

LP DUALITY

CSCI 4113/6101

INSTRUCTOR: NORBERT ZEH

SEPTEMBER 18, 2025

Linear programs always come in pairs. Every LP P has an associated LP D called its **dual**. If P is a maximization LP, then D is a minimization LP. These two LPs have the remarkable property that any feasible solution of P has an objective function value that is no greater than the objective function value of any feasible solution of D . What's more, any *optimal* solution of P has exactly the same objective function value as any optimal solution of D .

This relationship between an LP and its dual has many applications. First, we can use it to verify that a given solution of an LP is optimal. Second, when designing approximation algorithms, we can often consider the LP formulation of the problem we are trying to solve and construct a feasible solution of the dual LP whose objective function value differs from the solution produced by our algorithm by only a small factor. This then proves that the solution produced by our algorithm is a good approximation of an optimal solution for the given problem instance. This technique is called **dual fitting** and will be discussed in more detail later in this course. Third, and probably most importantly, there exists an algorithmic technique, called the **primal-dual schema**, which focuses on simultaneously finding optimal solutions to an LP and its dual. We will see numerous examples of this technique, for finding maximum flows, for finding minimum-weight perfect matchings, and for finding approximate solutions of NP-hard problems.

This chapter is organized as follows: Sec. 1 introduces the concept of the dual of an LP and proves that the optimal solutions of an LP and its dual always have the same objective function value. The concept of LP duality applies to LPs in canonical form. When using LP duality to reason about the correctness of numerous optimization algorithms, however, the LPs that arise are not in canonical form. A perfectly correct way to obtain the duals of such non-canonical LPs is to first convert them into canonical form—remember, this is always possible—construct the dual, and then “convert the dual back”. I encourage you to do this with a few non-canonical LPs until you are comfortable with duality, but overall, this is a rather tedious method to construct the dual of a non-canonical LP. Sec. 2 discusses how to obtain the dual of an arbitrary LP, one not necessarily in canonical form.

1 LP DUALITY

Consider a maximization LP P in canonical form:

$$\begin{aligned} &\text{Maximize } cx \\ &\text{s.t. } Ax \leq b \\ &\quad x \geq 0 \end{aligned} \tag{1}$$

In the context of LP duality, we call P the **primal LP**. Its corresponding **dual LP** is

$$\begin{aligned} &\text{Minimize } b^T y \\ &\text{s.t. } A^T y \geq c^T \\ &\quad y \geq 0. \end{aligned} \tag{2}$$

If you write both LPs in tabular form, similar to a tableau, then the dual LP is obtained by transposing the primal LP:

	c_1	c_2	$\cdots$	c_n			b_1	b_2	$\cdots$	b_m
b_1	a_{11}	a_{12}	$\cdots$	a_{1n}		c_1	a_{11}	a_{21}	$\cdots$	a_{m1}
b_2	a_{21}	a_{22}	$\cdots$	a_{2n}		c_2	a_{12}	a_{22}	$\cdots$	a_{m2}
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
b_m	a_{m1}	a_{m2}	$\cdots$	a_{mn}		c_n	a_{1n}	a_{2n}	$\cdots$	a_{mn}
	x_1	x_2	$\cdots$	x_n			y_1	y_2	$\cdots$	y_m

Primal

Dual

Given that every column of the dual is the transpose of a row of the primal and vice versa, it is often convenient to think about each dual variable y_i (corresponding to the i th column of the dual) as being associated with the i th constraint of the primal (corresponding to the i th row of the primal), and about each primal variable x_j (corresponding to the j th column of the primal) as being associated with the j th constraint of the dual (corresponding to the j th row of the dual). This association is important, for example, in the theory of complementary slackness, discussed in the next topic.

Something that often hampers the understanding of LP duality is to ask “what the dual means”. This is particularly tempting when expressing other optimization problems as LPs and then constructing the dual, for example, when using the primal-dual schema. It is sometimes possible to attach an intuitive interpretation to the dual, and this intuitive interpretation can be very helpful. However, as one of the leading researchers in approximation algorithms, Vijay Vazirani, once put it, much of the power of LP duality stems from the fact that it is a purely mechanical transformation of the given LP. Asking what the dual means is sometimes helpful but is often the wrong thing to do. Take this advice seriously.

