How to Take the Dual of a Linear Program

Sébastien Lahaie*

January 12, 2015

This is a revised version of notes I wrote several years ago on taking the dual of a linear program (LP), with some bug and typo fixes and new worked examples. The purpose of these notes is to: (1) explain how to recognize an LP, (2) show how to take the dual of an LP, and (3) list the fundamental results relating primal and dual LPs.

These notes do not provide any proofs and do not explain any of the deep geometric insight behind linear programming duality. For this there are several good references [1, 2, 3], as well as course lecture notes easily found online. The main goal here is simply to explain in detail the mechanical procedure of taking the dual. I believe that for students of computer science, operations research, and economics, taking the dual should be as natural as the process of taking derivatives and integrals.

I have found that the procedure given here is easily memorized after practicing on a few LPs, as opposed to the rule-based approaches given in certain textbooks. The intermediate steps also produce useful information as a by-product (complementary slackness conditions).

1 Formulation

A linear program (LP) is a formulation of an optimization problem: a minimization or maximization of an objective function over some domain. The objective function is linear, and the domain, or *feasible set*, is defined by linear constraints. Rather than give the generic form of an LP, here's a specific instance that we will work with throughout these notes. Once we've studied

*Microsoft Research, New York, NY 10011. `slahaie@microsoft.com`

this example, it should be quite clear what an LP looks like in general.

$$
\begin{aligned}
\max_{x_1 \geq 0,\, x_2 \leq 0,\, x_3} \quad & v_1 x_1 + v_2 x_2 + v_3 x_3 && (1) \\
\text{such that} \quad & a_1 x_1 + x_2 + x_3 \leq b_1 && (2) \\
& x_1 + a_2 x_2 = b_2 && (3) \\
& a_3 x_3 \geq b_3 && (4)
\end{aligned}
$$

The variables here are x_1, x_2, x_3. The remaining terms are constants (e.g., v_1, a_2, b_3). An LP consists of an objective function (1), and a set of inequality and equality constraints (2–4). The objective function

$$f(x) = v_1 x_1 + v_2 x_2 + v_3 x_3$$

is a *linear* function. Formally, this means that for two vectors x^1 and x^2 and real-valued constants c_1 and c_2 we have $f(c_1 x^1 + c_2 x^2) = c_1 f(x^1) + c_2 f(x^2)$. You can check that this indeed holds for our objective. The objective can either be a *maximization* or *minimization*. In our case, we have a maximization problem.

The left-hand side of each constraint is also a linear function. The right-hand sides are all constants. (You may encounter some LPs where the right-hand side for some constraint includes some variables, but the constraint can always be rearranged so that the right-hand side is constant.) In an LP, we can't have strict inequalities; for example, $x_1 + x_2 < 3$ wouldn't be allowed.

Constraints that specify whether each variable is non-negative, non-positive, or unrestricted are special and the convention in these notes is to list those under the max or min. This only holds for these kinds of constraints. A constraint of the form $x_1 \geq 2$, for example, is not special, whereas $x_1 \geq 0$ is special. In our LP, x_1 must be non-negative, x_2 must be non-positive, and x_3 is unrestricted (it helps to be explicit about this).

The dual of our LP turns out to be as follows.

$$
\begin{aligned}
\min_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad & b_1 \lambda_1 + b_2 \lambda_2 - b_3 \lambda_3 \\
\text{such that} \quad & a_1 \lambda_1 + \lambda_2 \geq v_1 && (5) \\
& \lambda_1 + a_2 \lambda_2 \leq v_2 && (6) \\
& \lambda_1 - a_3 \lambda_3 = v_3 && (7)
\end{aligned}
$$

This is also an LP with variables $\lambda_1, \lambda_2, \lambda_3$. The next section describes how this is derived.

2 Primal and Dual

The LP we start with is typically called the *primal.* To each LP there is associated another LP called its *dual.* The dual can be derived from the primal via a purely mechanical procedure. Let's see how to derive the dual of the LP of the previous section. There are seven steps. The first two steps put the primal in a 'standard form'.

