Homework #3
CMSC 35401: The Interplay of Learning and Game Theory
(Autumn'22)

Due Tue 12/06 2:00 pm

General Instructions The assignment is meant to be challenging. Feel free to discuss with fellow students, however please write up your solutions independently (e.g., start writing solutions after a few hours of any discussion) and acknowledge everyone you discussed the homework with on your writeup. The course materials are all on the the course webiste here www.haifeng-xu.com/cmsc35401fa22/index.htm. You may refer to any materials covered in our class. However, any attempt to consult outside sources, on the Internet or otherwise, for solutions to any of these homework problems is *not* allowed.

Whenever a question asks you to "show" or "prove" a claim, please provide a formal mathematical proof. These problems have been labeled based on their difficulties. `Short` problems are intended to take you 5-15 minutes each and `medium` problems are intended to take 15-30 minutes each. `Long` problems may take anywhere between 30 minutes to several hours depending on whether inspiration strikes.

Finally, please write your solutions in latex — hand written solutions will not be accepted. Hope you enjoy the homework!

Problem 1: Proper Scoring Rules

Recall that a proper *scoring rule* $S(i;p)$ is a reward function for eliciting an expert's prediction about a random event $E \in \{1, 2, \cdots, n\}$. Specifically, if the expert's reported distribution is $p \in \Delta_n$[1], then the expert receives reward $S(i;p)$ if the event E is realized to i in the future. A scoring rule is strictly proper if reporting her true belief $\lambda \in \Delta_n$ uniquely maximizes the expert's expected utility.

1. (`Short, 5 points`) The linear scoring rule is defined as $S(i;p) = p_i$. Show that linear scoring rule is not proper.

2. (`Short, 5 points`) We talked in class that log-scoring rule and quadratic scoring rule are both strictly proper. Please describe another strictly proper scoring rule and argue why it is strictly proper.

Problem 2: Linear Regression from Strategic Data Sources

Suppose you are an insurance company who is trying to design a linear regression function $w \cdot x + b$ (with parameter $w \in \mathbb{R}^d, b \in \mathbb{R}$) to determine the insurance payment for any customer with feature vector $x \in \mathbb{R}^d$

[1] $\Delta_n = \{p \in \mathbb{R}^n : \sum_{i=1}^n p_i = 1, p_i \geq 0 \ \forall i\}$ is the n-dimensional simplex

(e.g., riskier customers need to pay more). Naturally, any customer with true feature vector x would have incentives to misreport his feature as some $z \in \mathbb{R}^d$ to hopefully induce a lower payment, i.e., smaller $w \cdot z + b$. However, such manipulation of pretending to be z comes with a cost $c(z; x)$. Therefore, there is a tradeoff between the benefit and cost of manipulating their feature x. Formally, we assume any customer with true feature x will try to pick

$$z^*(x) = \arg\min_z [w \cdot z + b + c(z; x)].$$

We say a *regressor* (w, b) is *incentive compatible* if $z^*(x) = x$ for any $x \in \mathbb{R}^d$. Note that a regressor with $w = 0$ is trivially incentive compatible as its payment does not depend on the feature x at all. We call such a regressor *trivial*. However, there might be non-trivial regressors.

- (`Medium, 11 points`) Suppose $c(z; x) = \sqrt{\sum_{i=1}^d (z_i - x_i)^2}$ is the standard Euclidean distance. Please characterize the set of all incentive compatible regressors.
- (`Medium, 11 points`) Suppose $c(z; x) = \sum_{i=1}^d (z_i - x_i)^2$ is the *squared* Euclidean distance. Please characterize the set of all incentive compatible regressors. How many non-trivial regressors are there in this set?
- (`Long, 18 points`) In practice, the regressor is learned from training data by running a linear regression. Suppose you have n *un-manipulated* data $(x_1, y_1), \cdots, (x_n, y_n)$, where $y_j \in \mathbb{R}$ is the payment of customer j. Now you want to use this "pure" data to learn a regressor that will perform well when facing strategic customers with manipulation cost $c(z; x) = \sum_{i=1}^d (z_i - x_i)^2$, i.e., the *squared* Euclidean distance.

  Formulate the problem of learning the regressor by minimizing the empiric risks but taking into account strategic behaviors of future customers. Prove that your formulated problem is a convex optimization problem.

Problem 3: Boosting (`Long, 0 points`)

(This question is just for those who are interested, and has 0 point. So do it only if you like to)

A fundamental concept in learning theory is *boosting*, intuitively means that classifiers that perform only slightly better than random guess can be turned into a classifier that is never wrong. In this question, you will prove a basic version of this celebrated result using the minimax theorem for zero-sum games.

Let $\mathcal{X} = \{x_1, \cdots, x_n\}$ be any feature space and $\mathcal{H} = \{h : X \to \{-1, 1\}\}$ be a set of classifiers over $\mathcal{X}$ (a.k.a., hypothesis class). For example, $\mathcal{H}$ could be the set of all linear classifiers. However, for simplicity, in this question we will assume that $\mathcal{H} = \{h_1, \cdots, h_m\}$ is finite. Let $g : \mathcal{X} \to \{-1, 1\}$ be the ground truth, i.e., the true label of x_j is $g(x_j)$.

The *weak learnability assumption* on $\mathcal{H}$ says that $\mathcal{H}$ is good in the following sense: there exists $\epsilon > 0$ such that for any distribution $p(\in \Delta_n)$ over $\mathcal{X}$, there exists a classifier h_i such that h_i is correct with probability at least $\frac{1}{2} + \epsilon$ for point x drawn from p, or more formally,

$$\sum_{j=1}^n p_j \cdot \mathbb{I}[h_i(x_j) = g(x_j)] \geq \frac{1}{2} + \epsilon,$$

where $\mathbb{I}[h_i(x_j) = g(x_j)]$ is the indicator function. That is, $\mathbb{I}[h_i(x_j) = g(x_j)]$ equals 1 if $h_i(x_j) = g(x_j)$ and equals 0 otherwise.

It turns out that weak learnability implies something much stronger — we can combine classifiers in $\mathcal{H}$ to construct a classifier that is always correct (a.k.a., *strong learnability*), formally stated as follows.

If $\mathcal{H}$ satisfies the weak learnability assumption, then there always exists a distribution $q(\in \Delta_m)$ over $\mathcal{H}$ such that the following weighted classifier:

$$h_q(x) = \begin{cases} 1 & if \quad \sum_{i=1}^{m} q_i h_i(x) \geq 0 \\ -1 & otherwise \end{cases}$$

is always correct, that is, $h_q(x) = g(x)$ for any $x \in \mathcal{X}$.

Prove the above statement.

[Hint: The classification problem can be viewed as a zero-sum game played between a *classifier designer* whose pure strategy is to pick a classifier from $\mathcal{H}$ and an *adversary* whose pure strategy is to pick a data point from $\mathcal{X}$. Think about how to define the payoff matrix of this game and what weak learnability means in the zero-sum game context.]