Lecture 9 - Game Theory

Optimization and Approximation - ENS M1

Nicolas Bousquet

1 *Zero-sums games*

Game theory is a wide area in theoretical CS. Games are usually partitioned into two sets: cooperative and non-cooperative games, each of them are again divided into several categories. During this lecture, we will concentrate on a particular class of non-cooperative games called 2*-players zero-sum games*. A 2*-players zero-sum game* is a mathematical representation of a situation in which a participant's gain (or loss) is exactly balanced by the losses (or gains) of the other participant(s). In other words, a zero-game is a game where the payoff of player I is the opposite of the payoff of player II. In particular, there is only one winner (or differently if one player gain, the other loses, it may seem natural but there are contexts where both can win or both can lose, think at cooperative games).

The *strategic form*, or *normal form*, of a two-person zero-sum game is given by a triplet (X, Y, A), where:

- X is a nonempty set, the set of strategies of Player I.
- Y is a nonempty set, the set of strategies of Player II.
- A is a real-valued function defined on $X \times Y$. (Thus, $A(x, y)$ is a real number for every $x \in X$ and every $y \in Y$.)

The interpretation is as follows. Simultaneously, Player I chooses $x \in X$ and Player II chooses $y \in Y$, each does not know the choice of the other. Then their choices are made known and I wins the amount $A(x, y)$ from II. If $A(x, y)$ is negative, Player I pays the absolute value of this amount to player II. Thus, $A(x, y)$ represents the winnings of Player I and the losses of II. The function A is called the *payoff function*. Note that the payoff function can be represented as a matrix called the payoff matrix.

Example: Odd or Even. Let us illustrate these abstract concepts on a simple example. Players I and II simultaneously call out one of the numbers one or two. Player I's name is Odd; he gains if the sum of the numbers is odd. Player II's name is Even; she gains if the sum of the numbers is even. The amount paid to the winner by the loser is always the sum of the numbers in dollars. To put this game in strategic form we must specify X, Y and A. Here we may choose $X = \{1, 2\}, Y = \{1, 2\}$, and A as given in the following table.

$$A = \begin{pmatrix} -2 & +3 \\ +3 & -4 \end{pmatrix}$$

It turns out that one of the players has a distinct advantage in this game. Stop reading and try to guess which one.

If you have voted for "Player I", then you are right ! Let us try to understand why. Let us analyze this game from Player I's point of view. Suppose he calls strategy 'one' with probability $3/5$ and calls 'two' with probality $2/5$. In this case:

1. If II calls 'one', Player I loses 2 dollars with probability $3/5$ and gains 3 dollars $2/5$-ths of the time; on average, he gains $-2(3/5) + 3(2/5) = 0$ (he breaks even in the long run).
2. If II call 'two', Player I gains 3 dollars with probability $3/5$ and loses 4 dollars $2/5$-ths of the time; on average he gains $3(3/5) - 4(2/5) = 1/5$.

So Player I has a strategy which ensures that for every possible strategy of Plyaer II, he won't lose money (in average) ! But if Player II always plays its first strategy, the expected gain of Player I is 0. Can we find a strategy where the expected gain is positive for any possible strategy of Player II?

A *strategy* for Player I is a measure of probability on the set of strategies for Player I. Let p denote the probability that Player I calls 'one'. Since Player I only has two strategies, p completely characterizes the strategy of Player I. Let us try to choose p so that Player I gains the same amount on the average whether II calls 'one' or 'two'. Then since Player I's average winnings when Player II calls 'one' is $-2p+3(1-p)$, and his average winnings when II calls 'two' is $3p-4(1-p)$ Player I should choose p so that:

$$-2p+3(1-p)=3p-4(1-p) \Leftrightarrow 3-5p=7p-4 \Leftrightarrow 12p=7 \Leftrightarrow p=7/12.$$

We know that by calling $5/12$th of the time strategy 'one' and $7/12$ths of the time strategy 'two' then Player I gains the same amount of money. What is this amount?

$$-2(7/12)+3(5/12)=1/12.$$

Such a strategy that produces the same average gain no matter what the opponent does is called an *equalizing strategy*.

