COS 330: Great Ideas in Theoretical Computer Science

Fall 2025

Precept 7

My dual personality

Learn

Instead of introducing an extension to the lecture, today we are going to review the duality concepts from lecture in more depth. It's a very powerful tool, but it's not always intuitive at first glance, so it will be helpful to look at it again.

Duality is all about finding *upper bounds* (for maximization problems) or *lower bounds* (for minimization problems) on the optimal value of a linear program. The key idea is that any valid linear combination of the constraints of a linear program gives a valid bound on the objective function.

So let's say we have a maximization linear program (LP) that looks like the following:

$$\begin{aligned} \text{Maximize:} \quad & \sum_i c_i x_i \\ \text{Subject to:} \quad & \sum_i A_{ji} x_i \leq b_j \quad \forall j \\ & x_i \geq 0 \quad \forall i \end{aligned}$$

Our goal is to take linear combinations of the constraints to get an upper bound on the objective function $\sum_i c_i x_i$. Here's the four-step "recipe" for doing this:

Step 1: Create dual variables. For each constraint in the primal LP, create a dual variable w_j. These variables represent the "weights" we'll use to combine the constraints.

Step 2: Ensure validity. When we take a linear combination of the primal constraints using weights w_j, we get:

$$\sum_j w_j \left(\sum_i A_{ji} x_i \right) \leq \sum_j w_j b_j$$

For this combination to remain valid we need $w_j \geq 0$ (otherwise multiplying by a negative weight would flip the inequality). These conditions on w_j become the *validity constraints* in the dual.

Step 3: Ensure usefulness. We want the best upper bound on the objective possible, so we need the combined constraints to actually bound the objective function:

$$\sum_i c_i x_i \leq \sum_j w_j \left(\sum_i A_{ji} x_i \right) = \sum_i \left(\sum_j w_j A_{ji} \right) x_i$$

Above we swapped the order of the sums and rearranged the terms accordingly. By doing so, we can observe that we just need $c_i \leq \sum_j w_j A_{ji}$ (the coefficient of x_i on the right must dominate the left) to satisfy this for all i.

These conditions become the *usefulness constraints* in the dual.

Step 4: Optimize the bound. We now have:

$$\sum_i c_i x_i \leq \sum_i \left(\sum_j w_j A_{ji} \right) x_i \leq \sum_j w_j b_j$$

To get the tightest possible upper bound, we minimize $\sum_j w_j b_j$ over all valid choices of w_j. This becomes the dual LP! For completeness, here is the final dual LP:

$$\begin{aligned}
\text{Minimize:} \quad & \sum_j w_j b_j \\
\text{Subject to:} \quad & \sum_j w_j A_{ji} \geq c_i \quad \forall i \\
& w_j \geq 0 \quad \forall j
\end{aligned}$$

One important note that we slightly glossed over in lecture and above is that the steps we performed only work for a primal LP in the specific form given above (which is known as *standard form*). For example, if x_i were unconstrained (could be negative), then we can still perform similar steps, but the validity and usefulness constraints would be slightly different.

Example: From Matching LP to Vertex Cover LP

In lecture, we saw how to derive the dual of the matching LP for bipartite graphs, so let's revisit that example in detail using the four-step process above.

First, recall the problems. Consider a bipartite graph $G = (L \cup R, E)$ where L and R are the two parts of the bipartition. A *matching* is a set of edges with no shared endpoints, and a *vertex cover* is a set of vertices touching all edges. The maximum matching problem seeks the largest matching, while the minimum vertex cover problem seeks the smallest vertex cover.

Recall the LP relaxation for maximum matching:

$$\begin{aligned}
\text{Maximize:} \quad & \sum_{e \in E} x_e \\
\text{Subject to:} \quad & \sum_{e \text{ incident to } v} x_e \leq 1 \quad \forall v \in L \cup R \\
& x_e \geq 0 \quad \forall e \in E
\end{aligned}$$

The x_e represents whether edge e is in the matching, and the constraints ensure that each vertex is matched to at most one edge.

Note that the above is in the primal maximization form we discussed earlier, so we can apply the four-step process directly. Let's derive the dual step by step using the key principles:

Step 1: Create dual variables. For each constraint in the primal (one per vertex $v \in L \cup R$), we create a dual variable w_v. These will be the variables in our dual LP.

Step 2: Set up validity constraints. We take a linear combination of the primal constraints using weights w_v:

$$\sum_{v \in L \cup R} w_v \left(\sum_{e \text{ incident to } v} x_e \right) \leq \sum_{v \in L \cup R} w_v \cdot 1$$

We need $w_v \geq 0$ for all vertices v to ensure this combination remains valid. This gives us the *validity constraints*: $w_v \geq 0$ for all $v \in L \cup R$.