The first, useful and not all that difficult to prove, observation is that the objective function value of any solution of the dual LP is an upper bound on the objective function value of any solution of the primal LP.

LEMMA 1 (Weak Duality). *Let $\hat{x}$ be a feasible solution of the primal LP (1), and let $\hat{y}$ be a feasible solution of the dual LP (2). Then $c\hat{x} \leq b^T \hat{y}$.*

Proof. We have

$$\begin{aligned}
\sum_{j=1}^n c_j \hat{x}_j &\leq \sum_{j=1}^n \left(\sum_{i=1}^m a_{ij} \hat{y}_i \right) \hat{x}_j && \left(\text{by (2), } c_j \leq \sum_{i=1}^m a_{ij} \hat{y}_i \quad \forall 1 \leq j \leq n \right) \\
&= \sum_{i=1}^m \left(\sum_{j=1}^n a_{ij} \hat{x}_j \right) \hat{y}_i \\
&\leq \sum_{i=1}^m b_i \hat{y}_i && \left(\text{by (1), } \sum_{j=1}^n a_{ij} \hat{x}_j \leq b_i \quad \forall 1 \leq i \leq m \right).
\end{aligned}$$

□

An immediate consequence of Lem. 1 is the following corollary.

COROLLARY 2. *Let $\hat{x}$ be a feasible solution of the primal LP (1) and let $\hat{y}$ be a feasible solution of the dual LP (2). If $c\hat{x} = b^T \hat{y}$, then both $\hat{x}$ and $\hat{y}$ are optimal solutions of their respective LPs.*

Cor. 2 states a sufficient condition for a primal solution $\hat{x}$ and a dual solution $\hat{y}$ to be optimal. The following theorem states that this is also a necessary condition:

THEOREM 3 (Strong Duality). *Let $\hat{x}$ be a feasible solution of the primal LP (1) and let $\hat{y}$ be a feasible solution of the dual LP (2). $\hat{x}$ is an optimal solution of (1) and $\hat{y}$ is an optimal solution of (2) if and only if $c\hat{x} = b^T \hat{y}$.*

Proof. By Cor. 2, if $c\hat{x} = b^T \hat{y}$, then $\hat{x}$ and $\hat{y}$ are optimal solutions of (1) and (2), respectively. We need to prove that if $c\hat{x} < b^T \hat{y}$, then at least one of these two solutions is not optimal. To this end, we show that there exist solutions $\tilde{x}$ and $\tilde{y}$ of (1) and (2), respectively, such that $c\tilde{x} = b^T \tilde{y}$. By Lem. 1, $c\hat{x} \leq b^T \tilde{y} = c\tilde{x} \leq b^T \hat{y}$. If $c\hat{x} < b^T \hat{y}$, this implies that either $c\hat{x} < c\tilde{x}$ or $b^T \tilde{y} < b^T \hat{y}$ (or both). Thus, at least one of $\hat{x}$ and $\hat{y}$ is not an optimal solution, as required.

To find $\tilde{x}$ and $\tilde{y}$, we convert (1) into an LP in standard form, and solve this LP using the Simplex Algorithm. The primal solution $\tilde{x}$ will be the optimal primal solution found by the Simplex Algorithm (minus the slack variables introduced by the conversion to standard form). We will extract the dual solution $\tilde{y}$ from the final tableau produced by the Simplex Algorithm, which has the optimal primal solution as its BFS.

So consider the primal LP (1). Its corresponding LP in standard form is

$$\begin{aligned}
&\text{Maximize } cx \\
&\text{s.t. } Iz + Ax = b \\
&\quad x, z \geq 0,
\end{aligned} \tag{3}$$

where $z_1, \dots, z_m$ are slack variables corresponding to the m constraints of (1).

