MA3231 Linear Programming
W. J. Martin
September 21, 2018

Post-Optimality Analysis

Overview

We consider the situation where, after an optimal dictionary has been reached, a single piece of data in the original problem changes. The questions to ask are

1. What effect does this have on the optimal primal and/or dual solution?
2. Within what range of modification does the optimal basis remain optimal?
3. When changes to the data push us outside of this range, what steps should be taken in order to return to optimality?

We assume that we started with an LP in equality form

$$\begin{array}{rrcl} \mathbf{max} & c^\top x & & \\ & Ax & = & b \\ & x & \geq & 0 \end{array}$$

At optimality, the variables are partitioned into basic ($\mathcal{B}$) and non-basic ($\mathcal{N}$). So the matrix A, the vector x of variables, and cost vector c are partitioned accordingly

$$A = [\ B \ \vdots \ A_{\mathcal{N}}\ ], \qquad x^\top = [\ x_{\mathcal{B}}^\top \ \vdots \ x_{\mathcal{N}}^\top\ ], \qquad c^\top = [\ c_{\mathcal{B}}^\top \ \vdots \ c_{\mathcal{N}}^\top\ ]$$

and the same problem is now written

$$\begin{array}{rrcl} \mathbf{max} & c_{\mathcal{B}}^\top x_{\mathcal{B}} + c_{\mathcal{N}}^\top x_{\mathcal{N}} & & \\ & Bx_{\mathcal{B}} + A_{\mathcal{N}} x_{\mathcal{N}} & = & b \\ & x_{\mathcal{B}}, x_{\mathcal{N}} & \geq & 0 \end{array}$$

When we discuss the optimal dictionary, it is convenient to have both notation for its individual entries and expressions for these entries in terms of the original data. If we wish to refer to individual entries, we use the "bar" notation:

$$\begin{array}{rclcl} \zeta & = & \bar{\zeta} & + & \sum_{j \in \mathcal{N}} \bar{c}_j x_j \\ \hline x_i & = & \bar{b}_i & - & \sum_{j \in \mathcal{N}} \bar{a}_{ij} x_j \quad (i \in \mathcal{B}) \end{array}$$

(Note that the numbers $\bar{c}_j$ are called "reduced costs", a term that makes sense for minimization problems, but seems odd in our maximization setting.)

But we will also need to know how all of these values $\bar{b}_i$, $\bar{c}_j$, $\bar{a}_{ij}$ depend on the original data, so we keep track of the optimal dictionary in matrix form:

$$\begin{array}{ccccc} \zeta & = & c_{\mathcal{B}}^{\top} B^{-1} b & + & \left(c_{\mathcal{N}}^{\top} - c_{\mathcal{B}}^{\top} B^{-1} A_{\mathcal{N}}\right) x_{\mathcal{N}} \\ \hline x_{\mathcal{B}} & = & B^{-1} b & - & B^{-1} A_{\mathcal{N}} x_{\mathcal{N}} \end{array}$$

It is important here to know how to recover the matrix B^{-1} from the optimal dictionary, given the initial basis. For example, when the equality-form problem is obtained from an LP in standard form by introducing one slack variable for each constraint, then the slack variables become the initial basis and we recover B^{-1} column by column from the final dictionary — the i^{th} column of B^{-1} is the column of the final dictionary corresponding to the i^{th} slack variable, where we multiply by -1 in the case that the variable is non-basic.

We consider here only isolated changes in one datum. We will use δ to indicate the amount of change, $\hat{u}$ to indicate the modified value of datum u, and $\bar{\bar{u}}$ to indicate the resulting modification in an entry $\bar{u}$ in the optimal dictionary.

We begin with the objective coefficients, then consider the righthand side values, and then finally the entries of matrix A.

1. For $j \in \mathcal{N}$, c_j changes to $\hat{c}_j = c_j + \delta$.

   CONSEQUENCE: Reduced cost $\bar{c}_j$ will change.