Step 1. *If necessary, rewrite the objective as a minimization.*

In our case the objective (1) is a maximization, so we rewrite it as

$$\min \ -v_1x_1 - v_2x_2 - v_3x_3$$

A solution that maximizes an objective also minimizes the negative of that objective, so this doesn't change the set of optimal solutions to the LP.

Step 2. *Rewrite each inequality constraint as a "less than or equal", and rearrange each constraint so that the right-hand side is 0.*

After this step our LP now looks as follows.

$$\begin{aligned} \min_{x_1 \geq 0,\, x_2 \leq 0,\, x_3} \quad & -v_1x_1 - v_2x_2 - v_3x_3 \\ \text{such that} \quad & a_1x_1 + x_2 + x_3 - b_1 \leq 0 && (8) \\ & x_1 + a_2x_2 - b_2 = 0 && (9) \\ & -a_3x_3 + b_3 \leq 0 && (10) \end{aligned}$$

Note that the special constraints listed below the min do *not* change.

Step 3. *Define a non-negative dual variable for each inequality constraint, and an unrestricted dual variable for each equality constraint.*

To constraints (8) and (10) we associate variables $\lambda_1 \geq 0$ and $\lambda_3 \geq 0$ respectively. To constraint (9) we associate λ_2, which is unrestricted.

Step 4. *For each constraint, eliminate it and add the term*

$$\text{(dual variable)}\cdot\text{(left-hand side of constraint)}$$

to the objective. Maximize the result over the dual variables.

Concretely, for the first constraint (8) we would remove it and *add* the following term to the objective:

$$\lambda_1 (a_1x_1 + x_2 + x_3 - b_1)\,.$$

If we do this for each constraint (except of course special constraints), and maximize the result over the dual variables, we get

$$\begin{aligned}
\max_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \; \min_{x_1 \geq 0,\, x_2 \leq 0,\, x_3} \quad & -v_1 x_1 - v_2 x_2 - v_3 x_3 \\
+ \quad & \lambda_1 \left(a_1 x_1 + x_2 + x_3 - b_1\right) \qquad (11) \\
+ \quad & \lambda_2 \left(x_1 + a_2 x_2 - b_2\right) \qquad (12) \\
+ \quad & \lambda_3 \left(-a_3 x_3 + b_3\right) \qquad (13)
\end{aligned}$$

It helps to think of this as a two-player game, with an "outer player" and an "inner player". The outer player goes first and chooses some values for $\lambda_1, \lambda_2, \lambda_3$ (respecting the special constraints). With these values fixed, the inner player then chooses some values for x_1, x_2, x_3 so as to minimize the objective. Conscious that the inner player will do this, the outer player will choose values for $\lambda_1, \lambda_2, \lambda_3$ so that the minimum value in the inner player's minimization problem is as large as possible.

Step 5. *We now have an objective with several terms of the form*

(dual variable)·(expression with primal variables),

plus remaining terms involving only primal variables. *Rewrite the objective so that it consists of several terms of the form*

(primal variable)·(expression with dual variables),

plus remaining terms involving only dual variables.

If we do this to the objective of the previous step, we obtain

$$\begin{aligned}
\max_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \; \min_{x_1 \geq 0,\, x_2 \leq 0,\, x_3} \quad & -b_1 \lambda_1 - b_2 \lambda_2 + b_3 \lambda_3 \\
+ \quad & x_1(a_1 \lambda_1 + \lambda_2 - v_1) \qquad (14) \\
+ \quad & x_2(\lambda_1 + a_2 \lambda_2 - v_2) \qquad (15) \\
+ \quad & x_3(\lambda_1 - a_3 \lambda_3 - v_3) \qquad (16)
\end{aligned}$$

This step takes great care. If you get any sign wrong or forget any term, the resulting dual will of course be wrong, but also very misleading and confusing.

Step 6. *Remove each term of the form* (primal variable)·(expression with dual variables) *and replace it with a constraint of the form:*

- *expression* ≥ 0, *if the primal variable is* non-negative.

