CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

Allan Borodin
Tyrone Strangway and Young Wu (TAs)

September 19, 2016

Lecture 4

- Announcements
  - Tutorial this Friday in SS 1086
  - First 3 questions for assignment 1 have been posted.
  - Talk of possible interetst: This coming Tuesday, September 20, there will be a seminar "Preferences and Manipulative Actions in Elections" by Gabor Erdelyi. The talk will take place in Pratt 266 at 11AM.
  - Another talk of possible interest: Thursday, September 29, there will be a seminar "When Should an Expert Make a Prediction?" by Amir Ban. The talk will take place in the first floor conference room of the Fields Institute, 222 College Street. The time will either be at noon or at 1PM.
- Today's agenda
  - Briefly discuss some terminology
  - Zero-sum games; Read chapter 2 (excluding section 2.6) and section 3.2 (Hide and Seek Games)

Some terminology

- In lecture 2, we mentioned the concept of a *dominant strategy*. I have clarified that discussion. Here is how that now reads:

  A *(weakly) dominant strategy* D for a player P is one in which P will achieve the maximize payoff possible for any given strategy profile of the other players. (Usually one assumes that D is better for P for at least one strategy profile of the other players. If two strategies have the same value for all strategy profiles of the other players, then we do not have to distinguish these two strategies.)

  A *strictly dominant strategy* is one that yields a better value for all strategy profiles of the other players.

  Unless otherwise stated, dominate strategy will mean weakly dominant and we will assume that there do not exist two strategies for a player that have the same value for all strategy profiles of the other players.

Terminology continued

- In the critique of Nash equilibria, I mentioned that so far we have only been discussing "full information games". That terminology is not standard. I am now changing that to the following terminology that is consistent with the text and reasonably well (although perhaps not universally) accepted:

  There are two distinctions to be made, games of perfect vs imperfect information and complete vs incomplete information. The discussion thus far has been restricted to games of *perfect and complete information* where we know everyones strategies and precise values for all strategy profiles. When we do not know the precise values but have a probabilistic *prior* belief about the distributions of the values, then we have a Bayesian game of *perfect but incomplete information*.

  We return to these distinctions in Chapter 6 (Games in extensive form).

Zero-sum games

We now jump back to Chapters 2 and 3 (but only section 3.2) in the KP text where the topic is zero-sum games which are a restriction of general-sum games.

In two player general-sum games, each entry of the payoff matrix $U[i,j]$ is a pair $(a_{i,j}, b_{i,j})$ where $a_{i,j}$ (respectively, $b_{i,j}$)) is the payoff for the first (i.e. row) player (resp. the payoff for the second player) when the first player uses strategy i and the second player uses strategy j.

Note: In the chapters on zero-sum games, strategies are often called "actions".

Zero-sum games

We now jump back to Chapters 2 and 3 (but only section 3.2) in the KP text where the topic is zero-sum games which are a restriction of general-sum games.
In two player general-sum games, each entry of the payoff matrix $U[i,j]$ is a pair $(a_{i,j}, b_{i,j})$ where $a_{i,j}$ (respectively, $b_{i,j}$)) is the payoff for the first (i.e. row) player (resp. the payoff for the second player) when the first player uses strategy i and the second player uses strategy j.
Note: In the chapters on zero-sum games, strategies are often called "actions".

In a *zero-sum game*, we have $a_{i,j} = -b_{i,j}$ for all i, j. That is, what one player (say the row player) "gains" the other playeer loses. In particular, we can have two player games where each entry of the matrix is either $+1$ or -1; that is, when one player wins the game, the other loses and we often consider "which player has a winning strategy".

Zero-sum games

We now jump back to Chapters 2 and 3 (but only section 3.2) in the KP text where the topic is zero-sum games which are a restriction of general-sum games.
In two player general-sum games, each entry of the payoff matrix $U[i,j]$ is a pair $(a_{i,j}, b_{i,j})$ where $a_{i,j}$ (respectively, $b_{i,j}$)) is the payoff for the first (i.e. row) player (resp. the payoff for the second player) when the first player uses strategy i and the second player uses strategy j.
Note: In the chapters on zero-sum games, strategies are often called "actions".

In a *zero-sum game*, we have $a_{i,j} = -b_{i,j}$ for all i, j. That is, what one player (say the row player) "gains" the other playeer loses. In particular, we can have two player games where each entry of the matrix is either $+1$ or -1; that is, when one player wins the game, the other loses and we often consider "which player has a winning strategy".

