Lecture 4: Minimax theorem from strong duality

Rajat Mittal

IIT Kanpur

We have seen that the dual of a linear program gives tight bounds on the value of primal (strong duality) and vice versa. This theorem is useful in many situations, we start with Von Neumann's Minimax theorem.

It is a straight forward application of duality, but have far reaching applications. The standard context to introduce Min-Max theorem is through Game theory. First, let us look at the basics of game theory.

Game theory formally studies the interaction between various players in a strategic game. Many human interactions can be formulated as a game; starting from Chess, Soccer etc. to advertising, auctions and even political campaigns.

We will study one of the simple settings, called *normal form games*. It is a one round game. A game in this form is specified by number of players, their possible strategies and the payoff associated with every possible selection of strategies by each player. Specifically,

- Players: An N player game will have players from the set $[N] = \{1, 2, \cdots, N\}$. For this lecture note, we will only be concerned with $N = 2$.
- Strategies: Each player i will have a set of pure strategies S_i. A profile consists of one strategy per player. The set of all profiles is obtained by taking set product,

$$S = \prod_{i=1}^{N} S_i.$$

- Payoff: The outcome is represented by a *payoff* function $p : S \to \mathbb{R}^N$. In other words, p gives payoff for each player given a profile.

Let us take a simple example motivated from prisoner's dilemma. Suppose, there are two people in a project team. The project is worth 20 points. If they cooperate these marks will be divided equally between them (10 each).

Though, a student might feel that he/she has worked harder. In that case, instructor has kept a provision of complaining at the cost of 3 points. If both students complain, 3 marks are deducted from both the students. If one complains and not the other, 50% credit of other student is transferred to the complainant.

The payoff can be easily visualized as a matrix. Suppose rows represent first player's strategy and columns represent second player's strategies.

	Complain	*Do not complain*
Complain	$7, 7$	$12, 5$
Do not complain	$5, 12$	$10, 10$

This game is pretty interesting, for either player complaining is better than not complaining (if other person's strategy is kept fixed). Keeping this in mind, two rational players will complain. In this case, both complaining is an equilibrium, because no player would want to switch their strategy if other player's strategy is fixed. Even though the choice is not socially optimal, it is an equilibrium.

In any case, this way of describing a game is pretty useful. Notice that the matrix representation is a concise way to express payoffs, specially when there are two players involved. Many famous games like rock/paper/scissor and prisoner's dilemma can be described in this setting.

Exercise 1. Read about rock/paper/scissor and prisoner's dilemma; find n, S, p for these games.

1 Nash equilibrium in two player zero-sum games, and Minimax theorem

In general, a player can also be allowed to play a *mixed* strategy, a probability distribution over its set of strategies. We will be interested in expected payoffs for these mixed strategies.

One of the question in these games is to find a *Nash equilibrium.* A profile is a Nash equilibrium if no party benefits from deviating from the profile (keeping other strategies fixed). In the above game of project marks, both complaining is an equilibrium, though not a globally optimal strategy.

We will ignore the social question about whether Nash equilibrium is best for society/globally or not. Instead, our focus will be on the computational question, can we compute Nash equilibrium for a 2-player game?

Again, we will represent strategies of first player by rows and second player by columns. The payoff can be described by two matrices, A and B, for first and second player respectively. For both matrices, rows are indexed by the strategies for the first player and columns are indexed by the strategies for the second player. An entry in A (respectively B) gives the payoff for the first player (respectively second player) if they play with corresponding row and column strategy. For the project game above,

$$A = \begin{array}{lcc} & \textit{Complain} & \textit{Do not complain} \\ \textit{Complain} & 7 & 12 \\ \textit{Do not complain} & 5 & 10 \end{array}$$

Exercise 2. What is B in this case?

If $B = -A$, the 2-player game is called *zero-sum.* Assume that the game is zero-sum for the rest of this section, hence $B = -A$.

Let us assume that A is an $m \times n$ matrices. A mixed strategy for the first player is a vector $x \in \mathbb{R}^m$, such that,

$$\sum_{i=1}^{n} x_i = \mathbb{1}^T x = 1 \ \textit{and} \ x \geq 0.$$

In other words, x_i is the probability that the first player plays strategy i. Similarly, let $y \in \mathbb{R}^n$ represent a mixed strategy for the second player, i.e., $\mathbb{1}^T y = 1, \ y \geq 0$.