Step 3: Set up usefulness constraints. We can rewrite the left side by swapping the order of summation:

$$\sum_{v \in L \cup R} w_v \left(\sum_{e \text{ incident to } v} x_e \right) = \sum_{e \in E} \left(\sum_{v \text{ incident to } e} w_v \right) x_e$$

It's a bit subtle, but note how "e incident to v" becomes "v incident to e" when we swap the sums, and we go from summing over vertices to summing over edges. For this to be useful as an upper bound on the primal objective $\sum_{e \in E} x_e$, we need:

$$\sum_{e \in E} x_e \leq \sum_{e \in E} \left(\sum_{v \text{ incident to } e} w_v \right) x_e$$

Since each primal variable satisfies $x_e \geq 0$, this holds when:

$$1 \leq \sum_{v \text{ incident to } e} w_v \quad \forall e \in E$$

For an edge $e = (u, v)$ with endpoints u and v, this becomes:

$$w_u + w_v \geq 1 \quad \forall e = (u, v) \in E$$

These are the *usefulness constraints*.

Step 4: Minimize to get the tightest bound. Putting it all together, we have:

$$\sum_{e \in E} x_e \leq \sum_{e \in E} \left(\sum_{v \text{ incident to } e} w_v \right) x_e \leq \sum_{v \in L \cup R} w_v$$

To get the tightest upper bound on the primal objective, we minimize the right-hand side over all valid choices of w_v.

Thus, the resulting LP is:

$$\begin{aligned} \text{Minimize:} \quad & \sum_{v \in L \cup R} w_v \\ \text{Subject to:} \quad & w_u + w_v \geq 1 \quad \forall e = (u, v) \in E \\ & w_v \geq 0 \quad \forall v \in L \cup R \end{aligned}$$

Notice that this is exactly the LP relaxation for the minimum vertex cover problem! The dual variable w_v indicates whether vertex v is in the cover, and the constraints ensure every edge is covered.

In lecture, we then used strong duality to show that the optimal value of maximum matching is upper bounded by the optimal value of minimum vertex cover, but we'll stop here for today since our focus was on deriving the dual itself (you can revisit that part in the lecture notes if needed).

Practice

Problem 1

In the Learn section, we saw how to take the dual of an LP in the *standard maximization form* (maximize subject to $\leq$ constraints with $x_i \geq 0$). However, not all LPs come in this form. In this problem, you'll practice converting LPs to standard form.

For each of the following LPs, rewrite it in an equivalent form that matches the standard maximization form from the Learn section. You may need to introduce new variables or modify constraints.

(a) Minimization objective:

$$\begin{aligned} \text{Minimize:} \quad & 2x_1 + 3x_2 \\ \text{Subject to:} \quad & x_1 + x_2 \leq 5 \\ & x_1, x_2 \geq 0 \end{aligned}$$

(b) Greater-than-or-equal constraints:

$$\begin{aligned} \text{Maximize:} \quad & x_1 + 2x_2 \\ \text{Subject to:} \quad & 2x_1 + x_2 \geq 4 \\ & x_1 + 3x_2 \geq 6 \\ & x_1, x_2 \geq 0 \end{aligned}$$

(c) Equality constraints:

$$\begin{aligned} \text{Maximize:} \quad & 3x_1 + x_2 \\ \text{Subject to:} \quad & x_1 + x_2 = 5 \\ & 2x_1 + x_2 \leq 8 \\ & x_1, x_2 \geq 0 \end{aligned}$$

(d) Unconstrained variables:

$$\begin{aligned} \text{Maximize:} \quad & 2x_1 + x_2 \\ \text{Subject to:} \quad & x_1 + x_2 \leq 4 \\ & x_1 \geq 0, \quad x_2 \text{ unconstrained} \end{aligned}$$

Problem 2

In lecture, we saw that taking the dual of the dual gives you back the primal. In this problem, you'll prove this fact by adapting the four-step framework to work with minimization problems.

Consider the following primal LP in standard form:

$$\begin{aligned} \text{Maximize:} \quad & \sum_i c_i x_i \\ \text{Subject to:} \quad & \sum_i A_{ji} x_i \le b_j \quad \forall j \\ & x_i \ge 0 \quad \forall i \end{aligned} \tag{P}$$

Recall from the Learn section that the dual of this LP is:

$$\begin{aligned} \text{Minimize:} \quad & \sum_j w_j b_j \\ \text{Subject to:} \quad & \sum_j w_j A_{ji} \ge c_i \quad \forall i \\ & w_j \ge 0 \quad \forall j \end{aligned} \tag{D}$$

(a) Show that the dual of the dual (D) is the primal (P).

Hint: It might be tempting to first convert (D) to standard maximization form and then take the dual like in the Learn section, but this ends up being more complicated than necessary. Instead, directly apply the four-step framework from the Learn section, but adapt the direction of inequalities and the objective for minimization (it should be very similar to the Learn section, just with flipped inequalities).