Now recall that the Simplex Algorithm produces a sequence of LPs tableaux $T^{(0)}, \dots, T^{(t)}$ representing LPs equivalent to (3) and that the solution $(\tilde{z}, \tilde{x})$ it outputs is the BFS of $T^{(t)}$. Let the LP represented by

$T^{(t)}$ be

$$\begin{aligned} & \text{Maximize } c'z + c''x + d' \\ & \text{s.t. } A'z + A''x = b' \\ & \quad x, z \geq 0. \end{aligned}$$

Recall that since $(\tilde{z}, \tilde{x})$ is the BFS of $T^{(t)}$, its objective function value is d' . Since the LP represented by $T^{(t)}$ is equivalent to (3), the objective function of (3) assigns the same objective function value to $(\tilde{z}, \tilde{x})$:

$$c\tilde{x} = d'$$

Since $\tilde{z} \geq 0$, we have $A\tilde{x} \leq b$, that is, $\tilde{x}$ is a feasible solution of (1). Its objective function value as a solution of (1) is also d' , because (1) and (3) have the same objective function.

It remains to demonstrate that there exists a feasible dual solution $\tilde{y}$ with the same objective function value. We show that the solution

$$\tilde{y} = -c'$$

(the value of each dual variable y_i is the negation of the objective function coefficient c'_i of the slack variable z_i in $T^{(t)}$) is such a solution.

For any solution $(\dot{z}, \dot{x})$ of (3), we have

$$\begin{aligned} \sum_{j=1}^n c_j \dot{x}_j &= \sum_{i=1}^m c'_i \dot{z}_i + \sum_{j=1}^n c''_j \dot{x}_j + d' \\ &= d' + \sum_{i=1}^m (-\tilde{y}_i) \left(b_i - \sum_{j=1}^n a_{ij} \dot{x}_j \right) + \sum_{j=1}^n c''_j \dot{x}_j \\ &= \left(d' - \sum_{i=1}^m b_i \tilde{y}_i \right) + \sum_{j=1}^n \left(c''_j + \sum_{i=1}^m a_{ij} \tilde{y}_i \right) \dot{x}_j. \end{aligned} \tag{4}$$

The first equality holds because (3) and the LP represented by $T^{(t)}$ are equivalent, and thus assign the same objective function value to the solution $(\dot{z}, \dot{x})$. The second inequality follows by substituting $-\tilde{y}_i$ for c'_i , and $b_i - \sum_{j=1}^n a_{ij} \dot{x}_j$ for $\dot{z}_i$. This is valid because we defined $\tilde{y} = -c'$ and, being a feasible solution of (3), $(\dot{z}, \dot{x})$ satisfies $\dot{z}_i = b_i - \sum_{j=1}^n a_{ij} \dot{x}_j$. The last equality is obtained by rearranging terms.

By substituting $(\tilde{z}, \tilde{x})$ for $(\dot{z}, \dot{x})$ in (4), we obtain that

$$\sum_{j=1}^n c_j \tilde{x}_j = \left(d' - \sum_{i=1}^m b_i \tilde{y}_i \right) + \sum_{j=1}^n \left(c''_j + \sum_{i=1}^m a_{ij} \tilde{y}_i \right) \tilde{x}_j. \tag{5}$$

Next observe that, for any index $1 \leq k \leq n$ and any real value Δ , the vector $(\check{z}, \check{x})$ defined as

$$\begin{aligned} \check{x}_j &= \begin{cases} \tilde{x}_j + \Delta & \text{if } j = k \\ \tilde{x}_j & \text{otherwise,} \end{cases} \\ \check{z}_i &= b_i - \sum_{j=1}^n a_{ij} \check{x}_j = b_i - \sum_{j=1}^n a_{ij} \tilde{x}_j - a_{ik} \Delta = \tilde{z}_i - a_{ik} \Delta \end{aligned}$$

is also a (not necessarily feasible) solution of (3). Thus, this solution also satisfies (4), and we obtain

