Math 482: Linear Programming[1] — Mikhail Lavrov

Lecture 14: Duality and the Simplex Tableau

February 26, 2020 — *University of Illinois at Urbana-Champaign*

1 Finding the dual optimal solution

Suppose that we have a linear program (and a dual for it) as shown below:

$$(\mathbf{P}) \begin{cases} \underset{x,y\in\mathbb{R}}{\text{maximize}} & 2x+3y \\ \text{subject to} & -x+\ y \le 3 \quad (u) \\ & x-2y \le 2 \quad (v) \\ & x+\ y\le 7 \quad (w) \\ & x,y \ge 0 \end{cases} \qquad (\mathbf{D}) \begin{cases} \underset{u,v,w\in\mathbb{R}}{\text{minimize}} & 3u+2v+7w \\ \text{subject to} & -u+\ v+\ w \ge 2 \quad (x) \\ & u-2v+\ w\ge 3 \quad (y) \\ & u,v,w \ge 0 \end{cases}$$

We can solve this linear program using the simplex method:

	x	y	s_1	s_2	s_3	
s_1	-1	1	1	0	0	3
s_2	1	-2	0	1	0	2
s_3	1	1	0	0	1	7
$-z$	2	3	0	0	0	0

$\rightsquigarrow$

	x	y	s_1	s_2	s_3	
y	0	1	$1/2$	0	$1/2$	5
s_2	0	0	$3/2$	1	$1/2$	10
x	1	0	$-1/2$	0	$1/2$	2
$-z$	0	0	$-1/2$	0	$-5/2$	-19

Fine print: when we add slack variables to the primal program, the dual seems like it will change slightly. Actually, two different effects cancel out. When we change $-x+y \le 3$ into $-x+y+s_1 = 3$, we turn an inequality into an equation, which makes the dual variable u unrestricted (and the same for v and w). However, we also get three more primal variables s_1, s_2, s_3 and those give us three more dual constraints... which are exactly the constraints $u \ge 0$, $v \ge 0$, $w \ge 0$.

We can read off an optimal solution to $(\mathbf{P})$ from this pretty easily: it has $x = 2$ and $y = 5$. Today, we'll also learn to read off the optimal solution to $(\mathbf{D})$ from the simplex tableau.

First, recall some formulas we used for the revised simplex method. We computed the reduced cost of a variable in two steps:

$$r_j = c_j - \mathbf{u}^\mathsf{T} A_j \qquad \text{where} \qquad \mathbf{u}^\mathsf{T} = \mathbf{c}_\mathcal{B}{}^\mathsf{T} A_\mathcal{B}^{-1}.$$

It is not a coincidence that we denoted this vector with $\mathbf{u}$. When we finish solving the problem, $\mathbf{u}$ will become our optimal dual solution:

Lemma 1.1. *If $\mathcal{B}$ is an optimal basis for $(\mathbf{P})$, then $\mathbf{u}^\mathsf{T} = \mathbf{c}_\mathcal{B}{}^\mathsf{T} A_\mathcal{B}^{-1}$ is an optimal solution to $(\mathbf{D})$.*

[1] This document comes from the Math 482 course webpage: https://faculty.math.illinois.edu/~mlavrov/courses/482-spring-2020.html

Proof. Two steps go into checking that $\mathbf{u}^\mathsf{T}$ is an optimal dual solution.

First, we should check feasibility. Here, we make use of the formula for reduced costs: the reduced cost of the j^{th} variable is $r_j = c_j - \mathbf{u}^\mathsf{T} A_j$.

When we've found the optimal solution (assuming we're maximizing), all reduced costs should be nonpositive: $r_j \le 0$. That says that $c_j - \mathbf{u}^\mathsf{T} A_j \le 0$, or $\mathbf{u}^\mathsf{T} A_j \ge c_j$. But that's exactly the j^{th} dual constraint! Checking that every reduced cost is nonnegative tells us that every dual constraint is satisfied.