Exercise 3. What is the expected payoff for first player when she plays with strategy x and second player plays with pure strategy i?

If second player plays strategy i, then payoff's are given by the i-th column of A, say $A_{*,i}$. The expected payoff in that case will be $\sum_j x_j A_{j,i} = x^T A_{*,i}$.

Note 1. Check out the new notation. $A_{*,i}$ denotes the i-th column of A; similarly, $A_{i,*}$ will denote the i-th row.

Exercise 4. What is the expected payoff for first player when she plays with strategy x and second player plays with strategy y?

Applying the same idea, the expected payoff should be $\sum_{i,j} A_{i,j} x_i y_j = x^T A y$ for the first player and $-x^T A y$ for the second player (it is a zero-sum game). Clearly, first player wants to maximize payoff according to A and second player wants to minimize payoff according to A.

Let us look at the game from the perspective of the first player. She needs to decide x such that her payoff is maximized. Though, if she picks x, her payoff depends on the strategy chosen by the second player.

Assuming that the second player is also rational, it makes sense to maximize the worst possible payoff for any strategy x. What is the worst possible payoff for first player's strategy x? It is $\min_y x^T A y$. In other words, first player's task is to $\max_x \min_y x^T A y$.

Exercise 5. Similarly, can you write the optimization problem for the second player?

The famous minimax theorem by John von Neumann states that these two optimization problems have equal value.

Theorem 1 (Minimax). *For a 2-player zero-sum game, let first player have m strategies and second player have n strategies. Then,*

$$\max_x \min_y x^T Ay = \min_y \max_x x^T Ay,$$

where x, y are two probability distributions over $\mathbb{R}^m$ and $\mathbb{R}^n$ respectively.

What does it have to do with Nash equilibrium? Answer, everything!

Let us call a strategy pair (x_0, y_0) a *minimax strategy* if it optimizes the function in Minimax theorem. We first show that it is a Nash equilibrium. That means, fixing x_0, second player should not benefit from changing y and vice versa.

Suppose x_0 is fixed, then we know that y_0 minimizes $x_0^T Ay$ (why?). Hence for any strategy y,

$$x_0^T Ay_0 \leq x_0^T Ay.$$

This equation implies that second player loses by changing his strategy from y_0 to y. Similarly, first player will lose by changing x if y_0 is fixed. Hence (x_0, y_0) is a Nash equilibrium.

You will show in the assignment that a Nash equilibrium is a minimax strategy too. That means, given that fixing x second player does not benefit from changing y and vice versa, we get an optimal solution of minimax.

Proof of Von Neumann's Min-Max theorem, Theorem 1. We now show that this optimization can be done with the help of linear programming. Notice that the optimization problem mentioned in the theorem is not linear (x and y are both variables).

Let us look at the problem from the first player's perspective, optimize $\max_x \min_y x^T Ay$ (where x, y are two strategies). Notice that $x^T Ay$ is a convex combination of entries $x^T A_{*,i} = \sum_j x_j A_{j,i}$.

Exercise 6. Let x be fixed, show that $\min_{y \geq 0 : \mathbb{1}^T y = 1} x^T Ay$ is equal to $\min_i \sum_j x_j A_{j,i}$.

To help you with the exercise above, think of a situation when you went to a house to steal some stuff. You can only take away 1 kg of item. Given that the house has gold, silver, bronze in ample quantity; how will you partition your 1 kg among the three? Common sense assumption: Gold is more expensive than silver than bronze.

In other words, the task of first player is, $\max\limits_{x \geq 0 : \mathbb{1}^T x = 1} \min\limits_i \sum_j x_j A_{j,i}$. The inner function has become linear in the entries of x. Using a usual trick to get rid of minimization,

$$\begin{array}{rll} \max & v & \\ \text{subject to } v \leq & \sum_j x_j A_{j,i} & \forall i \in \{1, 2, \cdots, n\} \\ & x_j \geq 0 & \forall j \in \{1, 2, \cdots, m\} \\ & \sum_j x_j = 1 & \end{array}$$

The value of this program gives us the maximin strategy for the first player. Let us take a dual for this program.

Exercise 7. Calculate the dual of the above linear program.