$$\begin{aligned}\sum_{j=1}^n c_j \tilde{x}_j &= \left(d' - \sum_{i=1}^m b_i \tilde{y}_i \right) + \sum_{j=1}^n \left(c_j'' + \sum_{i=1}^m a_{ij} \tilde{y}_i \right) \tilde{x}_j \\ \sum_{j=1}^n c_j \tilde{x}_j + c_k \Delta &= \left(d' - \sum_{i=1}^m b_i \tilde{y}_i \right) + \sum_{j=1}^n \left(c_j'' + \sum_{i=1}^m a_{ij} \tilde{y}_i \right) \tilde{x}_j + \left(c_k'' + \sum_{i=1}^m a_{ik} \tilde{y}_i \right) \Delta.\end{aligned}$$

By subtracting (5), we obtain that

$$c_k \Delta = \left(c_k'' + \sum_{i=1}^m a_{ik} \tilde{y}_i \right) \Delta.$$

Since this is true for all $\Delta \in \mathbb{R}$, this implies that

$$c_k = c_k'' + \sum_{i=1}^m a_{ik} \tilde{y}_i. \quad (6)$$

Since this holds for all $1 \leq k \leq n$, we can substitute this equality into (5) and obtain

$$\begin{aligned}\sum_{j=1}^n c_j \tilde{x}_j &= \left(d' - \sum_{i=1}^m b_i \tilde{y}_i \right) + \sum_{j=1}^n \left(c_j'' + \sum_{i=1}^m a_{ij} \tilde{y}_i \right) \tilde{x}_j \\ &= \left(d' - \sum_{i=1}^m b_i \tilde{y}_i \right) + \sum_{j=1}^n c_j \tilde{x}_j\end{aligned}$$

and, hence,

$$d' - \sum_{i=1}^m b_i \tilde{y}_i = 0.$$

Thus, the objective function value of $\tilde{y}$ as a solution of (2) is

$$\sum_{i=1}^m b_i \tilde{y}_i = d'',$$

as claimed. To see that $\tilde{y}$ is a feasible solution of (2), observe that $T^{(t)}$ satisfies $c' \leq 0$ and $c'' \leq 0$ (otherwise, the Simplex Algorithm would continue pivoting). Thus, $\tilde{y} = -c' \geq 0$, and (6) together with $c'' \leq 0$ implies that

$$c \leq A^T \tilde{y}.$$

Therefore, $\tilde{y}$ is a feasible solution of (2). □

2 THE DUAL OF A NON-CANONICAL LP

We defined the dual of an LP in canonical form, a maximization LP

$$\begin{aligned} &\text{Maximize } cx \\ &\text{s.t. } Ax \leq b \\ &\quad x \geq 0, \end{aligned} \tag{7}$$

as the minimization LP

$$\begin{aligned} &\text{Minimize } b^T y \\ &\text{s.t. } A^T y \geq c^T \\ &\quad y \geq 0. \end{aligned} \tag{8}$$

What if the LP is not in canonical form? Can we still easily define its dual?

Let's start with a warm-up exercise. Every mathematician expects any duality construction to be its own inverse: if we take the dual of the dual, we should get back the primal. LP duality meets this expectation:

Observe that the dual LP (8) can easily be converted into canonical form by negating the objective function and all constraints:

$$\begin{aligned} &\text{Maximize } -b^T y \\ &\text{s.t. } -A^T y \leq -c^T \\ &\quad y \geq 0. \end{aligned}$$

The dual of this LP is

$$\begin{aligned} &\text{Minimize } -cx \\ &\text{s.t. } -Ax \geq -b \\ &\quad x \geq 0. \end{aligned}$$

But that's just our primal (7) written as a minimization LP via negation. So the dual of a maximization LP is a minimization LP, the dual of a minimization LP is a maximization LP, and LP duality is well-behaved: the dual of the dual of an LP is just the original LP.