   COMPUTE: $\bar{\bar{c}}_j = \bar{c}_j + \delta$

   ANALYSIS:
   case (i): If $\bar{c}_j + \delta \leq 0$ (e.g., when $\delta \leq 0$), then the current dictionary is still optimal. The primal solution does not change, nor does the dual solution.

   case (ii): If $\bar{c}_j + \delta > 0$, then we need to carry out a primal pivot. (Hopefully just one will do, but not always.)

   RANGE for δ under which optimality of the current dictionary is retained is $(-\infty, |\bar{c}_j|]$.

2. For $i \in \mathcal{B}$, c_i changes to $\hat{c}_i = c_i + \delta$.

   CONSEQUENCE: All reduced costs $\bar{c}_j$ can potentially change.

   COMPUTE: Let $\mathbf{e}_i$ denote the vector with rows indexed by basic variables, a one in position i and zeros elsewhere. With $\hat{c}_{\mathcal{B}} = c_{\mathcal{B}} + \delta \mathbf{e}_i$, we need

   $$\hat{c}_{\mathcal{B}}^{\top} B^{-1} A_{\mathcal{N}} = c_{\mathcal{B}}^{\top} B^{-1} A_{\mathcal{N}} + \delta(\mathbf{e}_i^{\top} B^{-1} A_{\mathcal{N}}).$$

   Ah, but all of the entries of $B^{-1} A_{\mathcal{N}}$ are already in the final dictionary! So we just update each non-basic reduced cost: $\bar{\bar{c}}_j = \bar{c}_j - \delta \bar{a}_{ij} \qquad (j \in \mathcal{N})$.

Analysis:
case (i): If all $\bar{\bar{c}}_j \leq 0$, then the current dictionary is still optimal. The primal solution is the same, but the dual optimal solution $y = \hat{c}_{\mathcal{B}}^\top B^{-1}$ and the optimal objective value $\bar{\zeta}$ will change.

case (ii): If some $\bar{\bar{c}}_j > 0$, then we need to carry out a primal pivot.

Range for δ under which optimality of the current dictionary is retained is

$$\delta \in \left(\max_{\bar{a}_{ij}>0} \frac{\bar{c}_j}{\bar{a}_{ij}} \ , \ \min_{\bar{a}_{ij}<0} \frac{\bar{c}_j}{\bar{a}_{ij}} \right).$$

(Please do not take my word for it; think about this!)

3. For $i \in \mathcal{B}$, b_i changes to $\hat{b}_i = b_i + \delta$.

   Consequence: Potentially, all right-hand side values change and feasibility is at stake.

   Compute: Let $\mathbf{e}_i$ denote the vector with rows indexed by basic variables, a one in position i and zeros elsewhere. Let v denote the i^{th} column of matrix B^{-1} (corresponding to the i^{th} slack variable in the standard form case). Then we have

$$\begin{aligned} B^{-1}\hat{b} &= B^{-1}(b + \delta \mathbf{e}_i) \\ &= B^{-1}b + \delta v \end{aligned}$$

   so that, for each $h \in \mathcal{B}$, the updated RHS is

$$\bar{\bar{b}}_h = \bar{b}_h + \delta v_h.$$

   Analysis:
   case (i): If all $\bar{\bar{b}}_h \geq 0$, then the current dictionary is still feasible. The primal solution will change accordingly and so will $\bar{\zeta}$, but the dual optimal solution y remains the same.

   case (ii): If some $\bar{\bar{b}}_h < 0$, then we need to carry out a dual simplex pivot. (Hopefully just one will suffice, but not always.)

   Range for δ under which optimality of the current dictionary is retained is

$$\delta \in \left(\max_{v_h>0} -\frac{\bar{b}_h}{v_h} \ , \ \min_{v_h<0} -\frac{\bar{b}_h}{v_h} \right).$$

   (Why?)

4. For some i, j, a_{ij} changes. First consider $j \in \mathcal{N}$ and assume a_{ij} changes to $\hat{a}_{ij} = a_{ij} + \delta$.

   Consequence: Since this is an entry of $A_{\mathcal{N}}$, and it's in column j, only $\bar{c}_j$ can change. We need to see if the dictionary is still optimal.