Since $a_{i,j}$ determines $b_{i,j}$, we only need to specify $a_{i,j}$ so that we will represent zero-sum games by a real-valued matrix A with $A[i,j] = a_{i,j}$

Zero-sum games: revisiting some concepts

As in general-sum games, we will again have pure and mixed strategies and for any choice of player strategies there is a well defined payoff to each player. We will also always have mixed Nash equilibria (and sometimes pure equilbria). Zero-sum games are, however, quite special and we will see that they will possess a special property. Namely,

- There is the concept of the *value* V of the game which is an amount that one player can guarantee as its minimum expected gain and the other can guarantee as its maximum expected loss.

Zero-sum games: revisiting some concepts

As in general-sum games, we will again have pure and mixed strategies and for any choice of player strategies there is a well defined payoff to each player. We will also always have mixed Nash equilibria (and sometimes pure equilbria). Zero-sum games are, however, quite special and we will see that they will possess a special property. Namely,

- There is the concept of the *value* V of the game which is an amount that one player can guarantee as its minimum expected gain and the other can guarantee as its maximum expected loss.
- Every Nash equlibrium has the same payoff to each player, the payoff being the value V of the game. Recall that in general-sum games there can be many NE with different payoffs to the individual players.

Von Neumann's Minimax Theorem and the value of a zero sum game

Let Δ_k be the set of probability distributions over a set of k possiibilities.

Suppose player 1 has m possible strategies and player 2 has n possible strategies. Then (as before) we can represent a mixed strategy by a pair of vectors $(\mathbf{x}, \mathbf{y})$ where $\mathbf{x} = (x_1, \ldots, x_m) \in \Delta_m$ with x_i being the probability that player 1 chooses strategy i and similarly $\mathbf{y} = (y_1, \ldots, y_n) \in \Delta_n$ with y_j the probability that player 2 chooses strategy j.

Von Neumann's Minimax Theorem and value V of a game

Let $A[a_{ij}]$ be the payoff matrix for a zero sum game. Then $V = \max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} = \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}$

Lets understand the meaning of this result using a few examples of zero-sum games; then we will return to discuss the significance of the result and a sketch of its proof.

A zero-sum game with a pure NE

		player II	
		action 1	action 2
player I	action 1	2	3
	action 2	1	0

Figure : Example in section 2.3.1 in KP text

In this game, the pair (action 1,action 1) is a pure NE (also called a *saddle point* in definition 2.3.1 of KP text). The value of this game is 2. That is, no matter what player II does, player one is guaranteed a gain of at least 2, and no matter what player I does, player II can lose at most 2.

Pick a hand: an example of a zero sum game with no pure NE

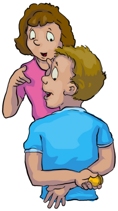

Figure : The "pick a hand" game; figure 2.1 in KP text

Chooser \ Hider	$L1$	$R2$
L	1	0
R	0	2

How can the Hider minimize his potential loss?

We will use the principle of indifference to compute the value of this game. Suppose the Hider plays L1 with probability y_1 and hence R1 with probability$1 - y_1$.

How can the Hider minimize his potential loss?

We will use the principle of indifference to compute the value of this game. Suppose the Hider plays L1 with probability y_1 and hence R1 with probability$1 - y_1$.

Then the Hider is trying to minimize $\max\{(y_1 \cdot 1, (1 - y_1) \cdot 2\}$. To achieve the minimum value, we must have

$$y_1 \cdot 1 = (1 - y_1) \cdot 2$$

so that $y_1 = \frac{2}{3}$ and Hider's loss (i.e. the value V of the game) is $\frac{2}{3}$.

How can the Hider minimize his potential loss?

We will use the principle of indifference to compute the value of this game. Suppose the Hider plays L1 with probability y_1 and hence R1 with probability$1 - y_1$.

Then the Hider is trying to minimize $\max\{(y_1 \cdot 1, (1 - y_1) \cdot 2\}$. To achieve the minimum value, we must have

$$y_1 \cdot 1 = (1 - y_1) \cdot 2$$

so that $y_1 = \frac{2}{3}$ and Hider's loss (i.e. the value V of the game) is $\frac{2}{3}$.

Then by the minimax theorem, the Chooser can always guarantee a gain of $V = \frac{2}{3}$ by an approriate mixed strategy.

How can the Hider minimize his potential loss?

We will use the principle of indifference to compute the value of this game. Suppose the Hider plays L1 with probability y_1 and hence R1 with probability$1 - y_1$.

