Theory Homework 3

CSE 493S/599S: Advanced Machine Learning
Instructor: Sewoong Oh
Due: Thursday, May 28$^{\text{th}}$ at 11:59pm

The goal of this homework is to help you better understand the ideas from theoretical machine learning we have covered in class.

Notes:

- You will be assigned a subset of the problems for each homework. Please submit the homework to Gradescope and link each page of your work to the corresponding problem.
- Please typeset your work using LaTeX.
- List every person with whom you discussed any problem in any depth, and every reference (outside of our course slides, lectures, and textbook) that you used.
- You may spend an arbitrary amount of time discussing and working out a solution with your listed collaborators, but **do not take notes, photos, or other artifacts of your collaboration**. Erase the board you were working on, and once you're alone, write up your answers yourself.
- The homework problems have been carefully chosen for their pedagogical value and hence might be similar or identical to those given out in similar courses at UW or other schools. Using any pre-existing solutions from these sources, from the Web or other textbooks constitutes a violation of the academic integrity expected of you and is strictly prohibited.

Version history:

V1 Initial version.

1 The perceptron (taught in the 6th theory lecture)

The perceptron is a classical algorithm for learning a separating hyperplane of a dataset $S = \{(x_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \{-1, 1\}$. It is presented in the following psuedocode:

Algorithm 1: Perceptron

1. Initialize $w_0 \leftarrow 0$.
2. While w_t does not separate the data:
   a. Select a random index $i_t \in \{1, \ldots, n\}$.
   b. If $y_{i_t} w_t^\top x_{i_t} < 1$, set $w_{t+1} \leftarrow w_t + y_{i_t} x_{i_t}$.
   c. Otherwise, set $w_{t+1} \leftarrow w_t$.
   d. Set $t \leftarrow t + 1$.
3. Return w_t.

The perceptron happens to be equivalent to learning a linear separator using SGD and the hinge loss! In this problem, we will prove some famous results about the perceptron.

1.1 Mistake bound

First, we will show that the perceptron performs well on the training data (denoted S) using a *mistake bound*. In particular, we will show that if there exists a linear separator of the training data, then the perceptron will find it provided the margin of S is not too small.

The margin is first defined for a particular hyperplane $\mathcal{H}_w = \{x : w^\top x = 0\}$ corresponding to a vector $w \in \mathbb{R}^d$. Supposing that $\mathcal{H}_w$ perfectly separates S, we define the margin $\gamma(S, w)$ as the smallest distance between a point in S and a point in $\mathcal{H}_w$:

$$\gamma(S, w) = \text{dist}(S, \mathcal{H}_w)$$

where $\text{dist}(A, B) = \min(\|a - b\| : a \in A, b \in B)$.

The margin of S is then defined as the largest margin achievable by any w:

$$\gamma(S) = \max_{\|w\|=1} \gamma(S, w).$$

Additionally, define the diameter of S to be $D(S) = \max_{(x,y)\in S} \|x\|$.

Our goal in this section is to prove the following theorem:

Theorem 1.1. algorithm 1 makes at most $(2 + D(S)^2)/\gamma(S)^2$ mistakes on any sequence of examples S that can be perfectly linearly separated.

The proof of this theorem can be broken into parts:

(1) First, we will upper bound $\|w_t\|$. In particular, show that

$$\|w_{t+1}\|^2 \leq \|w_t\|^2 + 2 + D(S)^2.$$

Now, let m_t be the total number of mistakes made by algorithm 1 during the first t iterations. Use the previous result to show that

$$\|w_t\| \leq \sqrt{m_t(2 + D(S)^2)}.$$

[8 points]

(2) Next we will lower bound $\|w_t\|$. Start by showing that for any unit vector w that perfectly separates S, we have

$$\langle w, w_{t+1} - w_t \rangle \geq \gamma(S, w).$$

when we make a mistake at iteration t.

Let unit vector $w^\star$ denote the hyperplane achieving the maximum margin $\gamma(S)$. Use the previous result to show that

$$\langle w^\star, w_t \rangle \geq m_t \gamma(S).$$

Use this to obtain a lower bound for $\|w_t\|$. **[8 points]**

(3) Combine the two bounds to obtain a bound on the number of mistakes m_t. **[4 points]**

[total 20 points]

1.2 Generalization bound

Let us assume that the data S was drawn i.i.d. from a fixed underlying distribution $\mathcal{D}$ which is linearly separable. In the previous section, we saw that algorithm 1 finds a linear predictor for S. Now we will show that this predictor also works on new data drawn from $\mathcal{D}$! In particular, use the result from the previous subsection to give a proof of the following theorem:

Theorem 1.2. Let S_n denote a set of n i.i.d. samples from $\mathcal{D}$. Let $w(S)$ be the output of algorithm 1 on dataset S. Let $Z = (X, Y)$ be an additional independent sample from $\mathcal{D}$. Then,

$$\mathbb{P}[Y w(S_n)^\top X < 1] \leq \frac{1}{n+1} \mathbb{E}_{S_{n+1}} \left[\frac{2 + D(S_{n+1})^2}{\gamma(S_{n+1})^2} \right].$$

Hint: Note that:

(i) Z can be swapped with any entry of S_n without changing the distribution of outcomes.

(ii) If for some S, the perceptron never makes a mistake on example $s \in S$, then $w(S) = w(S \setminus s)$. This implies that $w(S \setminus s)$ will predict s correctly.

Use the mistake bound from earlier to show that the reasoning in (ii) will apply to many examples.

[20 points]
