CS 3510 Algorithms 4/16/2024

Lecture 19: LP Duality

Lecturer: Abrahim Ladha *Scribe(s): Richard Zhang*

Last time, we discussed what linear programming (LP) was and how many diverse problems it can represent. Today, we do a far more in depth example and detail the Simplex method, as well as the duality theorems.

1 LP and Simplex In-Depth

Recall LP standard form. Given a $m \times n$ matrix A , a $m \times 1$ vector b, and an $n \times 1$ vector c, we find a solution x of size $n \times 1$ such that

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & c^T x \\ \text{subject to} \quad & Ax \leq b, \\ & x \geq 0 \end{aligned}$$

Consider the following linear program:

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & 2x_1 + 3x_2 + 5x_3 \\ \text{subject to} \quad & x_1 + 2x_2 + x_3 \leq 4, \\ & x_1 \geq 0, \\ & x_2 \geq 0, \\ & x_3 \geq 0 \end{aligned}$$

Each constraint adds a geometric figure to our n space (where n is the number of variables). The set of feasible solution to an LP forms a polyhedron in n-space. The constraint $x_1, x_2, x_3 \geq 0$ force us into the positive octant. This forms a triangle based pyramid (pictured below)!

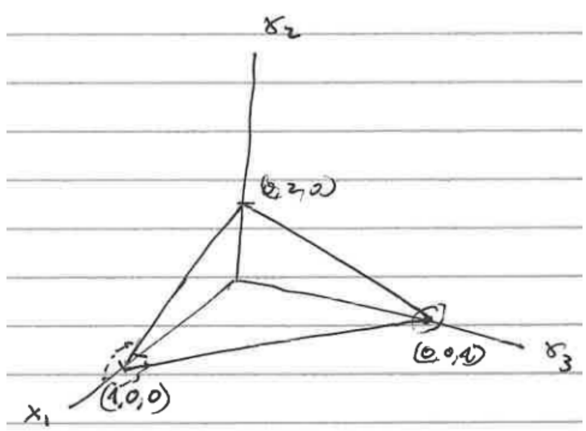

In order to maximize $c^T x$, we can treat it as a plane where it equals some value z that we control ($c^T x = z$). We take its plane and intersect it with this polyhedron. On some fixed plane with a fixed value z, the points will have the same objective function value, so we seek a maximal place to shift the plane given its slope. This shift is performed by changing the value of z and finding the maximum z that intersects with our polyhedron.

It is a guarantee that such a maximum exists at an extreme point (or equal to all points on a line). The simplex algorithm simply iterates over the polyhedron's external points, increasing the objective function each time until it cannot be increased any more. Here we see the max is at $(0, 0, 4)$ with a value of 20.

It is a feat we are searching for a real number solution but only need to check finitely many extreme points.

Let's now do a more complex example. Suppose the following LP in standard form was given:

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & 3x_1 + x_2 + 2x_3 \\ \text{subject to} \quad & x_1 + x_2 + 3x_3 \leq 30, \\ & 2x_1 + 2x_2 + 5x_3 \leq 24, \\ & 4x_1 + x_2 + 2x_3 \leq 36, \\ & x_1, x_2, x_3 \geq 0 \end{aligned}$$

First, we put this LP into "slack form" to convert inequalities to equalities.

$$\begin{aligned} z &= 3x_1 + x_2 + 2x_3 \\ x_4 &= 30 - x_1 - x_2 - 3x_3 \\ x_5 &= 24 - 2x_1 - 2x_2 - 5x_3 \\ x_6 &= 36 - 4x_1 - x_2 - 2x_3 \end{aligned}$$

Although not shown, all the variables are constrained to be nonnegative ($x_1, x_2, x_3, x_4, x_5, x_6 \geq 0$). If we have the constraints in the form $Ax \leq b$, then we can express such constaints to be $Ax + s = b$ with $s \geq 0$. The slack form would therefore be in the form $s = b - Ax$ and is shown above. In the example, $s = \begin{bmatrix} x_4 & x_5 & x_6 \end{bmatrix}^T$. z would be the value of the objective function which we are trying to maximize. It will become clear later why the objective function is expressed this way. Denote the variables on the left hand side as **basic** variables and the variables on the right hand side as **nonbasic** variables.

We need to start somewhere, so we start off with a basic feasible solution that sets all the nonbasic variables to be 0. To satisfying the constraints, the basic variables would be set accordingly, giving us an initial basic solution of $(x_1, x_2, x_3, x_4, x_5, x_6) = (0, 0, 0, 30, 24, 36)$. The current objective function value is 0.

As we run simplex, we will "pivot" to increase the current objective function value. This will rewrite our LP by swapping one basic variable for one nonbasic variable to maximize the objective function. Think of the basic variables as those which are loose. An equality constraint is **tight** if its nonbasic variables get its basic variable to be equal to 0.