Therefore, the game is clearly in Player I's favor. Can he do better than $1/12$ dollars per game on the average? The answer is: Not if II plays properly. In fact, II could use the same strategy:

- call 'one' with probability 7/12.
- call 'two' with probability 5/12.

If Player I calls 'one', Player II's average loss is $-2(7/12)+3(5/12)=1/12$. If Player I calls 'two', II's average loss is $3(7/12)-4(5/12)=1/12$.

Hence, Player I has a strategy that guarantees him at least $1/12$ on the average, and Player II has a strategy that keeps her average loss to at most $1/12$. The value $1/12$ is called the *value of the game*, and the procedure each uses to guarantee this gain (or loss) is called an *optimal strategy* or a *minimax strategy*. Strategies for Players I and II are *stable* if neither Players I nor II have interest to change their strategy. In other words, if Player I modify its strategy then its average payoff does not increase and if Player II modifies its strategy then its average payoff does not increase. The two stategies we just exhibited are stable. If both players choose a strategy such that none of them has any incentive to change their strategy, we are in a *Nash equilibrium*.

The main result concerning two-players zero-sum games ensures that a minimax strategy always exists.

2 Pure and Mixed strategies

A strategy is *pure* if the choice of players is not randomized. It means that Player I still plays the same strategy $x \in X$ then the strategy of Player I is pure. If the strategy of Player I is randomized we say that the strategy is *mixed* (note that a pure strategy also is mixed). We assumed that when a player uses a mixed strategy, he is only interested in his average return. He does not care about his – maximum possible gains or losses — only the average. This is actually a rather drastic assumption. We are evidently assuming that a player is indifferent between receiving 5 million dollars outright, and receiving 10 million dollars with probability 1/2 and nothing with probability 1/2. I think nearly everyone would prefer the $\$5,000,000$ outright. (we will come back to this "problem" a bit later)

This function can then be seen as a matrix where strategies of Player I corresponds to rows and strategies of player II corresponds to columns. Indeed assume that Player I chooses the i-th strategy while II chooses the j-th strategy. Let us denote by x_i the 0-vector except on coefficient i where its value equals 1 (and y_j has a similar value for the vector y). Then the payoff for Player I for this couple of strategies is:

$$x_i^t A y_j = a(i,j).$$

Indeed, the payoff $a(i,j)$ is counted in the final value with "probability" $x_i \cdot y_j$ which is the probability that I chooses strategy i and II chooses strategy II.

3 Minimax theorem

Theorem 1 (Von Neumann Minimax Theorem)**.** *For every finite two-person zero-sum game, there is a number V, called the value of the game, such that:*

- *there is a mixed strategy for Player I such that Player I's average gain is at least V no matter what II does, and*
- *there is a mixed strategy for Player II such that II's average loss is at most V no matter what Player I does.*

Strategies for Players I and II satisfying the minimax theorem are called *minimax strategies*. Note that minimax strategies are stable.

You can immediately check on the the Odd-or-Even game that this result does not hold for pure strategies. Note that this result is not straightforward and asserts a non-trivial statement: for every zero-sum game, there exists a strategy such that, even if we reveal our strategy before playing then we gain the maximum amount of money ! (so a two-players zero-sum games can be played without lies). To state in a more formal way this theorem, it ensures that:
[Minimax Theorem Restatement]

$$\max_x \min_y x^t Ay = \min_y \max_x x^t Ay$$

(where x and y are probability vectors). Indeed, if Player I reveals its strategy then II want to minimize the payoff of I and then choose y which minimizes $x^t Ay$. Since I want to play the best strategy, he selects x which maximizes this value. The second equation is corresponds to the second case.

Proof of the Minimax theorem. The Minimax Theorem says that if we allow the Row Player and Column Player to select probability distributions over the rows and columns of A, respectively, for their strategies, then maximization problem is equal to the minimization one. We can show this by considering a Linear Program The Minimax Theorem then follows immediately from the Strong LP-Duality Theorem. Consider the Linear Programs

$$\begin{aligned} \max t \\ \text{subject to:} \\ t - \sum_j A_{ij} y_j \leq 0 \text{ for all } i \\ \sum_j y_j = 1 \\ y_j \geq 0, t \text{ unconstrained} \end{aligned}$$

The first constraints ensures the following: for every i, if Player II decides to play the pure strategy i, then the average payoff of Player I is at most t. Indeed we want to be sure that this value is non-positive and then it means that $t \leq \sum_j A_{ij} y_j$ which is the average payoff of pure strategy i against the strategy y. The second constraint (and the non-negativity constraints) ensure that the vector y is a probability distribution.

