

Lecture 17: Linear Program Duality

Lecturer: Rong Ge

Scribe: Haoming Li

Overview

In this lecture, we will formulate the linear program (LP) for yet another graph problem. Then, we move on to talk about LP duality and its implication.

17.1 Using LP to Solve Graph Problems

In the previous lecture, we formed an LP instance for the maximum bipartite matching problem. In this section, we form an LP instance for the shortest path problem.

17.1.1 Example: Shortest Path

Given a weighted undirected graph $G = (V, E)$ and $s, t \in V$, we want to find the shortest path from s to t . We want to express the problem of finding (the weight of) such path as a linear program.

The variables of the LP will be the quantities that we would like to know. There are two possible choices: 1) $x_{uv} = 1$ if edge (u, v) is in the shortest s - t path, or 2) $x_u = d[u] =$ length of the shortest path from s to u . In this lecture, we will work with 1).

Let $w(u, v)$ be the weight on edge (u, v) . The constraint then follows:

$$\begin{aligned} x_s &= 0 \\ x_v &\leq x_u + w(u, v) \quad \forall (u, v) \in E \end{aligned}$$

For the object, a naive guess would be to minimize x_t since we are finding the shortest path. However, this is incorrect because setting $x_u = 0$ for all u would then be a feasible solution. The correct objective is, in fact, to maximize x_t . We will prove by induction that, with this objective, it is indeed true that the solution to the LP $x_t^* = d[t]$.

Proof: For every edge (u, v) , clearly $d[v] \leq d[u] + w(u, v)$. Therefore the $d[u]$'s form a feasible solution to the LP. Since x_t^* maximizes the value of variable x_t over all the feasible solutions, we have $x_t^* \geq d[t]$.

It remains to show that $x_t^* \leq d[t]$. Let v_1, v_2, \dots, v_k be a shortest s - t path. Thus $v_1 = s$ and $v_k = t$. We will show by induction that $x_{v_i}^* \leq d[v_i]$, $\forall i \in [1, k]$.

Base case: The constraint forces $x_s^* = 0$ and clearly $d[s] = 0$.

Induction hypothesis: Suppose $x_{v_i}^* \leq d[v_i]$ for some $i < k$.

Inductive step: since the solution satisfy the constraint, we have $x_{v_{i+1}}^* \leq x_{v_i}^* + w(v_i, v_{i+1})$. By the IH, we have $x_{v_i}^* + w(v_i, v_{i+1}) \leq d[v_i] + w(v_i, v_{i+1}) = d[v_{i+1}]$. Therefore, $x_{v_{i+1}}^* \leq d[v_{i+1}]$

In particular, since $v_k = t$, $x_t^* \leq d[t]$. Therefore, $x_t^* = d[t]$

■

17.2 LP Duality

Now we move on to introduce LP duality. We will begin with a motivating example from game theory.

17.2.1 Two-Player Zero-Sum Games

In game theory, two-player zero-sum game is the game with two competing players; when one wins the other loses. An example would be rock-paper-scissors. It can be specified by the *payoff matrix* illustrated here. There are two players, called Row and Column, and they each pick a move from $\{r, p, s\}$. They then look up the matrix entry corresponding to their moves, and Column pays Row this amount. It is Row's gain and Column's loss.

		Column		
		r	p	s
Row	r	0	-1	1
	p	1	0	-1
	s	-1	1	0

A player can choose to use a single strategy at all time. This is called a *pure strategy*. This is obviously not a good way to play the rock-paper-scissors game. It makes sense to consider using a *mixed strategy*: play r with probability p_1 , etc. The expected payoff of a player, therefore, would be taken over the 9 possible outcomes. Pure strategy, hence, is a special case of mixed strategy in which all the probability is put on one of the actions.

17.2.2 Solving Two-Player Zero-Sum Games by LP

Consider the game specified by the payoff matrix (of the row player) below. We will use LP to try to find a good strategy for the row player.

	A	B	C
A	3	1	-1
B	-2	3	2
C	1	-2	4

Let x_1, x_2, x_3 denote the probability the row player plays A, B and C in her mixed strategy. What makes a good strategy is that no matter what the opponent does, she gets a good payoff. Let x_4 denote the payoff. Consider the following LP:

$$\begin{aligned}
 & \max x_4 \\
 & x_1 + x_2 + x_3 = 1 \\
 & 3x_1 - 2x_2 + x_3 \geq x_4 \\
 & x_1 + 3x_2 - 2x_3 \geq x_4 \\
 & -x_1 + 2x_2 + 4x_3 \geq x_4 \\
 & x_1, x_2, x_3 \geq 0
 \end{aligned}$$

The solution to this LP turns out to be setting $x_1 = 9/19$, $x_2 = 6/19$ and $x_3 = 4/19$, which achieves an expected payoff of $x_4 = 1$.

17.2.3 Duality in Two-Player Zero-Sum Games

Since what the row player loses is what the column player gains, and vice versa, another way to look at this optimization is that the column player wants to make sure the row player's payoff is low no matter what the row player does.

