3.10 THE LU FACTORIZATION

We have now come full circle, and we are back to where the text began—solving a nonsingular system of linear equations using Gaussian elimination with back substitution. This time, however, the goal is to describe and understand the process in the context of matrices.

If $\mathbf{Ax} = \mathbf{b}$ is a nonsingular system, then the object of Gaussian elimination is to reduce $\mathbf{A}$ to an upper-triangular matrix using elementary row operations. If no zero pivots are encountered, then row interchanges are not necessary, and the reduction can be accomplished by using only elementary row operations of Type III. For example, consider reducing the matrix

$$\mathbf{A} = \begin{pmatrix} 2 & 2 & 2 \\ 4 & 7 & 7 \\ 6 & 18 & 22 \end{pmatrix}$$

to upper-triangular form as shown below:

$$\begin{aligned} \begin{pmatrix} 2 & 2 & 2 \\ 4 & 7 & 7 \\ 6 & 18 & 22 \end{pmatrix} \begin{matrix} \\ R_2 - 2R_1 \\ R_3 - 3R_1 \end{matrix} &\longrightarrow \begin{pmatrix} 2 & 2 & 2 \\ 0 & 3 & 3 \\ 0 & 12 & 16 \end{pmatrix} \begin{matrix} \\ \\ R_3 - 4R_2 \end{matrix} \\ &\longrightarrow \begin{pmatrix} 2 & 2 & 2 \\ 0 & 3 & 3 \\ 0 & 0 & 4 \end{pmatrix} = \mathbf{U}. \end{aligned} \tag{3.10.1}$$

We learned in the previous section that each of these Type III operations can be executed by means of a left-hand multiplication with the corresponding elementary matrix $\mathbf{G}_i$, and the product of all of these $\mathbf{G}_i$'s is

$$\mathbf{G}_3\mathbf{G}_2\mathbf{G}_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -4 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -3 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 5 & -4 & 1 \end{pmatrix}.$$

In other words, $\mathbf{G}_3\mathbf{G}_2\mathbf{G}_1\mathbf{A} = \mathbf{U}$, so that $\mathbf{A} = \mathbf{G}_1^{-1}\mathbf{G}_2^{-1}\mathbf{G}_3^{-1}\mathbf{U} = \mathbf{LU}$, where $\mathbf{L}$ is the lower-triangular matrix

$$\mathbf{L} = \mathbf{G}_1^{-1}\mathbf{G}_2^{-1}\mathbf{G}_3^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 4 & 1 \end{pmatrix}.$$

Thus $\mathbf{A} = \mathbf{LU}$ is a product of a lower-triangular matrix $\mathbf{L}$ and an upper-triangular matrix $\mathbf{U}$. Naturally, this is called an **LU *factorization*** of $\mathbf{A}$.

Observe that $\mathbf{U}$ is the end product of Gaussian elimination and has the pivots on its diagonal, while $\mathbf{L}$ has 1's on its diagonal. Moreover, $\mathbf{L}$ has the remarkable property that below its diagonal, *each entry* ℓ_{ij} *is precisely the multiplier used in the elimination* (3.10.1) *to annihilate the* (i,j)*-position.*

This is characteristic of what happens in general. To develop the general theory, it's convenient to introduce the concept of an ***elementary lower-triangular matrix,*** which is defined to be an $n \times n$ triangular matrix of the form

$$\mathbf{T}_k = \mathbf{I} - \mathbf{c}_k \mathbf{e}_k^T,$$

where $\mathbf{c}_k$ is a column with zeros in the first k positions. In particular, if

$$\mathbf{c}_k = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ \mu_{k+1} \\ \vdots \\ \mu_n \end{pmatrix}, \quad \text{then} \quad \mathbf{T}_k = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & -\mu_{k+1} & 1 & \cdots & 0 \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -\mu_n & 0 & \cdots & 1 \end{pmatrix}. \tag{3.10.2}$$

