Math 482: Linear Programming[1] — Mikhail Lavrov

Lecture 19: Zero-Sum Games II

March 9, 2020 — *University of Illinois at Urbana-Champaign*

1 The "best worst-case" (maximin) strategy

Last time, we described a zero-sum game between Alice and Bob by an $a \times b$ payoff matrix A listing Alice's payoffs from each outcome (which are the negative of Bob's). A mixed strategy for Alice, corresponding to choosing between her strategies at random with some probabilities, is given by a vector $\mathbf{x} \in \mathbb{R}^a$ with $x_1 + x_2 + \cdots + x_a = 1$ and $\mathbf{x} \geq \mathbf{0}$. A mixed strategy for Bob is given by a vector $\mathbf{y} \in \mathbb{R}^b$ with $y_1 + y_2 + \cdots + y_b = 1$ and $\mathbf{y} \geq \mathbf{0}$.

We saw that if Alice and Bob play these mixed strategies against each other, the expected payoff for Alice is given by $\mathbf{x}^\mathsf{T} A\mathbf{y}$.

Alice would like to choose the best mixed strategy $\mathbf{x}$. But she can't use the formula $\mathbf{x}^\mathsf{T} A\mathbf{y}$ to evaluate how good a strategy is directly, because she doesn't know which $\mathbf{y}$ is Bob's strategy. Instead, one thing Alice might try is to choose the strategy $\mathbf{x}$ with the best performance against the worst possible counter to $\mathbf{x}$. This is called Alice's *maximin* strategy.

For all we know, this could be a terrible idea! In fact, we can cook up lots of examples of games that *aren't* zero-sum, in which the maximin strategy does terribly.

Consider the "Win, Lose, and Copy" game, defined as follows. Bob has two options: "Win \$100" and "Lose \$100". Alice also has two options: "Don't play" and "Copy Bob". In this game, Alice's maximin strategy is not to play: copying Bob has the risk that Bob will pick "Lose \$100", and not playing can't lose any money. But Bob isn't stupid and will never pick "Lose \$100", so "Copy Bob" is guaranteed to earn Alice \$100 as well.

We will see that in the case of zero-sum strategies, the maximin strategy is reasonable, but it will take us some time to get there.

2 Linear programming for zero-sum games

First, here is a claim to simplify the analysis.

Lemma 2.1. *If Alice's mixed strategy $\mathbf{x} \in \mathbb{R}^a$ is fixed, Bob's best response is a pure strategy (once which always picks the same option).*

Proof. If Alice is playing mixed strategy $\mathbf{x}$, then $\mathbf{x}^\mathsf{T} A$ is the vector of her possible payoffs, depending on Bob's choices. If Bob plays the pure strategy "always pick option i" for some $i \in \{1, 2, \ldots, b\}$, then Alice's payoff is going to be $(\mathbf{x}^\mathsf{T} A)_i$: the i^{th} component of this vector.

[1] This document comes from the Math 482 course webpage: https://faculty.math.illinois.edu/~mlavrov/courses/482-spring-2020.html

If Bob plays a mixed strategy $\mathbf{y} \in \mathbb{R}^b$, then Alice's expected payoff $\mathbf{x}^\mathsf{T} A\mathbf{y}$ is a weighted average of the payoffs above, where payoff $(\mathbf{x}^\mathsf{T} A)_i$ is multiplied by weight y_i. The weighted average can't be lower than the smallest of the payoffs $(\mathbf{x}^\mathsf{T} A)_1, \dots, (\mathbf{x}^\mathsf{T} A)_b$. So the smallest of those payoffs is Alice's worst case.

Or in other words: if option $i \in \{1, 2, \dots, b\}$ is the best response for Bob, then playing a mixed strategy $\mathbf{y}$ instead could be described as "with probability y_i, do the best thing; the rest of the time, do something worse." That's obviously suboptimal.

(Technically, multiple options could be tied for Bob's best response, in which case choosing randomly between them is just as good as choosing one of them; but choosing randomly will never be strictly better.) □

Based on this claim, the worst-case payoff for Alice's mixed strategy $\mathbf{x}$ is given by

$$\min \left\{ (\mathbf{x}^\mathsf{T} A)_1, (\mathbf{x}^\mathsf{T} A)_2, \dots, (\mathbf{x}^\mathsf{T} A)_b \right\}.$$

Alice's maximin strategy is given by the solution to the following optimization problem:

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^a}{\text{maximize}} \quad & \min \left\{ (\mathbf{x}^\mathsf{T} A)_1, (\mathbf{x}^\mathsf{T} A)_2, \dots, (\mathbf{x}^\mathsf{T} A)_b \right\} \\ \text{subject to} \quad & x_1 + x_2 + \dots + x_a = 1 \\ & \mathbf{x} \geq \mathbf{0} \end{aligned}$$

This is not a linear program. But we can simplify it into one with a standard trick. To maximize the minimum of multiple options, we can maximize an auxiliary variable u, subject to the constraint that u is smaller than each option. In other words, we maximize u, adding the constraints $u \leq (\mathbf{x}^\mathsf{T} A)_1$, $u \leq (\mathbf{x}^\mathsf{T} A)_2$, $\dots$, $u \leq (\mathbf{x}^\mathsf{T} A)_b$.

