

Lecture notes 4: Duality

Vincent Conitzer

1 Introduction

Let us again consider the linear program for our original painting problem instance:

$$\begin{array}{ll}\text{maximize} & 3x_1 + 2x_2 \\ \text{subject to} & \\ & 4x_1 + 2x_2 \leq 16 \\ & x_1 + 2x_2 \leq 8 \\ & x_1 + x_2 \leq 5 \\ & x_1 \geq 0; x_2 \geq 0\end{array}$$

We already know that the optimal solution is to set $x_1 = 3, x_2 = 2$, for a solution value of $9 + 4 = 13$. But how do we *prove* that this is optimal? Originally, we proved this using a picture: from the picture it was easy to see that the optimal solution is where the lines corresponding to the blue and red (first and third) constraints intersect, and it is easy to show that this is at the point $x_1 = 3, x_2 = 2$. However, this approach does not generalize to more variables—perhaps convincing pictures can still be drawn for three variables, but how would you draw one for four variables? Moreover, it is somewhat informal and relies on our ability to draw straight lines. Once we introduce an algorithm that always finds the optimal solution to a linear program, we can simply give the run of the algorithm as a proof of optimality. Still, such a proof would be quite long and unpleasant to read.

We can use *duality* to give a much nicer proof of optimality. To introduce duality, let us first pursue a more modest objective: we want to find a simple upper bound on the objective value (revenue) that can be obtained. The lower the upper bound, the better; and if we can give an upper bound of 13, then we will have proven that our solution $x_1 = 3, x_2 = 2$ is optimal, since it in fact produces a revenue of 13.

Here is a very simple upper bound: we know that the blue constraint $4x_1 + 2x_2 \leq 16$ must hold. Because x_1 is nonnegative, we must have $3x_1 + 2x_2 \leq 4x_1 + 2x_2 \leq 16$. Hence, we know that 16 is an upper bound on the revenue $3x_1 + 2x_2$ that we can obtain.

We can get a better upper bound by multiplying the red constraint $x_1 + x_2 \leq 5$ by 3. This gives us $3x_1 + 3x_2 \leq 15$, and because x_2 is nonnegative, $3x_1 + 2x_2 \leq 3x_1 + 3x_2$. So 15 is also an upper bound.

We can also multiply constraints by fractions; moreover, we can add them to each other. If we multiply the blue constraint by $1/2$, we get $2x_1 + x_2 \leq 8$. If we then add this to the red constraint, we get $3x_1 + 2x_2 \leq 13$. Hence, we get an upper bound of 13—the optimal solution value. We cannot get a lower upper bound than that, since then, it would not be a lower bound, because a revenue of 13 can actually be obtained. We note that in finding this optimal upper bound, we used only the constraints that are *binding* in the optimal solution—that is, they are satisfied without any slack. This is intuitively sensible because these are the constraints that matter. We will see later on in these lecture notes that this is a special case of something known as *complementary slackness*.

All of the above upper bounds were obtained by multiplying the constraints with various non-negative numbers, and then adding the results together. If we multiply by a negative number, the inequality in the constraint flips, so that would not be helpful. Let y_1, y_2, y_3 be the numbers by which we multiply the blue, green, and red constraints, respectively. Hence, the first upper bound that we obtained consisted of setting $y_1 = 1, y_2 = 0, y_3 = 0$; the second upper bound consisted of setting $y_1 = 0, y_2 = 0, y_3 = 3$; and the third, optimal upper bound consisted of setting $y_1 = 0.5, y_2 = 0, y_3 = 1$.

Now, let us think about how we might find the optimal upper bound. As it turns out, we can model this as a linear program! As we already noted, all the y_i should be nonnegative. In addition, we need to make sure that, once we add up all the multiplied constraints, the coefficient on x_1 is at least 3, because the coefficient on x_1 in the objective is 3. That is, $4y_1 + y_2 + y_3 \geq 3$. Similarly, for x_2 , we must have $2y_1 + 2y_2 + y_3 \geq 2$. These inequalities guarantee that we will in fact have a correct upper bound. Finally, given that we satisfy all of these inequalities, we want to minimize the upper bound that this gives us, since smaller upper bounds are better. The upper bound that we obtain is a linear combination of the right-hand sides of the original inequalities: $16y_1 + 8y_2 + 5y_3$. Putting it together, we get:

$$\begin{aligned}
 &\text{minimize } 16y_1 + 8y_2 + 5y_3 \\
 &\text{subject to} \\
 &4y_1 + y_2 + y_3 \geq 3 \\
 &2y_1 + 2y_2 + y_3 \geq 2 \\
 &y_1 \geq 0; y_2 \geq 0; y_3 \geq 0
 \end{aligned}$$

