CS295 Introduction to Algorithmic Game Theory

Instructor: Ioannis Panageas Scribed by: Nadia Etemadi, Yotam Golan

Lecture 3. Linear Programming Duality and Zero-sum Games.

1 Linear Programming

1.1 Feasibility and Optimization

A linear program takes the standard form

$$Ax \leq b$$
$$x \geq 0.$$

A is a matrix of size $n \times m$. When we are looking for a *feasible* solution, we are looking for an existing solution x^* of our linear program, if one exists. The size of our matrix A represents the fact that we have n constraints and m variables in our linear program.

Now suppose that we want to find the most *optimal* solution for our linear program–that is, a solution that both exists and maximizes (or minimizes) a metric of our choosing. Our linear program can now take on the following form:

$$\max\ c^\intercal x$$
$$s.t Ax \leq b$$
$$x \geq 0.$$

With this form, we have the added caveat of wanting to maximize (or minimize) a metric of our choosing, in this case $c^\intercal x$, subject to some rule that we can define, which in this case is $Ax \leq b$. Our goal here now is to find an optimized solution x^*, or if we aren't able to, designate the program as *infeasible.*

Lemma 1.1 *Equivalence: These two problems are polynomial-time equivalent.*

1.2 Dual and Primal Formulations

Often when trying to create a solvable algorithm for an LP-definable problem, the problem needs to be decomposed into a standard formulation. We can demonstrate that every LP-definable problem can be transformed into this standard form, otherwise known as a primal formulation.

Suppose then that we are given a LP in the standard form:

$$\begin{gathered} \max\ c^\intercal x \\ s.t Ax \leq b \\ x \geq 0. \end{gathered}$$

The goal of a primal formulation is to then return either a optimal answer, or that is infeasible. In different terms, the primal formulation's goal is to identify the player strategy x which maximizes some criteria c, while obeying the given constraints in a zero-sum game.

As the primal formulation is an Linear Program, it also has a inverse. This is the *Dual Formulation*.

$$\begin{array}{ccc} \max\ c^\intercal x & & \text{-min } (-c^\intercal)x \\ Ax \leq b & = & -Ax \geq -b \\ x \geq 0. & & x \geq 0. \end{array}$$

The problem on the right is in standard form, and we can then take its dual to get the desired LP

$$\begin{array}{ccc} \max\ (-b^\intercal)y & & \min\ (b^\intercal)y \\ (-A)y \leq b & = & Ay \geq b \\ y \geq 0. & & y \geq 0. \end{array}$$

This primal-dual pairing of two LP's A and B is related via the ***Weak Duality Theorem***.
Let us take the two Linear Programs,
U:

$$\begin{gathered} \max\ c^\intercal x \\ s.t Ax \leq b \\ x \geq 0. \end{gathered}$$

and V:

$$\begin{aligned} &\min b^\intercal y \\ &s.t Ay \geq b \\ &y \geq 0. \end{aligned}$$

A primal-dual formulation, as demonstrated above. Given m constraints and n variables, if $x \in R^n$ for U is feasible, and $y \in R^m$ for V is feasible then we can state that $x^\intercal A^\intercal y \geq x^\intercal c$ and $y^\intercal Ax \leq y^\intercal b$. Given these two inequalities, we can then state

$$c^\intercal x \leq y^\intercal Ax \leq b^\intercal y$$

From this we can make several conclusions. It follows, that if U is unbounded, then V is infeasible. Similarly, if V is unbounded, then U is also infeasible. Finally, if $c^\intercal x = b^\intercal y$ with x being feasible for U and y being feasible for V, then both must solve with their respective variable.

Indeed, this primal-dual pairing of Linear Programs possess four key characteristics which describe their relation. These are the only four possible states the primal-dual pairing can exist in:
1. The Primal is bounded and feasible $\rightarrow$ The Dual is bounded and feasible.
2. The Primal is unbounded and feasible $\rightarrow$ The Dual is infeasible.
3. The Primal is infeasible $\rightarrow$ The Dual is unbounded and feasible.
4. The Primal is infeasible $\rightarrow$ The Dual is infeasible.

Using the *Weak Duality Theorem* we can prove that If either U or V has a finite optimal value, then so does the other. Furthermore, the optimal values will coincide, and the optimal solutions to both U and V will exist. This is known as the ***Strong Duality Theorem***. Here is an example:

Primal

max $0 \cdot x_1 + 0 \cdot x_2 + 1 \cdot z$

s.t $\begin{pmatrix} -3 & 2 & 1 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ z \end{pmatrix} \leq 0$

$x_1 + x_2 = 1$

$x_1, x_2 \geq 0$

Dual

min $0 \cdot y_1 + 0 \cdot y_2 + 1 \cdot w$

s.t $\begin{pmatrix} -3 & 1 & 1 \\ 2 & -1 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ w \end{pmatrix} \geq 0$

$y_1 + y_2 = 1$

$y_1, y_2 \geq 0$

These primal and dual forms are equivalent by the Strong Duality Theorem.

