicit label set or any kind of supervision. But the agent receives a reward or punishment for any action. The agent's goal is to maximize reward in the long run.

**Active learning:** In this scenario, the algorithm (student) carefully chooses which data points to include during learning and then queries a teacher about its label. The goal is to get results comparable to the supervised learning scenario with fewer data points because sometimes getting labels can be expensive.