Primal-Dual Algorithm

CS 149 Staff

November 18, 2006

Introduction

The primal dual algorithm, while not a good general purpose LP solution technique, is valuable because it is easy to customize for a particular problem. Variants are frequently used in approximation algorithms for NP hard problems that can be formulated as integer programs.

Recall the complementary slackness conditions stating that for optimal primal and dual solutions, a dual variable can only be nonzero if the corresponding constraint is tight in the primal, and a primal variable can only be nonzero if the corresponding constraint is tight in the dual. Note that for a primal LP in standard form all primal constraints are equality constraints so any feasible point has all constraints tight. This implies that primal feasibility is sufficient to satisfy $\pi^T(Ax - b) = 0$, so the harder complementary slackness condition to satisfy is $x^T(c - \pi^T A) = 0$.

The primal dual algorithm proceeds by iteratively improving a dual feasible point π until a primal feasible point x can be found satisfying $x^T(c - \pi^T A) = 0$. Given a current π, the complementary slackness condition is satisfied if and only if for every i such that $x_i \neq 0$, $(c - \pi^T A -)_i == 0$. Denote the set of i satisfying the second equation by $J(\pi)$. We need to look for a feasible primal point x that has nonzero components only within $J(\pi)$.

We are looking for a feasible point, so the two phase simplex algorithm is in order. Introduce artificial variables and attempt to minimize their sum in the same way we did before. To ensure that prohibited components of x stay 0, simply refuse to consider them when choosing a column to enter the basis.

Here's a summary of the primal dual algorithm:

1. Choose a dual feasible π, usually $\pi = 0$.

2. Search for a primal feasible point x satisfying the complementary slackness condition. If the restricted primal has optimal value of zero, terminate with the now optimal x and π. Otherwise take note of $\overline{\pi}$, the optimum for the dual of the restricted primal.

3. Use $\overline{\pi}$ to adjust π by setting $\pi \leftarrow \pi + \lambda\overline{\pi}$, taking λ as large as possible without violating dual feasibility ($\pi^T A - c \leq 0$). If λ can be set arbitrarily large, then the dual problem is unbounded and hence the primal problem is infeasible.

The following theorem shows that we make progress in the sense that the dual objective function keeps increasing. This is insufficient to prove termination by itself, because there exist infinite, bounded and increasing sequences such as $1-1/n$. Termination proves need to be done separately for each problem.

The primal dual algorithm suggests reasonable algorithms for many problems, including min cost perfect matching and shortest path.

Theorem

If a restricted problem has an optimal objective greater than zero, then that solution cannot be optimal. Let $\overline{\pi}$ denote the optimal solution of (DRP). There exists a $\lambda > 0$ such that $\pi^* := \pi + \lambda\overline{\pi}$ is feasible for (D) with $b^T\pi^* > b^T\pi$.

Proof

Let's examine the dual of the restricted problem:

$$\begin{aligned} \max \quad & b^T\overline{\pi} \\ \text{s.t.} \quad & \overline{\pi} \leq 1 \\ & A_J^T\overline{\pi} \leq 0 \end{aligned} \tag{1}$$

Say $\overline{\pi}$ is optimal for the DRP. DRP must have the same objective value as the RP, so $b^T\overline{\pi} = g_0 > 0$ when we are not optimal yet.

We wish to find a $\lambda > 0$ and consider the solution $\pi^* = \pi + \lambda\overline{\pi}$, improving our objective:

$$b^T\pi^* = b^T\pi + \lambda b^T\overline{\pi} = b^T\pi + \lambda g_0 > b^T\pi$$

We want to choose $\lambda > 0$ maximal such that π^* is dual feasible (for the original problem (D)):

$$\begin{aligned} A^T\pi^* &\leq c \\ \Leftrightarrow A^T + \lambda A^T\overline{\pi} &\leq c, \end{aligned}$$

so it suffices to show that

$$A_J^T\pi + \lambda A_J^T\overline{\pi} \leq c_J \text{ and } A_{J^c}^T\pi + \lambda A_{J^c}^T\overline{\pi} \leq c_{J^c}$$

Now, we know that $A_{J^c}^T\pi < c_{J^c}$ by the definition of J.

So, we can choose

$$\lambda_0 = \min_{\substack{j \notin J \\ \text{such that} \\ \overline{\pi}^T A_j > 0}} \left\{ \frac{c_j - \pi^T A_j}{\overline{\pi}^T A_j} \right\}$$

which will imply that $\lambda_0 > 0$.

Notes

Recall that the reduced costs for the restricted primal are defined by $\bar{c} = c^T - \bar{\pi}^T A$. Let A_a denote the portion of the A matrix that corresponds to the artificial variables (and similar definition of c_a). Note that $A_a = I$ an identity matrix, and $c_a = 1$, a vector of all ones. Therefore $\bar{\pi}^T = 1^T - \bar{c}^T$, so the dual variables for the restricted primal can be read out of the final tableau.