We can multiply each inequality with y_i and the equality with u. It will give us the dual.

$$\begin{array}{rll} \min & u & \\ \text{subject to } u \geq & \sum_i y_i A_{i,j} & \forall j \in \{1, 2, \cdots, m\} \\ & y_i \geq 0 & \forall i \in \{1, 2, \cdots, n\} \\ & \sum_i y_i = 1 & \end{array}$$

A keen observation of this equation reveals that it is the minimax strategy for the second player. This is the same argument as the one which was used to convert maximin strategy of the first player to the linear program above.

Now, strong duality gives us the minimax theorem. □

In other words, there is always a Nash equilibrium for a 2-player zero-sum game and it can be computed using any LP solver. Now you know why John Nash was *a beautiful mind*.

The proof for existence of Nash equilibrium was an easy implication of strong duality. Surprisingly, it turns out that even strong duality can be proven assuming the existence of Nash equilibrium. The proof will not be covered in this course, interested readers can refer to [3]

2 Extra reading: *Non* zero-sum games

When the game is not zero sum, still there is a linear programming based approach to find Nash equilibrium. Most of the content of this section is inspired from the survey article by Bernhard von Stengel [2].

What about Nash equilibrium of 2-player games which are not zero-sum? We will see a linear programming inspired approach to find equilibrium in such games.

Remember that the payoff will be given by two matrices, A for player 1 and B for player 2. Again, fixing a strategy y_0 for player 2, player 1 wants to maximize $x^T Ay_0$. This can be written as a linear program (y_0 is not a variable).

$$\begin{aligned} \max \quad & x^T Ay_0 \\ \text{subject to} \quad & x_j \geq 0 \quad \forall j \in \{1, 2, \cdots, m\} \\ & \textstyle\sum_j x_j = 1 \end{aligned}$$

The dual of this program is pretty simple.

$$\begin{aligned} \min \quad & u \\ \text{subject to} \quad & u \geq \textstyle\sum_j (y_0)_j A_{i,j} = A_{i,*}^T y_0 \ \forall i \in \{1, 2, \cdots, m\} \end{aligned}$$

Exercise 8. Show that the maximum value of this program is $\max_i A_{i,*}^T y_0$.

Notice that $A_{i,*}^T y_0$ is the expected payoff for player 1 if he plays pure strategy i. If the strategy of column player is fixed, it is not surprising that player 1 should play with the pure strategy which maximizes her expected payoff. Any pure strategy i such that $A_{i,*}^T y_0$ attains the maximum, will be considered best response strategy of player 1 against y_0.

The dual LP simply calculates that maximum.

Exercise 9. What will be the complementary slackness condition for this primal-dual pair?

The complementary slackness condition states that x is optimal for this y_0 iff x, u are feasible and $x^T(u\mathbb{1} - Ay_0) = 0$. In other words, x, u is an optimal pair if and only if

- $\mathbb{1}^T x = 1$ and $x \geq 0$,
- $u\mathbb{1} - Ay_0 \geq 0$,
- $x^T(u\mathbb{1} - Ay_0) = 0$.

Notice that the last condition can be interpreted as: player 1 plays a strategy with positive probability only if that strategy is a best response strategy against y_0.

We can have similar conditions from the perspective of player 2. y, v is an optimal pair for a given strategy x_0 iff

- $\mathbb{1}^T y = 1$ and $y \geq 0$,

– $v\mathbb{1}^T - x_0^T B \leq 0$,
– $(v\mathbb{1}^T - x_0^T B)y = 0$.

For a Nash equilibrium, x_0 should be the best strategy for y_0 and y_0 should be the best strategy for x_0. Combining, x_0, y_0 is a Nash equilibrium if and only if both sets of complementary slackness conditions are satisfied.

We get a characterization of Nash equilibrium: a profile x, y is a Nash equilibrium if and only if,

1. $\mathbb{1}^T x = 1, x \geq 0$ and $\mathbb{1}^T y = 1, y \geq 0$.
2. $u\mathbb{1} - Ay \geq 0$ and $v\mathbb{1}^T - x^T B \leq 0$.
3. $x^T(u\mathbb{1} - Ay) = 0$ and $(v\mathbb{1}^T - x^T B)y = 0$.

Exercise 10. Convince yourself that this is a characterization of Nash equilibrium in 2-player games.

Note 2. Do not confuse this characterization with complementary slackness condition for a single pair of linear programs. We get this by combining complementary slackness for two different primal-dual pairs and the definition of Nash equilibrium.

We know from the earlier discussion that $u = \max_i A_{i,*}^T y$ and $v = \min_i x^T A_{*,i}$. This showed that the third condition is equivalent to: $x_i \neq 0$ means i is the best response strategy for y and vice versa.

