Notes on Row Reduction

Francis J. Narcowich
Department of Mathematics
Texas A&M University

September 2003

1 The Row-Reduction Algorithm

The row-reduced form of a matrix contains a great deal of information, both about the matrix itself and about systems of equations that may be associated with it. To talk about row reduction, we need to define several terms and introduce some notation.

1. *Notation for row operations.* Row reduction is made easier by having notation for various row operations. This is the notation that we will use here.

   (a) *Elementary transposition.* $R \leftrightarrow R'$ means interchange rows R and R'.

   (b) *Elementary multiplication.* $R = cR$ means multiply the current row R by $c \neq 0$ and make the result the new row R.

   (c) *Elementary modification.* $R = R + cR'$ means multiply R' by c, then add cR' to R, and make the result the new row R.

2. *Row equivalence.* A matrix A is *row equivalent* to a matrix B if A can be transformed into B using a finite number of elementary row operations. Since such operations are reversible, B is also row-equivalent to A, and we simply say that A and B are row equivalent; we write $A \Leftrightarrow B$.

3. *Leading entry.* The *leading entry* in a row is the first non-zero entry in a row. The leading entries in each row of M are in boldface type.

$$M = \begin{pmatrix} 0 & \mathbf{1} & 3 & 2 \\ \mathbf{2} & 4 & 0 & -1 \\ 0 & 0 & \mathbf{6} & 5 \end{pmatrix}$$

4. *Row-reduced form of a matrix.* A matrix is in *row-reduced form* if these hold:

   (a) All rows with only zeros for entries are at the bottom.

   (b) In any non-zero row, the leading entry is 1.

   (c) Any column containing a leading entry for some row has zeros in all its other entries.

The matrix M above is *not* in row-reduced form. Here are two examples of matrices that *are* in row-reduced form.

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 3 \\ 0 & 0 & 1 & 0 & -4 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } \begin{pmatrix} 1 & 0 & 0 & 7 \\ 0 & 1 & 0 & 4 \\ 0 & 0 & 1 & 3 \end{pmatrix}.$$

Although we will not prove it here, every matrix is row equivalent to a matrix that is in row-reduced form. We call the process of finding the row-reduced matrix equivalent to a given matrix *row reduction.* Below is an algorithm for row-reducing a matrix; it uses two sets of steps, *forward* ones (Gaussian elimination) and *backward* ones (Jordan reduction).

Forward Steps

1. Starting on the left, find the first column that contains a leading entry. If there are several leading entries in this column, choose a convenient one—for example, an entry that is 1.

2. Interchange rows until the row containing the leading entry that you have chosen is the top one.

3. Use the top row and elementary modification to zero-out the other entries in the column that you are working with.

4. Repeat the first three steps with the submatrix comprising all columns to the right of the one you were working with, and all rows below the top. Stop when either there are no more rows left or the next submatrix consists of zeros. When this step is complete, all rows with only zeros for entries are at the bottom of the matrix.

Backward Steps

1. Find the leading entry in the last non-zero row. If this entry is not already 1, use elementary multiplication to make it 1. Use elementary modification to zero-out all of the entries above this leading entry.

2. Repeat step one for the leading entry in the second to the last non-zero row. Once this is done, do the same thing for the third to the last row, fourth to last, and so on. When there are no more rows left, the matrix is in row-reduced form. (Actually, it is in a special kind of row-reduced form, called *reduced echelon* form.)

Rank of a matrix. Before we give an example, we want to mention a very important quantity that one can obtain from the row-reduced form of a matrix; namely, the *rank* of a matrix. The *rank* of a matrix A, which we denote by rank$[A]$, is defined to be the number of leading entries in the row reduced form of A. For future reference, we point out that this is equal to the number of non-zero rows in the row-reduced form of A. One more thing: If, after row-reducing a matrix, a column contains a leading entry, then that column (which is a column in the original matrix) is called a *leading column*. The rank of A is obviously also equal to the number of leading columns in A.

