CSE 5401 Lec. 7

Games

1. Definitions of games and equilibria
2. Solving Zero-Sum games & Min-Max Theorem

Today we'll consider the problem of finding optimal moves in some very simple games: two players simultaneously take a single move, which determines the outcome. For example, rock-paper-scissors: each player has three possible moves, {rock, paper, scissors} If both players pick the same move, the game is a draw. Otherwise, paper beats rock; scissors beats paper; and rock beats scissors. We can represent the game with a pair of payoff matrices

"row" player:

	rock	paper	scissors
rock	0	-1	1
paper	1	0	-1
scissors	-1	1	0

"column" player:

	rock	paper	scissors
rock	0	1	-1
paper	-1	0	-1
scissors	1	-1	0

In general, if the row player plays move (index) r & the column player plays index c, then the row player receives the score in entry (r,c) of his matrix and the column player receives the score in entry (r,c) of her matrix. Both players want to maximize their score, but in rock-paper-scissors, the corresp entries all sum to zero — thus, equivalently, the column player wishes to minimize the row player's score and vice-versa.

Not all games are "zero sum." Consider the following game, known as "prisoner's dilemma"

	rat	hold
rat	-1	5
hold	-2	4

	rat	hold
rat	-1	-2
hold	5	4

Interrogating the prisoners in separate rooms...

Now the players would like to cooperate, and both play "hold out", netting 4 points. Unfortunately for these two, no matter what the other player does, they will net an additional point by ratting out their accomplace. Formally, it is a "dominating strategy". So, if both players are rational, they will rat, getting 1. Most games do not have dominating strategies — rock-paper-scissors doesn't. Indeed, if we commit to a single move, we lose. The situation is entirely different if we allow the players to choose moves randomly:

Defn A "mixed strategy" is a probability distribution on moves. We think of the pleyers as committing to mixed strategies, and then when it comes time to play the game, sampling their move from this distribution, independent of the other player. We focus on the expected score of a player. Given a mixed strategy for the opponent, a best response is any move that maximizes the player's expected score.

Lemma Suppose we fix a mixed strategy for one player. Then any mixed strategy that maximizes the opponent's score is only supported on best responses of the opponent.

Proof? If the player gives probability $p>0$ to some move with expected score less than a best response by $\delta>0$, then the mixed strategy has expected score that is less than simply playing the best response by at least $\delta p>0$. Hence it cannot be optimal. ▨

Indeed, the expected score of the mixed strategy is the same as simply playing any one of the best-

response moves in the support. Therefore, it is computationally easy to find one: e.g., given a mixed strategy $\vec{p}$ for the column player, a best response of the row player is $i^* = \arg\max_i (M\vec{p})_i$, where M is the row player's matrix. (Similarly for the column player.) The running time is dominated by the matrix-vector product, $O(n^2)$.

The players would love for the other player to "go first," so they can play a best response, maximizing their scores. Focusing on zero-sum games for now, the player who "goes first" has to solve the problem of finding a distribution $\vec{p}$ solving $\arg\max_{\vec{p}\in\Delta_m} \min_{\vec{q}\in\Delta_n} \vec{p}^T M \vec{q}$ (where the row player has m moves and the column player has n, and $\Delta_n = \{\vec{x} \in [0,1]^n : \|\vec{x}\|_1 = 1\}$ denotes the set of probability distributions.). Perhaps surprisingly, it does not matter who "chooses the distribution first":

Theorem ("Min-Max", von Neumann): for any $M \in \mathbb{R}^{m\times n}$,

$$\max_{\vec{p}\in\Delta_m} \min_{\vec{q}\in\Delta_n} \vec{p}^T M \vec{q} = \min_{\vec{q}\in\Delta_n} \max_{\vec{p}\in\Delta_n} \vec{p}^T M \vec{q}$$

Consider rock-paper-scissors: if the row player plays uniformly at random, the column player's expected scores are $[1/3\ 1/3\ 1/3]\begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix} = [0\ 0\ 0]$, so the column player cannot obtain a positive expected score. But the column player can also play uniformly at random, and then the row player cannot obtain a positive expected score. Thus, optimal

play in rock-paper-scissors yields an expected score of 0 for both players: they can guarantee they obtain this score, and cannot obtain a higher score. Such a pair of mixed strategies $(\vec{p}^*, \vec{q}^*)$ s.t. $\vec{p}^{*T} M \vec{q}^* = \max_{\vec{p}} \min_{\vec{q}} \vec{p}^T M \vec{q} = \min_{\vec{q}} \max_{\vec{p}} \vec{p}^T M \vec{q}$ is called an equilibrium: Given $\vec{p}^*$, no $\vec{q}'$ obtains a higher expected score than $\vec{q}^*$ and vice-versa. Beyond zero-sum games such equilibria still always exist

Theorem (Nash) For any pair $M \in \mathbb{R}^{m \times n}$ and $N \in \mathbb{R}^{m \times n}$, there exist $(\vec{p}^*, \vec{q}^*) \in \Delta_m \times \Delta_n$ such that $\vec{p}^* = \arg\max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q}^*$ and $\vec{q}^* = \arg\max_{\vec{q} \in \Delta_n} \vec{p}^{*T} N \vec{q}$.

Unfortunately, no polynomial-time algorithm for finding such Nash equilibria given the pair of matrices (M, N) is known, and there is some evidence that no such algorithm exists. The problem is believed not to be NP-hard because solutions always exist, but it is complete for a family of problems that includes, e.g., finding fixed points for discrete approximations to continuous functions. By contrast, we will give an algorithm for solving zero-sum games, and obtain a proof of the Min-Max Theorem as a corollary of its guarantee.

