

Lecture 9: Formulas for the Simplex Tableau

February 10, 2020

University of Illinois at Urbana-Champaign

1 Row reduction and inverse matrices

This section is a summary of some relevant facts from linear algebra.

When we row-reduce a system of equations $A\mathbf{x} = \mathbf{b}$, every step is an elementary row operation:

- Multiplying the i^{th} row by a constant $c \neq 0$,
- Adding c times the i^{th} row to the j^{th} row, or
- Swapping the i^{th} row and the j^{th} row. (Though we won't ever need this third operation in this class.)

These can all be described by taking a corresponding elementary matrix E , and multiplying both sides of the equation $A\mathbf{x} = \mathbf{b}$ by E , from the left, getting $E A \mathbf{x} = E \mathbf{b}$. Here is an example of this in action, where the elementary row operation is “subtract twice the first row from the third row”:

$$\begin{array}{ccc}
 E & A\mathbf{x} = \mathbf{b} & E A \mathbf{x} = E \mathbf{b} \\
 \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -2 & 0 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 7 \\ 8 \\ 9 \end{bmatrix} & \begin{bmatrix} 1 & 4 \\ 2 & 5 \\ 1 & -2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 7 \\ 8 \\ -5 \end{bmatrix}
 \end{array}$$

The product of many of these elementary matrices is some other, more complicated matrix. As a result, *every* system of equations we get from $A\mathbf{x} = \mathbf{b}$ by row operations will have the form $M A \mathbf{x} = M \mathbf{b}$, for some matrix M . A single matrix can summarize the composition of arbitrarily many row operations.

In particular, we might care about which matrix M corresponds to the row operations that row-reduce a matrix. Well, which matrix is it?

To answer this question, let's first split up the columns of our system into the basic variables \mathcal{B} and the nonbasic variables \mathcal{N} , as we did in the previous lecture. When we row-reduce the matrix A with respect to \mathcal{B} , our goal is to have the i^{th} basic variable have a 1 in the i^{th} row, and 0 in all the other rows. In other words, in the matrix MA of the resulting row-reduced system $M A \mathbf{x} = M \mathbf{b}$, the columns $(MA)_{\mathcal{B}}$ corresponding to the basic variables just form the identity matrix, I .

This is enough to tell us what M must be. If left-multiplying all of M by A gives us a matrix MA whose submatrix $(MA)_{\mathcal{B}}$ is equal to I , then left-multiplying just the submatrix $A_{\mathcal{B}}$ by M should give us just I . But there is exactly one matrix that has this effect: it is the inverse matrix $A_{\mathcal{B}}^{-1}$.

¹This document comes from the Math 482 course webpage: <https://faculty.math.illinois.edu/~mlavrov/courses/482-spring-2020.html>

So, when we row-reduce the system $A\mathbf{x} = \mathbf{b}$ with respect to the basic variables \mathcal{B} , the new system of equations we get (equivalent to the old, but in a nicer form) is $A_{\mathcal{B}}^{-1}A\mathbf{x} = A_{\mathcal{B}}^{-1}\mathbf{b}$.

For another way to think about what's going, consider splitting up the system $A\mathbf{x} = \mathbf{b}$ into

$$A_{\mathcal{B}}\mathbf{x}_{\mathcal{B}} + A_{\mathcal{N}}\mathbf{x}_{\mathcal{N}} = \mathbf{b}.$$

Then, we can solve for $\mathbf{x}_{\mathcal{B}}$ in terms of $\mathbf{x}_{\mathcal{N}}$ by multiplying by $A_{\mathcal{B}}^{-1}$ and rearranging:

$$\mathbf{x}_{\mathcal{B}} = A_{\mathcal{B}}^{-1}\mathbf{b} - A_{\mathcal{B}}^{-1}A_{\mathcal{N}}\mathbf{x}_{\mathcal{N}}.$$

This gives us a formula for *all* the solutions to $A\mathbf{x} = \mathbf{b}$: we can pick any value for $\mathbf{x}_{\mathcal{N}}$ that we like, and then we must use the formula above to find the value of $\mathbf{x}_{\mathcal{B}}$ that goes with it. In particular, we can set $\mathbf{x}_{\mathcal{N}} = \mathbf{0}$, and then the formula gives us $\mathbf{x}_{\mathcal{B}} = A_{\mathcal{B}}^{-1}\mathbf{b}$: this is what we called the basic solution.

2 Formulas for a simplex tableau

How does this look in the simplex tableau?