We will now row-reduce a matrix using the algorithm given above. Consider the matrix A shown below.

$$\text{(Starting matrix) } A = \begin{pmatrix} 1 & 1 & 1 \\ 2 & 2 & 2 \\ 1 & -1 & 2 \end{pmatrix}$$

This is the matrix we will row reduce. We will begin by carrying out the forward steps in the algorithm. Leading entries are in boldface type.

$$R_2 = R_2 - 2R_1 \, , \ R_3 = R_3 - R_1 : \quad A \Leftrightarrow \begin{pmatrix} \mathbf{1} & 1 & 1 \\ 0 & 0 & 0 \\ 0 & -\mathbf{2} & 1 \end{pmatrix}$$

$$R_2 \leftrightarrow R_3 : \quad A \Leftrightarrow \begin{pmatrix} \mathbf{1} & 1 & 1 \\ 0 & -\mathbf{2} & 1 \\ 0 & 0 & 0 \end{pmatrix}$$

$$R_2 = -\frac{1}{2}R_2 : \quad A \Leftrightarrow \begin{pmatrix} \mathbf{1} & 1 & 1 \\ 0 & \mathbf{1} & -\frac{1}{2} \\ 0 & 0 & 0 \end{pmatrix}$$

This ends the set of forward steps. We will move on to the backward steps.

$$R_1 = R_1 - R_2: \quad A \Leftrightarrow \begin{pmatrix} \mathbf{1} & 0 & \frac{3}{2} \\ 0 & \mathbf{1} & -\frac{1}{2} \\ 0 & 0 & 0 \end{pmatrix}$$

This completes the algorithm. The last matrix on the right above is the row- reduced form of A. Note that this implies that rank$[A] = 2$ and that the leading columns of A are the first and second.

2   Applications to Solving Systems of Equations

One of the most important applications of the row-reduction algorithm is solving a system of linear equations. The procedure begins by converting the system

$$\mathbf{S}: \begin{cases} a_{11}x_1 + \cdots + a_{1n}x_n & = & b_1 \\ a_{21}x_1 + \cdots + a_{2n}x_n & = & b_2 \\ \quad\quad\quad \vdots & & \vdots \quad \vdots \\ a_{m1}x_1 + \cdots + a_{mn}x_n & = & b_m \end{cases}$$

into the matrix form $Ax = b$. Here A is the $m \times n$ coefficient matrix for S, x is the $n \times 1$ vector of unknowns, and b being the $m \times 1$ vector of b_j's.

There is an important matrix associated with the system S. We call the $m \times (n+1)$ matrix

$$[A|b] = \begin{pmatrix} a_{11} & \ldots & a_{1n} & b_1 \\ \vdots & \ldots & \vdots & \vdots \\ a_{m1} & \ldots & a_{mn} & b_m \end{pmatrix}$$

the *augmented* matrix for the system S. In view of the connection between row operations and operations on the individual equations comprising the system S, any matrix $[A'|b']$ that is row equivalent to the original system $[A|b]$ is the augmented matrix for a system S' equivalent (i.e., having the same solution set) to S. Let us state this formally.

Theorem 2.1 *If $[A|b]$ and $[A'|b']$ are augmented matrices for two linear systems of equations S and S', and if $[A|b]$ and $[A'|b']$ are row equivalent, then S and S' are equivalent systems.*

By examining the possible row-reduced matrices corresponding to the augmented matrix for the system S, one can use Theorem 2.1 to obtain the following result, which we state without proof.

Theorem 2.2 *Consider the system S with coefficient matrix A and augmented matrix $[A|b]$. As above, the sizes of b, A, and $[A|b]$ are $m \times 1$, $m \times n$, and $m \times (n+1)$, respectively; in addition, n is the number of unknowns. We have these possibilities:*

1. *S is inconsistent if and only if $\operatorname{rank}[A] < \operatorname{rank}[A|b]$.*

2. *S has a unique solution if and only if $\operatorname{rank}[A] = \operatorname{rank}[A|b] = n$.*

3. *S has infinitely many solutions if and only if $\operatorname{rank}[A] = \operatorname{rank}[A|b] < n$.*

We need to illustrate the use of this theorem. To do that, look at the simple systems below.

$$
\begin{array}{lll}
x_1 + 2x_2 = 1 & 3x_1 + 2x_2 = 3 & 3x_1 + 2x_2 = 3 \\
3x_1 + x_2 = -2 & -6x_1 - 4x_2 = 0 & -6x_1 - 4x_2 = -6
\end{array}
$$

The augmented matrices for these systems are, respectively,

$$
\left(\begin{array}{cc|c} 1 & 2 & 1 \\ 3 & 1 & -2 \end{array}\right)
\qquad
\left(\begin{array}{cc|c} 3 & 2 & 3 \\ -6 & -4 & 0 \end{array}\right)
\qquad
\left(\begin{array}{cc|c} 3 & 2 & 3 \\ -6 & -4 & -6 \end{array}\right)
$$