2 Zero-Sum Games as Linear Programs

2.1 Representing Zero-Sum Games with Linear Programming

Linear programs can seem very overly theoretical or difficult to understand without any context. Why do linear programs matter, and more importantly, how can we use them to our benefit in algorithmic game theory?

Consider a zero-sum game. We have an outcome matrix R_{ij}, where the row player x chooses a strategy $x \in \Delta_n$, and the column player picks a strategy $y \in \Delta_m$.

The row player's payout is $x^\intercal Ry$ received, and the column player's payout is $x^\intercal Ry$ paid, thus designating the scenario as a zero-sum game.

If we assume that player x plays first, and that they want to receive a payout of at least z, then for all pure strategies of player y we can assume that player x receives a payout of at least z. As a more formal definition, we can say that

$$x^\intercal R \geq z * \mathbf{1}^\intercal$$

$$-x^\intercal R + z * \mathbf{1}^\intercal \leq 0.$$

In addition, x should be a randomized strategy. More formally, we can define this in the following manner:

$$x^\intercal \mathbf{1} = 1$$
$$x \geq 0$$

where $\mathbf{1}$ represents the identity matrix.

2.2 Dual Formulations of Zero-Sum Games

Using the logic from above, we can write a linear program that represents the maximal payout of player x, the row player:

$$\max z$$
$$x^\intercal R \geq z * \mathbf{1}^\intercal$$
$$x^\intercal \mathbf{1} = 1$$
$$x \geq 0$$

We want to maximize z, subject to the constraints that player x receives a payout of at least z using a randomized mixed strategy. Note that the above linear program is the same as expressing the following:

$$\max_{x\in\Delta_n}\min_{y\in\Delta_m} x^\intercal R y$$

If we consider the dual form of the previous linear program, we can find some interesting behavior.

$$\begin{gathered}\min z' \\ -y^\intercal R^\intercal + z' * \mathbf{1}^\intercal \geq 0 \\ y^\intercal \mathbf{1} = 1 \\ y \geq 0\end{gathered}$$

Now, if we substitute $z'' = -z'$, we get the following:

$$\begin{gathered}-\max z'' \\ y^\intercal * (-R)^\intercal \geq z'' * \mathbf{1}^\intercal \\ y^\intercal \mathbf{1} = 1 \\ y \geq 0\end{gathered}$$

If we flip the signs, we get a linear program that represents what would occur if y, the column player, played first. The property of LP duality allows us to make some interesting conclusions regarding Nash equilibrium as well.

3 Nash Equilibrium and Linear Programs

3.1 Equilibrium Between Linear Programs

A zero-sum game can be represented fully as a series of n Linear Programs, each representing a player. In this series, each player will have an optimal x and z such that *(x*,z*)* will be optimal for their respective LP. In this series, let x^1, z^1 and x^2, z^2 represent the optimal values for their respective LP's 1 and 2. If this is true, then x^1, x^2 is the Nash Equilibrium of the zero sum game on their reward matrix R with the payoffs being z^1 and z^2 respectively.

Since (x^1, z^1) is optimal, it is also by definition feasible. We can expand it out to $x^{1\intercal} R x^2 \geq z^1$.
Similarly, x^2, z^2 is optimal and feasible, so -$x^{2\intercal} R^\intercal x^1 \geq z^2$.
Using Strong Duality, we can now say that -z = z^2 and thus $x^{1\intercal} R x^2 = z^1$.

This means that no matter what player 2 does, they will pay at least z as long as player 1 chooses x^1. And no matter what player 1 does, they will get at most z, as long as player 2 chooses x^2. This is a Nash Equilibrium!

4 Corollaries

Corollary 4.1 *Von Neuman minmax theorem: it holds that*

$$max_{x\in\Delta_n} min_{y\in\Delta_m} x^\intercal R y \;=\; min_{y\in\Delta_m} max_{x\in\Delta_n} x^\intercal R y$$

Corollary 4.2 *Uniqueness of payout: In all Nash equilibrium points of a zero-sum game, the payouts of the row player are always the same. The same holds for the column player–their payout will always be the same.*

Corollary 4.3 *Convexity of Nash equilibrium: The set of Nash equilibrium points in a game is convex.*

References

[1] James V. Burke *https://sites.math.washington.edu/ burke/crs/407/notes/section4.pdf*. University Of Washington.

[2] James King *https://sites.math.washington.edu/ king/coursedir/m308a01/Projects/m308a01-pdf/peydaye-saheli.pdf*. University Of Washington.