The Primal-Dual Algorithm

P&S Chapter 5

Last Revised – October 30, 2006

Simplex solves LP by starting at a Basic Feasible Solution (BFS) and moving from BFS to BFS, always improving the objective function, until no more improvement is possible.

Recall that x, π are jointly optimal solutions to primal and dual iff they jointly satisfy the complementary slackness conditions (CSC).

The Primal Dual Algorithm start with a feasible π and then iterates the following operations

1. Search for an x that jointly satisfies CSC with π.

2. If such an x exists, optimality has been achieved. Stop.

3. Otherwise *improve* π and go to 1.

Primal			Dual		
$\min z$	$=$	$c'x$	$\max w$	$=$	$\pi' b$
Ax	$=$	$b \geq 0$	$\pi' A$	$\leq$	c'
x	$\geq$	0	π'	$\gtrless$	0

We may always assume that $b \geq 0$ since, if not, we can multiply appropriate equalities by -1.

We always assume that we know feasible π of dual. If $c \geq 0$ we may take $\pi = 0$.

If not, we can use a trick due to Beale and introduce
(a) a new variable x_{n+1} with cost $c_{n+1} = 0$.
(b) constraint $x_1 + x_2 + \cdots + x_n + x_{n+1} = b_{m+1}$
where b_{m+1} is large
(c) new dual variable π_{m+1}.

New dual is

$$\begin{aligned} &\max w = \pi' b + \pi_{m+1} b_{m+1} \\ &\pi' A_j + \pi_{m+1} \leq c_j \qquad j = 1, \ldots, n \\ &\qquad \pi_{m+1} \leq 0 \end{aligned}$$

which has feasible solution

$$\begin{aligned} \pi_i &= 0 \qquad i = 1, \ldots, m \\ \pi_{m+1} &= \min_{1 \leq j \leq n} \{c_j\} < 0 \end{aligned}$$

Primal			Dual		
$\min z$	$=$	$c'x$	$\max w$	$=$	$\pi' b$
Ax	$=$	$b \geq 0$	$\pi' A$	$\leq$	c'
x	$\geq$	0	π'	$\gtrless$	0

Assume $b \geq 0$ and we know dual feasible π.

Recall that x, π are jointly optimal iff they satisfy

$$\forall i,\ \pi_i(a_i' x - b_i) = 0 \quad \text{and} \quad \forall j,\ (c_j - \pi' A_j) x_j = 0.$$

The Primal-Dual algorithm maintains a feasible π. At each step it solves a Restricted Primal (RP) trying to find a jointly optimal primal solution x. If it doesn't succeed, that gives the Dual of RP (DRP) enough information to "improve" π, while keeping it feasible. This procedure iterates and converges to optimum in a finite number of steps.

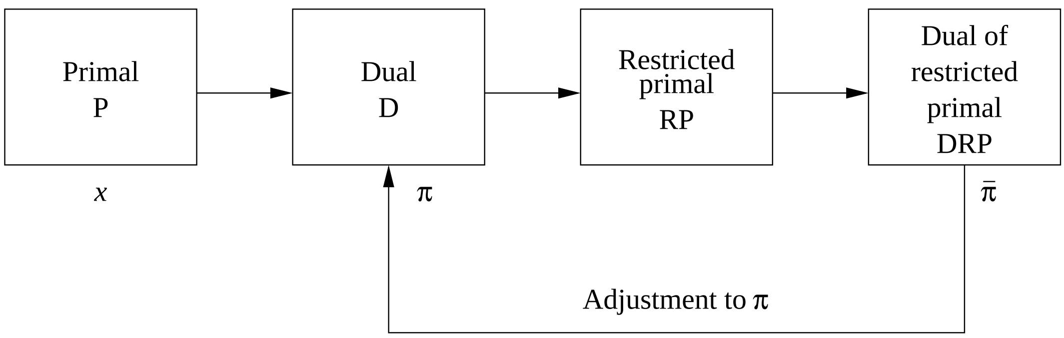

Primal			Dual		
$\min z$	$=$	$c'x$	$\max w$	$=$	$\pi' b$
Ax	$=$	$b \geq 0$	$\pi' A$	$\leq$	c'
x	$\geq$	0	π'	$\gtrless$	0

Complementary slackness conditions are

$\forall i,\ \pi_i(a_i' x - b_i) = 0$ and $\forall j,\ (c_j - \pi' A_j) x_j = 0.$

Define set of admissible columns
$J = \{j : \pi' A_j = c_j\}$.