What if we're minimizing, not maximizing? No problem! In that case, we are asking for $r_j \ge 0$ for all j, so the conclusion changes: we conclude that $\mathbf{u}^\mathsf{T} A_j \le c_j$ for all j. But the dual constraints also change: if (**P**) is a minimization problem, then (**D**) is a maximization problem, and its constraints will be $\le$ constraints. Either way, the constraints we get are exactly the constraints we want.

Once we know that $\mathbf{u}^\mathsf{T}$ is dual feasible, we want to check that it's optimal. Here, we remember the formula $z_0 = \mathbf{c}_\mathcal{B}{}^\mathsf{T} A_\mathcal{B}^{-1} \mathbf{b}$. We can rewrite this in two ways:

1. From (**P**)'s point of view, it is $z_0 = \mathbf{c}_\mathcal{B}{}^\mathsf{T} \mathbf{p}$. But our optimal solution $\mathbf{x}$ has $\mathbf{x}_\mathcal{B} = \mathbf{p}$ and $\mathbf{x}_\mathcal{N} = \mathbf{0}$. Following this line of reasoning, we just get $z_0 = \mathbf{c}^\mathsf{T}\mathbf{x}$, which is not surprising: that's just the objective function of (**P**).
2. From (**D**)'s point of view, it's better to write z_0 as $\mathbf{u}^\mathsf{T}\mathbf{b}$, which is the dual objective function.

But either way, we get the same value z_0. So we have $\mathbf{u}^\mathsf{T}\mathbf{b} = \mathbf{c}^\mathsf{T}\mathbf{x}$. When a feasible solution to (**P**) has the same objective value as a feasible solution to (**D**), we know that both are optimal. □

This is a very convenient formula if we're using the revised simplex method, because then we've already computed $\mathbf{u}^\mathsf{T} = \mathbf{c}_\mathcal{B}{}^\mathsf{T} A_\mathcal{B}^{-1}$ when we were checking all the reduced costs. What if we just have the simplex tableau above?

It turns out that it is *also* easy to use the formula if we started with a linear program in the form $A\mathbf{x} \le \mathbf{b}$, and added slack variables. Then the new matrix for the program is the augmented matrix $\begin{bmatrix} A & I \end{bmatrix}$. If we compute the reduced costs of all the slack variables, we get

$$\mathbf{r}_\mathcal{S}{}^\mathsf{T} = \mathbf{c}_\mathcal{S}{}^\mathsf{T} - \mathbf{u}^\mathsf{T} A_\mathcal{S}.$$

Here, $\mathbf{c}_\mathcal{S}{}^\mathsf{T} = \mathbf{0}^\mathsf{T}$, because the slack variables do not appear in the objective function. Meanwhile, $A_\mathcal{S}$ (the submatrix corresponding to the slack variables) is just the identity matrix, I. So we get

$$\mathbf{r}_\mathcal{S}{}^\mathsf{T} = \mathbf{0}^\mathsf{T} - \mathbf{u}^\mathsf{T} I = -\mathbf{u}^\mathsf{T}.$$

This means that we can compute the dual solution by taking the negatives of the reduced costs of the slack variables, in the final tableau. In our example, the reduced costs are $\begin{bmatrix} -1/2 & 0 & -5/2 \end{bmatrix}$, and therefore the dual solution is

$$\begin{bmatrix} u & v & w \end{bmatrix} = -\mathbf{r}_\mathcal{S}{}^\mathsf{T} = \begin{bmatrix} 1/2 & 0 & 5/2 \end{bmatrix}.$$

We can read it off without having to do any arithmetic!

A word of caution: this will give you slightly misleading answers when (**P**) contains $\ge$ constraints. If we started off with a constraint such as $x + y \ge 1$, we would turn it into an equation by *first* multiplying it by -1 (to get $-x-y \leq -1$) and *then* adding a slack variable (to get $-x-y+s = -1$).

When we apply this process, we get the correct dual variable for the dual of the equation constraint $-x-y+s=-1$. This happens to be the same as the correct dual variable for the dual of the inequality constraint $-x-y \leq -1$. However, multiplying the constraint by -1 also has the effect of multiplying the dual variable by -1.

