Chapter 7

Linear Programming II: LP Duality

In this chapter, we go over the important concept of *duality* in linear programming.

7.1 A Motivating Example for Duality

In the last chapter, we introduced linear programming and three of its main applications: (a) solving "day-to-day life" (continuous optimization) problems, (b) designing algorithms for combinatorial optimization problems via *rounding*, and (c) proving different structural results. However, we only saw examples of the first two types of applications.

In this chapter, we will see an example of a type (c) application and then switch back to type (b) ones. To do so, we need to introduce one of the most important notions in LPs: *duality*. We use this section to present a motivating application of this concept and then formalize it in the next section.

Consider the following LP:

$$\begin{aligned} \max_{x_1,x_2} \quad & 2x_1 + 3x_2 \\ \text{subject to} \quad & 4x_1 + 8x_2 \leqslant 12 \\ & 2x_1 + x_2 \leqslant 3 \\ & 3x_1 + 2x_2 \leqslant 4 \\ & x_1, x_2 \geqslant 0. \end{aligned}$$

While the optimal solution to this LP may not be immediately clear, it is easy to see that it is at most 12 just given the first constraint:

$$\underbrace{2x_1 + 3x_2}_{\text{objective function}} \leqslant \underbrace{4x_1 + 8x_2 \leqslant 12}_{\text{first constraint}}.$$

In fact, we can divide both sides of the same constraint by two and obtain an even better upper bound of 6 on the objective function:

$$\underbrace{2x_1 + 3x_2}_{\text{objective function}} \leqslant 2x_1 + 4x_2 = \underbrace{\frac{1}{2} \cdot (4x_1 + 8x_2) \leqslant 6}_{\text{first constraint}}.$$

Yet another way is to sum up the first two constraints and then divide by three to obtain an upper bound of 5 as follows:

$$\underbrace{2x_1 + 3x_2}_{\text{objective function}} = \frac{1}{3} \cdot (6x_1 + 9x_2) = \frac{1}{3} \cdot (\underbrace{4x_1 + 8x_2}_{\text{LHS of first constraint}} + \underbrace{2x_1 + x_2}_{\text{LHS of second constraint}}) \leqslant \frac{1}{3} \cdot (12 + 3) = 5.$$

We can continue playing this game to find better and better upper bounds. But let us instead formulate this game entirely as follows:

1. Multiply *new* variables $y_1, y_2, y_3 \geqslant 0$ separately on the three inequalities, which leads to

$$\begin{aligned} &y_1 \cdot (4x_1 + 8x_2) \leqslant 12 \cdot y_1 \\ &y_2 \cdot (2x_1 + x_2) \quad \leqslant 3 \cdot y_2 \\ &y_3 \cdot (3x_1 + 2x_2) \leqslant 4 \cdot y_3 \\ &x_1, x_2 \geqslant 0 \\ &y_1, y_2, y_3 \geqslant 0. \end{aligned}$$

   (We ensure the order of inequalities does not change by making y_1, y_2, y_3 non-negative).

2. Sum up the above inequality constraints to obtain the following linear combination of the constraints:

$$(4y_1 + 2y_2 + 3y_3) \cdot x_1 + (8y_1 + y_2 + 2y_3) \cdot x_2 \leqslant 12y_1 + 3y_2 + 4y_3. \tag{7.1}$$

3. Ensure that the coefficient of x_1 is at least 2 and the coefficient of x_2 is at least 3 in the above inequality:

$$\begin{aligned} &4y_1 + 2y_2 + 3y_3 \geqslant 2 \\ &8y_1 + y_2 + 2y_3 \quad \geqslant 3. \end{aligned}$$

4. With these constraints, we now have that RHS of Eq (7.1) will be an *upper bound* on the objective value of the original program (as $x_1, x_2 \geqslant 0$); thus, to obtain the best upper bound, we should minimize this expression as much as possible. Putting all these together gives us the following formulation:

$$\begin{aligned} \min_{y_1,y_2,y_3} \quad & 12y_1 + 3y_2 + 4y_3 \\ \text{subject to} \quad & 4y_1 + 2y_2 + 3y_3 \geqslant 2 \\ & 8y_1 + y_2 + 2y_3 \quad \geqslant 3 \\ & y_1, y_2, y_3 \geqslant 0. \end{aligned}$$

   This is another LP that is called the **dual** of the original LP, which is called the **primal**. It is worth taking the dual of this LP again and seeing that we get back to the primal LP.

