

Topic 6: Game Theory and Lower Bounds for Randomized Algorithms

David Woodruff

Outline

- 2-player zero-sum games and minimax optimal strategies
- Connection to randomized algorithms
- General sum games, Nash equilibria

Game Theory

- How people make decisions in social and economic interactions
 - Applications to computer science
- Users interacting with each other in large systems
 - Routing in large networks
 - Auctions on Ebay

Definitions

- A game has
 - Participants, called **players**
 - Each player has a set of choices, called **actions**
 - Combined actions of players leads to **payoffs** for each player

Shooter-Goalie Game

- 2 players: shooter and goalie
- Shooter has 2 actions: shoot to her left or shoot to her right
- Goalie has two actions: dive to shooter's left or to shooter's right
 - left and right are defined with respect to shooter's actions
- Set of actions for both Shooter and Goalie is {L, R}
- If shooter and goalie each choose L, or each choose R, then goalie makes a save
- If shooter and goalie choose different actions, then the shooter makes a goal

Payoff Matrix

- If goalie makes a save, goalie has payoff +1, shooter has payoff -1
- If shooter makes a goal, goalie has payoff -1, shooter has payoff +1

payoff matrix M	goalie	
	L	R
shooter L	(-1, 1)	(1, -1)
R	(1, -1)	(-1, 1)

- Payoff is (r,c) , where r is payoff to row player, and c is payoff to the column player
- For each entry (r,c) , $r+c = 0$. This is called a **zero-sum game**
- Zero-sum game does not imply "fairness". If all entries are $(1,-1)$ it is still zero-sum

An Aside

- Row-payoff matrix R consists of the payoffs to the row player
- C is the column-payoff matrix
- $M_{i,j} = (R_{i,j}, C_{i,j})$ for all i and j

payoff matrix M	goalie	
	L	R
shooter L	(-1, 1)	(1, -1)
R	(1, -1)	(-1, 1)

Row payoff matrix R	goalie	
	L	R
shooter L	-1	1
R	1	-1

- $R + C = 0$ for zero-sum games

Pure and Mixed Strategies

- How should the players play?
- **Pure strategy:**
 - Row player chooses a deterministic action I
 - Column player chooses a deterministic action J
 - Payoff is $R_{I,J}$ for row player, and $C_{I,J}$ for column player
- Pure strategies are deterministic, what about randomized strategies?
 - Players have a distribution over their actions
 - Row player decides on a $p_i \in [0,1]$ for each row, with $\sum_{\text{actions } i} p_i = 1$
 - Column player decides on a $q_j \in [0,1]$ for each column, with $\sum_{\text{actions } j} q_j = 1$
- Distributions p and q are **mixed strategies**

How to define payoff for mixed strategies?

Expected Payoff

- Assume players have independent randomness
- $V_R(p, q) = \sum_{i,j} \Pr[\text{row player plays } i, \text{column player plays } j] \cdot R_{i,j} = \sum_{i,j} p_i q_j R_{i,j}$
- $V_C(p, q) = \sum_{i,j} \Pr[\text{row player plays } i, \text{column player plays } j] \cdot C_{i,j} = \sum_{i,j} p_i q_j C_{i,j}$
- What is $V_R(p, q) + V_C(p, q)$?
 - 0, since zero-sum game

payoff matrix M	goalie	
	L	R
shooter L	(-1, 1)	(1, -1)
R	(1, -1)	(-1, 1)

If $p = (.5, .5)$ and $q = (.5, .5)$ what is V_R ?
 $V_R = .25 \cdot (-1) + .25 \cdot 1 + .25 \cdot 1 + .25 \cdot (-1)$

If $p = (.75, .25)$ and $q = (.6, .4)$ what is V_R ?
 $V_R = -0.1$

Minimax Optimal Strategies

- Row player wants a distribution p^* maximizing her expected payoff over all strategies q of her opponent
- p^* achieves lower bound $lb = \max_p \min_q V_R(p, q)$

$$lb := \max_p \overbrace{\min_q V_R(p, q)}^{\text{mixed strategy that maximizes the minimum expected payoff}}$$

payoff when opponent plays
optimal strategy against our choice p

- The row player can guarantee this payoff no matter what the column player does. lb is a lower bound on the row-player's payoff

Minimax Optimal Strategies

- Column player wants distribution q^* maximizing his expected payoff over all strategies p of his opponent
 - q^* achieving $\max_q \min_p V_C(p, q)$

Claim: $\max_q \min_p V_C(p, q) = - \min_q \max_p V_R(p, q)$

Proof: $\max_q \min_p V_C(p, q) = \max_q \min_p -V_R(p, q)$
 $= \max_q (- \max_p V_R(p, q))$
 $= - \min_q \max_p V_R(p, q)$

Payoff to row player if column player plays q^* is $ub = \min_q \max_p V_R(p, q)$

Column player can guarantee the row player does not achieve a larger expected payoff, so this is an upper bound ub on row player's expected payoff

Lower and Upper Bounds

- Row player guarantees she has expected payoff at least

$$lb = \max_p \min_q V_R(p, q)$$

- Column player guarantees row player has expected payoff at most

$$ub = \min_q \max_p V_R(p, q)$$

$lb \leq ub$, but how close is lb to ub ?