- *expression* ≤ 0*, if the primal variable is* non-positive*.*
- *expression* $= 0$*, if the primal variable is* unrestricted*.*

This step may seem hard to memorize but there's an intuitive reason for the rules. Let's consider term (14):

$$x_1(a_1\lambda_1 + \lambda_2 - v_1).$$

Because $x_1 \geq 0$, we introduce constraint $a_1\lambda_1 + \lambda_2 - v_1 \geq 0$. Why must this hold? If we had $a_1\lambda_1 + \lambda_2 - v_1 < 0$, then note that the inner player could choose $x_1 \to +\infty$ (i.e., arbitrarily large), and thus the objective value would be $-\infty$. Since the outer player wants to maximize the value of the inner player's minimization problem, it should therefore choose the values of $\lambda_1, \lambda_2, \lambda_3$ so that $a_1\lambda_1 + \lambda_2 - v_1 \geq 0$.

The same reasoning applies to the other two terms. The outer player must choose its values so that $\lambda_1 + a_2\lambda_2 - v_2 \leq 0$ otherwise the inner player can make the term

$$x_2(\lambda_1 + a_2\lambda_2 - v_2)$$

tend to $-\infty$ by choosing $x_2 \to -\infty$. Finally, since x_3 is unrestricted, the only way to make sure the term

$$x_3(\lambda_1 - a_3\lambda_3 - v_3)$$

cannot approach $-\infty$ is to set the dual variables such that $\lambda_1 - a_3\lambda_3 - v_3 = 0$. After these changes we have a new LP in which the primal variables no longer appear.

$$\begin{aligned}
\max_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad & -b_1\lambda_1 - b_2\lambda_2 + b_3\lambda_3 \\
\text{such that} \quad & a_1\lambda_1 + \lambda_2 - v_1 \geq 0 && (17)\\
& \lambda_1 + a_2\lambda_2 - v_2 \leq 0 && (18)\\
& \lambda_1 - a_3\lambda_3 - v_3 = 0 && (19)
\end{aligned}$$

Step 7. *If the original LP was a maximization rewritten as a minimization in Step 1, rewrite the result of the previous step as a minimization.*

The result looks as follows. Optionally, the constraints can be also be rearranged in whichever form is most natural.

$$\begin{aligned}
\min_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad & b_1\lambda_1 + b_2\lambda_2 - b_3\lambda_3 \\
\text{such that} \quad & a_1\lambda_1 + \lambda_2 \geq v_1 && (20)\\
& \lambda_1 + a_2\lambda_2 \leq v_2 && (21)\\
& \lambda_1 - a_3\lambda_3 = v_3 && (22)
\end{aligned}$$

Note that, according to these instructions, the dual will only have either non-negative or unrestricted variables. It may be more intuitive to redefine some variables via the substitution $\lambda \leftarrow -\lambda$, depending on their interpretation, and as a result the dual may have non-negative, non-positive, or unrestricted variables just like the primal.

This completes the process of taking the dual. By this process, we see that there is a variable in the dual for each constraint in the primal, and vice-versa. As an exercise, you can check for yourself that by taking the dual of this new LP, you recover the primal.

3 Key Results

An LP may be *infeasible*, *unbounded*, or have a *finite optimum*. An LP is infeasible if there is no solution that satisfies all the given constraints. For instance, suppose that in our original primal program we have $a_1 = a_2 = a_3 = 1$, and $b_1 = -1$ whereas $b_2 + b_3 \geq 1$. Then it is not hard to check that there is no solution to constraints (2–4). (Feasibility only has to do with the constraints, not with the objective.)

An LP can alternatively be unbounded. This means that for any feasible solution, there is another feasible solution with strictly higher objective value, in the case of a maximization program, or strictly lower objective value, in the case of a minimization program. For instance, suppose that in our original primal we have $a_1 = a_2 = 1$ and $a_3 = -1$, whereas $b_1 = b_2 = b_3 = 0$. The coefficients satisfy $v_3 < 0$ (the others don't matter). Then note that $(0, 0, c)$ is a feasible solution for any $c \leq 0$. The objective value of this solution is cv_3, which we can make arbitrarily large by letting $c \to -\infty$. Technically, this means there is no optimal solution, even though there are feasible solutions.