Next let us tackle the problem of constructing the dual of an arbitrary maximization LP. We can partition its constraints into three groups: upper bound constraints, lower bound constraints, and equality constraints. We can also partition its variables into three groups: non-negative variables, non-positive variables, and completely unconstrained variables. Thus, a general maximization LP can be written like this:

$$\begin{aligned} &\text{Maximize } c_1 x_+ + c_2 x_- + c_3 x_{\pm} \\ &\text{s.t. } A_1 x_+ + A_2 x_- + A_3 x_{\pm} \leq b_1 \\ &\quad A_4 x_+ + A_5 x_- + A_6 x_{\pm} \geq b_2 \\ &\quad A_7 x_+ + A_8 x_- + A_9 x_{\pm} = b_3 \\ &\quad x_+ \geq 0 \\ &\quad x_- \leq 0. \end{aligned} \tag{9}$$

Note that c_1, c_2, c_3 are row vectors and $b_1, b_2, b_3, x_+, x_-, x_{\pm}$ are column vectors, not numbers.

The canonical form of (9) is

$$\begin{aligned}
& \text{Maximize } c_1 x_+ - c_2 x'_- + c_3 x'_\pm - c_3 x''_\pm \\
& \text{s.t. } A_1 x_+ - A_2 x'_- + A_3 x'_\pm - A_3 x''_\pm \leq b_1 \\
& \quad -A_4 x_+ + A_5 x'_- - A_6 x'_\pm + A_6 x''_\pm \leq -b_2 \\
& \quad A_7 x_+ - A_8 x'_- + A_9 x'_\pm - A_9 x''_\pm \leq b_3 \\
& \quad -A_7 x_+ + A_8 x'_- - A_9 x'_\pm + A_9 x''_\pm \leq -b_3 \\
& \quad x_+, x'_-, x'_\pm, x''_\pm \geq 0.
\end{aligned}$$

Here, $x'_- = -x_-$ and $x'_\pm - x''_\pm = x_\pm$.

The dual of this LP is

$$\begin{aligned}
& \text{Minimize } b_1^T y_+ - b_2^T y'_- + b_3^T y'_\pm - b_3^T y''_\pm \\
& \text{s.t. } A_1^T y_+ - A_4^T y'_- + A_7^T y'_\pm - A_7^T y''_\pm \geq c_1^T \\
& \quad -A_2^T y_+ + A_5^T y'_- - A_8^T y'_\pm + A_8^T y''_\pm \geq -c_2^T \\
& \quad A_3^T y_+ - A_6^T y'_- + A_9^T y'_\pm - A_9^T y''_\pm \geq c_3^T \\
& \quad -A_3^T y_+ + A_6^T y'_- - A_9^T y'_\pm + A_9^T y''_\pm \geq -c_3^T \\
& \quad y_+, y'_-, y'_\pm, y''_\pm \geq 0.
\end{aligned} \tag{10}$$

Here, y_+ is the vector of dual variables corresponding to the original upper bound constraints, y'_- is the vector of dual variables corresponding to the original lower bound constraints, $y'_\pm$ is the vector of dual variables corresponding to the upper bound constraints derived from the equality constraints, and $y''_\pm$ is the vector of dual variables corresponding to the lower bound constraints derived from the equality constraints.

Here comes the cool part. First, we can flip the signs of the second and last groups of constraints in (10):

$$\begin{aligned}
& \text{Minimize } b_1^T y_+ - b_2^T y'_- + b_3^T y'_\pm - b_3^T y''_\pm \\
& \text{s.t. } A_1^T y_+ - A_4^T y'_- + A_7^T y'_\pm - A_7^T y''_\pm \geq c_1^T \\
& \quad A_2^T y_+ - A_5^T y'_- + A_8^T y'_\pm - A_8^T y''_\pm \leq c_2^T \\
& \quad A_3^T y_+ - A_6^T y'_- + A_9^T y'_\pm - A_9^T y''_\pm \geq c_3^T \\
& \quad A_3^T y_+ - A_6^T y'_- + A_9^T y'_\pm - A_9^T y''_\pm \leq c_3^T \\
& \quad y_+, y'_-, y'_\pm, y''_\pm \geq 0.
\end{aligned}$$