Indeed, the algorithm is an instantiation of the Multiplicative Weights Algorithm: We first rescale M so its entries lie in $[-1, 1]$. Now, the algorithm is (again)

Initialize $w_1^{(1)}, \ldots, w_m^{(1)} \leftarrow 1$

For $t = 1, \ldots, T$

| Put $\vec{p}^{(t)} = \frac{\vec{w}^{(t)}}{\sum_i w_i^{(t)}}$. Put $j^{(t)} \leftarrow \arg\min_j \left(\vec{p}^{(t)T} M\right)_j$

Put $\vec{r}^{(t)} \leftarrow M(\cdot, j^{(t)})$ [column $j^{(t)}$ of M]

$$w_i^{(t+1)} \leftarrow \begin{cases} w_i^{(t)} (1+\varepsilon)^{r_i^{(t)}} & \text{if } r_i^{(t)} \geq 0 \\ w_i^{(t)} (1-\varepsilon)^{-r_i^{(t)}} & \text{if } r_i^{(t)} < 0 \end{cases}$$

Return $\frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)}$.

Where we will take $\varepsilon = \frac{\delta}{2}$ and $T = \left\lceil \frac{4 \ln m}{\delta^2} \right\rceil$.

Theorem The above algorithm returns $\vec{p}^* \in \Delta_m$ such that $\min_{\vec{q} \in \Delta_n} \vec{p}^{*T} M \vec{q} \geq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q} - \delta$. It runs in time $O\left(\frac{1}{\delta^2} mn \ln m\right)$.

Proof: The running time is immediate for our choice of T.

Recall the Multiplicative Weight Algorithm's guarantee:

$\frac{1}{T}\sum_{t=1}^{T} \langle \vec{p}^{(t)}, \vec{r}^{(t)} \rangle \geq \frac{1}{T}\sum_{t=1}^{T} r_i^{(t)} - \delta$ for this ε & T, for all i.

which for us is $\frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)T} M(\cdot, j^{(t)}) \geq \frac{1}{T}\sum_{t=1}^{T} M_{ij^{(t)}} - \delta$.

Consider any $\vec{p}' \in \Delta_n$. If we add the i^{th} ineq. w/ wt. p_i', we get $\frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)T} M(\cdot, j^{(t)}) \geq \frac{1}{T}\sum_{t=1}^{T} (\vec{p}'^T M)_{j^{(t)}} - \delta = \vec{p}'^T M \left(\sum_{t=1}^{T} \frac{1_{j^{(t)}}}{T}\right) - \delta$

In particular, $\vec{p}' = \arg\max_{\vec{p}} \vec{p}^T M \left(\sum_{t=1}^{T} \frac{1_{j^{(t)}}}{T}\right)$ gives the RHS is at least $\max_{\vec{p} \in \Delta_m} \vec{p}^T M \left(\sum_{t=1}^{T} \frac{1_{j^{(t)}}}{T}\right) \geq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q} - \delta$. Now, how does $\frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)} M(\cdot, j^{(t)})$ compare to $\min_{\vec{q} \in \Delta_n} \left(\frac{1}{T}\sum_{t=1}^{T} p^{(t)}\right)^T M \vec{q}$? $j^{(t)}$ was a best response to $\vec{p}^{(t)}$, so for all $\vec{q}$ $\vec{p}^{(t)} M(\cdot, j^{(t)}) \leq \vec{p}^{(t)} M \vec{q}$. For any fixed $\vec{q}$, averaging over $t = 1, \cdots, T$ gives $\frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)} M(\cdot, j^{(t)}) \leq \frac{1}{T}\sum_{t=1}^{T} p^{(t)} M \vec{q}$.

Thus: $\min_{\vec{q}} \left(\frac{1}{T}\sum_{t=1}^{T} p^{(t)}\right) M \vec{q} \geq \frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)} M(\cdot, j^{(t)}) \geq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q} - \delta$.

We therefore see $\frac{1}{T}\sum_{t=1}^{T} \vec{p}^{(t)}$ is as claimed ▨

With the construction of mixed strategies provided by the Multiplicative Weights Algorithm, it is easy to prove the Min-Max Theorem:

Proof of Min-Max Theorem For every $\delta > 0$, we have a strategy $\vec{p}^*$ s.t. $\min_{\vec{q} \in \Delta_n} \vec{p}^{*T} M \vec{q} \geq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q} - \delta$. It is immediate therefore that $\max_{\vec{p} \in \Delta_m} \min_{\vec{q} \in \Delta_n} \vec{p}^T M \vec{q} \geq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q}$ by letting $\delta \to 0$. At the same time, $\min_{\vec{q} \in \Delta_n} \vec{p}' M \vec{q} \leq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q}$ for any $\vec{p}' \in \Delta_m$ so $\max_{\vec{p} \in \Delta_m} \min_{\vec{q} \in \Delta_n} \vec{p}^T M \vec{q} \leq \min_{\vec{q} \in \Delta_n} \max_{\vec{p} \in \Delta_m} \vec{p}^T M \vec{q}$. Hence, they are equal. ▨

The min-max theorem is also useful in the analysis of algorithms, specifically lower bounds for randomized algs. Let's consider the following zero-sum game: the strategies for the row player are deterministic algorithms running in t steps, and the strategies for the column player are inputs of size n. (Both are effectively finite.) If the alg. is correct on the input, the row player gets payoff 1, and otherwise 0. A <u>randomized</u> alg. corresponds to a mixed strategy for the row player: and the expected payoff on any input is the probability the alg. is correct. The worst-case input is a best response, so the value of the game is the worst-case probability that the algorithm is correct. But now we can invoke the Min-Max Theorem, and we find that this is the same as the best success probability of a deterministic algorithm for a worst case distribution on inputs. Deterministic algorithms on average-case inputs are much easier to reason about.