	Finite optimum	**Unbounded**	**Infeasible**
Finite optimum	✓	×	×
Unbounded	×	×	✓
Infeasible	×	✓	✓

Table 1: Possible combinations for the primal and its dual.

If a program is feasible and bounded, it has a finite optimum (though it may not be unique). Table 1 lists the possible relationships between feasibility of a primal program and its dual. In particular, note that if the primal program is unbounded, the dual is infeasible, and if the dual program is unbounded, the primal is infeasible. But it may be the case that both are infeasible.

The value of the objective function of the primal program at an optimum is denoted V_P and the optimum value of the dual is denoted V_D. The central result in the theory of linear programming is the following.

Theorem 1. (Strong duality) *If an LP has an optimal solution, so does its dual, and* $V_P = V_D$.

This result relates the values of the primal and dual programs, but not their solutions. The following result is written in reference to our running example. It should be straightforward to adapt it to other LPs.

Theorem 2. (Complementary slackness) *Let* (x_1, x_2, x_3) *and* $(\lambda_1, \lambda_2, \lambda_3)$ *be feasible solutions to the primal and dual problems, respectively. They are optimal solutions for their respective problems if and only if*

$$
\begin{aligned}
\lambda_1 (a_1 x_1 + x_2 + x_3 - b_1) &= 0 \\
\lambda_2 (x_1 + a_2 x_2 - b_2) &= 0 \\
\lambda_3 (-a_3 x_3 + b_3) &= 0
\end{aligned}
$$

and

$$
\begin{aligned}
x_1(a_1\lambda_1 + \lambda_2 - v_1) &= 0 \\
x_2(\lambda_1 + a_2\lambda_2 - v_2) &= 0 \\
x_3(\lambda_1 - a_3\lambda_3 - v_3) &= 0
\end{aligned}
$$

These constraints are known as the *complementary slackness conditions.* Note that the first three constraints can be read off from (11–13), and the final three from (14–16). So we get the complementary slackness conditions as a by-product of our procedure for taking the dual.

The complementary slackness conditions are sometimes written in alternative forms. For instance, we can write the condition

$$\lambda_1 (a_1 x_1 + x_2 + x_3 - b_1) = 0$$

instead as

$$\lambda_1 > 0 \implies a_1 x_1 + x_2 + x_3 - b_1 = 0,$$

or equivalently (taking the contrapositive)

$$a_1 x_1 + x_2 + x_3 - b_1 < 0 \implies \lambda_1 = 0.$$

Given an optimal primal solution, the complementary slackness conditions define a system of equalities that can be solved to identify an optimal dual solution (the latter also has to satisfy the dual feasibility constraints), and vice-versa.

4 Worked Examples

Here are two more worked examples. The first demonstrates the procedure using vector notation, which is often used to formulate LPs. The second gives an instance of an LP where the variables in the primal and dual have concrete interpretations.

Vector Notation

The constants of the problem are as follows. Let $c = (c_1, \ldots, c_n)$ be a vector with n components, which we call an n-vector. Let b be an m_1-vector and d an m_2-vector. Let A be an $m_1 \times n$ matrix and C be an $m_2 \times n$ matrix. The variables of the problem are captured in x, an n-vector. Consider the following LP.

$$
\begin{aligned}
\min_{x \geq 0} \quad & c^\top x \\
\text{s.t.} \quad & Ax = b \\
& Cx \leq d
\end{aligned}
$$

When applied to vectors, inequalities are understood component-wise. Let us now go through the steps. For Step 1, the objective is already a minimization. For Step 2, we rewrite the constraints in the canonical form.

$$
\begin{aligned}
\min_{x \geq 0} \quad & c^\top x \\
\text{s.t.} \quad & Ax - b = 0 && (23) \\
& Cx - d \leq 0 && (24)
\end{aligned}
$$