Then the Hider is trying to minimize $\max\{(y_1 \cdot 1, (1 - y_1) \cdot 2\}$. To achieve the minimum value, we must have

$$y_1 \cdot 1 = (1 - y_1) \cdot 2$$

so that $y_1 = \frac{2}{3}$ and Hider's loss (i.e. the value V of the game) is $\frac{2}{3}$.

Then by the minimax theorem, the Chooser can always guarantee a gain of $V = \frac{2}{3}$ by an approriate mixed strategy.

The Chooser's mixed strategy is also (L,R) $= (\frac{2}{3}, \frac{1}{3})$.

The (simplified) soccer penalty kick game

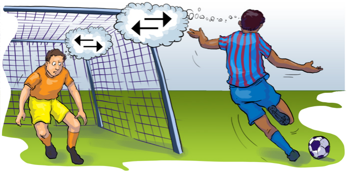

FIGURE 4. The game of Penalty Kicks.

In soccer, a player taking a penalty shot generally tries for the right or left hand side of the net. Given the size of the goal, the goalie has to make a decision whether to leap left or right in order to have a chance of stopping a kick.

The (simplified) soccer penalty kick game continued

In the introductory chapter, the KP text provides a matrix (for data collected by Palacios-Huerta) representing the success probabilty of scoring on a penalty kick based on 1417 penalty kicks in European professional games. Here 'R' represents the dominant side (i.e. right-footed or left-footed) of the kicker (which is known to the goalie).

		goalie	
		L	R
kicker	L	0.58	0.95
	R	0.93	0.70

Figure : Matrix for probability of scoring based on data collected from professional games

The (simplified) soccer penalty kick game continued

Given these probabilities, the minimax strategy for the kicker is to kick to the dominant side R with probability $p = .62$ and the goalie's minimax strategy is to leap to the dominant side R with probability $q = .58$.

The (simplified) soccer penalty kick game continued

Given these probabilities, the minimax strategy for the kicker is to kick to the dominant side R with probability $p = .62$ and the goalie's minimax strategy is to leap to the dominant side R with probability $q = .58$.

You should verify these probabilities and compute the value of the game

The (simplified) soccer penalty kick game continued

Given these probabilities, the minimax strategy for the kicker is to kick to the dominant side R with probability $p = .62$ and the goalie's minimax strategy is to leap to the dominant side R with probability $q = .58$.

You should verify these probabilities and compute the value of the game

Palacios-Huerta reports that the actual strategy fractions were $p' = .60$ and $q' = .577$. This provides some evidence that the players (as a community) have learned to use minimax strategies.

The easy direction of the minimax theorem

An equality $A = B$ is equivalent to the two inequalities, $A \leq B$ and $B \leq A$.

The easy direction of the minimax theorem

$\max_{\mathbf{x} \in \Delta_m} \min_{\mathbf{y} \in \Delta_n} \mathbf{x}^T A \mathbf{y} \leq \min_{\mathbf{y} \in \Delta_n} \max_{\mathbf{x} \in \Delta_m} \mathbf{x}^T A \mathbf{y}$

Proof: We are only considering the case of finitely many strategies so the proof in Lemma 2.5.3 suffices since $X = \Delta_m$ and $Y = \Delta_n$ are closed and bounded sets and $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y}$ is a continuous function guaranteeing the existence of the max and min.
Consider any $(\mathbf{x}', \mathbf{y}')$ such that $\mathbf{x}'$ (resp. $\mathbf{y}'$) is in the support X of the row (resp. column) player's distribution Y).

$$\min_{\mathbf{y} \in Y} f(\mathbf{x}', y) \leq f\mathbf{x}', \mathbf{y}') \leq \max_{\mathbf{x} \in X} f(\mathbf{x}, \mathbf{y}')$$

Since this holds for arbitrary $\mathbf{x}'$,

$$\max_{\mathbf{x}' \in X} \min_{\mathbf{y} \in Y} f(\mathbf{x}', \mathbf{y}) \leq \max_{\mathbf{x} \in X} f(\mathbf{x}, \mathbf{y}')$$

And then minimizing over $\mathbf{y} \in Y$, the desired inequality is obtained.

The other direction of the minimax theorem

The harder direction of the minimax theorem is derived in the KP text using a result called the separating hyperplane theorem.

The separating hyperplane theorem

Consider a region K in $\mathbb{R}^d$ that is closed (i.e. contains all of its limit points) and convex (i.e. any point on the line between two points in K is also in K).
Then if $\mathbf{0} \notin K, \exists \mathbf{z} \in \mathbb{R}^d, c \in \mathbb{R} : 0 < c < \mathbf{z}^T \mathbf{v}$ for all $\mathbf{v} \in K$.
That is, the hyperplane $\{\mathbf{x} \in \mathbb{R}^d : \mathbf{z}^T \mathbf{x} = c\}$ separates K from 0.