If $\forall j \notin J,\ x_j = 0$ then x is optimal.
This is equivalent to searching for x that satisfies

$$\begin{aligned} \sum_{j \in J} a_{ij} x_j = b_i \qquad & i = 1, \ldots, m \\ x_j \geq 0 \qquad & j \in J \\ x_j = 0 \qquad & j \notin J \end{aligned}$$

If we can find x that satisfies equalities below then x is optimal:

$$\begin{aligned} \sum_{j \in J} a_{ij} x_j = b_i \quad & i = 1, \ldots, m \\ x_j \geq 0 \quad & j \in J \\ x_j = 0 \quad & j \notin J \end{aligned}$$

We therefore introduce m new variables $x_i^a, i = 1, \ldots, m$, and the Restricted Primal (RP)

$$\begin{aligned} \min \xi &= \sum_{i=1}^{m} x_i^a \\ \sum_{j \in J} a_{ij} x_j + x_i^a &= b_i \quad i = 1, \ldots, m \\ x_j &\geq 0 \quad j \in J \\ (x_j &= 0 \quad j \notin J) \\ x_i^a &\geq 0 \end{aligned}$$

Solve RP , e.g., using simplex. If optimal solution has $\xi = 0$ then have found optimal x for original problem.

If optimal has $\xi > 0$, consider dual DRP of RP.

$$
\begin{array}{rcll}
\multicolumn{4}{l}{\text{RP}} \\
\min \xi & = & \sum_{i=1}^{m} x_i^a & \\
\sum_{j \in J} a_{ij} x_j + x_i^a & = & b_i & i \le m \\
x_j & \ge & 0 & j \in J \\
(x_j & = & 0 & j \notin J) \\
x_i^a & \ge & 0 & i \le m
\end{array}
$$

Assume that RP has $\xi > 0$.
Consider DRP, the dual of RP:

$$
\begin{array}{rcll}
\multicolumn{4}{l}{\text{DRP}} \\
\max w & = & \pi' b & \\
\pi' A_j & \le & 0 & j \in J \\
\pi_i & \le & 1 & i \le m \\
\pi_i & \gtrless & 0 & i \le m
\end{array}
$$

Let $\bar{\pi}$ be optimal solution to DRP derived from optimal solution to RP; $\bar{\pi} = \hat{c}'_B B^{-1}$ where B is basis columns of optimal BFS of RP and $\hat{c}$ is cost function of RP.

RP

$$\begin{aligned}
\min \xi &= \sum_{i=1}^{m} x_i^a \\
\sum_{j \in J} a_{ij} x_j + x_i^a &= b_i \quad i \le m \\
x_j &\ge 0 \quad j \in J \\
(x_j &= 0 \quad j \notin J) \\
x_i^a &\ge 0
\end{aligned}$$

DRP

$$\begin{aligned}
\max w &= \pi' b \\
\pi' A_j &\le 0 \quad j \in J \\
\pi_i &\le 1 \quad i \le m \\
\pi_i &\gtrless 0
\end{aligned}$$

$J = \{j : \pi' A_j = c_j\}$
In original Dual

Started with feasible π. Using RP, tried to find x that jointly satisfied CSC with π.

Optimum $\xi_0 > 0$, so this didn't exist, but we can find $\bar{\pi}$, optimum of DRP. Idea is to try and improve π to π^* by finding "good" θ to set

$$\pi^* = \pi + \theta \bar{\pi}.$$

Cost of π^* is

$$\pi^{*\prime} b = \pi' b + \theta \bar{\pi}' b.$$

Since RP and DRP are a primal-dual pair we have $\bar{\pi}' b = \xi_{opt} > 0$. Therefore, to *improve* π to π^*, we must have $\theta > 0$.

Dual

$$\begin{aligned} \max w &= \pi' b \\ \pi' A &\leq c' \\ \pi' &\gtrless 0 \end{aligned}$$

$$J = \{j : \pi' A_j = c_j\}$$

π is feasible

DRP

$$\begin{aligned} \max w &= \pi' b & \\ \pi' A_j &\leq 0 & j \in J \\ \pi_i &\leq 1 & i \leq m \\ \pi_i &\gtrless 0 & \end{aligned}$$

$\bar{\pi}$ is optimal

We "improve" cost of π by setting

$$\pi^* = \pi + \theta\bar{\pi}, \quad \theta > 0.$$

In order to maintain *feasibilty* of π we need

$$\forall j, \quad \pi^{*\prime} A_j = \pi' A_j + \theta\bar{\pi}' A_j \leq c_j.$$

If $\bar{\pi}' A_j \leq 0$ this is not a problem.