   Compute: We have access to the dual optimal solution y. So we just compute

   $$\bar{\bar{c}}_j = c_j - y^\top (\hat{A}_{\mathcal{N}})_j = \bar{c}_j - \delta y_i$$

   where $\hat{A}$ is the updated matrix (with only the (i,j)-entry changed and $(\hat{A}_{\mathcal{N}})_j$ is shorthand for the j^{th} column of this matrix. (Why does this simplify? Think about it.)

   Analysis:
   case (i): If $\bar{\bar{c}}_j = \bar{c}_j - \delta y_i \leq 0$, then the current dictionary is still optimal. The primal solution does not change, nor does the objective value. But if j was a slack variable, then the corresponding dual variable will change.

   case (ii): If we have $\bar{\bar{c}}_j > 0$, then we need to carry out one or more pivots of the simplex method.

   Range for δ under which optimality of the current dictionary is retained is

   $$\delta \in \left(\frac{\bar{c}_j}{y_i}, \infty \right)$$

   provided $y_i > 0$.

5. Some a_{ij} changes where j is basic.

   We Will Not Deal With This Case on Assignments or Tests.

   Consequence: Since this is an entry of B, the matrix B^{-1} changes. So both feasibility and dual feasibility are at risk.

   Compute: Let v denote the j^{th} column of B and let w denote the i^{th} column of B^{-1}. We have

   $$\begin{aligned} B^{-1} v &= \mathbf{e}_j \\ B^{-1}(v + \delta \mathbf{e}_i) &= \mathbf{e}_j + \delta w \end{aligned}$$

   so the new B^{-1} is obtained by adding, for each h, $\delta w_h / (1 + \delta w_j)$ times row j of the old B^{-1} to row h of the old B^{-1}. (This is just Gaussian elimination, applied to compute the inverse of a matrix; in our scenario, most of the work is already done, except for one column.) Call this new inverse matrix B'.

   Now we must re-compute

   $$\bar{\bar{b}} = B' b, \qquad \bar{\bar{y}}^\top = c_{\mathcal{B}}^\top B', \qquad c_{\mathcal{N}}^\top - \bar{\bar{y}}^\top A_{\mathcal{N}}.$$

   The analysis involves checking both feasibility and dual feasibility. If one is compromised, then we proceed to pivot with either the primal or dual method. If both are compromised, we need to appeal to some sort of Phase I algorithm.

Adding and Removing Constraints and Variables

One more important component of post-optimality analysis is the addition and deletion of primal and dual variables. In short,

- if after optimizing, you find a constraint that needs to be removed, then you can make that slack variable basic and subsequently delete that row from the dictionary; you don't care anymore whether that variable is positive or negative or zero.

- if after optimizing, you are forced to add a new constraint, just give it a new slack variable, insert this slack variable into the basis and the equation for it gets appended to the bottom of the dictionary. Typically, the right-hand side of this equation will involve other basic variables, so these need to be eliminated before you can proceed to analyze this dictionary for optimality. Often, an infeasibility results and the dual simplex method is employed to return to feasibility.

- if after optimizing, a variable is to be eliminated (i.e., forced to zero), just make it non-basic and delete that column from the dictionary.

- on the other hand, if a new variable is introduced, we have to introduce it carefully into the set $\mathcal{N}$. Suppose this variable, x', has objective coefficient c' (in terms of the original data) and the corresponding column of the constraint matrix is a'. Then put x' into $\mathcal{N}$, give it coefficients $\bar{a}' = -B^{-1}a'$ in the current optimal dictionary and give it coefficient $\bar{c}' = c' + y^\top \bar{a}'$ in the objective row. If $\bar{c}' \leq 0$, the the dictionary is still optimal and the newfangled variable does not enter into the problem. If $\bar{c}' > 0$, then proceed with the simplex method to return to optimality. A key point here is that the computation $c' - y^\top a'$ quickly decides if this new variable is worth considering (vis a vis the current dictionary).

That's all I feel like writing about for now. I hope this helps!