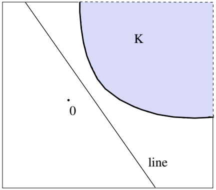

Figure : Illustration (figure 2.4) of separating hyperplane for $d = 2$

The separating hyperplane theorem in machine learning

The separating hyperplane theorem is the basis for a binary classification method in machine learning called support vector machines.
CSC 2545 - Kernel Methods and Support Vector Machines

The idea is that vectors represent various feature values of examples that need to be classified as being "good examples" (e.g. movies we like) or "bad examples" (movies we don't like). Given a training set, we compute this hyperplane (ignoring misclassified points and how to use kernal maps so as to apply the method in general). So if K are the good examples, the support vector is the vector that determines the separating hyperplane that will be used to classify new points.

Furthermore the more positive (resp. negative) $\mathbf{z}^T\mathbf{v}$, the more confidant (or the better the evidence) that the example represented by $\mathbf{v}$ is a good (resp. bad) example.

Sketch of the harder direction of the minimax theorem: see Theorem 2.5.4 in KP text

We will assume the separating hyperplane theorem.
Suppose by way of contradiction that

$$\max_{\mathbf{x}\in\Delta_m}\min_{\mathbf{y}\in\Delta_n}\mathbf{x}^T A\mathbf{y} < \lambda < \min_{\mathbf{y}\in\Delta_n}\max_{\mathbf{x}\in\Delta_m}\mathbf{x}^T A\mathbf{y}$$

for some λ.

By defining a new game with $\tilde{a}_{ij} = a_{ij} - \lambda$, we have reduced the payoff in each entry of the matrix by λ and hence the expected payoff of every mixed strategy (for the new game) is also reduced by λ. Therefore

$$\max_{\mathbf{x}\in\Delta_m}\min_{\mathbf{y}\in\Delta_n}\mathbf{x}^T \tilde{A}\mathbf{y} < 0 < \min_{\mathbf{y}\in\Delta_n}\max_{\mathbf{x}\in\Delta_m}\mathbf{x}^T \tilde{A}\mathbf{y}$$

Sketch of harder direction continued

To establish the contradiction, one shows that

1. The set $K = \{\tilde{A}\mathbf{y} + \mathbf{v}, \mathbf{y} \in \Delta_n, \mathbf{v} \geq \mathbf{0}$ satisfies the conditions of the separating hyperplane theorem. That is,
   - K is convex and closed
   - $\mathbf{0} \notin K$
2. Therefore $\exists \mathbf{z}$ and c such that $0 < c < \mathbf{z}^T(\tilde{A}\mathbf{y} + \mathbf{v})$
3. Furthermore $\mathbf{z} \geq \mathbf{0}$ and $\mathbf{z} \neq \mathbf{0}$.
4. The mixed strategy $\mathbf{x}' = \mathbf{z}/\sum z_i$ gives a positive expected gain against any mixed strategy $\mathbf{y}$ establishing the contradiction.

Two important computational applications of the minimax theorem

We will discuss two important applications of the minimax theorem.

1. The first application is actually an equivalent result, namely the (LP) **duality theorem** of linear programming (LP). Linear programming is one of the most important concepts in combinatorial optimization (leading to efficient optimal and approximation algorithms) and LP duality is arguably the central theorem of linear programmming. Since LPs can be solved optimally in polynomial time, this will imply that (unlike general-sum NEs), we can always solve zero-sum games (i.e. find the mixed strategies that yield the value of the game).
2. The second application is The Yao Principle which is a direct consequence of the minimax theorem and is a basic tool in proving "negative" results for randomized algorithms.

Integer programming, linear programming and LP duality

An integer program (IP) (resp. LP) formulates an optimization problem as the maximization or minimization of a linear function of integral (resp. real) variables subject to a set of linear inequalities and equalities.

Most combinatorial optimization have reasonably natural IP (and sometimes LP) formulations. But solving an IP is an NP complete problem so that one does not *in general* expect efficient algorithms for optimally solving IPs. However, many IPs are solved "in practice" by IPs and there are classes of IPs that do have worst case efficient algorithms.

Another important idea in approximately solving IPs is to relax the integral constraints to real valued constraints (e.g. $x_i \in \{0, 1\}$ is relaxed to $x_i \in [0, 1]$) and then "rounding" the fractional solution to an integral solution.
Note: if the objective function and all constraints have rational coefficients, then we can assume rational solutions for an LP.