For Step 3, we define an m_1-vector λ of unrestricted variables associated with constraints (23), and an m_2-vector μ of non-negative variables associated with constraints (24). For Step 4, we can now eliminate the constraints and replace them with terms in the objective, and maximize over the newly defined dual variables.

$$\begin{aligned}\max_{\lambda,\mu\geq 0}\ \min_{x\geq 0} &\quad c^\top x \\ &+\ \lambda^\top(Ax-b) \\ &+\ \mu^\top(Cx-d)\end{aligned}$$

Step 5 is the most delicate. We collect all terms that apply to the primal variables.

$$\begin{aligned}\max_{\lambda,\mu\geq 0}\ \min_{x\geq 0} &\quad -\lambda^\top b - \mu^\top d \\ &+\ \left(c + A^\top\lambda + C^\top\mu\right)^\top x\end{aligned} \tag{25}$$

For Step 6, we replace the terms involving primal variables with constraints. To reason about this, note first that

$$c + A^\top\lambda + C^\top\mu \tag{26}$$

is an n-vector. If any of its components is negative, say component i, then the inner minimization could make the objective arbitrarily small by taking $x_i \to +\infty$. (Recall that $x \geq 0$, and we take λ and μ as fixed for the inner minimization.) To avoid this, the outer maximization must choose variables λ and $\mu \geq 0$ such that (26) has non-negative components. We therefore obtain the following LP, which only involves dual variables.

$$\begin{aligned}\max_{\lambda,\mu\geq 0} &\quad -\lambda^\top b - \mu^\top d \\ \text{s.t.} &\quad c + A^\top\lambda + C^\top\mu \geq 0\end{aligned} \tag{27}$$

Finally, Step 7 does not apply because we originally had a minimization program. We can optionally rearrange some of the constraints and switch the signs of some variables if it makes the dual clearer. For instance, the variables tend to appear with negative signs, so we might consider the substitutions $\lambda \leftarrow -\lambda$ and $\mu \leftarrow -\mu$. The final form is

$$\begin{aligned}\max_{\lambda,\mu\leq 0} &\quad \lambda^\top b + \mu^\top d \\ \text{s.t.} &\quad A^\top\lambda + C^\top\mu \leq c\end{aligned} \tag{28}$$

To be extremely careful it is always a good exercise to then recover the primal from the derived dual.

Assignment Problem

In this example we have n agents and m items. The value of item j to agent i is denoted by v_{ij}. There is 1 unit of each item. It is possible to assign fractional units of items, but each agent can obtain at most 1 unit in total. We would like to assign items to agents in order to maximize the total value. This can be formulated as an LP. Let $x_{ij} \in [0, 1]$ be a variable corresponding to the number of units of item j assigned to agent i.

$$\begin{aligned} \max_{x \geq 0} \quad & \sum_{i=1}^{n} \sum_{j=1}^{m} v_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_{i=1}^{n} x_{ij} \leq 1 \qquad (j = 1, \ldots, m) && (29) \\ & \sum_{j=1}^{m} x_{ij} \leq 1 \qquad (i = 1, \ldots, n) && (30) \end{aligned}$$

The objective is to maximize the total value of the assigned items to the agents, as stated. Constraints (29) ensure that at most 1 unit of each item j is assigned. Constraints (30) ensure that each agent i obtains at most 1 unit in total from the items. Note that we have the special constraints $x \geq 0$. We haven't explicitly constrained $x \leq 1$ because it is implied from (29–30) combined with $x \geq 0$. Let us now take the dual.

For Step 1 we first rewrite the objective as a minimization by negating each term. For Step 2 we then rewrite the constraints in canonical form, ensuring that each constraint is a "less than or equal" and that the right-hand sides are 0. We obtain the following form of the LP.

$$\begin{aligned} \min_{x \geq 0} \quad & \sum_{i=1}^{n} \sum_{j=1}^{m} -v_{ij} x_{ij} \\ \text{s.t.} \quad & \sum_{i=1}^{n} x_{ij} - 1 \leq 0 \quad (j = 1, \ldots, m) && (31) \\ & \sum_{j=1}^{m} x_{ij} - 1 \leq 0 \quad (i = 1, \ldots, n) && (32) \end{aligned}$$