We choose the variable in the right side of z whose coefficient is largest and is nonbasic. This is the **entering variable**. Then, choose the "tightest" constraint which we will elaborate more later. The basic variable associated with that constraint will be the **leaving variable**. In this example, x_1 is the entering variable since its coefficient in the objective function is the highest and is a nonbasic variable. As we try to increase the value x_1 and increase the objective function value, we find that x_6 will be the first basic variable that will become zero, specifically when $x_1 = 9$. This makes the corresponding constraint for x_6 the tightest constraint. Another way of seeing it is that each of the constraints provide an upper bound on the value of x_1. The upper bounds would be $30, 12, 9$ and the constraint with the lowest upper bound (constraint with basic variable x_6) is chosen. We then solve for x_1 on this constraint and obtain

$$x_1 = 9 - \frac{x_2}{4} - \frac{x_3}{2} - \frac{x_6}{4}$$

Finally, we rewrite the LP by substituting this (basically doing Gaussian elimination) for any occurrence of x_1 on the right side of the LP. What we get is this:

$$\begin{aligned} z &= 27 + \frac{x_2}{4} + \frac{x_3}{2} - \frac{3x_6}{4} \\ x_1 &= 9 - \frac{x_2}{4} - \frac{x_3}{2} - \frac{x_6}{4} \\ x_4 &= 21 - \frac{3x_2}{4} - \frac{5x_3}{2} + \frac{x_6}{4} \\ x_5 &= 6 - \frac{3x_2}{2} - 4x_3 + \frac{x_6}{2} \end{aligned}$$

Notice that x_1 is now a basic variable and x_6 is now a nonbasic variable. x_1 has "entered" the group of basic variables, while x_6 left that group. We then create a new solution that again make the current nonbasic variables 0. The current solution is now $(x_1, x_2, x_3, x_4, x_5, x_6) = (9, 0, 0, 21, 6, 0)$ and has an objective function value of 27.

Now let's do another "pivot" operation. What is the next variable we choose as the entering variable? It would not be x_6 since it has a negative coefficient in the objective function. x_3 has the largest coefficient in the objective function and is nonbasic, so it will become the entering variable. The constraint with basic variable x_5 limits the increase of x_3 the most, so x_5 is the leaving variable and we substitute the following:

$$x_3 = \frac{3}{2} - \frac{3x_2}{8} - \frac{x_5}{4} + \frac{x_6}{8}$$

The new LP would now be the following after the substitution:

$$\begin{aligned}
z &= \frac{111}{4} + \frac{x_2}{16} - \frac{x_5}{8} - \frac{11x_6}{16} \\
x_1 &= \frac{33}{4} - \frac{x_2}{16} + \frac{x_5}{8} - \frac{5x_6}{16} \\
x_3 &= \frac{3}{2} - \frac{3x_2}{8} - \frac{x_5}{4} + \frac{x_6}{8} \\
x_4 &= \frac{69}{4} + \frac{3x_2}{16} + \frac{5x_5}{8} - \frac{x_6}{16}
\end{aligned}$$

x_3 is now a basic variable and x_5 is now a nonbasic variable. Again, after setting the current nonbasic variables to be all zero, we have a current solution of $(x_1, x_2, x_3, x_4, x_5, x_6) = (\frac{33}{4}, 0, \frac{3}{2}, \frac{69}{4}, 0, 0)$ and a current objective function value of $z = \frac{111}{4}$.

We can only increase x_2 to increase the objective function value. The upper bounds would be $132, 4, \infty$. The constraint with x_4 as the basic variable has an upper bound of ∞ for x_2 since x_4 increases as x_2 increases. The constraint for the basic variable x_3 constraints x_2 the most, so x_2 will enter and x_3 will leave. Performing the same substitution as the previous "pivot" operations, we get the following LP:

$$\begin{aligned}
z &= 28 - \frac{x_3}{6} - \frac{x_5}{6} - \frac{2x_6}{3} \\
x_1 &= 8 + \frac{x_3}{6} + \frac{x_5}{6} - \frac{x_6}{3} \\
x_2 &= 4 - \frac{8x_3}{3} - \frac{2x_5}{3} + \frac{x_6}{3} \\
x_4 &= 18 - \frac{x_3}{2} + \frac{x_5}{2}
\end{aligned}$$

Since all the coefficients in z are now all negative, we cannot continue pivoting and increasing the objective function value. We have reached the optimal solution and maximum objective function value, which is $(x_1, x_2, x_3, x_4, x_5, x_6) = (8, 4, 0, 18, 0, 0)$ and $z = 28$. Observe the values of the original slack variables x_4, x_5, x_6. These values determine how much **slack**, or difference, between the light and right hand sides of their corresponding constraints.

We have now reached the end of the simplex algorithm for this example. Simplex at a high level converts an LP into slack form and keeps pivoting to increase the objective function value until it cannot increase it anymore. The last solution found becomes the final solution returned.