Again, we assume that our linear program has the form

$$\begin{aligned} & \underset{\mathbf{x} \in \mathbb{R}^n}{\text{maximize}} && \mathbf{c}^T \mathbf{x} \\ & \text{subject to} && A\mathbf{x} = \mathbf{b} \\ & && \mathbf{x} \geq \mathbf{0} \end{aligned}$$

where A is an $m \times n$ matrix, $\mathbf{b} \in \mathbb{R}^m$, and $\mathbf{c} \in \mathbb{R}^n$; the rows of A are linearly independent.

Every choice of basic variables \mathcal{B} produces a valid simplex tableau if two things happen:

- $A_{\mathcal{B}}$ is invertible: then by setting $\mathbf{x}_{\mathcal{N}} = \mathbf{0}$ and $\mathbf{x}_{\mathcal{B}} = A_{\mathcal{B}}^{-1}\mathbf{b}$, we get a basic solution.
- $A_{\mathcal{B}}^{-1}\mathbf{b} \geq \mathbf{0}$; then, the basic solution we get is feasible.

To make the simplex tableau look “cleaner”, we can sort the columns to put the basic variables $\mathbf{x}_{\mathcal{B}}$ first and $\mathbf{x}_{\mathcal{N}}$ last: then, the tableau will look like

	$\mathbf{x}_{\mathcal{B}}$	$\mathbf{x}_{\mathcal{N}}$	
$\mathbf{x}_{\mathcal{B}}$	I	Q	\mathbf{p}
$-z$	$\mathbf{0}^T$	\mathbf{r}^T	$-z_0$

for some matrix Q , vectors \mathbf{p} and \mathbf{r} , and scalar z_0 . (Note: I've used the notation from section 5.5 of Gärtner and Matoušek, with one deviation: their definition of Q is the negative of ours. This is because their simplex tableaux are written differently.)

Usually, in practice, we don't sort the columns, and so things look slightly different. But swapping the columns doesn't hurt anything, as long as we keep them labeled. For example, for the linear program

$$\begin{aligned} & \underset{x,y \in \mathbb{R}}{\text{maximize}} && 2x + 3y \\ & \text{subject to} && -x + y + s_1 & & = 3 \\ & && x - 2y & + s_2 & = 2 \\ & && x + y & & + s_3 = 7 \\ & && x, y, s_1, s_2, s_3 \geq 0 \end{aligned}$$

that we looked at in Lecture 5, when we were done solving it, we got the tableau on the left below, but we could have written it in "cleaned-up form" as on the right.

	x	y	s_1	s_2	s_3	
y	0	1	1/2	0	1/2	5
s_2	0	0	3/2	1	1/2	10
x	1	0	-1/2	0	1/2	2
$-z$	0	0	-1/2	0	-5/2	-19

	y	s_2	x	s_1	s_3	
y	1	0	0	1/2	1/2	5
s_2	0	1	0	3/2	1/2	10
x	0	0	1	-1/2	1/2	2
$-z$	0	0	0	-1/2	-5/2	-19

How can we find the values of \mathbf{p} , Q , \mathbf{r} , and z_0 just from knowing the problem and which basic variables \mathcal{B} we chose?

Let's start with \mathbf{p} . This is the right-hand side of the tableau, and holds the values of the basic variables in the current basic feasible solution. In other words, it's exactly what we get from the formula $A_{\mathcal{B}}^{-1}\mathbf{b}$. In the problem above, if we first rearrange A to a matrix A' that puts the columns in y, s_2, x, s_1, s_3 order, we get:

$$A' = \begin{bmatrix} 1 & 0 & -1 & 1 & 0 \\ -2 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 \end{bmatrix}, \quad A_{\mathcal{B}} = \begin{bmatrix} 1 & 0 & -1 \\ -2 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix}, \quad A_{\mathcal{B}}^{-1}\mathbf{b} = \begin{bmatrix} 1 & 0 & -1 \\ -2 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 3 \\ 2 \\ 7 \end{bmatrix} = \begin{bmatrix} 5 \\ 10 \\ 2 \end{bmatrix}.$$

We also know that the row-reduction steps needed to start at $A\mathbf{x} = \mathbf{b}$ and end up at the tableau above can be replaced by left-multiplying by $A_{\mathcal{B}}^{-1}$. When we multiply the $A_{\mathcal{B}}$ part of the tableau by $A_{\mathcal{B}}^{-1}$, we get the big identity matrix. When we multiply the $A_{\mathcal{N}}$ part of the tableau by $A_{\mathcal{B}}^{-1}$, we get the matrix Q we wanted a formula for. So $Q = A_{\mathcal{B}}^{-1}A_{\mathcal{N}}$; In our example,

$$Q = A_{\mathcal{B}}^{-1}A_{\mathcal{N}} = \begin{bmatrix} 1 & 0 & -1 \\ -2 & 1 & 1 \\ 1 & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1/2 & 1/2 \\ 3/2 & 1/2 \\ -1/2 & 1/2 \end{bmatrix}.$$

What about the reduced costs row, consisting of $\mathbf{0}^T$, \mathbf{r}^T , and $-z_0$?