In Step 3 we define the dual variables. For each constraint j in (31), associate a dual variable $p_j \geq 0$. For each constraint i in (32), associate a dual variable $\pi_i \geq 0$. For Step 4, we now eliminate constraints (31–32) and replace them with objective terms involving the dual variables.

$$
\begin{aligned}
\max_{p\geq 0,\pi\geq 0} \ \min_{x\geq 0} \quad & \sum_{i=1}^{n}\sum_{j=1}^{m} -v_{ij}x_{ij} \\
+ \quad & \sum_{j=1}^{m} p_j \left(\sum_{i=1}^{n} x_{ij} - 1\right) \\
+ \quad & \sum_{i=1}^{n} \pi_i \left(\sum_{j=1}^{m} x_{ij} - 1\right)
\end{aligned}
$$

For Step 5, we collect all terms for each primal variable x_{ij}.

$$
\begin{aligned}
\max_{p\geq 0,\pi\geq 0} \ \min_{x\geq 0} \quad & -\sum_{j=1}^{m} p_j - \sum_{i=1}^{n} \pi_i \\
+ \quad & \sum_{i=1}^{n}\sum_{j=1}^{m} x_{ij}(-v_{ij} + p_j + \pi_i)
\end{aligned} \tag{33}
$$

Now for Step 6, we have $x_{ij} \geq 0$, so to make sure the inner maximization cannot take the terms in (33) to $-\infty$, the outer maximization must ensure $-v_{ij} + p_j + \pi_i \geq 0$ for each i and j. This leads to the following LP, purely in terms of the dual variables.

$$
\begin{aligned}
\max_{p\geq 0,\pi\geq 0} \quad & -\sum_{j=1}^{m} p_j - \sum_{i=1}^{n} \pi_i \\
\text{s.t.} \quad & -v_{ij} + p_j + \pi_i \geq 0 \qquad \begin{array}{c}(i = 1,\ldots,n)\\(j = 1,\ldots,m)\end{array}
\end{aligned}
$$

For the final Step 8, we have to switch the objective to a minimization, because the primal objective was originally switched in Step 1. After slight re-arranging, the dual is

$$
\begin{aligned}
\min_{p\geq 0,\pi\geq 0} \quad & \sum_{j=1}^{m} p_j + \sum_{i=1}^{n} \pi_i \\
\text{s.t.} \quad & \pi_i \geq v_{ij} - p_j \qquad \begin{array}{c}(i = 1,\ldots,n)\\(j = 1,\ldots,m)\end{array}
\end{aligned}
$$

There is in fact a great deal of structure and economic intuition associated with these relatively simple LPs. Without going too much into the details, dual variables p_j can be interpreted as item *prices*. In that case the expression $v_{ij} - p_j$ can then be interpreted as the surplus or *utility* (i.e., value minus price) of item j to agent i. Note that at an optimal solution, we will necessarily have

$$\pi_i = \max_j \{v_{ij} - p_j\}.$$

Therefore π_i should be interpreted as a utility variable, which at an optimal solution records the maximum utility agent i can obtain by selecting a single item under prices p. The complementary slackness conditions may also have useful interpretations. For instance, from (33) and invoking Theorem 2, we obtain that optimal primal and dual solutions must satisfy the relationship

$$x_{ij}(-v_{ij} + p_j + \pi_i) = 0,$$

or more intuitively in this case,

$$x_{ij} > 0 \implies \pi_i = v_{ij} - p_j.$$

As just mentioned, at an optimal solution π_i is the maximum utility agent i can achieve over the items, under prices p. Thus the condition says that if some units of item j are assigned to agent i at the optimal primal solution x, then item j maximizes the agent's utility under optimal dual solution p.

References

[1] Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to linear optimization*. Athena Scientific, 1997.

[2] Vašek Chvátal. *Linear programming*. WH Freeman and Company, New York, 1983.

[3] David G. Luenberger and Yinyu Ye. *Linear and nonlinear programming*. Springer Science & Business Media, 2008.