Let $\mathbf{1}$ denote the vector in which every component is 1. (In this case, we'll want to have $\mathbf{1} \in \mathbb{R}^b$, but in general, we'll abuse notation and write $\mathbf{1}$ for the all-ones vector of whichever dimension we need.) Then a quick way to write down these constraints on u is $u\mathbf{1}^\mathsf{T} \leq \mathbf{x}^\mathsf{T} A$, or $u\mathbf{1}^\mathsf{T} - \mathbf{x}^\mathsf{T} A \leq \mathbf{0}^\mathsf{T}$. Similarly, the constraint $x_1 + x_2 + \dots + x_a = 1$ can be written as $\mathbf{x}^\mathsf{T}\mathbf{1} = 1$, where $\mathbf{1} \in \mathbb{R}^a$.

So we get the following linear program:

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^a, u \in \mathbb{R}}{\text{maximize}} \quad & u \\ \text{subject to} \quad & u\mathbf{1}^\mathsf{T} - \mathbf{x}^\mathsf{T} A \leq \mathbf{0}^\mathsf{T} \\ & \mathbf{x}^\mathsf{T}\mathbf{1} = 1 \\ & \mathbf{x} \geq \mathbf{0} \end{aligned}$$

3 Zero-sum games and duality

Similarly, we can talk about Bob's *minimax* strategy, the strategy with the worst performance against Alice's best counter. If Bob is playing strategy $\mathbf{y} \in \mathbb{R}^b$, then Alice's best payoff in response is given by

$$\max \{ (A\mathbf{y})_1, (A\mathbf{y})_2, \dots, (A\mathbf{y})_a \}$$

and we get the following linear program for Bob:

$$\begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^b, v\in\mathbb{R}}{\text{minimize}} \quad & v \\ \text{subject to} \quad & \mathbf{1}v - A\mathbf{y} \ge \mathbf{0} \\ & \mathbf{1}^\mathsf{T}\mathbf{y} = 1 \\ & \mathbf{y} \ge \mathbf{0} \end{aligned}$$

Now observe that Alice and Bob's linear programs are duals of each other!

$$\begin{aligned} \underset{\mathbf{x}\in\mathbb{R}^a, u\in\mathbb{R}}{\text{maximize}} \quad & u & & \\ \text{subject to} \quad & u\mathbf{1}^\mathsf{T} - \mathbf{x}^\mathsf{T}A \le \mathbf{0}^\mathsf{T} & (\mathbf{y}) \\ & \mathbf{x}^\mathsf{T}\mathbf{1} = 1 & (v) \\ & \mathbf{x} \ge \mathbf{0} & \\ & u \text{ unrestricted} & \end{aligned} \qquad \begin{aligned} \underset{\mathbf{y}\in\mathbb{R}^b, v\in\mathbb{R}}{\text{minimize}} \quad & v & & \\ \text{subject to} \quad & \mathbf{1}v - A\mathbf{y} \ge 0 & (\mathbf{x}) \\ & \mathbf{1}^\mathsf{T}\mathbf{y} = 1 & (u) \\ & \mathbf{y} \ge \mathbf{0} & \\ & v \text{ unrestricted} & \end{aligned}$$

This doesn't tell us anything about how to solve these linear programs that we didn't already know.

However, strong duality tells us that the two linear programs have the same objective value z^*. This means that Alice has a mixed strategy $\mathbf{x}^*$ that guarantees her a payoff of z^*, and Bob has a mixed strategy $\mathbf{y}^*$ requiring him to pay at most z^* to Alice (equivalently, guaranteeing him a payoff of $-z^*$).

Strong duality justifies the decision to analyze zero-sum games by looking at the maximin and minimax strategies. They are, in fact, the optimal strategies for Alice and Bob.

The reasoning is the same as our reasoning in the previous lecture about saddle points. Alice can't have any other strategy that guarantees her a payoff of more than z^*, because it would be thwarted by Bob's mixed strategy $\mathbf{y}^*$. So Alice's mixed strategy $\mathbf{x}^*$ is the best. Similarly, Bob can't have any other strategy that guarantees him a payoff of more than $-z^*$, because it would be thwarted by Alice's mixed strategy $\mathbf{x}^*$. So Bob's mixed strategy $\mathbf{y}^*$ is the best.

