

ASSIGNMENT 2

Exercise 1 (Finite classes are learnable). In this exercise we will show that any finite class is learnable. More specifically we will establish the following result:

Theorem: Let \mathcal{H} be a finite set of functions. Then, for any empirical risk minimization (ERM) function \hat{f}_{S^m} we have

$$\mathbb{P}_{S^m}(\{S^m : \mathbb{P}_X(\hat{f}_{S^m}(X) \neq f(X)) > \varepsilon\}) \leq |\mathcal{H}|e^{-\varepsilon m}$$

for any data distribution and for any $f \in \mathcal{H}$.

1. Define the set of bad hypothesis as

$$\mathcal{H}_B = \{h \in \mathcal{H} : \mathbb{P}_X(h(X) \neq f(X)) > \varepsilon\}$$

and define the set of misleading samples as

$$\mathcal{M} = \{S^m : \exists h \in \mathcal{H}_B \text{ such that } L_{S^m}(h) = 0\}$$

Use realizability to show that (and this is the main step of the proof)

$$\mathbb{P}_{S^m}(\{S^m : \mathbb{P}_X(\hat{f}_{S^m}(X) \neq f(X)) \geq \varepsilon\}) \leq \mathbb{P}_{S^m}(\mathcal{M})$$

2. Argue that

$$\mathbb{P}_{S^m}(\mathcal{M}) \leq \sum_{h \in \mathcal{H}_B} \mathbb{P}_{S^m}(S^m : L_{S^m}(h) = 0)$$

3. Argue that

$$\mathbb{P}_{S^m}(S^m : L_{S^m}(h) = 0) \leq (1 - \varepsilon)^m$$

4. Conclude the proof.

Exercise 2 (No Free Lunch Theorem). In this exercise we show that if \mathcal{H} is unrestricted, that is if \mathcal{H} contains all the functions from \mathcal{X} to \mathcal{Y} then PAC learning \mathcal{H} requires to have “huge” number of samples, of the order of $|\mathcal{X}|$:

Theorem: Let \mathcal{H} denote the set of all functions from \mathcal{X} to \mathcal{Y} . Then the number of samples m needed to PAC learn \mathcal{H} with accuracy $\varepsilon = 1/8$ and confidence $\delta = 1/8$, $m_{\mathcal{H}}(1/8, 1/8)$, is at least $|\mathcal{X}|/2$. Hence, learning a “rich class” of functions necessarily requires seeing lots of samples.

1. Suppose first that $|\mathcal{X}| < \infty$ and that $\mathcal{Y} = \{0, 1\}$ —at the end we show that the bound we obtain holds without these assumptions.

Let $\hat{f}_{S^m}(\cdot)$ be a predictor algorithm for the class \mathcal{H} . Let the test symbol $X \in \mathcal{X}$ and the training sequence $S^m = S_1, S_2 \dots S_m \in \mathcal{X}^m$ be i.i.d. $\sim P$ for some distribution P . Justify the following inequalities:

$$\begin{aligned} & \sup_{h \in \mathcal{H}} \mathbb{E}_{S^m} [P_X(\hat{f}_{S^m}(X) \neq h(X))] \\ & \geq \mathbb{E}_h \mathbb{E}_{S^m} (P_X(\hat{f}_{S^m}(X) \neq h(X))) \quad \text{for any distribution over } h \end{aligned} \quad (1)$$

$$= \mathbb{E}_h \mathbb{E}_{S^m} \mathbb{E}_X (\mathbf{1}\{\hat{f}_{S^m}(X) \neq h(X)\}) \quad (2)$$

$$= \mathbb{P}_{h, S^m, X}(\hat{f}_{S^m}(X) \neq h(X)) \quad (3)$$

$$\geq \mathbb{P}_{h, S^m, X}(\hat{f}_{S^m}(X) \neq h(X) | X \notin \{S_1, \dots, S_m\}) \mathbb{P}_{S^m, X}(X \notin \{S_1, \dots, S_m\}) \quad (4)$$

2. Suppose h is uniformly distributed over the set of functions from \mathcal{X} to $\{0, 1\}$. Show that

$$\mathbb{P}_{h, S^m, X}(\hat{f}_{S^m}(X) \neq h(X) | X \notin \{S_1, \dots, S_m\}) = 1/2.$$

3. Let P_X be the uniform distribution over \mathcal{X} . Show that

$$\mathbb{P}_{S^m, X}(X \notin \{S_1, \dots, S_m\}) \geq \frac{|\mathcal{X}| - m}{|\mathcal{X}|}.$$

4. Argue that there exists h and a data distribution such that

$$\mathbb{P}_{S^m, X}(\hat{f}_{S^m}(X) \neq h(X)) \geq \frac{1}{4}$$

whenever $m \leq |\mathcal{X}|/2$.

5. Define events (for notational convenience)

$$\mathcal{E} = \{\hat{f}_{S^m}(X) \neq h(X)\}$$

$$\mathcal{E}_\varepsilon = \{s^m : \mathbb{P}_X(\hat{f}_{s^m}(X) \neq h(X)) \geq \varepsilon\}$$

Argue that there exists a function h and a data distribution such that

$$\frac{1}{4} \leq \mathbb{P}_{S^m, X}(\mathcal{E}) \leq Pr(S^m \in \mathcal{E}_\varepsilon) + \varepsilon.$$

6. Conclude that $m_{\mathcal{H}}(1/8, 1/8) \geq |\mathcal{X}|/2$.

7. Show that the previous bound also holds if $|\mathcal{Y}| > 2$.