An immediate consequence of our dual construction is that *any* feasible solution to the dual LP already gives us a valid upper bound on the optimal objective of the primal LP (this is often referred to as **weak duality**). In other words, finding an upper bound for the primal LP reduces to finding a feasible solution in the dual. But how tight is this approach, i.e., how well can we upper bound the objective of the primal LP if we manage to minimize the dual? The answer is *completely tight*! Minimum value of the dual LP coincides with the maximum value of the primal LP (this is often referred to as **strong duality**).

In the rest of this chapter, we formalize these notions and then see some simple applications of duality.

7.2 Duality in Linear Programming

7.2.1 Primal and Dual LPs

Consider a primal LP as follows:

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & c^\top x \\ \text{subject to} \quad & A \cdot x \leqslant b \\ & x \geqslant 0, \end{aligned}$$

where $A \in \mathbb{R}^{m \times n}$, i.e., we have m constraints (besides non-negativity constraints). Recall that any LP can be written in this form.

1. Let $y = [y_1, \ldots, y_m]$ be the vector of *non-negative* coefficients in the four-step approach of the previous section for writing the dual. Let A_i denote the i-th row of matrix A and A^i be its i-th column. So, we have

$$A = \begin{bmatrix} a_{11} & \ldots & a_{1n} \\ \vdots & \ddots & \\ a_{m1} & & a_{mn} \end{bmatrix} = \begin{bmatrix} A_1 \\ \vdots \\ A_m \end{bmatrix} = \begin{bmatrix} A^1 & \ldots & A^n \end{bmatrix}.$$

   We can thus write:

$$\begin{aligned} y_1 \cdot \langle A_1 \, , \, x \rangle &\leqslant b_1 \cdot y_1 \\ y_2 \cdot \langle A_2 \, , \, x \rangle &\leqslant b_2 \cdot y_2 \\ &\vdots \\ y_m \cdot \langle A_m \, , \, x \rangle &\leqslant b_m \cdot y_m. \end{aligned}$$

   In the matrix form, this can be written as:

$$\begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} \odot \left(\begin{bmatrix} a_{11} & \ldots & a_{1n} \\ \vdots & \ddots & \\ a_{m1} & & a_{mn} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \right) \leqslant \begin{bmatrix} y_1 \\ \vdots \\ y_m \end{bmatrix} \odot \begin{bmatrix} b_1 \\ \vdots \\ b_m \end{bmatrix},$$

   where $\odot$ stands for the element-wise multiplication.

2. Now, by adding all these constraints together and factoring out $x_1, \ldots, x_n$, we obtain that:

$$\langle y \, , \, A^1 \rangle \cdot x_1 + \langle y \, , \, A^2 \rangle \cdot x_2 + \ldots + \langle y \, , \, A^n \rangle \cdot x_n \leqslant \langle b \, , \, y \rangle.$$

3. In order for the LHS to form an upper bound on the objective value of the primal, we need to have,

$$\begin{aligned} \langle y \, , \, A^1 \rangle &\geqslant c_1 \\ \langle y \, , \, A^2 \rangle &\geqslant c_2 \\ &\vdots \\ \langle y \, , \, A^n \rangle &\geqslant c_n, \end{aligned}$$

   so that for $x \geqslant 0$, we will get $c^\top x \leqslant \sum_{j=1}^n \langle y \, , \, A^j \rangle \cdot x_j$. In the matrix form, this is simply

$$A^\top y \geqslant c.$$

4. Finally, we would like to minimize the RHS of the equation in step (2), which gives us the following:

$$\begin{aligned} \min_{y \in \mathbb{R}^m} \quad & b^\top y \\ \text{subject to} \quad & A^\top \cdot y \geqslant c \\ & y \geqslant 0. \end{aligned}$$

This way, we get:

Definition 7.1. For any primal LP (P) as below, the dual LP (D) is defined as:

Primal LP:	**Dual LP:**
$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & c^\top x \\ \text{subject to} \quad & A \cdot x \leqslant b \\ & x \geqslant 0, \end{aligned}$	$\begin{aligned} \min_{y \in \mathbb{R}^m} \quad & b^\top y \\ \text{subject to} \quad & A^\top \cdot y \geqslant c \\ & y \geqslant 0. \end{aligned}$

An immediate observation given the above formula is that:

Observation 7.2 ("Dual of the dual is primal")**.** *Given an LP (P) and its dual (D), the dual of (D) is (P).*

7.2.2 Weak and Strong Duality Theorems

Weak duality. The weak duality theorem states that for any maximization primal LP (P), any feasible solution to its dual LP (D) gives an upper bound on the objective value of (P). Formally,

Theorem 7.3 ("Weak Duality")**.** *Let (P) be any maximization LP and (D) be its dual in the form of Definition 7.1. For any feasible point x in (P) and any feasible point y in (D):*

$$c^\top \cdot x \leqslant b^\top \cdot y.$$

Proof. Proof of this theorem is straightforward basically by our construction of the dual (and the whole motivation behind it). For any $i \in [m]$ and $j \in [n]$, let A_i and A^j denote the i-th row and the j-th column of the matrix $A \in \mathbb{R}^{m \times n}$, respectively. This way, the primal and dual constraints will be

$$\forall i \in [m] \quad \langle A_i \,, x \rangle \leqslant b_i \qquad \text{and} \qquad \forall j \in [n] \quad \langle A^j \,, y \rangle \geqslant c_j. \tag{7.2}$$

We now have

$$\begin{aligned} c^\top \cdot x = \sum_{j=1}^{n} c_j \cdot x_j &\leqslant \sum_{j=1}^{n} \langle A^j \,, y \rangle \cdot x_j && \text{(by the right inequality of Eq (7.2))} \\ &= \sum_{j=1}^{n} \sum_{i=1}^{m} a_{i,j} \cdot y_i \cdot x_j = \sum_{i=1}^{m} \langle A_i \,, x \rangle \cdot y_i \\ &\leqslant \sum_{i=1}^{m} b_i \cdot y_i = b^\top \cdot y. && \text{(by the left inequality of Eq (7.2))} \end{aligned}$$

□

We have the following immediate corollary of Theorem 7.3.

Observation 7.4. *Let (P) be any maximization LP and (D) be its dual in Definition 7.1. Then,*

1. *If the objective value of (P) can go to $+\infty$ (is unbounded from above), then (D) has no feasible solution.*
2. *If the objective value of (D) can go to $-\infty$ (is unbounded from below), then (P) has no feasible solution.*

Proof. For the first one, suppose (D) has a feasible solution y. Then, by weak duality, we need $b^\top \cdot y$ to be larger than any feasible objective value of (P), but the latter is going to $+\infty$, a contradiction. The second part holds via the same argument. $\square$

We also have a simple example that shows it is possible for both primal and dual LP to be infeasible.

Observation 7.5. *The following is an example of a primal (P) and dual (D) pair which are both infeasible*[1]*:*

(P)

$$\begin{aligned} \max_{(x_1,x_2)\in\mathbb{R}^2} \quad & x_1 - x_2 \\ \text{subject to} \quad & x_1 + x_2 = 1 \\ & 2x_1 + 2x_2 = -2 \end{aligned}$$

(D)

$$\begin{aligned} \min_{(y_1,y_2)\in\mathbb{R}^2} \quad & y_1 - 2y_2 \\ \text{subject to} \quad & y_1 + 2y_2 = 1 \\ & y_1 + 2y_2 = -1. \end{aligned}$$

Just using weak duality and the above observations, we can infer the following possibilities between the primal and dual LPs:

(P)/(D)	Unbounded	Infeasible	Feasible
Unbounded	no	yes	no
Infeasible	yes	yes	???
Feasible	no	???	???

Table 7.1: Here, for the primal (P), being unbounded means having objective value going to $+\infty$ (first row) and for the dual (D), means having objective value going to $-\infty$ (first column).

What happens in the remaining cases? What can we say about primal and dual when they are both feasible and bounded? Strong duality handles these cases (and more!).