Let y_1, y_2, y_3 denote the probability the column player plays A, B and C in her mixed strategy. She wants to make sure that the row player's payoff is lower than y_4 . Consider the following LP:

$$\begin{aligned} \min y_4 \\ y_1 + y_2 + y_3 &= 1 \\ 3y_1 + y_2 - y_3 &\leq x_4 \\ -2y_1 + 3y_2 + 2y_3 &\leq y_4 \\ y_1 - 2y_2 + 4y_3 &\leq y_4 \\ y_1, y_2, y_3 &\geq 0 \end{aligned}$$

This second LP turns out to be the *dual* of the first LP. The solution to this LP turns out to be setting $y_1 = y_2 = y_3 = 1/3$ and achieves an expected payoff of $y_4 = 1$.

Let us summarize. By solving the LP (the first LP), the row player (the maximizer) can determine a strategy for herself that guarantees an expected outcome of at least x_4 no matter what the column player does. And by solving the dual LP (the second LP), the column player (the minimizer) can guarantee an expected outcome of at most y_4 , no matter what the row player does.

This example is easily generalized to arbitrary games and shows the existence of mixed strategies that are optimal for both players and achieve the same value a fundamental result of game theory called the *min-max theorem*, which says the following:

Theorem 17.1 (Min-Max Theorem) *For any two-player zero-sum game, there is always a pair of optimal strategies and a single value V . If the row player plays its optimal strategy, then it can guarantee a payoff of at least V . If the column player plays its optimal strategy, then it can guarantee a payoff of at most V .*

This explains why, in our previous example, $x_4 = y_4 = 1$.

Corollary 17.2 (Strong Duality of LP) *The solution to an LP and its dual must be equal.*

17.2.4 Duality for LP, An Example

To demonstrate the duality for LP in general, let's step away from our previous two-player game example, and consider the following LP:

$$\begin{aligned} \min 2x_1 - 3x_2 + x_3 \\ x_1 - x_2 &\geq 1 \quad (1) \\ x_2 - 2x_3 &\geq 2 \quad (2) \\ -x_1 - x_2 - x_3 &\geq -7 \quad (3) \\ x_1, x_2, x_3 &\geq 0 \end{aligned}$$

The solution to this LP turns out to be setting $(x_1, x_2, x_3) = (4, 3, 0)$ and achieves an optimal solution of -1 . It follows directly, from the fact that the objective is to minimize, that the optimal solution is at most -1 . However, with the goal of driving the dual of this LP, we want to answer the question: how to prove that the optimal is *at least* -1 ?

Proof: Our strategy is to use a linear combination of the three constraints.

$2.5 \times (1) + 0.5 \times (3)$, we get $2x_1 - 3x_2 - 0.5x_3 \geq -1$. Since $x_3 \geq 0$, we have $2x_1 - 3x_2 + x_3 \geq 2x_1 - 3x_2 - 0.5x_3 \geq -1$

The numbers 2.5 and 0.5 appears magical at the moment; in fact, they are the solution to the dual LP! To see this, let's generalize the proof idea above.

Let $y_1, y_2, y_3 \geq 0$. $y_1 \times (1) + y_2 \times (2) + y_3 \times (3)$, we get:

$$\begin{aligned} y_1(x_1 - x_2) + y_2(x_2 - 2x_3) + y_3(-x_1 - x_2 - x_3) &\geq y_1 + 2y_2 - 7y_3 \\ (y_1 - y_3)x_1 + (-y_1 + y_2 - y_3)x_2 + (-2y_2 - y_3)x_3 &\geq y_1 + 2y_2 - 7y_3 \end{aligned}$$

To lower-bound the optimal solution, we want:

$$2x_1 - 3x_2 + x_3 \geq (y_1 - y_3)x_1 + (-y_1 + y_2 - y_3)x_2 + (-2y_2 - y_3)x_3$$

Since $x_1, x_2, x_3 \geq 0$, this can be true only if the coefficient on LHS is greater than the coefficient on RHS, that is:

$$\begin{aligned} y_1 - y_3 &\leq 2 \\ -y_1 + y_2 - y_3 &\leq -3 \\ -2y_2 - y_3 &\leq 1 \end{aligned}$$

Hence, for any y_1, y_2, y_3 that satisfy these three inequalities, we have:

$$2x_1 - 3x_2 + x_3 \geq y_1 + 2y_2 - 7y_3$$

In order for the proof to be strong, we want to maximize $y_1 + 2y_2 - 7y_3$. Hence, we have derived the following LP:

$$\begin{aligned} \max y_1 + 2y_2 - 7y_3 \\ y_1 - y_3 &\leq 2 \\ -y_1 + y_2 - y_3 &\leq -3 \\ -2y_2 - y_3 &\leq 1 \\ y_1, y_2, y_3 &\geq 0 \end{aligned}$$

This LP is, in fact, the dual of the original LP. The optimal solution to this LP is to set $(y_1, y_2, y_3) = (2.5, 0, 0.5)$ and achieves value -1 .

17.2.5 Duality for LP, In general

For any (primal) LP of the form:

$$\begin{aligned} \min \langle c, x \rangle \\ Ax &\geq b \\ x &\geq 0 \end{aligned}$$

Its dual LP is defined by:

$$\begin{aligned} \max \langle b, y \rangle \\ A^T y &\leq c \\ y &\geq 0 \end{aligned}$$

Each variable in the dual LP corresponds to a constraint in the primal LP, and each constraints in the dual LP corresponds to a variable in the primal LP. The strong duality of LP, introduces in Corollary 17.2, says that the two LP has the same optimal value.