A Pure Strategy Observation

- Suppose we want to find row player's optimal strategy p^*
- **Claim:** can assume column player plays a pure strategy. **Why?**
 - For any strategy p of the row player, $V_R(p, q) = \sum_{i,j} p_i q_j R_{i,j} = \sum_j q_j \cdot (\sum_i p_i R_{i,j})$
 - Column player can choose q to be the j for which $\sum_i p_i R_{i,j}$ is minimal
- $lb = \max_p \min_q V_R(p, q) = \max_p \min_j \sum_i p_i R_{i,j}$
- $ub = \min_q \max_p V_R(p, q) = \min_q \max_i \sum_j q_j R_{i,j}$

Shooter-Goalie Example

payoff matrix M	goalie	
	L	R
shooter L	(-1, 1)	(1, -1)
R	(1, -1)	(-1, 1)

Claim: minimax-optimal strategy for both players is (.5, .5)

Proof: For the shooter (row-player), let $\mathbf{p} = (p_1, p_2)$ be the minimax optimal strategy

$p_1 \geq 0, p_2 \geq 0$, and $p_1 + p_2 = 1$. Write $\mathbf{p} = (p, 1-p)$ with p in $[0,1]$

Suppose goalie (column-player) plays L

Shooter's payoff is $p \cdot (-1) + (1-p) \cdot (1) = 1 - 2p$

Suppose goalie plays R


Shooter's payoff is $p \cdot (1) + (1-p) \cdot (-1) = 2p - 1$

Choose $p \in [0,1]$ to maximize $lb = \max_p \min(1 - 2p, 2p - 1)$

$p = \frac{1}{2}$ realizes this, and $lb = 0$

Similarly show optimal strategy $\mathbf{q} = (q_1, q_2)$ of goalie is $(1/2, 1/2)$ and $ub = 0$

$ub = lb = 0$, which is the **value of the game**



Asymmetric Shooter-Goalie

	L	R
shooter L	$(-\frac{1}{2}, \frac{1}{2})$	(1, -1)
R	(1, -1)	(-1, 1)

Goalie is now weaker on the left

Let $\mathbf{p} = (p_1, p_2)$ be the minimax optimal shooter (row-player) strategy

Suppose goalie (column player) plays L

Shooter's payoff is $p \cdot (-\frac{1}{2}) + (1-p) \cdot (1) = 1 - (\frac{3}{2})p$

Suppose goalie plays R

Shooter's payoff is $p \cdot (1) + (1-p) \cdot (-1) = 2p - 1$

Choose $p \in [0,1]$ to maximize $lb = \max_p \min(1 - (\frac{3}{2})p, 2p - 1)$

Maximized when $1 - (\frac{3}{2})p = 2p - 1$, so $p = 4/7$, and $lb = 1/7$

What is the column player's minimax strategy?

Asymmetric Shooter-Goalie

	L	R
shooter L	$(-\frac{1}{2}, \frac{1}{2})$	(1, -1)
R	(1, -1)	(-1, 1)

Let $\mathbf{q} = (q, 1 - q)$ be the minimax optimal goalie (column-player) strategy

Suppose shooter (row player) plays L

Goalie's payoff is $q \cdot (\frac{1}{2}) + (1-q) \cdot (-1) = \frac{3q}{2} - 1$

Suppose shooter plays R

Goalie's payoff is $q \cdot (-1) + (1-q) \cdot (1) = 1 - 2q$

Choose $q \in [0,1]$ to realize $\max_q \min(\frac{3q}{2} - 1, 1 - 2q)$

$\frac{3q}{2} - 1 = 1 - 2q$ implies $q = 4/7$, and expected payoff at least $-1/7$

Remember: this means row player's ub at most $1/7$


Uhh... $lb = ub$ again... Value of the game is $1/7$

Another Example

- Suppose in a zero-sum game, Row player's payoffs are:
 - 1 -2
 - 1 2
- What is row player's minimax strategy? **Why?**
- Suppose her distribution is $(p, 1-p)$
- Expected payoff if column player plays first action is:

$$p \cdot (-1) + (1-p) \cdot 1 = 1 - 2p$$
- Expected payoff if column player plays second action is:

$$p \cdot (-2) + (1-p) \cdot 2 = 2 - 4p$$
- These lines both have a negative slope
- Should play $p = 0$
- Can show column player should always play first action and value of game is 1



Exercise 1: What if both players have somewhat different weaknesses? What if the payoffs are:

$$\begin{array}{cc} (-1/2, 1/2) & (3/4, -3/4) \\ (1, -1) & (-3/2, 3/2) \end{array}$$

Show that minimax-optimal strategies are $\mathbf{p} = (2/3, 1/3)$, $\mathbf{q} = (3/5, 2/5)$ and value of game is 0.