Exercise 11. Show that i is the best response strategy against y_0 iff $A_{i,*}^T y \geq A_{j,*}^T y$ for all j.

So, getting rid of u, v, a profile x, y is a Nash equilibrium if and only if

1. $\mathbb{1}^T x = 1, x \geq 0$ and $\mathbb{1}^T y = 1, y \geq 0$.
2. For all $i \in I = \{1, 2 \cdots, m\}$, either $x_i = 0$ or $A_{i,*}^T y \geq A_{k,*}^T y$ for all k. What is the range of k?
3. For all $j \in J = \{1, 2 \cdots, n\}$, either $y_j = 0$ or $x^T A_{*,j} \leq x^T A_{*,k}$ for all k. What is the range of k in this case?

Looking at the characterization, first condition just states that x, y are strategies. Observe that there are m conditions on y from second part of OR in the second condition and n conditions because of the first part of OR in the third condition. Similarly, we have $m + n$ conditions on x. The above characterization states: for every entry in $I \cup J$, at least x satisfies its corresponding condition or y does.

For notational ease, change $J = \{1, 2, \cdots, n\}$ to $J = \{m+1, m+2, \cdots, m+n\}$ where i is switched to $m + i$.

Given two strategies x, y, we can label the entries in I, J satisfied by them. So, define

$$L(x) = \{i \in I : \ x_i = 0\} \cup \{j \in J : x^T A_{*,j} \leq x^T A_{*,k} \ \forall k \in \{1, 2, \cdots, n\}.$$

and

$$L(y) = \{j \in J : \ y_j = 0\} \cup \{i \in I : A_{i,*}^T y \geq A_{k,*}^T y \ \forall k \in \{1, 2, \cdots, m\}.$$

So, x, y are in Nash equilibrium if and only if $L(x) \cup L(y) = \{1, 2, \cdots, m+n\}$. Lemke-Howson algorithm is a way to find such equilibrium and is motivated by the simplex algorithm. Interested readers can refer to the survey by Bernhard von Stengel [2].

3 Communication complexity

There is another application of Von Neumann's Minimax theorem 1 in a completely different domain, *communication complexity.* Communication complexity is a part of complexity theory where two parties, Alice and Bob, want to compute a function $f : \{0,1\}^n \rightarrow \{0,1\}$ such that both parties *only* have a part of the input. The cost is measured in terms of bits communicated between the two parties, any computation they do on their side is *for free.*

This model of computation has been extremely useful with applications in VLSI circuits, data structures, streaming algorithms and many more. In particular, showing a lower bound on communication for computing a function allows us to show hardness of varied tasks in the mentioned areas. We will see that the Minimax theorem gives a generic strategy to put lower bound on communication. With this picture in mind, let us dive into the details.

Formally, we want to compute a function $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$, where Alice will get an input $x \in \mathcal{X}$ and Bob will get an input $y \in \mathcal{Y}$. Let us stick with deterministic protocols to start with, where player's protocol only depends on their input and bits communicated. (As opposed to a randomized protocol where it can depend upon some random coin tosses too.)

The cost of the protocol on x, y is the number of bits communicated to get $f(x, y)$. The cost of the protocol is

$$\max_{x,y} \textit{ Cost of the protocol on } x, y.$$

The deterministic communication complexity of f is the cost of the best deterministic protocol for f.

Stuff becomes more interesting with randomized protocols. Alice and Bob can toss some random coins and decide the protocol on the basis of these coins. We assume that the coin tosses are shared (they are known to both parties). This setting is called the *public coin model*. As is the case in any randomized model, some error proability is allowed. What do we mean by error probability?

We can see two sources of randomness; R for the random coin tosses, and say μ to be a distribution over inputs $\mathcal{X} \times \mathcal{Y}$. There are two ways in which we can define the error of the protocol.

- Worst case: In this setting our protocol should work on any input (worst case over input). Protocol has error less than ϵ, if for every x, y:
$$P_r[\textit{Protocol is wrong}] \leq \epsilon.$$
- Average case: In this case we have a distribution μ on $\mathcal{X} \times \mathcal{Y}$ in mind. The error is calculated assuming inputs come from this distribution.
$$P_{r,\mu}[\textit{Protocol is wrong}] \leq \epsilon.$$

Note 3. A successful randomized protocol incurs only a constant error less than 1/2 (we fix it to be 1/3). The actual constant does not matter, we can go from one constant to another using constant repetitions. Since we care only about asymptotics, constant factor can be ignored.