Applying the row-reduction algorithm yields the row-reduced form of each of these augmented matrices. The result is, again respectively,

$$
\left(\begin{array}{cc|c} 1 & 0 & -1 \\ 0 & 1 & 1 \end{array}\right)
\qquad
\left(\begin{array}{cc|c} 1 & \frac{2}{3} & 1 \\ 0 & 0 & 6 \end{array}\right)
\qquad
\left(\begin{array}{cc|c} 1 & \frac{2}{3} & 1 \\ 0 & 0 & 0 \end{array}\right)
$$

From each of these row-reduced versions of the augmented matrices, one can read off the rank of the coefficient matrix as well as the rank of the augmented matrix. Applying Theorem 2.2 to each of these tells us the number of solutions to expect for each of the corresponding systems. We summarize our findings in the table below.

System	rank$[A]$	rank$[A|b]$	n	# of solutions
First	2	2	2	1
Second	1	2	2	0 *(inconsistent)*
Third	1	1	2	∞

We now turn to the discussion of a very important class of systems, homogeneous ones. A homogeneous system is one in which the vector $b = 0$.

5

By simply plugging $X = 0$ into the equation $AX = 0$, we see that every homogeneous system has at least *one* solution, the *trivial* solution $X = 0$. Are there any others? To answer this, we need the following result:

Corollary 2.3 *Let A be an $m \times n$ matrix. A homogeneous system of equations $AX = 0$ will have a unique solution, the trivial solution $X = 0$, if and only if rank$[A] = n$. In all other cases, it will have infinitely many solutions. As a consequence, if $n > m$—i.e., if the number of unknowns is larger than the number of equations—, then the system will have infinitely many solutions.*

Proof: Since $X = 0$ is always a solution, case (i) of Theorem 2.2 is eliminated as a possibility. Therefore, we must always have rank$[A] = $ rank$[A|0] \leq n$. By Theorem 2.2, case (ii), equality will hold if and only if $X = 0$ is the only solution. When it does not hold, we are always in case (iii) of Theorem 2.2; there are thus infinitely many solutions for the system. If $n > m$, then we need only note that rank$[A] \leq m < n$ to see that the system has to have infinitely solutions. $\square$

Thus far we have only discussed *how many* solutions a system of linear equations has. We have said nothing about *how to obtain* these solutions. The key to the whole process is row-reducing the augmented matrix for the original system, S.

Let us look at an example. Suppose that we have found that our system has an augmented matrix $[A|b]$ that is row equivalent to the matrix

$$[A'|b'] = \left(\begin{array}{ccccc|c} \mathbf{1} & 2 & 0 & 0 & 0 & -1 \\ 0 & 0 & \mathbf{1} & 0 & -3 & 7 \\ 0 & 0 & 0 & \mathbf{1} & 3 & 4 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right).$$

We now convert this back to a system, one that is of course equivalent to whatever one we started with. The result is the following system of equations:

$$\begin{aligned} x_1 + 2x_2 &= -1 \\ x_3 - 3x_5 &= 7 \\ x_4 + 3x_5 &= 4. \end{aligned}$$

Notice that the variables corresponding to the leading columns, which are called *leading variables*, appear in this set only once. That means that they

can be solved for in terms of the other variables. Solving for these leading variables results in the system

$$\begin{aligned} x_1 &= -2x_2 - 1 \\ x_3 &= 3x_5 + 7 \\ x_4 &= -3x_5 + 4. \end{aligned}$$

It turns out that by assigning arbitrary values to the *non*-leading variables gives us all possible solutions to the system. It is customary to show that this assignment has been made by assigning new letters to the non-leading variables. In our example, we could set $x_2 = s$, $x_5 = t$, and then rewrite the whole solution in the column form shown below.

$$X = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{pmatrix} = \begin{pmatrix} -2s - 1 \\ s \\ 3t + 7 \\ -3t + 4 \\ t \end{pmatrix} = \begin{pmatrix} -1 \\ 0 \\ 7 \\ 4 \\ 0 \end{pmatrix} + s \begin{pmatrix} -2 \\ 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + t \begin{pmatrix} 0 \\ 0 \\ 3 \\ -3 \\ 1 \end{pmatrix}$$

Written in this way, we see that if we set $s = t = 0$, we get a *particular* solution to the original system. When this column is subtracted off, what is left is a solution to the corresponding *homogeneous* system. This happens in every case: A solution to $AX = b$ may be written $X_p + X_h$, where X_p is a fixed column vector satisfying $AX_p = b$, and X_h runs over all solutions to $AX_h = 0$. This is exactly analogous to what happens in the case of linear differential equations.