LU-Factorization

Isaiah Lankham, Bruno Nachtergaele, Anne Schilling
(March 12, 2007)

1 Introduction

Given a system of linear equations, a complete reduction of the coefficient matrix to Reduced Row Echelon (RRE) form is far from the most efficient algorithm if one is only interested in finding a solution to the system. However, the Elementary Row Operations (EROs) that constitute such a reduction are themselves at the heart of many frequently used numerical (i.e., computer-calculated) applications of Linear Algebra. In the Sections that follow, we will see how EROs can be used to produce a so-called *LU-factorization* of a matrix into a product of two significantly simpler matrices. Unlike Diagonalization and the Polar Decomposition for Matrices that we've already encountered in this course, these LU Decompositions can be computed reasonably quickly for many matrices. LU-factorizations are also an important tool for solving linear systems of equations.

You should note that the factorization of complicated objects into simpler components is an extremely common problem solving technique in mathematics. E.g., we will often factor a polynomial into several polynomials of lower degree, and one can similarly use the prime factorization for an integer in order to simply certain numerical computations.

2 Upper and Lower Triangular Matrices

We begin by recalling the terminology for several special forms of matrices.

A square matrix $A = (a_{ij}) \in \mathbb{F}^{n\times n}$ is called *upper triangular* if $a_{ij} = 0$ for each pair of integers $i, j \in \{1, \ldots, n\}$ such that $i > j$. In other words, A has the form

$$A = \begin{bmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ 0 & a_{22} & a_{23} & \cdots & a_{2n} \\ 0 & 0 & a_{33} & \cdots & a_{3n} \\ \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & a_{nn} \end{bmatrix}.$$

Similarly, $A = (a_{ij}) \in \mathbb{F}^{n\times n}$ is called *lower triangular* if $a_{ij} = 0$ for each pair of integers $i, j \in \{1, \ldots, n\}$ such that $i < j$. In other words, A has the form

$$A = \begin{bmatrix} a_{11} & 0 & 0 & \cdots & 0 \\ a_{21} & a_{22} & 0 & \cdots & 0 \\ a_{31} & a_{32} & a_{33} & \cdots & 0 \\ \vdots & \vdots & & \ddots & \vdots \\ a_{n1} & a_{n2} & a_{n3} & \cdots & a_{nn} \end{bmatrix}.$$

Finally, we call $A = (a_{ij})$ a *diagonal matrix* when $a_{ij} = 0$ for each $i, j \in \{1, \ldots, n\}$ such that $i \neq j$. In other words, A is both an upper triangular matrix and a lower triangular matrix so that the only non-zero entries of A are along the diagonal (i.e., when $i = j$).

Before applying triangularity to the solving of linear systems, we first derive several useful algebraic properties of triangular matrices.

Theorem 2.1. *Let $A, B \in \mathbb{F}^{n\times n}$ be square matrices and $c \in \mathbb{R}$ be any real scalar. Then*

1. *A is upper triangular if and only if A^T is lower triangular.*
2. *if A and B are upper triangular,*
   (a) *cA is upper triangular,*
   (b) *$A + B$ is upper triangular, and*
   (c) *AB is upper triangular.*
3. *all properties of Part 2 hold when* upper triangular *is replaced by* lower triangular.

Proof. The proof of Part 1 follows directly from the definitions of upper and lower triangularity in addition to the definition of transpose. The proofs of Parts 2(a) and 2(b) are equally straightforward. Moreover, the proof of Part 3 follows by combining Part 1 with Part 2 but when A is replaced by A^T and B by B^T. Thus, we need only prove Part 2(c).

To prove Part 2(c), we start from the definition of the matrix product. Denoting $A = (a_{ij})$ and $B = (b_{ij})$, note that $AB = ((ab)_{ij})$ is an $n \times n$ matrix having "i-j entry" given by

$$(ab)_{ij} = \sum_{k=1}^{n} a_{ik} b_{kj}.$$

Since A and B are upper triangular, we have that $a_{ik} = 0$ when $i > k$ and that $b_{kj} = 0$ when $k > j$. Thus, to obtain a non-zero summand $a_{ik}b_{kj} \neq 0$, we must have both $a_{ik} \neq 0$, which implies that $i \leq k$, and $b_{kj} \neq 0$, which implies that $k \leq j$. In particular, these two conditions are simultaneously satisfiable only when $i \leq j$. Therefore, $(ab)_{ij} = 0$ when $i > j$, from which AB is upper triangular. $\square$

As a side remark, we mention that Parts 2(a) and 2(b) of Theorem 2.1 imply that the set of all $n \times n$ upper triangular matrices forms a proper subspace of the vector space of all $n \times n$ matrices. (One can furthermore see from Part 2(c) that the set of all $n \times n$ upper triangular matrices actually forms a subalgebra of the algebra of all $n \times n$ matrices.)

