Homework #2
CMSC 35401: The Interplay of Economics and Machine Learning (Winter'24)

Due Saturday 02/03, 9:00 pm

General Instructions The assignment is meant to practice your understanding of course materials, and some of them are challenging. You are allowed 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, equally importantly, acknowledge everyone you discussed the homework with on your writeup. All course materials are available on the course webiste here https://www.haifeng-xu.com/cmsc35401win24. 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. Note that, the total score is meant to *not* be normalized to 100 (for instance, this HW has 30 in total for regular students and additional 15 points for those who take it as elective).

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

Problem 1: Rock-Paper-Scissor

In this problem, you will learn to master the rock-paper-scissor game. Recall that the game has the following payoff structure where each utility (x, y) means the row player receives x and the column player receives y.

	Rock	Paper	Scissor
Rock	(0, 0)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0, 0)	(-1, 1)
Scissor	(-1, 1)	(1, -1)	(0, 0)

Table 1: Payoffs of the Standard Rock-Paper-Scissor Game

1. (`Short,5 points`) Prove that the above rock-paper-scissor has a *unique* Nash equilibrium, which is that each player picks one of $\{Rock, Paper, Scissor\}$ uniformly at random.

2. (`Medium, 6 points`) Consider the situation where the column player is forbidden to play $Scissor$ (equivalently, the last column of the above payoff matrix is deleted). What is the Nash equilibrium of this new variant of the game.

3. (`Short, 5 points`) Consider the situation where the two players are encouraged to collaborate. In particular, if they play the same action, each will receive 0.5. This results in the following game variant. What is the Nash equilibrium of this new game?

	Rock	Paper	Scissor
Rock	(0.5, 0.5)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0.5, 0.5)	(-1, 1)
Scissor	(-1, 1)	(1, -1)	(0.5, 0.5)

Table 2: Payoffs of the Rock-Paper-Scissor Game with Encouraged Collaboration

Problem 2: Stackelberg Games

In this problem, you will learn another type of games called **Stackelberg games**. A Stackelberg game is a two-player game but with *sequential* player moves. In particular, a normal-form Stackelberg game is described by two matrices $A, B \in \mathbb{R}^{n \times m}$ where A is the payoff matrix of the row player who has action set $[n] = \{1, \cdots, n\}$ and B is the payoff matrix of the column player who has action set $[m] = \{1, \cdots, m\}$. The row player moves first (call *her* the *leader*) and the column player (call *him* the *follower*) moves second and thus can see the row player's strategy and then responds with his best action. Similar to the argument we saw in class, such a best response can without loss of generality be a pure best response. Sometimes there may be multiple best responses. In this case we assume that the follower is a benign player so that he will always pick the one that is the best for the leader, i.e., the follower breaks ties in favor of the leader.

It is not difficult to see that, after seeing the leader's strategy — either pure strategy or mixed strategy — the follower's best response action is easy to compute. That is, simply check the utility of each follower action $j \in [m]$ and then pick the best one. Therefore, research in Stackelberg games mainly focuses on computing the optimal leader strategy, which is also called the leader's Strong Stackelberg Equilibrium (SSE) strategy.

Answer the following questions about Stackelberg games.

1. (`Short, 5 points`) **A warm-up example.** Recall the traffic light game from Lecture 4, as follows. Assume that the row player is the leader and she can only play a pure strategy[1], what is the leader's SSE strategy?

	STOP	GO
STOP	(-3, -2)	(-3, 0)
GO	(0, -2)	(-100, -100)

Table 3: Payoffs of the Traffic Light Game

[1] For example, maybe because the follower can observe whatever pure action the leader takes.

2. (`Short, 5 points`) Consider the normal-form Stackelberg game and assume that the leader can only play a pure strategy. Show that there is a $\mathcal{O}(nm)$ time algorithm that computes the leader's pure SSE strategy.

3. (`Medium, 7 points`) We now consider the case where the leader can play a mixed strategy. To compute the leader's SSE (mixed) strategy, consider the following simpler *SSE with promise* problem. That is, imagine that there is an oracle who promises us that when the leader plays the mixed SSE strategy, the follower's best response action will be j^*. Show that given this credible promise, the leader's SSE strategy can be computed by a linear program.

   Explain how we can still compute the leader's SSE strategy efficiently even without the oracle's promise, by solving m different linear programs.

4. (`Medium, 7 points`) Prove that the leader's utility by playing the SSE mixed strategy (and the follower will best respond) is at least her utility in any Nash equilibrium of the game when players move simultaneously.

Additional Problems for those who take the course as Elective

These following problems are designed specifically for the very few *UChicago CS PhD* student who want to take the course as elective, hence will need these to fulfill their elective requirement.

Notes for regular students: *if you are interested*, you can try to solve these problems as well, but this is *not* a requirement for you. We will grade your solution just *by courtesy*, but your grades on these problems will NOT count towards your final grade.

Problem 3: Correlated Equilibrium (`Medium, 7 points`)

Let us come back to the collaborative version of rock-paper-scissor as in Problem 1(3), with payoff matrix as in Table 2.

Now suppose that you are an outsider who watches two players playing the above game variant with encouraged collaboration, and you can recommend actions to the two players using a correlated equilibrium. If you want to *minimize* the sum of the two players' expected utilities, which correlated equilibrium should you use? If you want to *maximize* the sum of their expected utilities, which correlated equilibrium should you use?

Hint: you can write code to compute solution for this problem, using any linear program solver, e.g., cvx /www.cvxpy.org/.

Problem 4: Boosting (`Long, 13 points`)

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.]