By observing that $\mathbf{e}_k^T \mathbf{c}_k = 0$, the formula for the inverse of an elementary matrix given in (3.9.1) produces

$$\mathbf{T}_k^{-1} = \mathbf{I} + \mathbf{c}_k \mathbf{e}_k^T = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & \mu_{k+1} & 1 & \cdots & 0 \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mu_n & 0 & \cdots & 1 \end{pmatrix}, \tag{3.10.3}$$

which is also an elementary lower-triangular matrix. The utility of elementary lower-triangular matrices lies in the fact that all of the Type III row operations needed to annihilate the entries below the k^{th} pivot can be accomplished with one multiplication by $\mathbf{T}_k$. If

$$\mathbf{A}_{k-1} = \begin{pmatrix} * & * & \cdots & \alpha_1 & * & \cdots & * \\ 0 & * & \cdots & \alpha_2 & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \alpha_k & * & \cdots & * \\ 0 & 0 & \cdots & \alpha_{k+1} & * & \cdots & * \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \alpha_n & * & \cdots & * \end{pmatrix}$$

is the partially triangularized result after $k-1$ elimination steps, then

$$\mathbf{T}_k\mathbf{A}_{k-1} = \left(\mathbf{I} - \mathbf{c}_k\mathbf{e}_k^T\right)\mathbf{A}_{k-1} = \mathbf{A}_{k-1} - \mathbf{c}_k\mathbf{e}_k^T\mathbf{A}_{k-1}$$

$$= \begin{pmatrix} * & * & \cdots & \alpha_1 & * & \cdots & * \\ 0 & * & \cdots & \alpha_2 & * & \cdots & * \\ \vdots & \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \alpha_k & * & \cdots & * \\ 0 & 0 & \cdots & 0 & * & \cdots & * \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & * & \cdots & * \end{pmatrix}, \qquad \text{where} \quad \mathbf{c}_k = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \alpha_{k+1}/\alpha_k \\ \vdots \\ \alpha_n/\alpha_k \end{pmatrix}$$

contains the multipliers used to annihilate those entries below α_k. Notice that $\mathbf{T}_k$ does not alter the first $k-1$ columns of $\mathbf{A}_{k-1}$ because $\mathbf{e}_k^T[\mathbf{A}_{k-1}]_{*j} = 0$ whenever $j \le k-1$. Therefore, if no row interchanges are required, then reducing $\mathbf{A}$ to an upper-triangular matrix $\mathbf{U}$ by Gaussian elimination is equivalent to executing a sequence of $n-1$ left-hand multiplications with elementary lower-triangular matrices. That is, $\mathbf{T}_{n-1}\cdots\mathbf{T}_2\mathbf{T}_1\mathbf{A} = \mathbf{U}$, and hence

$$\mathbf{A} = \mathbf{T}_1^{-1}\mathbf{T}_2^{-1}\cdots\mathbf{T}_{n-1}^{-1}\mathbf{U}. \tag{3.10.4}$$

Making use of the fact that $\mathbf{e}_j^T\mathbf{c}_k = 0$ whenever $j \le k$ and applying (3.10.3) reveals that

$$\begin{aligned} \mathbf{T}_1^{-1}\mathbf{T}_2^{-1}\cdots\mathbf{T}_{n-1}^{-1} &= \left(\mathbf{I} + \mathbf{c}_1\mathbf{e}_1^T\right)\left(\mathbf{I} + \mathbf{c}_2\mathbf{e}_2^T\right)\cdots\left(\mathbf{I} + \mathbf{c}_{n-1}\mathbf{e}_{n-1}^T\right) \\ &= \mathbf{I} + \mathbf{c}_1\mathbf{e}_1^T + \mathbf{c}_2\mathbf{e}_2^T + \cdots + \mathbf{c}_{n-1}\mathbf{e}_{n-1}^T. \end{aligned} \tag{3.10.5}$$