3 Back and Forward Substitution

Consider an $n \times n$ linear system of the form

$$Ax = b$$

with A an upper triangular matrix. To solve such a linear system, it is not at all necessary to first reduce it to Reduced Row Echelon (RRE) form. To see this, note that the last equation in such a system can only involve the single unknown x_n and that, moreover,

$$x_n = \frac{b_n}{a_{nn}}$$

as long as $a_{nn} \neq 0$. If $a_{nn} = 0$, then we must be careful to distinguish the two cases where $b_n = 0$ and $b_n \neq 0$. Thus, for concreteness, we assume that the diagonal elements of A are all nonzero, and so we can next substitute the solution for x_n into the second-to-last equation. Since A is upper triangular, the resulting equation will involve only the single unknown x_{n-1}, and, moreover,

$$x_{n-1} = \frac{b_{n-1} - a_{n-1,n}x_n}{a_{n-1,n-1}}.$$

One can then similarly substitute the solutions for x_n and x_{n-1} into the ante penultimate equation in order to solve for x_{n-2}, and so on until the complete solution is found. We call this process *back substitution*. Note that, as an intermediate step in our algorithm for reduction to RRE form, we obtain an upper triangular matrix that is row equivalent to A. Back substitution allows one to stop the reduction at that point and solve the linear system.

A similar procedure can be applied when A is lower triangular. In this case, the first equation contains only x_1, so

$$x_1 = \frac{b_1}{a_{11}},$$

where we are again assuming that the diagonal entries of A are all nonzero. Then, as above, we can substitute the solution for x_1 into the second equation to obtain

$$x_2 = \frac{b_2 - a_{21}x_1}{a_{22}}.$$

Continuing this process, we obtain the *forward substitution* procedure. In particular,

$$x_n = \frac{b_n - \sum_{k=1}^{n-1} a_{nk}x_k}{a_{nn}}.$$

We next consider a more general linear system $Ax = b$, for which we assume that there is a lower triangular matrix L and an upper triangular matrix U such that $A = LU$. Such a system is more general since it clearly includes the special cases of A being either lower or upper triangular. In order to solve such a system, we can again exploit triangularity in order to produce a solution without applying a single Elementary Row Operation. To accomplish this, we first set $y = Ux$, where x is the as yet unknown solution of $Ax = b$. Then note that y must satisfy

$$Ly = b.$$

As L is lower triangular, we can solve for y via forward substitution (assuming that the diagonal entries of L are all nonzero). Then, once we have y, we can apply back substitution to solve for x in the system

$$Ux = y$$

since U is upper triangular. (As with L, we must also assume that every diagonal entry of U is nonzero.)

In summary, we have given an algorithm for solving any linear system $Ax = b$ in which we can factor $A = LU$, where L is lower triangular, U is upper triangular, and both L and U have all non-zero diagonal entries. Moreover, the solution is found entirely through simple forward and back substitution, and one can easily verify that the solution obtained is unique.

We note in closing that the simple procedures of back and forward substitution can also be regarded as a straightforward method for computing the inverses of lower and upper triangular matrices. The only condition we imposed on our triangular matrices above was that all diagonal entries were non-zero. It should be clear to you that this non-zero diagonal restriction is a necessary and sufficient condition for a triangular matrix to be non-singular. Moreover, once the inverses of L and U have been obtained, then we can immediately calculate the inverse for A by noting that

$$A^{-1} = (LU)^{-1} = U^{-1}L^{-1}.$$

4 LU factorization

Based upon the discussion in the previous Section, it should be clear that one can find many uses for the factorization of a matrix $A = LU$ into the product of a lower triangular matrix L and an upper triangular matrix U. This form of decomposition of a matrix is called an *LU-factorization* (or sometimes *LU-decomposition*). One can prove that such a factorization, with L and U satisfying the condition that all diagonal entries are non-zero, is equivalent to either A or some permutation of A being non-singular. For simplicity, we will now explain how such an LU-factorization of A may be obtained in the most common case that A can be reduced to RRE form without requiring any row swapping operations. (Row-exchanges can, however, be included in the discussion below with only a small additional effort.)