2 Duality

How do we know simplex returns the optimum? Recall how a max-flow min-cut solution was the optimum. If a flow for a network was equal to the min-cut of that same network, that flow was the maximum flow. Finding the max flow was a maximization problem and finding the min cut was a minimization problem, but they both have the same optimal objective function value. Similarly, take any LP in standard form and call it the Primal. It has an "evil" LP called a Dual which minimizes but has the same optimal objective function value.

Primal

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & c^T x \\ \text{subject to} \quad & Ax \leq b, \\ & x \geq 0 \end{aligned}$$

Dual

$$\begin{aligned} \underset{y}{\text{minimize}} \quad & b^T y \\ \text{subject to} \quad & A^T y \geq c, \\ & y \geq 0 \end{aligned}$$

Note that the Dual of the Dual is the Primal. They are like brothers. If the Primal has m constraints, the Dual will have m variables. If the Dual has n constraints, the Primal will have n variables. We will prove that their optimal objective values are the same with the Weak Duality Theorems.

Theorem (Weak Duality Theorem)**.** *Given a primal LP (A, b, c) and its dual, let x be a feasible solution for the primal LP and let y be a feasible solution the dual LP. Then, we have*

$$c^T x \leq b^T y$$

The Weak Duality Theorem can show that the max-flow $\leq$ min-cut in a flow network.

Proof: Let n be the number of variables in the Primal and m be the number of variables in the Dual.

$$c^T x = \sum_{j=1}^{n} c_j x_j \leq \sum_{j=1}^{n} \left[\sum_{i=1}^{m} (a_{ij} y_i) x_j \right] = \sum_{j=1}^{m} \left[\sum_{i=1}^{n} (a_{ij} x_j) y_j \right] \leq \sum_{i=1}^{m} b_i y_j = b^T y$$

Two important results come from the Weak Duality Theorem.

1. If $c^T x = b^T y$, then x, y are the optimal solutions for their corresponding LP's. If a solution x for the Primal has an objective function value equal to $b^T y$, then this must be the highest. Otherwise, a solution with a higher objective function value would be greater than $b^T y$ and thus contradict the Weak Duality Theorem. A similar explanation can be made for a solution y for the Dual.

2. If the Primal has a solution that is unbounded, then the Dual has no feasible solutions. If the Dual has a solution that is unbounded, then the Primal has no feasible solutions. Say that the Primal was unbounded, but the Dual had a feasible solution. This feasible solution would have a finite objective function value. However, the Weak Duality Theorem states that this Dual feasible solution should have a greater objective function value than any solution of the Primal, which is impossible since the possible Primal objective function values are unbounded! Therefore, no feasible solution exists for the Dual. A similar explanation can be made for the second case.

Why is simplex optimal? It actually solves both the primal and the dual simultaneously, just like how the max-flow algorithm also finds the min cut.

Recall our solution to the example problem was $(x_1, x_2, x_3) = (8, 4, 0)$ with $z = 28$. The final z equation representing the objective function was in the form $v' + \sum c'_j x_j$. With n as the number of original variables in the Primal, each entry in y for the Dual would then be the following:

$$y_i = \begin{cases} -c'_{n+i} & \text{if } x_{n+i} \text{ is nonbasic} \\ 0 & \text{otherwise} \end{cases}$$

Applying this to the final z equation in the example ($z = 28 - \frac{x_3}{6} - \frac{x_5}{6} - \frac{2x_6}{3}$), we get that $y_1 = 0, y_2 = \frac{1}{6}, y_3 = \frac{2}{3}$. $y_1 = 0$ since $x_{1+3} = x_4$ and x_4 become a basic variable at the end. Plugging these values into $b^T y$ gives us 28 which is the same as the maximum objective function value for the Primal.

3 A quick way for certifying optimality

We can certify that a solution from simplex is the best by applying Gaussian elimination and the rules of algebra. Let's look at the following LP:

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & x_1 + 6x_2 \\ \text{subject to} \quad & x_1 \leq 100 \ (a), \\ & x_2 \leq 300 \ (b), \\ & x_1 + x_2 \leq 400 \ (c) \end{aligned}$$

Simplex returns $x_1 = 100, x_2 = 300$ with the objective function value as 1900. We can use Gaussian elimination by multiplying inequality (b) by 5 and adding it to inequality (c) $(5b + c = x_1 + 6x_2 \leq 1500 + 400 = 1900$. With these constraints, we have shown that the objective function value cannot exceed 1900. Since the simplex solution has a objective function value of 1900, we know that this solution is the optimum.

4 Quick Remark about the Runtime of Simplex

The runtime for the average case is polynomial. However, with some inputs, simplex may run for $\binom{m+n}{n}$ iterations. Since each iteration runs in $O(mn)$ time, the worst case runtime would be exponential.