By observing that

$$\mathbf{c}_k\mathbf{e}_k^T = \begin{pmatrix} 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & \ell_{k+1,k} & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \ell_{nk} & 0 & \cdots & 0 \end{pmatrix},$$

where the ℓ_{ik}'s are the multipliers used at the k^{th} stage to annihilate the entries below the k^{th} pivot, it now follows from (3.10.4) and (3.10.5) that

$$\mathbf{A} = \mathbf{LU},$$

where

$$\mathbf{L} = \mathbf{I} + \mathbf{c}_1\mathbf{e}_1^T + \mathbf{c}_2\mathbf{e}_2^T + \cdots + \mathbf{c}_{n-1}\mathbf{e}_{n-1}^T = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ \ell_{21} & 1 & 0 & \cdots & 0 \\ \ell_{31} & \ell_{32} & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \ell_{n1} & \ell_{n2} & \ell_{n3} & \cdots & 1 \end{pmatrix} \tag{3.10.6}$$

is the lower-triangular matrix with 1's on the diagonal, and where ℓ_{ij} is precisely the multiplier used to annihilate the (i, j)-position during Gaussian elimination. Thus the factorization $\mathbf{A} = \mathbf{LU}$ can be viewed as the matrix formulation of Gaussian elimination, with the understanding that no row interchanges are used.

LU Factorization

If $\mathbf{A}$ is an $n \times n$ matrix such that a zero pivot is never encountered when applying Gaussian elimination with Type III operations, then $\mathbf{A}$ can be factored as the product $\mathbf{A} = \mathbf{LU}$, where the following hold.

- $\mathbf{L}$ is lower triangular and $\mathbf{U}$ is upper triangular. (3.10.7)
- $\ell_{ii} = 1$ and $u_{ii} \neq 0$ for each $i = 1, 2, \ldots, n$. (3.10.8)
- Below the diagonal of $\mathbf{L}$, the entry ℓ_{ij} is the multiple of row j that is subtracted from row i in order to annihilate the (i, j)-position during Gaussian elimination.
- $\mathbf{U}$ is the final result of Gaussian elimination applied to $\mathbf{A}$.
- The matrices $\mathbf{L}$ and $\mathbf{U}$ are uniquely determined by properties (3.10.7) and (3.10.8).

The decomposition of $\mathbf{A}$ into $\mathbf{A} = \mathbf{LU}$ is called the ***LU factorization of*** $\mathbf{A}$, and the matrices $\mathbf{L}$ and $\mathbf{U}$ are called the ***LU factors of*** $\mathbf{A}$.

Proof. Except for the statement concerning the uniqueness of the LU factors, each point has already been established. To prove uniqueness, observe that LU factors must be nonsingular because they have nonzero diagonals. If $\mathbf{L}_1\mathbf{U}_1 = \mathbf{A} = \mathbf{L}_2\mathbf{U}_2$ are two LU factorizations for $\mathbf{A}$, then

$$\mathbf{L}_2^{-1}\mathbf{L}_1 = \mathbf{U}_2\mathbf{U}_1^{-1}. \tag{3.10.9}$$

Notice that $\mathbf{L}_2^{-1}\mathbf{L}_1$ is lower triangular, while $\mathbf{U}_2\mathbf{U}_1^{-1}$ is upper triangular because the inverse of a matrix that is upper (lower) triangular is again upper (lower) triangular, and because the product of two upper (lower) triangular matrices is also upper (lower) triangular. Consequently, (3.10.9) implies $\mathbf{L}_2^{-1}\mathbf{L}_1 = \mathbf{D} = \mathbf{U}_2\mathbf{U}_1^{-1}$ must be a diagonal matrix. However, $[\mathbf{L}_2]_{ii} = 1 = [\mathbf{L}_2^{-1}]_{ii}$, so it must be the case that $\mathbf{L}_2^{-1}\mathbf{L}_1 = \mathbf{I} = \mathbf{U}_2\mathbf{U}_1^{-1}$, and thus $\mathbf{L}_1 = \mathbf{L}_2$ and $\mathbf{U}_1 = \mathbf{U}_2$. ■