4 Example

Recall the *even-odd game*, which we used as an example in the previous lecture. In this game, Alice and Bob each hold up either 1 or 2 fingers. Let N be the total. Alice wants an odd total: if N is odd, Bob gives Alice N dollars. Bob wants an even total: if N is even, Alice gives Bob N dollars. The payoff matrix is

	Bob: 1	Bob: 2
Alice: 1	-2	3
Alice: 2	3	-4

Alice's linear program, in this case, is

$$\begin{aligned} \underset{x_1,x_2,u\in\mathbb{R}}{\text{maximize}} \quad & u \\ \text{subject to} \quad & u \le -2x_1 + 3x_2 \\ & u \le 3x_1 - 4x_2 \\ & x_1 + x_2 = 1 \\ & x_1, x_2 \ge 0 \\ & u \text{ unrestricted} \end{aligned}$$

This is kind of obnoxious to solve (we'd end up with 6 variables and 3 equations) but the optimal solution is $(x_1, x_2, u) = (\frac{7}{12}, \frac{5}{12}, \frac{1}{12})$. In other words, by holding up 1 finger $\frac{7}{12}$ of the time, and 2 fingers $\frac{5}{12}$ of the time, Alice guarantees herself an average profit of $\frac{1}{12}$, no matter what Bob does.

In this case, because the game is small, we can take a shortcut and simplify our analysis considerably. Because we didn't find any dominating strategies for this 2×2 game, we know that Bob's optimal strategy is going to be a mixed strategy, with $y_1, y_2 > 0$. Complementary slackness tells us that in Alice's linear program, the corresponding constraints are tight: $u = -2x_1 + 3x_2$ and $u = 3x_1 - 4x_2$. So we can find x_1, x_2 by solving the system

$$\begin{cases} -2x_1 + 3x_2 = 3x_1 - 4x_2 \\ x_1 + x_2 = 1 \end{cases} \implies \begin{cases} x_1 = \frac{7}{12} \\ x_2 = \frac{5}{12} \end{cases}$$

With more complicated games, this strategy might not work, because it required knowing that *all* of Bob's choices are played with positive probability. If Bob has 3 or more options, it's possible that, for example, one option is dominated by a mixed strategy that combines the other options. This is not easy to see just by looking at the payoff matrix, so we'll have to solve the linear program to succeed.

That being said, there is a way (as described in the next, optional section of the notes) to simplify the linear program in general.

5 A simplified linear program (optional)

Suppose that we know Alice's optimal payoff is positive. Then with $u \geq 0$, we can take this linear program

$$\begin{aligned} \underset{x_1,x_2,u\in\mathbb{R}}{\text{maximize}} \quad & u \\ \text{subject to} \quad & u \leq -2x_1 + 3x_2 \\ & u \leq \quad 3x_1 - 4x_2 \\ & x_1 + x_2 = 1 \\ & x_1, x_2, u \geq 0 \end{aligned}$$

and divide all three constraints by u, getting the linear program

$$\begin{aligned} \underset{x_1,x_2,u\in\mathbb{R}}{\text{maximize}} \quad & u \\ \text{subject to} \quad & 1 \leq -2x_1/u + 3x_2/u \\ & 1 \leq \quad 3x_1/u - 4x_2/u \\ & x_1/u + x_2/u = 1/u \\ & x_1, x_2, u \geq 0 \end{aligned}$$

Let $x_1' = x_1/u$ and let $x_2' = x_2/u$. Maximizing u is equivalent to minimizing $1/u$, and we have a formula for $1/u$: $1/u = x_1/u + x_2/u = x_1' + x_2'$. So we get the linear program

$$\begin{aligned} \underset{x_1',x_2'\in\mathbb{R}}{\text{minimize}} \quad & x_1' + x_2' \\ \text{subject to} \quad & -2x_1' + 3x_2' \geq 1 \\ & \quad 3x_1' - 4x_2' \geq 1 \\ & x_1', x_2' \geq 0 \end{aligned}$$

(We have dropped the variable u: it is now unnecessary.) The resulting linear program is easier to solve: it gets rid of one unrestricted variable and one equation. In general, this gives us the linear program

$$\begin{aligned} \underset{\mathbf{x}'\in\mathbb{R}^a}{\text{minimize}} \quad & x_1' + \cdots + x_a' \\ \text{subject to} \quad & \mathbf{x}'^{\mathsf{T}} A \geq \mathbf{1}^{\mathsf{T}} \\ & \mathbf{x}' \geq \mathbf{0} \end{aligned}$$

Note that the optimal values of $\mathbf{x}'$ and the optimal objective value are a bit harder to interpret: $\mathbf{x}'$ is a vector *proportional* to the optimal mixed strategy (but needs to be scaled down to sum to 1), and the objective value is the reciprocal of Alice's optimal payoff.

There's one more thing to watch out for. This LP only works if Alice's optimal payoff is positive. If it's negative or zero, we'll get an infeasible linear program this way.

We can get around this problem by modifying Alice's payoff matrix A, adding the same number to every entry. It's enough to add a number large enough to make all entries of A positive. This changes Alice's payoffs, but doesn't change her optimal strategy (essentially, it represents Bob paying Alice a large fee to play the game). So we can find the optimal strategy using this method, and then change back.