A natural question is, how is the communication complexity of a function differ in the two settings above (worst case and average case). Yao's minimax principle provide a very useful theorem in this regard.

Theorem 2 (Yao's minimax theorem). *The worst case complexity of a function f (randomized communication complexity) is equal to the maximum possible average case complexity for that function. Here maximum is taken over all possible distributions μ on $\mathcal{X} \times \mathcal{Y}$.*

Exercise 12. Show that worst case complexity is an upper bound on the average case complexity for any μ.

Proof of Yao's mimimax theorem 2. This seems like a case where proving lower bound was easy. Though how do we prove that there always exists a distribution μ such that the average case complexity with respect to μ is same as the worst case complexity. This might remind you of duality. Indeed, we prove the theorem using minimax theorem (strong duality).

Note 4. An important observation is that a randomized protocol is a distribution over deterministic protocols (assignment question).

We just need to construct a matrix M such that maximin and minimax strategies correspond to the average case and worst case behavior respectively. Suppose all average case strategies with error 1/3 take cost at most c. We need to show that there exists a randomized protocol with cost c which works for every input.

From now on, a protocol will mean a protocol with cost less than or equal to c. The matrix M is very natural, its columns are indexed by protocols and rows by inputs to the protocols.

$$M_{(x,y),P} = \mathbb{1}_{\textit{ Protocol } P \textit{ is correct on } (x,y)}.$$

We will use variable i for distribution on inputs and p for distributions on protocols.

Exercise 13. What does $i^T Mp$ denote for a deterministic protocol p and input i? What if i and p are distributions?

You can convince yourself that $i^T Mp = \sum_{(x,y),P} M_{(x,y),P} \Pr_i[(x,y)] \Pr_p[P]$. In other words it is the success probablity of randomized protocol p on input distribution i. Consider the minimax quantity $\min_i \max_p i^T Mp$. We know that for any distribution there is a protocol with success probability 2/3; in other words, for any i, $\max_p i^T Mp$ is at least 2/3. We get that

$$\min_i \max_p i^T MP \geq 2/3.$$

From minimax theorem 1, $\max_p \min_i i^T MP \geq 2/3$. So, there exists a p (randomized protocol) such that $\min_i i^T Mp \geq 2/3$. Since i goes over all input distributions (even point distributions), p succeeds with worst case probability 2/3 proving the theorem.

□

Even though the above theorem (Theorem 2) seems to only give equivalence between average case and worst case hardness. The most important application of it is in proving lower bounds on worst case complexity. Notice that the hardness on any distribution of inputs is a lower bound, a good lower bound amounts to coming up with a distribution such that no deterministic protocol works for it.

Exercise 14. Why??

Lower bound for disjointness Disjointness is one of the standard problems in communication complexity. Both Alice and Bob are given strings in $\{0,1\}^n$ (say x and y). They can be interpreted as the indicator of a subset of $[n]$ in a canonical way, say S_x and S_y.

Exercise 15. What is the natural way to map a string in $\{0,1\}^n$ to a subset of $[n]$?

The problems is to find if S_x and S_y are disjoint or not. It turns out that a lower bound can be given using Yao's minimax [1, Chapter 4.3].

4 Assignment

Exercise 16. Consider a two player zero-sum game.

- Show that a pair of strategies (x,y) is a Nash equilibrium if and only if

$$x_i > 0 \Rightarrow A_{i,*}^T y = \max_j A_{j,*}^T y \text{ and } y_i > 0 \Rightarrow A_{*,i}^T y = \min_j A_{*,j}^T y.$$

- Show that all Nash equilibriums in a 2-player zero-sum game optimize the function in Minimax theorem (Thm. 1).

Exercise 17. Read about Lemke-Howson algorithm.

Exercise 18. Show that a randomized protocol is a distribution over deterministic protocols.

References

1. T. Roughgarden. Communication complexity (for algorithm designers). `https://theory.stanford.edu/~tim/w15/l/w15.pdf`.
2. B. von Stengel. Computing equilibria for two-person games. `http://maths.lse.ac.uk/Personal/stengel/TEXTE/nashsurvey.pdf`.
3. B. von Stengel. Zero-sum games and linear programming duality. `https://arxiv.org/pdf/2205.11196.pdf`.