In particular, if $\bar{\pi}' A_j \leq 0$ for all j, then θ can be made arbitrarily large so original dual is **unbounded** and original primal is infeasible.

Dual

$$\begin{aligned} \max w &= \pi' b \\ \pi' A &\leq c' \\ \pi' &\gtreqless 0 \end{aligned}$$

$$J = \{j : \pi' A_j = c_j\}$$

DRP

$$\begin{aligned} \max w &= \pi' b & \\ \pi' A_j &\leq 0 & j \in J \\ \pi_i &\leq 1 & i \leq m \\ \pi_i &\gtreqless 0 & \end{aligned}$$

$\bar{\pi}$ is optimal

We just saw that if, $\forall j,\ \bar{\pi}' A_j \leq 0$, then original primal is infeasible.

We know that $\forall j \in J,\ \bar{\pi}' A_j \leq 0$ since $\bar{\pi}$ is optimal and thus feasible. Then

Theorem If $\xi_{opt} > 0$ in RP and the optimal dual (in DRP) satisfies

$$\bar{\pi}' A_j \leq 0 \quad \text{for } j \notin J$$

then P is infeasible.

Dual

$$\begin{aligned} \max w &= \pi' b \\ \pi' A &\leq c' \\ \pi' &\gtrless 0 \end{aligned}$$

$J = \{j : \pi' A_j = c_j\}$

π is feasible

DRP

$$\begin{aligned} \max w &= \pi' b \\ \pi' A_j &\leq 0 \quad j \in J \\ \pi_i &\leq 1 \quad i \leq m \\ \pi_i &\gtrless 0 \end{aligned}$$

$\bar{\pi}$ is optimal

In order to maintain *feasibilty* of π we need

$$\forall j, \quad \pi^{*\prime} A_j = \pi' A_j + \theta \bar{\pi}' A_j \leq c_j.$$

From previous slide, when maintaining feasibility, we only worry about $\bar{\pi}' A_j > 0$ for some $j \notin J$. i.e.,

$$\pi^{*\prime} A_j = \pi' A_j + \theta \bar{\pi}' A_j \leq c_j \quad j \notin J \text{ and } \bar{\pi}' A_j > 0$$

Theorem: When $\xi_{opt} > 0$ in RP and there is a $j \notin J$ with $\bar{\pi}' A_j > 0$, the largest θ that maintains the feasibility of $\pi^* = \pi + \theta \bar{\pi}$ is

$$\theta_1 = \min_{\substack{j \notin J \\ s.t.\ \bar{\pi}' A_j > 0}} \left[\frac{c_j - \pi' A_j}{\bar{\pi}' A_j} \right]$$

The new cost is

$$w^* = \pi' b + \theta_1 \bar{\pi}' b = w + \theta_1 \bar{\pi}' b > w.$$

procedure primal-dual
begin

infeasible := 'no', opt := 'no';
let π be feasible in D
while infeasible ='no' **and** opt ='no' **do**
begin set $J = \{j : \pi' A_j = c_j\}$;
  solve RP by the simplex algorithm;
  if $\xi_{opt} = 0$ **then** opt:= 'yes'
  else if $\bar{\pi}' A_j \leq 0$ for all $j \notin J$
    then infeasible := 'yes'
    else $\pi := \pi + \theta_1 \bar{\pi}$
    (**comment:** θ_1 from last slide)
end

end

Note: recall that $\bar{\pi} = \hat{c}'_B B^{-1}$ is optimal solution to DRP derived from optimal solution to RP, where B is composed of basis columns of optimal BFS of RP and $\hat{c}$ is cost function of RP.

Quick Review of Relative Cost

Recall that given an LP in standard form and its dual D then

1. Let B be a BFS of the LP $Ax = b$, $x \geq 0$ and c'_B the associated cost vector. For all j, the relative cost of x_j is

$$\bar{c}_j = c_j - z_j = c_j - c'_B B^{-1} A_j$$

2. If B is an optimal BFS then we can choose $\pi' = c'_B B^{-1}$ as an optimal solution to Dual.

3. If x is an optimal BFS and $\pi' = c'_B B^{-1}$ is the associated optimal dual solution, then

$$\bar{c}_j = c_j - z_j = c_j - \pi' A_j.$$

Recall that $J = \{j : \pi' A_j = c_j\}$. We now claim:

Theorem: Every admissible column in the optimal basis of RP remains admissible at the start of the next iteration.