This LP is equivalent to:

$$\begin{aligned}
& \text{Minimize } b_1^T y_+ - b_2^T y'_- + b_3^T y'_\pm - b_3^T y''_\pm \\
& \text{s.t. } A_1^T y_+ - A_4^T y'_- + A_7^T y'_\pm - A_7^T y''_\pm \geq c_1^T \\
& \quad A_2^T y_+ - A_5^T y'_- + A_8^T y'_\pm - A_8^T y''_\pm \leq c_2^T \\
& \quad A_3^T y_+ - A_6^T y'_- + A_9^T y'_\pm - A_9^T y''_\pm = c_3^T \\
& \quad y_+, y'_-, y'_\pm, y''_\pm \geq 0.
\end{aligned}$$

Second, we can replace every occurrence of $y'_\pm - y''_\pm$ in the objective function or in any constraint with $y_\pm = y'_\pm - y''_\pm$. While $y'_\pm, y''_\pm \geq 0$, $y_\pm$ is unconstrained; any value we assign to $y_\pm$ can be written as the difference of two non-negative vectors assigned to $y'_\pm$ and $y''_\pm$. We can also replace every occurrence of

y'_- with $-y_- = y'_-$. Since $y'_- \geq 0$, we have $y_- \leq 0$. This gives our final dual LP:

$$\begin{aligned}
& \text{Minimize } b_1^T y_+ + b_2^T y_- + b_3^T y_{\pm} \\
& \text{s.t. } A_1^T y_+ + A_4^T y_- + A_7^T y_{\pm} \geq c_1^T \\
& \quad A_2^T y_+ + A_5^T y_- + A_8^T y_{\pm} \leq c_2^T \\
& \quad A_3^T y_+ + A_6^T y_- + A_9^T y_{\pm} = c_3^T \\
& \quad y_+ \geq 0 \\
& \quad y_- \leq 0.
\end{aligned} \tag{11}$$

If you refer back to the primal LP (9) that we started with, you can observe that, just as in the dual of an LP in canonical form, the coefficients of the dual objective function are the right-hand sides of the primal constraints, the right-hand sides of the dual constraints are the coefficients of the primal objective function, and the matrix of coefficients in the dual constraints is the transpose of the matrix of coefficients in the primal constraints. What differs is that some dual variable are now constrained to be non-positive or completely unconstrained, and some constraints are upper bound constraints or equality constraints instead of lower bound constraints. By looking at (9) and (11), we can easily discern the rules that govern which constraint should be an upper bound, lower bound or equality constraint:

- The dual constraint corresponding to any non-negative primal variable is a lower bound constraint (the constraints corresponding to variables x_+ in our example).
- The dual constraint corresponding to any non-positive primal variable is an upper bound constraint (the constraints corresponding to variables x_- in our example).
- The dual constraint corresponding to any unconstrained primal variable is an equality constraint (the constraints corresponding to variables $x_{\pm}$ in our example).

Similarly, our example tells us which dual variables are non-negative, non-positive or unconstrained:

- Every dual variable corresponding to an upper bound constraint in the primal is non-negative (the vector y_+ in our example).
- Every dual variable corresponding to a lower bound constraint in the primal is non-positive (the vector y_- in our example).
- Every dual variable corresponding to an equality constraint in the primal is unconstrained (the vector $y_{\pm}$ in our example).

These rules apply when the primal is a maximization LP and, consequently, the dual is a minimization LP. We observed before that

- The dual of a maximization LP is a minimization LP and the dual of a minimization LP is a maximization LP.

If the primal is a minimization LP and we convert it into a dual that is a maximization LP, however, the rules above are reversed: Unconstrained primal variables still translate into equality constraints in the dual and equality constraints in the primal still translate into unconstrained dual variables. However, upper bound and lower bound constraints in the primal now translate into non-positive and non-negative dual variables, respectively, and non-negative and non-positive primal variables now translate into upper bound and lower bound constraints in the dual, respectively.