Strong duality. Strong duality states that whenever at least one of the primal or dual is feasible and bounded, then so is the other one – more importantly, their objective values are equal. Formally,

Theorem 7.6 ("Strong Duality")**.** *Let (P) be any LP and (D) be its dual in Definition 7.1. Suppose at least one of these LPs is both feasible and bounded. Then, both LPs are feasible and bounded and have optimal solutions x^* for (P) and y^* for (D) such that*

$$c^\top \cdot x^* = b^\top \cdot y^*.$$

The proof of this theorem is beyond the scope of our book and is thus omitted.

But we can now use it to complete Table 7.1 as follows:

[1]You should first convince yourself that these two LPs are indeed primal and dual of each other. Also recall from Exercise 6.1 how these LPs can be written in the "usual" form of LPs we use in Definition 7.1.

(P)/(D)	Unbounded	Infeasible	Feasible
Unbounded	no	yes	no
Infeasible	yes	yes	no
Feasible	no	no	yes†

Table 7.2: All possibilities for primal and dual LPs.
†: In this case, both optimal objective value of (P) and (D) are the same.

7.3 LP Duality Applications

Let us now visit some basic applications of LP duality.

7.3.1 Optimizing vs Feasibility Checking LPs

We can use strong duality to say that checking feasibility of LPs in general is as hard as solving general LPs. In other words, if we have an algorithm that given a polyhedron can decide if it is feasible or not, then we can use the algorithm to also find optimal solutions of any given LP.

Concretely, suppose we are given an LP of the following form

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & c^\top x \\ \text{subject to} \quad & A \cdot x \leqslant b \\ & x \geqslant 0, \end{aligned}$$

where A is in $\mathbb{R}^{m \times n}$. We build an $(m+n)$-dimensional polyhedron as below by combining the LP and its dual and add one more constraint that requires the objective function of the LP and its dual LP to be equal.

Polyhedron (P1):

$$\begin{aligned} A \cdot x &\leqslant b \\ A^\top \cdot y &\geqslant c \\ c^\top \cdot x &= b^\top \cdot y \\ x, y &\geqslant 0. \end{aligned}$$

Then we give this polyhedron to the black box algorithm as input. If the algorithm decides that the polyhedron has a feasible point (x^*, y^*), then we know the original LP has an optimal solution x^* by the strong duality theorem (Theorem 7.6), and if this polyhedron is infeasible, then the original LP should also be either unbounded or infeasible.

7.3.2 $\mathsf{LP} \in \mathbf{NP} \cap \mathbf{coNP}$

Consider the decision version of the LP problem as the set

$$\mathsf{LP} := \{(P, q) \mid \text{(P) is an LP of the form in Definition 7.1 with objective value at least } q\}.$$

For simplicity, we further assume that in this problem we are promised that (P) is bounded and feasible (this assumption is not necessary). What complexity class this problem belongs to?

An immediate consequence of strong duality is that $\mathsf{LP} \in \mathbf{NP} \cap \mathbf{coNP}$, meaning that both the *Yes* instances and *No* instances of the LP problem can be *verified* in polynomial time:

- LP $\in$ **NP**: The certificate is a feasible solution x' for (P) such that $c^\top x' \geqslant q$. The algorithm simply verifies that x' is a feasible solution of (P) by checking each constraint respectively and verifying that $c^\top x' \geqslant q$. Then, we can conclude the optimal value of (P) is at least q and the algorithm only takes polynomial time[2].

- LP $\in$ **coNP**: We need to prove that the objective value of (P) is less than the given parameter q. To do this, we take the dual (D) and the certificate is a feasible solution y' for (D) such that $b^\top y' < q$. By strong duality (Theorem 7.6), such a certificate must always exist and by weak duality (Theorem 7.3), given y', we have the proof that $c^\top x \leqslant b^\top y' < q$ for all feasible x in (P). We can verify y' is a feasible solution in (D) as in the first part.

Of course recall that by Theorem 6.3, we have the stronger result that LP $\in$ **P** $\subseteq$ **NP** $\cap$ **coNP**.