Exercise 2: For the game with payoffs:

$$\begin{array}{cc} (-1/2, 1/2) & (3/4, -3/4) \\ (1, -1) & (-2/3, 2/3) \end{array}$$

Show that minimax-optimal strategies are $\mathbf{p} = (\frac{4}{7}, \frac{3}{7})$, $\mathbf{q} = (\frac{17}{35}, \frac{18}{35})$ and value of game is $\frac{1}{7}$.

Exercise 3: For the game with payoffs:

$$\begin{array}{cc} (-1/2, 1/2) & (-1, 1) \\ (1, -1) & (2/3, -2/3) \end{array}$$

Show that minimax-optimal strategies are $\mathbf{p} = (0, 1)$, $\mathbf{q} = (0, 1)$ and value of game is $\frac{2}{3}$.

Von Neumann's Minimax Theorem

- In each example,
 - row player has a strategy \mathbf{p}^* guaranteeing a payoff of lb for him
 - column player has a strategy \mathbf{q}^* guaranteeing row player's payoff is at most ub
 - $lb = ub!$
- Von Neumann:** Given a finite 2-player zero-sum game,

$$lb = \max_p \min_q V_R(p, q) = \min_q \max_p V_R(p, q) = ub$$
 Common value is the **value of the game**
- In a zero-sum game, the row and column players can tell their strategy to each other and it doesn't affect their expected performance!
 - Don't tell each other your randomness

Lower Bounds for Randomized Algorithms


- A randomized algorithm is a zero-sum game
- Create a row-payoff matrix R :
 - Rows are possible inputs (for sorting, $n!$)
 - Columns are possible algorithms (e.g. every algorithm for sorting)
 - $R_{i,j}$ is cost of algorithm j on input i (e.g. number of comparisons)
- A deterministic algorithm with good worst-case guarantee is a column with entries that are all small
- A randomized algorithm with good expected guarantee is a distribution \mathbf{q} on columns so the expected cost in each row is small

Lower Bounds for Randomized Algorithms

- Minimax-optimal strategy for column player is best randomized algorithm
- A lower bound for a randomized algorithm is a distribution \mathbf{p} on inputs so for every algorithm j , expected cost of running j on input distribution \mathbf{p} is large
 - show lb is large for the game
 - give strategy for the row player (distribution on inputs) such that every column (deterministic algorithm) has high cost

Lower Bounds for Randomized Sorting

- **Theorem:** Let A be a randomized comparison-based sorting algorithm. There's an input on which A makes an expected $\Omega(\lg n!)$ comparisons
- **Proof:** consider uniform distribution on $n!$ permutations of n distinct numbers
 - $n!$ leaves
 - No two inputs go to same leaf
 - How many leaves at depth $\lg(n!) - 10$?
 - $\leq 1+2+4+\dots + 2^{(\lg n!) - 10} \leq \frac{n!}{512}$
 - $511/512 > .99$ fraction of inputs are at depth $> \lg(n!) - 10$
 - Expected depth $> .99(\lg(n!) - 10) = \Omega(\lg n!)$



General-Sum Two-Player Games

- Many games are not zero-sum, have “win-win” or “lose-lose” payoffs
- Game of “chicken”
- Suppose two drivers facing each other each drive on their left (L) or right (R)

payoff matrix M		Bob	
		L	R
Alice	L	(1, 1)	(-1, -1)
	R	(-1, -1)	(1, 1)

- What is a good notion of optimality to look at?

Nash Equilibria

- (\mathbf{p}, \mathbf{q}) is stable if no player has incentive to individually switch strategy
 - For any other strategy \mathbf{p}' of row player,
 - row player's new payoff $= \sum_{i,j} p'_i q_j R_{i,j} \leq \sum_{i,j} p_i q_j R_{i,j} =$ row player's old payoff
 - For any other strategy \mathbf{q}' of column player,
 - column player's new payoff $= \sum_{i,j} p_i q'_j C_{i,j} \leq \sum_{i,j} p_i q_j C_{i,j} =$ column player's old payoff
- For chicken, $((1,0),(1,0))$ and $((0,1),(0,1))$ and $((1/2,1/2),(1/2,1/2))$ are Nash Equilibria
- **Theorem (Nash):** Every finite player game (with a finite number of strategies) has a Nash equilibrium