Not surprisingly, the factorization procedure itself involves Elementary Row Operations (EROs). Suppose that $A = (a_{ij})$ is an $n \times n$ matrix with $a_{11} \neq 0$. The first step in our algorithm for reducing A to RRE form is then the ERO that replaces the second row by the second minus the first row after it has been multiplied (i.e., rescaled) by the factor a_{21}/a_{11}. This produced the matrix

$$A_1 = \begin{bmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ 0 & a_{22} - a_{12}a_{21}/a_{11} & a_{23} - a_{13}a_{21}/a_{11} & \cdots & a_{2n} - a_{1n}a_{21}/a_{11} \\ a_{31} & a_{32} & a_{33} & \cdots & a_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \end{bmatrix}.$$

The new trick comes now. The ERO we just performed on A can be viewed as the multiplication of A from the left with the lower triangular matrix E_{12} defined by

$$E_{12} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ -a_{21}/a_{11} & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}.$$

In particular, you should verify that

$$A_1 = E_{12}A$$

The next ERO that we would perform replaces the third row by the third row minus an appropriate rescaling of the first row, such that the left-most entry of the third row is made to vanish. In other words, we want to left multiply A_1 by the lower triangular matrix

$$E_{13} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ -a_{31}/a_{11} & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}.$$

to obtain $A_2 = E_{13}A_1$. Proceeding in this fashion, is should be clear how to define lower triangular matrices E_{1j} such that

$$A_{n-1} = E_{1n}E_{1,n-1}\cdots E_{13}E_{12}A$$

is the result of the usual initial $n-1$ EROs applied to A. Moreover, A_{n-1} has all zeroes in the first column except for the top element.

Note that the next sequence of EROs at this point would cause the entries of the 2nd column under the diagonal to vanish. This can be achieved by further left multiplication, this time using lower triangular matrices of the form

$$E_{2j} = \begin{bmatrix} 1 & 0 & 0 & \cdots & & & 0 \\ 0 & 1 & 0 & \cdots & & & 0 \\ 0 & 0 & 1 & \cdots & & & 0 \\ \vdots & \vdots & \vdots & \ddots & & & \vdots \\ 0 & -a_{j2}/a_{22}^{(1)} & 0 & \cdots & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & & & 1 \end{bmatrix},$$

where $a_{22}^{(1)}$ denotes the "2-2 entry" of A_1 and we again assume that $a_{22}^{(1)} \neq 0$.

Continuing in the fashion with EROs that zero out the entries in the third column under the diagonal, then the fourth column, and so on, we will eventually obtain an upper triangular matrix U that is row-equivalent to A. The above procedure thus results in

$$U = E_{n-1,n}E_{n-2,n-1}E_{n-2,n} \cdots E_{23}E_{1n} \cdots E_{12}A\,, \tag{1}$$

where each E_{ij} was specifically formed to be lower triangular.

The next step is to notice that the inverse of each E_{ij} is again a lower triangular matrix of the same form, which follows by changing the sign of the only non-vanishing off-diagonal element. E.g.,

$$E_{12}^{-1} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ a_{21}/a_{11} & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}.$$

Thus, using these inverses for each E_{ij}, one can solve for A in Equation (1) to obtain

$$A = LU$$

with

$$L = E_{12}^{-1}E_{13}^{-1} \cdots E_{1n}^{-1}E_{23}^{-1} \cdots E_{n-1,n}^{-1}\,.$$

Thus, appealing to Theorem 2.1 in Section 2 above, it follows that L is lower triangular since it is a product of lower triangular matrices. We have therefore obtained an LU-factorization for the matrix A.

When this procedure is implemented in a computer program, one does not actually perform such a large number of matrix multiplications with matrices whose entries are mostly zero. Instead, it is easy to implement this procedure so that it only modifies one row of the matrix at a time. Then, by thinking about how the matrix product can be defined row by row, this leads to an even more efficient method for computing LU-factorizations.