This LP ensures that for every pure strategy x, the average profit is at most t for Player I. Thus the same holds for any mixed strategies.

Exercise 1. *Prove formally that if the average payoff of any pure strategy is at most α then the average payoff of any mixed strategy is at most α.*

So it means that there exists some strategy for II (min) such that for every possible strategy of Player I (max) the average profit of Player I is at most t. Since it holds for every possible pure strategy, it also holds for mixed strategies (BTW, show that when we know the strategy of the other player, there always exists some pure strategy which is optimal, why?). So this t satisfies

$$t = \max_x \min_y x^t Ay$$

Let us now compute the dual of this problem and show it corresponds to the other case of the minimax theorem:

$$\min s$$
$$\text{subject to:}$$
$$s - \sum_i A_{ji} x_i \geq 0 \text{ for all } j$$
$$\sum_i x_i = 1$$
$$x_i \geq 0, s \text{ unconstrained}$$

Exercise 2. • *Show that it means that there exists some strategy for Player I (max) such that for every possible strategy of II (min) the average profit of Player I is at least t.*

- *Show that $s = \min_y \max_x x^t Ay$.*

Now we can conclude using Strong Duality Theorem. □

Note that the proof of the minimax theorem gives a way to compute this value.

3.1 Limits of this model: St Petersburg paradox

In the previous Sections, we have assumed that each player plays its best strategy, and we defined "best strategy" as the strategy which maximized the average payoff. But in practice, maximizing the payoff is not necessarily the best strategy (from a human point of view at least). Consider the following game:

> A con is flipped repeatedly until the first head. If the first head appears on the n-th toss, you get paid $\$2^n$. Now a casino proposes this game for $\$1,000$, do you play?

You answer is probably **no**. Indeed, your argument is the following: with probability $\frac{1}{2}$ Player I gain only $\$2$, with probability $\frac{3}{4}$, Player I gain less than $\$4$, with probability $\frac{7}{8}$, Player I gain less than $\$8$, it is not interesting doing it. Let us now make the calculation: what is our expected payoff?

$$\frac{1}{2} \cdot 2 + \frac{1}{4} \cdot 4 + \frac{1}{8} \cdot 8 \cdots + \frac{1}{2^n} \cdot 2^n \cdots = 1 + 1 + 1 + 1 \cdots = +\infty$$

Thus it is interesting to play against the casino: our expected payoff is infinite. So why don't we play?

We don't play because our utility is small. What do we mean by that? Let us give another example. Player I give you the choice between two choices:

- You gain $\$1,000,000$!
- You gain $\$3,000,000$ with probability $\frac{1}{2}$ and 0 with probability $\frac{1}{2}$.

The usual choice consists in choosing the first case. What is the reason? The expected value in the second case is larger but the risk of this second choice implies that it is worth the risk. In other words, we have to distinguish two notions: the expected value, which is the theoretical payoff and the utility which is what is the "real" gain for the player.

Utility functions The are several possible models of utility. Bernouilli represents it with a logarithmic function. In other words, when the payoff is $\$1,000,000$, Bernouilli considered that the payoff only is $\log(1,000,000)$. The intuitive reason is the following: each additional dollar is less valuable than the previous one. If you already gain $\$1,000,000$, you don't care about one dollar. But if you gain $\$5$ then you care. And obviously:

$$\log(1,000,001) - \log(1,000,000) < log(6) - log(5)$$

Bernouilli assumed that we are playing a game is our utility is larger than the utility of entering into the game. What is the utility of the game?

$$\frac{1}{2}\log 2 + \frac{1}{4}\log 4 + \cdots + \frac{1}{2^n}\log 2^n + \cdots = \frac{1}{2} + \frac{2}{4} + \frac{3}{8} + \cdots + \frac{n}{2^n} + \cdots \longrightarrow 4$$

So you may accept to play this game up to an entering cost of $\log c = 4$, *i.e.* $c = 16$ (and you would you play for 16 dollars?).