7.3.3 Bipartite Matching and Bipartite Vertex Covers

We defined the bipartite matching and bipartite vertex cover problems in the previous chapter: in the former problem we want to pick the maximum number of edges that do not share any vertices, and in the latter the goal is to pick the minimum number of vertices so that every edge has at least one chosen endpoint. We prove that for both LPs, the optimal solution can be rounded to a matching and a vertex cover of the same size. I.e., the optimal value of each LP is equal to the value of the original graph problem. These results however did not say anything about the connections between these two problems. We use strong duality to address this.

Recall that the LP for the fractional matching in a bipartite graph $G = (L \sqcup R, E)$ is:

$$\begin{aligned} \max_{x \in \mathbb{R}^E} \quad & \sum x_e \\ \text{subject to} \quad & \sum_{e \ni v} x_e \leqslant 1 && \forall\, v \in L \sqcup R, \\ & x_e \geqslant 0 && \forall\, e \in E. \end{aligned}$$

Let (P) denote this LP. We can obtain the dual (D) of (P) using the same approach as before:

$$\begin{aligned} \min_{y \in \mathbb{R}^V} \quad & \sum y_v \\ \text{subject to} \quad & y_u + y_v \geqslant 1 && \forall\, (u,v) \in E, \\ & y_v \geqslant 0 && \forall\, v \in L \sqcup R. \end{aligned}$$

As before, *if* we assume the values in this dual (D) are integral, we get a program for minimum vertex cover, because $y_v = 1$ can be interpreted as picking the vertex v in the cover, and $y_u + y_v \geqslant 1$ constraints ensure that from every edge we are picking at least one vertex. So:

Dual of the maximum bipartite matching LP is the minimum bipartite vertex cover LP.

By the results we have proven for fractional matching and vertex cover, plus strong duality (Theorem 7.6), we have that for any *bipartite* graph:

$$\begin{aligned} \text{maximum matching size} &= \text{optimal value of (P)} \\ &= \text{optimal value of (D)} = \text{minimum vertex cover size}; \end{aligned}$$

[2] We have skipped a step here by ignoring the *bit-complexity* of x', namely, how large the proof itself needs to be. But it is true that the bit-complexity of the optimal solution of the LP is polynomial in the input size.

here, the first equality is by Proposition 6.5, the second one is by strong duality (Theorem 7.6), and the third one is by Proposition 6.6.

Figure 7.1 gives an illustration of this equality on a small bipartite graph, where a maximum matching and a minimum vertex cover both have size three.

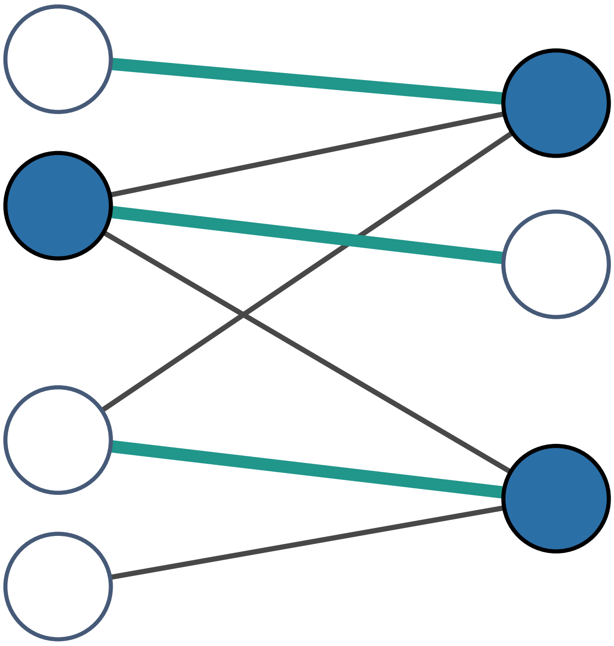

Figure 7.1: Duality of maximum matching and minimum vertex cover on a small bipartite graph: the three thick (green) edges form a maximum matching and the three shaded (blue) vertices form a minimum vertex cover, both of size three.

This duality result is known as *Kőnig's theorem*, which was first proved by Kőnig in 1931 combinatorially, without using any connection to LPs [Kő31] (see [Szá20] for an English translation). This is a great example of type (c) applications of LP: proving a structural result without necessarily any algorithmic aspects.