Example 3.10.1

Once $\mathbf{L}$ and $\mathbf{U}$ are known, there is usually no need to manipulate with $\mathbf{A}$. This together with the fact that the multipliers used in Gaussian elimination occur in just the right places in $\mathbf{L}$ means that $\mathbf{A}$ can be successively overwritten with the information in $\mathbf{L}$ and $\mathbf{U}$ as Gaussian elimination evolves. The rule is to store the multiplier ℓ_{ij} in the position it annihilates—namely, the (i,j)-position of the array. For a 3×3 matrix, the result looks like this:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \xrightarrow{\text{Type III operations}} \begin{pmatrix} u_{11} & u_{12} & u_{13} \\ \ell_{21} & u_{22} & u_{23} \\ \ell_{31} & \ell_{32} & u_{33} \end{pmatrix}.$$

For example, generating the LU factorization of

$$\mathbf{A} = \begin{pmatrix} 2 & 2 & 2 \\ 4 & 7 & 7 \\ 6 & 18 & 22 \end{pmatrix}$$

by successively overwriting a single 3×3 array would evolve as shown below:

$$\begin{pmatrix} 2 & 2 & 2 \\ 4 & 7 & 7 \\ 6 & 18 & 22 \end{pmatrix} \begin{matrix} R_2 - 2R_1 \\ R_3 - 3R_1 \end{matrix} \longrightarrow \begin{pmatrix} 2 & 2 & 2 \\ \textcircled{2} & 3 & 3 \\ \textcircled{3} & 12 & 16 \end{pmatrix} \begin{matrix} \\ R_3 - 4R_2 \end{matrix} \longrightarrow \begin{pmatrix} 2 & 2 & 2 \\ \textcircled{2} & 3 & 3 \\ \textcircled{3} & \textcircled{4} & 4 \end{pmatrix}.$$

Thus

$$\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 4 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{U} = \begin{pmatrix} 2 & 2 & 2 \\ 0 & 3 & 3 \\ 0 & 0 & 4 \end{pmatrix}.$$

This is an important feature in practical computation because it guarantees that an LU factorization requires no more computer memory than that required to store the original matrix $\mathbf{A}$.

Once the LU factors for a nonsingular matrix $\mathbf{A}_{n\times n}$ have been obtained, it's relatively easy to solve a linear system $\mathbf{Ax} = \mathbf{b}$. By rewriting $\mathbf{Ax} = \mathbf{b}$ as

$$\mathbf{L}(\mathbf{Ux}) = \mathbf{b} \quad \text{and setting} \quad \mathbf{y} = \mathbf{Ux},$$

we see that $\mathbf{Ax} = \mathbf{b}$ is equivalent to the two triangular systems

$$\mathbf{Ly} = \mathbf{b} \quad \text{and} \quad \mathbf{Ux} = \mathbf{y}.$$

First, the lower-triangular system $\mathbf{Ly} = \mathbf{b}$ is solved for $\mathbf{y}$ by ***forward substitution.*** That is, if

$$\begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ \ell_{21} & 1 & 0 & \cdots & 0 \\ \ell_{31} & \ell_{32} & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \ell_{n1} & \ell_{n2} & \ell_{n3} & \cdots & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ \vdots \\ b_n \end{pmatrix},$$

set

$$y_1 = b_1, \quad y_2 = b_2 - \ell_{21}y_1, \quad y_3 = b_3 - \ell_{31}y_1 - \ell_{32}y_2, \quad \text{etc.}$$

The forward substitution algorithm can be written more concisely as

$$y_1 = b_1 \quad \text{and} \quad y_i = b_i - \sum_{k=1}^{i-1} \ell_{ik}y_k \quad \text{for} \quad i = 2, 3, \ldots, n. \tag{3.10.10}$$

After $\mathbf{y}$ is known, the upper-triangular system $\mathbf{Ux} = \mathbf{y}$ is solved using the standard back substitution procedure by starting with $x_n = y_n/u_{nn}$, and setting

$$x_i = \frac{1}{u_{ii}}\left(y_i - \sum_{k=i+1}^{n} u_{ik}x_k\right) \quad \text{for} \quad i = n-1,\ n-2,\ \ldots, 1. \tag{3.10.11}$$

It can be verified that only n^2 multiplications/divisions and $n^2 - n$ additions/subtractions are required when (3.10.10) and (3.10.11) are used to solve the two triangular systems $\mathbf{Ly} = \mathbf{b}$ and $\mathbf{Ux} = \mathbf{y}$, so it's relatively cheap to solve $\mathbf{Ax} = \mathbf{b}$ once $\mathbf{L}$ and $\mathbf{U}$ are known—recall from §1.2 that these operation counts are about $n^3/3$ when we start from scratch.