This linear program is the *dual* of the original (also called *primal*) linear program. Any feasible solution to the dual corresponds to an upper bound on any solution to the primal (this is known as the *weak duality* property). Also, we saw that in this case, there is a feasible solution to the dual whose value is equal to the optimal value for the primal (and therefore, we know that this feasible solution to the dual is in fact optimal for the dual, because if there were a better solution for the dual, it would no longer be an upper bound for the primal). It turns out that this was no accident: in fact, the values of the optimal solutions to the primal and dual are always equal. This property is known as *strong duality*. It should be emphasized that we are only talking about linear programs here. If we are considering a (mixed) integer program (say, a maximization problem), we can take its linear program relaxation, and then take the dual of that. Because the optimal value of the LP relaxation is no lower than that of the original MIP, feasible solutions to the dual will in fact correspond to upper bounds for the original MIP. However, in general, the optimal value for the dual of the LP relaxation will not be equal to the optimal value for the original MIP.

Finally, let us consider what happens if we take the dual of the dual. We have not yet shown how to take the dual of a minimization problem. One way to do this is to convert it to maximization form first, but it is not difficult to reason about the minimization problem directly. Again, we want to put nonnegative weights on the constraints; let us call the weights x_1 and x_2 for the above minimization problem. This time, we want to find *lower* bounds. For example, if we set $x_1 = 3, x_2 = 2$, we obtain the relationship $16y_1 + 8y_2 + 5y_3 \geq (3 \cdot 4 + 2 \cdot 2)y_1 + (3 \cdot 1 + 2 \cdot 2)y_2 + (3 \cdot 1 + 2 \cdot 1)y_3 \geq 3 \cdot 3 + 2 \cdot 2 = 13$. That is, we now need to make sure that the coefficients resulting from our linear combination of constraints are *less than* or equal to those in the objective, and subject to this we want to *maximize* the linear combination of the right-hand side. This results in the following linear program:

$$\begin{array}{ll}
\text{maximize} & 3x_1 + 2x_2 \\
\text{subject to} & \\
& 4x_1 + 2x_2 \leq 16 \\
& x_1 + 2x_2 \leq 8 \\
& x_1 + x_2 \leq 5 \\
& x_1 \geq 0; x_2 \geq 0
\end{array}$$

This is just the primal problem again! Thus, the dual of the dual is the primal (this is true in general).

2 The dual: general form and results

We recall the following general form of a linear program:

$$\begin{array}{ll}
\text{maximize} & c_1x_1 + \dots + c_nx_n \\
\text{subject to} & \\
& a_{11}x_1 + \dots + a_{1n}x_n \leq b_1 \\
& \vdots \\
& a_{m1}x_1 + \dots + a_{mn}x_n \leq b_m \\
& x_1 \geq 0; \dots; x_n \geq 0
\end{array}$$

The dual of this linear program is:

$$\begin{array}{ll}
\text{minimize} & b_1y_1 + \dots + b_my_m \\
\text{subject to} & \\
& a_{11}y_1 + \dots + a_{m1}y_m \geq c_1 \\
& \vdots \\
& a_{1n}y_1 + \dots + a_{mn}y_m \geq c_n \\
& y_1 \geq 0; \dots; y_m \geq 0
\end{array}$$

Note that in the dual, we have m variables and n constraints. This is because variables in the dual correspond to constraints in the primal, and vice versa.

We have the following results.

