ECE595 / STAT598: Machine Learning I
Lecture 24.2: Probably Approximately Correct - PAC Framework

Spring 2020

Stanley Chan

School of Electrical and Computer Engineering
Purdue University

Outline

- Lecture 22 Is Learning Feasible?
- Lecture 23 Probability Inequality
- Lecture 24 Probably Approximate Correct

Today's Lecture:

- Two ingredients of generalization
  - Training and testing error
  - Hoeffding inequality
  - Interpreting the bound
- PAC Framework
  - PAC learnable
  - Confidence and accuracy
  - Example

Accuracy and Confidence

Recall the equation

$$\mathbb{P}\Big[\,|E_{\text{in}}(h) - E_{\text{out}}(h)| > \epsilon\Big] \le 2e^{-2\epsilon^2 N}$$

- $\delta = 2e^{-2\epsilon^2 N}$. **confidence**: $1 - \delta$.
- $\epsilon = \sqrt{\frac{1}{2N}\log\frac{2}{\delta}}$. **accuracy**: $1 - \epsilon$.
- Then the equation becomes

$$\mathbb{P}\Big[\,|E_{\text{in}}(h) - E_{\text{out}}(h)| > \epsilon\Big] \le \delta$$

- which is equivalent to

$$\mathbb{P}\Big[\,|E_{\text{in}}(h) - E_{\text{out}}(h)| \le \epsilon\Big] > 1 - \delta$$

Probably Approximately Correct

- **Probably**: Quantify error using probability:

$$\mathbb{P}\left[\, |E_{\text{in}}(h) - E_{\text{out}}(h)| \leq \epsilon \right] \geq 1 - \delta$$

- **Approximately Correct**: In-sample error is an approximation of the out-sample error:

$$\mathbb{P}\left[|E_{\text{in}}(h) - E_{\text{out}}(h)| \leq \epsilon\right] \geq 1 - \delta$$

- If you can find an algorithm $\mathcal{A}$ such that for any ϵ and δ, there exists an N which can make the above inequality holds, then we say that the target function is **PAC-learnable**.
- The following example is taken from Mohri et al. Foundation of Machine Learning, Example 2.4.

Example: Rectangle Classifier

Consider a set of 2D data points.

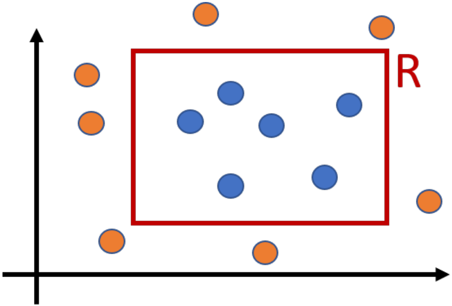

- The target function is a rectangle R
- Inside R: blue. Outside R: orange. Data is intrinsically separable.
- Goal: Pick a hypothesis rectangle R' using the available data point
- Question: Is this problem PAC learnable?

What Shall We Do?

- This question is very general.
- It asks about the nature of the problem.
- We want to show that this problem is indeed PAC learnable.
- To do so, we need to propose an **algorithm** $\mathcal{A}$ which takes the training data and returns an R', such that for any $\epsilon > 0$ and $\delta > 0$, there exists an N (which is a function of ϵ and δ) with
$$\mathbb{P}[|E_{\text{in}}(R') - E_{\text{out}}(R')| > \epsilon] \leq \delta.$$
- If we find such algorithm, then the problem is PAC learnable.

Proposed Algorithm

- $\mathcal{A}$: Give me the data point points, find the **tightest rectangle** that covers the blue circles.

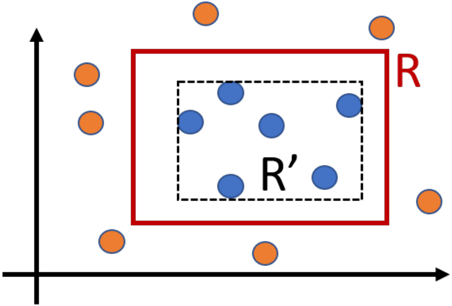

Intuition

- As N grows, we can find a R' which is getting closer and closer to R.
- So for any $\epsilon > 0$ and $\delta > 0$, it seems possible that as long as N is large enough we will be able to make training error close to testing error.
- See Appendix for proof.

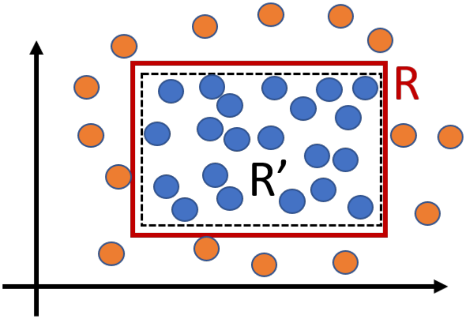

Summary

- Not all problems are learnable.
- Those that are learnable require **training and testing** samples are correlated.
- Then **Hoeffding inequality** applies

$$\mathbb{P}[|E_{\text{out}}(R') - E_{\text{in}}(R')| > \epsilon] \leq \delta.$$

- For any accuracy ϵ and any confidence δ, if you can find an algorithm $\mathcal{A}$ such that as long as N is large enough the above inequality can be proved, then the target function is PAC learnable.
- Next time: Look at the hypothesis set $\mathcal{H}$.

Reading List

- Yasar Abu-Mustafa, Learning from Data, Chapter 1.3, 2.1.
- Mehryar Mohri, Foundations of Machine Learning, Chapter 2.1.
- Martin Wainwright, High Dimensional Statistics, Cambridge University Press 2019. (Chapter 2)
- CMU Note https://www.cs.cmu.edu/~mgormley/courses/10601-s17/slides/lecture28-pac.pdf
- Iowa State Note http://web.cs.iastate.edu/~honavar/pac.pdf
- Princeton Note https://www.cs.princeton.edu/courses/archive/spring08/cos511/scribe_notes/0211.pdf
- Stanford Note http://cs229.stanford.edu/notes/cs229-notes4.pdf