Gems of TCS

PAC Learning

Sasha Golovnev

November 29, 2021

Classes of Learning Problems

- Classification

Classes of Learning Problems

- Classification
- Ranking

Classes of Learning Problems

- Classification
- Ranking
- Regression

Classes of Learning Problems

- Classification
- Ranking
- Regression
- Clustering

Classes of Learning Problems

- Classification
- Ranking
- Regression
- Clustering
- ...

Example

- Given a set of <u>labeled</u> emails

Example

- Given a set of labeled emails
- Build a classifier that predicts spam/non-spam labels for incoming emails

Setup

- Partition labeled data into three sets:
  - training sample
  - validation sample
  - test sample

Setup

- Partition labeled data into three sets:
    - training sample
    - validation sample
    - test sample
- Identify relevant features

Setup

- Partition labeled data into three sets:
  - training sample
  - validation sample
  - test sample
- Identify relevant features
- Train on training sample

Setup

- Partition labeled data into three sets:
  - training sample
  - validation sample
  - test sample
- Identify relevant features
- Train on training sample
- Tune parameters using validation sample

Setup

- Partition labeled data into three sets:
    - training sample
    - validation sample
    - test sample
- Identify relevant features
- Train on training sample
- Tune parameters using validation sample
- Evaluate using test sample

What can be learned?

- What can be learned?
- What cannot be learned?
- How many samples do we need to learn?

What can be learned?

- What can be learned?
- What cannot be learned?
- How many samples do we need to learn?
- Framework of PAC learning (L. Valiant, 1984)

Definitions

- X—set of all possible instances/examples

Definitions

- X—set of all possible instances/examples
- $\mathcal{D}$—target distribution over X

Definitions

- X—set of all possible instances/examples
- $\mathcal{D}$—target distribution over X
- c—target concept

DEFINITIONS

- X—set of all possible instances/examples
- $\mathcal{D}$—target distribution over X
- c—target concept
- C—concept class

Definitions

- X—set of all possible instances/examples
- $\mathcal{D}$—target distribution over X
- c—target concept
- C—concept class
- Goal: given training set, select h that approximates c well

Errors

Generalization Error

For hypothesis h, target concept c, and target distribution D:

$$R(h) = \Pr_{x \sim D}[h(x) \neq c(x)] \,.$$

Errors

Generalization Error

For hypothesis h, target concept c, and target distribution D:

$$R(h) = \Pr_{x \sim D}[h(x) \neq c(x)] \,.$$

Empirical Error

For hypothesis h, target concept c, and sample $S = (x_1, \ldots, x_m)$:

$$\widehat{R}(h) = \frac{|\{x_i : h(x_i) \neq c(x_i)\}|}{m} \,.$$

Average Error

$$\mathbb{E}_S[\widehat{R}(h)] = R(h)\,.$$

PAC Learning

PAC (Probably Approximately Correct)

Concept class C is PAC-learnable if there exists learning algorithm s.t.

- for all $c \in C, \varepsilon > 0, \delta > 0$, all distributions D:

$$\Pr_{S \sim D^m}[R(h_S) \leq \varepsilon] \geq 1 - \delta\,,$$

PAC Learning

PAC (Probably Approximately Correct)

Concept class C is PAC-learnable if there exists learning algorithm s.t.

- for all $c \in C, \varepsilon > 0, \delta > 0$, all distributions D:

$$\Pr_{S \sim D^m}[R(h_S) \leq \varepsilon] \geq 1 - \delta\,,$$

- for random samples of size

$$m \leq \mathsf{poly}(1/\varepsilon, 1/\delta, n)\,.$$

- Probably: confidence $1 - \delta$
- Approximately correct: accuracy $1 - \varepsilon$

Example