Theorem 1 (Weak duality) *If x_1, \dots, x_n and y_1, \dots, y_m are feasible solutions for the primal and dual, respectively, and the primal problem is a maximization problem, then $c_1x_1 + \dots + c_nx_n \leq b_1y_1 + \dots + b_my_m$.*

This is merely saying that any feasible solution to the dual gives an upper bound to any feasible solution to the primal—which should be clear from the way we introduced the dual. The strong duality theorem is more difficult to prove:

Theorem 2 (Strong duality) *If x_1, \dots, x_n and y_1, \dots, y_m are optimal solutions for the primal and dual, respectively, then $c_1x_1 + \dots + c_nx_n = b_1y_1 + \dots + b_my_m$. Moreover, if either the primal or the dual has an optimal solution, then so does the other.*

We will give a proof of the strong duality theorem later. It should be noted that this theorem assumes that one of the primal and dual has an optimal solution. If the primal (dual) is unbounded, then the dual (primal) has no feasible solution—because such a feasible solution would give a bound on the former problem, which contradicts its unboundedness. It is possible for the primal and dual to be simultaneously infeasible.

The following result refers to the slack variables of the primal and dual, which indicate “by how much” a constraint is met. That is, the primal slack variables are given by $w_i = b_i - a_{i1}x_1 - \dots - a_{in}x_n$, and the dual slack variables are given by $z_j = c_j - a_{1j}y_1 - \dots - a_{mj}y_m$.

Theorem 3 (Complementary slackness) *Suppose that x_1, \dots, x_n and y_1, \dots, y_m are feasible solutions for the primal and dual, respectively, and let w_1, \dots, w_m and z_1, \dots, z_n be the corresponding slacks for the primal and dual, respectively. Then x_1, \dots, x_n and y_1, \dots, y_m are both optimal if and only if the following conditions hold:*

1. for all i , $w_i y_i = 0$, and
2. for all j , $z_j x_j = 0$.

That is, for the optimal solutions to the primal and the dual, for any variable that is set to a positive value in the primal (dual), the corresponding slack variable in the dual (primal) must be set to zero. Conversely, if all of these constraints are satisfied for a pair of feasible solutions, then these solutions must be optimal.

Using complementary slackness, we can quickly convert an optimal solution to the primal to an optimal solution to the dual (and vice versa). Let us consider the painting problem instance again, and suppose that somehow we know that the optimal solution to the primal is $x_1 = 3, x_2 = 2$. If we consider the second (green) constraint $x_1 + 2x_2 \leq 8$, we see that there is a slack of 1 on this constraint for the optimal solution ($x_1 + 2x_2 = 3 + 4 = 7 = 8 - 1$). This means that the corresponding dual variable y_2 must be zero in an optimal solution to the dual. Also, because both x_1 and x_2 are set to nonzero values, the dual constraints $4y_1 + y_2 + y_3 \geq 3$ and $2y_1 + 2y_2 + y_3 \geq 2$ must both be binding in an optimal solution to the dual, that is, there can be no slack on these constraints. This leaves us to solve the system of equalities $4y_1 + y_3 = 3$ and $2y_1 + y_3 = 2$, which has as its unique solution $y_1 = 0.5, y_3 = 1$.

3 Equality constraints and duals

Suppose that our linear program has an equality constraint $a_{i1}x_1 + \dots + a_{in}x_n = b_i$. We can still take the dual of this linear program by simply converting it back to the standard inequality format. It is instructive to see what happens if we do this. The equality constraint is equivalent to the following two constraints: $a_{i1}x_1 + \dots + a_{in}x_n \leq b_i$ and $a_{i1}x_1 + \dots + a_{in}x_n \geq b_i$. The latter can be converted to $-a_{i1}x_1 - \dots - a_{in}x_n \leq -b_i$, giving us two less-than-or-equal-to constraints. Now, in the dual, let y'_i be the variable associated with the former constraint, and y''_i the variable associated with the latter constraint. Then, as part of the left-hand side of the j th constraint of the dual, we have $a_{ij}y'_i - a_{ij}y''_i = a_{ij}(y'_i - y''_i)$. Similarly, as part of the dual’s objective, we have $b_i y'_i - b_i y''_i = b_i(y'_i - y''_i)$. That is, y'_i and y''_i always occur together as $(y'_i - y''_i)$. Hence, we can just replace them with a single new variable $y_i = (y'_i - y''_i)$, and then all the constraints and the objective of the dual look the same as they would have if the original constraint had been a less-than-or-equal-to constraint. There is one difference: y_i can take any negative value (by setting $y''_i > y'_i$) as well as any nonnegative value (by setting $y'_i \geq y''_i$). That is, y_i is a *free variable*.