Of these, z_0 is easier to find: it is the objective value. In general, our objective value is $\mathbf{c}^T\mathbf{x}$, which we can split up into two parts: $\mathbf{c}_{\mathcal{B}}^T\mathbf{x}_{\mathcal{B}} + \mathbf{c}_{\mathcal{N}}^T\mathbf{x}_{\mathcal{N}}$.

At our current basic solution, $\mathbf{x}_{\mathcal{B}} = A_{\mathcal{B}}^{-1}\mathbf{b}$ and $\mathbf{x}_{\mathcal{N}} = \mathbf{0}$. Therefore $\mathbf{c}^T\mathbf{x} = \mathbf{c}_{\mathcal{B}}^T A_{\mathcal{B}}^{-1}\mathbf{b}$ is the current objective value; that's the formula for $-z_0$. In terms of \mathbf{p} , we have $z_0 = \mathbf{c}_{\mathcal{B}}^T\mathbf{p}$.

To find \mathbf{r}^T , let's think about where the reduced cost row comes from. Generally, we compute it one row operation at a time, as we move from one basic feasible solution to another. But we can

always compute the reduced cost row from scratch: start with the original row of costs (which is the vector \mathbf{c}^\top) and row-reduce it to eliminate the entries below the pivots.

The big matrix replacing A in the tableau (consisting of both I and Q) is given by $A_B^{-1}A$. (Or, $A_B^{-1}A'$, if we decided to rearrange the columns to put the basic variables first.) To eliminate the costs below the basic variables, we proceed in two steps:

- Take this matrix $A_B^{-1}A$ and left-multiply it by \mathbf{c}_B^\top , to get a vector that matches \mathbf{c}^\top in the basic-variable components.
- Subtract this vector from \mathbf{c}^\top , getting a vector with zeroes in the basic-variable components, and the reduced costs in the nonbasic-variable components.

In our example, the first step gives us

$$\mathbf{c}_B^\top(A_B^{-1}A') = [3 \ 0 \ 2] \begin{bmatrix} 1 & 0 & 0 & 1/2 & 1/2 \\ 0 & 1 & 0 & 3/2 & 1/2 \\ 0 & 0 & 1 & -1/2 & 1/2 \end{bmatrix} = [3 \ 0 \ 2 \ 1/2 \ 5/2]$$

and when we subtract this from \mathbf{c}^\top (also reordered to match our new order of variables), we get

$$[3 \ 0 \ 2 \ 0 \ 0] - [3 \ 0 \ 2 \ 1/2 \ 5/2] = [0 \ 0 \ 0 \ -1/2 \ -5/2] = [\mathbf{0}^\top \ \mathbf{r}^\top].$$

So our formula for the bottom row of the tableau is $\mathbf{c}^\top - \mathbf{c}_B^\top A_B^{-1}A$.

There are two ways to use this. First, we could get the reduced cost of a specific variable: if A_i is the column of A corresponding to the variable x_i , then its reduced cost r_i is given by

$$r_i = c_i - \mathbf{c}_B^\top A_B^{-1}A_i.$$

Second, we could get all the nonbasic reduced costs \mathbf{r}^\top by the formula

$$\mathbf{r}^\top = \mathbf{c}_N^\top - \mathbf{c}_B^\top A_B^{-1}A_N.$$

You may have noticed that the vector $\mathbf{c}_B^\top A_B^{-1}$ appears here a lot. It can be convenient to give it a name, \mathbf{u}^\top , and to compute it when we're finding reduced costs. We'll see more properties of the vector \mathbf{u}^\top later on in this course!

To summarize, here are the formulas for the parts of the simplex tableau:

- $\mathbf{p} = A_B^{-1}\mathbf{b}$.
- $Q = A_B^{-1}A_N$.
- $\mathbf{r}^\top = \mathbf{c}_N^\top - \mathbf{c}_B^\top A_B^{-1}A_N$.

We can write this as $\mathbf{c}_N^\top - \mathbf{c}_B^\top Q$ if we already found Q , or as $\mathbf{c}_N^\top - \mathbf{u}^\top A_N$ in terms of $\mathbf{u}^\top = \mathbf{c}_B^\top A_B^{-1}$.

- $z_0 = \mathbf{c}_B^\top A_B^{-1}\mathbf{b}$.

We can write this as $\mathbf{c}_B^\top \mathbf{p}$ if we already found \mathbf{p} , or as $\mathbf{u}^\top \mathbf{b}$ in terms of $\mathbf{u}^\top = \mathbf{c}_B^\top A_B^{-1}$.