If only one system $\mathbf{Ax} = \mathbf{b}$ is to be solved, then there is no significant difference between the technique of reducing the augmented matrix $[\mathbf{A}|\mathbf{b}]$ to a row echelon form and the LU factorization method presented here. However, suppose it becomes necessary to later solve other systems $\mathbf{Ax} = \tilde{\mathbf{b}}$ with the same coefficient matrix but with different right-hand sides, which is frequently the case in applied work. If the LU factors of $\mathbf{A}$ were computed and saved when the original system was solved, then they need not be recomputed, and the solutions to all subsequent systems $\mathbf{Ax} = \tilde{\mathbf{b}}$ are therefore relatively cheap to obtain. That is, the operation counts for each subsequent system are on the order of n^2, whereas these counts would be on the order of $n^3/3$ if we would start from scratch each time.

Summary

- To solve a nonsingular system $\mathbf{Ax} = \mathbf{b}$ using the LU factorization $\mathbf{A} = \mathbf{LU}$, first solve $\mathbf{Ly} = \mathbf{b}$ for $\mathbf{y}$ with the forward substitution algorithm (3.10.10), and then solve $\mathbf{Ux} = \mathbf{y}$ for $\mathbf{x}$ with the back substitution procedure (3.10.11).

- The advantage of this approach is that once the LU factors for $\mathbf{A}$ have been computed, any other linear system $\mathbf{Ax} = \tilde{\mathbf{b}}$ can be solved with only n^2 multiplications/divisions and $n^2 - n$ additions/subtractions.

Example 3.10.2

Problem 1: Use the LU factorization of $\mathbf{A}$ to solve $\mathbf{Ax} = \mathbf{b}$, where

$$\mathbf{A} = \begin{pmatrix} 2 & 2 & 2 \\ 4 & 7 & 7 \\ 6 & 18 & 22 \end{pmatrix} \quad \text{and} \quad \mathbf{b} = \begin{pmatrix} 12 \\ 24 \\ 12 \end{pmatrix}.$$

Problem 2: Suppose that after solving the original system new information is received that changes $\mathbf{b}$ to

$$\tilde{\mathbf{b}} = \begin{pmatrix} 6 \\ 24 \\ 70 \end{pmatrix}.$$

Use the LU factors of $\mathbf{A}$ to solve the updated system $\mathbf{Ax} = \tilde{\mathbf{b}}$.

Solution 1: The LU factors of the coefficient matrix were determined in Example 3.10.1 to be

$$\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 4 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{U} = \begin{pmatrix} 2 & 2 & 2 \\ 0 & 3 & 3 \\ 0 & 0 & 4 \end{pmatrix}.$$

The strategy is to set $\mathbf{Ux} = \mathbf{y}$ and solve $\mathbf{Ax} = \mathbf{L}(\mathbf{Ux}) = \mathbf{b}$ by solving the two triangular systems

$$\mathbf{Ly} = \mathbf{b} \quad \text{and} \quad \mathbf{Ux} = \mathbf{y}.$$

First solve the lower-triangular system $\mathbf{Ly} = \mathbf{b}$ by using forward substitution:

$$\begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 4 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} 12 \\ 24 \\ 12 \end{pmatrix} \quad \Longrightarrow \quad \begin{aligned} y_1 &= 12, \\ y_2 &= 24 - 2y_1 = 0, \\ y_3 &= 12 - 3y_1 - 4y_2 = -24. \end{aligned}$$

Now use back substitution to solve the upper-triangular system $\mathbf{Ux} = \mathbf{y}$:

$$\begin{pmatrix} 2 & 2 & 2 \\ 0 & 3 & 3 \\ 0 & 0 & 4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 12 \\ 0 \\ -24 \end{pmatrix} \quad \Longrightarrow \quad \begin{aligned} x_1 &= (12 - 2x_2 - 2x_3)/2 = 6, \\ x_2 &= (0 - 3x_3)/3 = 6, \\ x_3 &= -24/4 = -6. \end{aligned}$$

