The Primal-Dual Algorithm

Let A be an $m \times n$ matrix, and consider the following primal problem P in standard form and its dual D.

$$
\begin{array}{llll}
\text{P:} & \min & \mathbf{c}'\mathbf{x} & \\
 & \text{subject to} & A\mathbf{x} & = \mathbf{b} \\
 & & \mathbf{x} & \geq \mathbf{0}
\end{array}
\qquad
\begin{array}{llll}
\text{D:} & \max & \mathbf{p}'\mathbf{b} & \\
 & \text{subject to} & \mathbf{p}'A & \leq \mathbf{c}' \\
 & & \mathbf{p} & \text{free}
\end{array}
$$

The primal-dual algorithm is initialized with a dual feasible solution $\mathbf{p}$. Define the index set J of admissible columns to be $J = \{j : \mathbf{p}'A_j = c_j\}$, where A_j denotes the j^{th} column of A. The set J is the set of dual constraints which are tight at $\mathbf{p}$.

Suppose that $\mathbf{x}$ is a primal feasible solution. Then complementary slackness states that $\mathbf{x}$ and $\mathbf{p}$ are optimal for their respective problems if and only if

$$
\begin{aligned}
p_i(\mathbf{a}_i\mathbf{x} - b_i) &= 0, \text{ for } i = 1, \ldots, m, \text{ and} \\
x_j(\mathbf{p}'A_j - c_j) &= 0, \text{ for } j = 1, \ldots, n.
\end{aligned}
$$

(Here, $\mathbf{a}_i$ denotes the i^{th} row of A.) The first set of equations is always satisfied since $\mathbf{x}$ is feasible and hence $A\mathbf{x} = \mathbf{b}$. The second set of equations is equivalent to $x_j = 0$ for $j \notin J$. Thus, we seek a primal feasible $\mathbf{x}$ where $x_j = 0$ for $j \notin J$. If such an $\mathbf{x}$ exists, then by complementary slackness, $\mathbf{x}$ and $\mathbf{p}$ are optimal for their respective problems.

We determine if such an $\mathbf{x}$ exists by testing whether the linear program constraints

$$
\begin{array}{lrl}
\text{FP:} \quad \text{subject to} & A\mathbf{x} & = \mathbf{b} \\
 & \mathbf{x} & \geq \mathbf{0} \\
 & x_j & = 0 \text{ if } j \notin J
\end{array}
$$

are feasible using Phase I of the two phase simplex algorithm. In Phase I, we solve an auxiliary linear program formed by adding m artificial variables y_i and changing the objective function. We call the auxiliary linear program the *restricted primal* RP.

$$
\begin{array}{llcccl}
\text{RP:} & \min & & \mathbf{1y} & & \\
 & \text{subject to} & A\mathbf{x} & +I\mathbf{y} & = & \mathbf{b} \\
 & & \mathbf{x} & & \geq & \mathbf{0} \\
 & & & \mathbf{y} & \geq & \mathbf{0} \\
 & & x_j & & = & 0, \text{ if } j \notin J.
\end{array}
$$

(Here $\mathbf{1}$ denotes the vector of all 1s.)

The auxiliary linear program of Phase I always has a finite optimum. Let ξ be the optimal cost of RP. If $\xi = 0$, then there exists a feasible solution $\mathbf{x}$ to the primal P such that $\mathbf{x}$ and $\mathbf{p}$ satisfy the complementary slackness conditions. Hence, $\mathbf{x}$ and $\mathbf{p}$ are optimal. Note that $\mathbf{x}$ is obtained by solving RP (and driving artificial variables out of the basis if necessary).

If $\xi > 0$, then we use a solution to the dual of the restricted primal to improve our solution $\mathbf{p}$ to the dual. The dual of the restricted primal is

$$
\begin{array}{llll}
\text{DRP:} & \max & \mathbf{p}'\mathbf{b} & \\
 & \text{subject to} & \mathbf{p}'A_j & \leq 0, \text{ for } j \in J, \\
 & & \mathbf{p} & \leq \mathbf{1} \\
 & & \mathbf{p} & \text{free.}
\end{array}
$$

Let $\bar{\mathbf{p}}$ be an optimal solution to DRP with cost ξ. We update our dual feasible vector by setting $\hat{\mathbf{p}} = \mathbf{p} + \theta\bar{\mathbf{p}}$ for some $\theta > 0$. Note that the cost of $\hat{\mathbf{p}}$ is

$$\hat{\mathbf{p}}'\mathbf{b} = \mathbf{p}'\mathbf{b} + \theta\bar{\mathbf{p}}'\mathbf{b} > \mathbf{p}'\mathbf{b}$$

since $\bar{\mathbf{p}}'\mathbf{b} = \xi > 0$, and hence we have improved the cost of our dual feasible solution.