Proof: Suppose column j is in the optimal basis of RP at the end of an iteration. Then its relative cost (in RP) is

$$0 = \bar{c_j} = -\bar{\pi}' A_j.$$

This means that

$$\pi^{*\prime} A_j = \pi' A_j + \theta_1 \bar{\pi}' A_j = \pi' A_j = c_j$$

so j remains in J.

One consequence is that if at some iteration RP has an optimal BFS $\widehat{x}_0$ then, at the start of the next iteration, $\widehat{x}_0$ remains a BFS (although probably no longer optimal) *in the new RP.* We may therefore start simplex in the new RP at the old optimal solution.

Recall the definition of

$$\theta_1 = \min_{\substack{j \notin J \\ s.t.\ \bar{\pi}' A_j > 0}} \left[\frac{c_j - \pi' A_j}{\bar{\pi}' A_j} \right]$$

Let $j = j_0$ be value at which minimum occurs. Then

$$\pi^{*\prime} A_{j_0} = \pi' A_{j_0} + \theta_1 \bar{\pi}' A_{j_0} = c_{j_0}$$

so j_0 *enters* J.

At the end of the previous iteration it is possible that j_0 could not enter the BFS because it was not in J and was therefore not considered.

Now that $j_0 \in J$, it might be able to enter a BFS. Let $\widehat{x}_0$ be current optimal BFS of RP at end of last iteration. It remains a BFS in the new RP.

Since $\bar{\pi}' A_j > 0$ we have (from page 13 (1,2)) and fact that $\widehat{c}_j = 0$) that, for BFS $\widehat{x}_0$, the *relative cost* of x_{j_0} in the new RP is $-\bar{\pi}' A_{j_0} < 0$.

We can therefore *pivot* on x_{j_0} and (if the BFS is not degenerate) we will improve the cost.

$$
\begin{array}{rcl}
 & \text{RP} & \\
\min \xi & = & \sum_{i=1}^{m} x_i^a \\
\sum_{j \in J} a_{ij} x_j + x_i^a & = & b_i \quad i \leq m \\
x_j & \geq & 0 \quad j \in J \\
x_j & = & 0 \quad j \notin J \\
x_i^a & \geq & 0
\end{array}
$$

Consider a primal P' where the words (i) $j \in J$ and (ii) $x_j = 0,\ j \notin J$ were deleted. Any BFS of RP would be a BFS of P', so there are only a finite number of BFSs shared among all of the RPs.

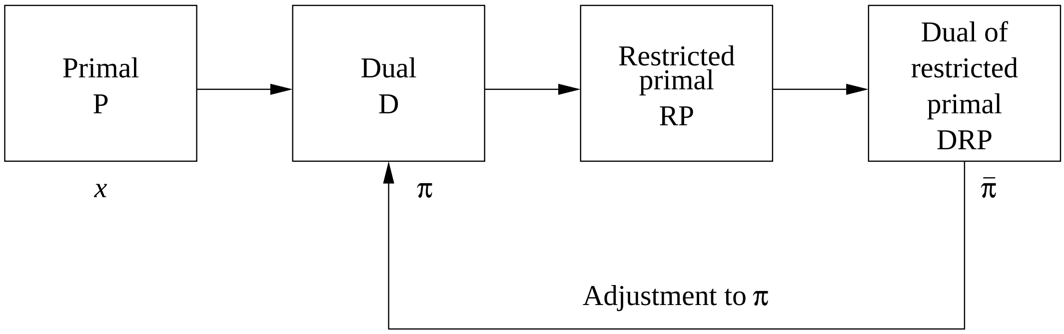

We can consider an iteration of the primal-dual algorithm as starting from some BFS of P' which is a BFS of our current RP and finding a pivot j_0 to start moving to another BFS of P' which is also BFS of our current RP. If the pivot is not degenerate then our cost will decrease. (If the pivot is degenerate we use an anti-cycling rule to guarantee that we will not cycle). Then, since our algorithm moves from BFS to BFS without ever repeating a BFS, it must terminate. When it terminates, it either shows infeasibility of original problem P or reaches optimality of P.

Next: The reason that this is an interesting technique is that we will soon see that solving RP and/or DRP can often be done using other combinatorial algorithms.