So, if your original program had $\geq$ constraints, then the values you get from the formula $\mathbf{u}^\mathsf{T} = -\mathbf{r}_\mathcal{S}{}^\mathsf{T}$ should be negated if you want the correct dual solution for the original program.

2 Strong duality

We now have a formula for an optimal dual solution with the same objective value as the optimal primal solution. This is a proof of strong duality: that the primal and dual linear programs have the same optimal objective value!

(This is somehow not a very well-respected proof, because in order for it to work, we need to know that the simplex method always works, which has its own convoluted proof. But we've already done all that hard work, so we might as well use it.)

Specifically:

- From the lemma in the previous section, we show that whenever the primal program has an optimal solution, so does the dual (because we found one) and the objective values agree (because that's how we proved it was optimal).
- Because duality is symmetric, we get the converse for free: whenever the dual program has an optimal solution, so does the primal, and the objective values agree.
- Our proof only works for linear programs in equational form, but we can put all linear programs in this form, and that plays well with duality.

As a bonus, we can learn what happens when the primal or the dual program is infeasible or unbounded. Specifically, there are only four possible cases:

1. Both the primal program and the dual program have an optimal solution. (Most of the examples we've seen have been like this.)
2. The primal is infeasible, and the dual is unbounded.
3. The primal is unbounded, and the dual is infeasible.
4. Both the primal and the dual are infeasible.

Why can't other things happen? Well, there's two arguments ruling out the other cases.

First, the primal program and the dual program can't both be unbounded. If the primal is unbounded, that means in particular that it has a feasible solution. The objective value of that feasible solution limits the possible objective values in the dual, so the dual can't be unbounded.

Second, strong duality says that if one of the two programs has an optimal solution, the other can't be infeasible or unbounded.

Here are some examples of the cases that are possible. In the pair

$$(\mathbf{P}) \begin{cases} \underset{x\in\mathbb{R}}{\text{maximize}} & x \\ \text{subject to} & x \le -1 \quad (u) \\ & x \ge 0 \end{cases} \qquad (\mathbf{D}) \begin{cases} \underset{u\in\mathbb{R}}{\text{minimize}} & -u \\ \text{subject to} & u \ge 1 \quad (x) \\ & u \ge 0 \end{cases}$$

the primal program is infeasible (we can't have $x \le -1$ and $x \ge 0$ at the same time) and the dual program is unbounded (by setting u to be very large, we make $-u$ very small).

We can get an example where the primal program is infeasible and the dual program is unbounded simply by reversing the roles of the two programs. Or, if we want to keep things as a max and a min, we could do a slight variant of the example above:

$$(\mathbf{P}) \begin{cases} \underset{y\in\mathbb{R}}{\text{maximize}} & y \\ \text{subject to} & -y \le 1 \quad (v) \\ & y \ge 0 \end{cases} \qquad (\mathbf{D}) \begin{cases} \underset{v\in\mathbb{R}}{\text{minimize}} & v \\ \text{subject to} & -v \ge 1 \quad (y) \\ & v \ge 0 \end{cases}$$

Here, any nonnegative y is primal feasible, but no v is dual feasible.

To get an example where both linear programs are infeasible, just take these two examples and stick them together:

$$(\mathbf{P}) \begin{cases} \underset{x,y\in\mathbb{R}}{\text{maximize}} & x+y \\ \text{subject to} & x \le -1 \quad (u) \\ & -y \le 1 \quad (v) \\ & x, y \ge 0 \end{cases} \qquad (\mathbf{D}) \begin{cases} \underset{u,v\in\mathbb{R}}{\text{minimize}} & -u+v \\ \text{subject to} & u \ge 1 \quad (x) \\ & -v \ge 1 \quad (y) \\ & u, v \ge 0 \end{cases}$$

Here, the primal is infeasible because we can't choose a value of x, and the dual is infeasible because we can't choose a value of v.