We wish to make θ as large as possible to increase the cost of $\hat{\mathbf{p}}$ while still maintaining dual feasibility. If $j \in J$, or $j \notin J$ but $\bar{\mathbf{p}}'A_j \leq 0$, then

$$\hat{\mathbf{p}}'A_j = \mathbf{p}'A_j + \theta\bar{\mathbf{p}}'A_j \leq c_j.$$

If $j \notin J$ and $\bar{\mathbf{p}}'A_j > 0$, then we require that

$$\hat{\mathbf{p}}'A_j = \mathbf{p}'A_j + \theta\bar{\mathbf{p}}'A_j \leq c_j,$$

$$\text{and so} \qquad \theta \leq \frac{c_j - \mathbf{p}'A_j}{\bar{\mathbf{p}}'A_j}.$$

Since these are the only restrictions on θ, we choose the most restrictive bound, setting

$$\theta^* = \min_{\substack{j\notin J \\ \bar{\mathbf{p}}'A_j>0}} \left[\frac{c_j - \mathbf{p}'A_j}{\bar{\mathbf{p}}'A_j}\right] \qquad \text{and} \quad \hat{\mathbf{p}} = \mathbf{p} + \theta^*\bar{\mathbf{p}}.$$

Note that $c_j - \mathbf{p}'A_j > 0$ for $j \notin J$, so the value of θ^* is strictly positive. If $\bar{\mathbf{p}}'A_j \leq 0$ for all $j \notin J$, then θ can be chosen arbitrarily large. Hence, the dual has optimal cost $+\infty$ and the primal is infeasible.

We repeat this process until obtaining an optimal primal solution $\mathbf{x}$ or proving that the primal is infeasible. We summarize an iteration of the primal-dual algorithm:

1. Given a dual feasible solution $\mathbf{p}$, find the set $J = \{j : \mathbf{p}'A_j = c_j\}$ of admissible columns. J is found by considering which dual constraints are tight at $\mathbf{p}$.

2. Solve the restricted primal problem RP using the simplex method. If the optimal cost ξ of RP is 0, the optimal solution $\mathbf{x}$ to RP is also an optimal solution to the primal problem, and the algorithm terminates.

3. If the optimal cost ξ of RP is positive, find an optimal solution $\bar{\mathbf{p}}$ to the dual of the restricted problem DRP. If $\bar{\mathbf{p}}'A_j \leq 0$ for all $j \notin J$, then the dual has optimal cost $+\infty$, the primal is infeasible, and the algorithm terminates. Otherwise, set $\hat{\mathbf{p}} = \mathbf{p} + \theta^*\bar{\mathbf{p}}$, where

$$\theta^* = \min_{\substack{j\notin J \\ \bar{\mathbf{p}}'A_j>0}} \left[\frac{c_j - \mathbf{p}'A_j}{\bar{\mathbf{p}}'A_j}\right].$$

   Repeat Step 1 with the new dual feasible solution $\hat{\mathbf{p}}$.

To discuss finiteness and the advantages of the primal-dual algorithm, we have the following

Lemma 1. *Every admissible column in the optimal basis of the restricted primal RP remains admissible at the start of the next iteration.*

Proof. If A_j is in the optimal basis of RP, then the reduced cost of x_j in the optimal tableau is $\bar{c}_j = -\bar{\mathbf{p}}'A_j = 0$. (Note that the cost c_j in RP of x_j is 0). Hence,

$$\hat{\mathbf{p}}'A_j = \mathbf{p}'A_j + \theta^*\bar{\mathbf{p}}'A_j = \mathbf{p}'A_j = c_j, \text{ since } \bar{\mathbf{p}}'A_j = 0,$$

and so A_j is an admissible column at the start of the next iteration. □

Thus, the primal-dual algorithm can be thought of as pivoting in a primal tableau for Phase I with restrictions on which columns may enter the basis. The final tableau simultaneously satisfies primal and dual feasibility and hence is optimal. This perspective shows that the primal-dual algorithm is finite.

Theorem 2. *The primal-dual algorithm terminates in a finite number of iterations.*

Proof. We view the primal-dual algorithm as providing a sequence of bases for the Phase I auxiliary linear program AP for the original primal problem. Each basis corresponds to a basic feasible solution with strictly decreasing costs ξ. No basic feasible solution of AP is degenerate, then no basis will be repeated, and the primal-dual algorithm terminates using only a finite number of pivots solving restricted primal problems. In the presence of degeneracy, the use of anticycling rules guarantees that the primal-dual algorithm terminates in a finite number of iterations. □

We conclude by discussing three advantages of the primal-dual algorithm.

1. The restricted primal problems have the same number of rows as the primal problem, but they have significantly fewer variables.

2. In forming the restricted primal problems, we "combinatorialize the cost"; that is, we replace the original cost function $\mathbf{c}$ with the simpler cost function prescribed by Phase I of the simplex method, assigning a cost of 1 to each artificial variable y_i and a cost of 0 to every other variable x_j.

3. Lemma 1 states that every column in an optimal basis of the restricted primal remains admissible at the start of the next iteration. Hence, that optimal basis can be used as the initial basic feasible solution when solving the subsequent restricted primal problem. The tableau needs to be adjusted by adding columns for newly admissible columns and removing the columns that are no longer admissible. Suppose that A_j is newly admissible. Then the reduced cost of x_j is $\bar{c}_j = 0 - \mathbf{p}'A_j$. Note that if A_j is newly admissible, then from the previous iteration we have $\mathbf{p}'A_j > 0$, and so $\bar{c}_j < 0$. The rest of column j in the tableau can be computed using $B^{-1}A_j$.

References

[1] C. H. Papadimitriou and K. Steiglitz. Combinatorial optimization: algorithms and complexity. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1982.