Solution 2: To solve the updated system $\mathbf{Ax} = \tilde{\mathbf{b}}$, simply repeat the forward and backward substitution steps with $\mathbf{b}$ replaced by $\tilde{\mathbf{b}}$. Solving $\mathbf{Ly} = \tilde{\mathbf{b}}$ with forward substitution gives the following:

$$\begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 3 & 4 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \end{pmatrix} = \begin{pmatrix} 6 \\ 24 \\ 70 \end{pmatrix} \quad \Longrightarrow \quad \begin{aligned} y_1 &= 6, \\ y_2 &= 24 - 2y_1 = 12, \\ y_3 &= 70 - 3y_1 - 4y_2 = 4. \end{aligned}$$

Using back substitution to solve $\mathbf{Ux} = \mathbf{y}$ gives the following updated solution:

$$\begin{pmatrix} 2 & 2 & 2 \\ 0 & 3 & 3 \\ 0 & 0 & 4 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} 6 \\ 12 \\ 4 \end{pmatrix} \quad \Longrightarrow \quad \begin{aligned} x_1 &= (6 - 2x_2 - 2x_3)/2 = -1, \\ x_2 &= (12 - 3x_3)/3 = 3, \\ x_3 &= 4/4 = 1. \end{aligned}$$

Example 3.10.3

Computing $\mathbf{A}^{-1}$. Although matrix inversion is not used for solving $\mathbf{Ax} = \mathbf{b}$, there are a few applications where explicit knowledge of $\mathbf{A}^{-1}$ is desirable.

Problem: Explain how to use the LU factors of a nonsingular matrix $\mathbf{A}_{n\times n}$ to compute $\mathbf{A}^{-1}$ efficiently.

Solution: The strategy is to solve the matrix equation $\mathbf{AX} = \mathbf{I}$. Recall from (3.5.5) that $\mathbf{AA}^{-1} = \mathbf{I}$ implies $\mathbf{A}[\mathbf{A}^{-1}]_{*j} = \mathbf{e}_j$, so the j^{th} column of $\mathbf{A}^{-1}$ is the solution of a system $\mathbf{Ax}_j = \mathbf{e}_j$. Each of these n systems has the same coefficient matrix, so, once the LU factors for $\mathbf{A}$ are known, each system $\mathbf{Ax}_j = \mathbf{LUx}_j = \mathbf{e}_j$ can be solved by the standard two-step process.

(1) Set $\mathbf{y}_j = \mathbf{Ux}_j$, and solve $\mathbf{Ly}_j = \mathbf{e}_j$ for $\mathbf{y}_j$ by forward substitution.

(2) Solve $\mathbf{Ux}_j = \mathbf{y}_j$ for $\mathbf{x}_j = [\mathbf{A}^{-1}]_{*j}$ by back substitution.

This method has at least two advantages: it's efficient, and any code written to solve $\mathbf{Ax} = \mathbf{b}$ can also be used to compute $\mathbf{A}^{-1}$.

Note: A tempting alternate solution might be to use the fact $\mathbf{A}^{-1} = (\mathbf{LU})^{-1} = \mathbf{U}^{-1}\mathbf{L}^{-1}$. But computing $\mathbf{U}^{-1}$ and $\mathbf{L}^{-1}$ explicitly and then multiplying the results is not as computationally efficient as the method just described.

Not all nonsingular matrices possess an LU factorization. For example, there is clearly no nonzero value of u_{11} that will satisfy

$$\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \ell_{21} & 1 \end{pmatrix} \begin{pmatrix} u_{11} & u_{12} \\ 0 & u_{22} \end{pmatrix}.$$

The problem here is the zero pivot in the (1,1)-position. Our development of the LU factorization using elementary lower-triangular matrices shows that if no zero pivots emerge, then no row interchanges are necessary, and the LU factorization can indeed be carried to completion. The converse is also true (its proof is left as an exercise), so we can say that *a nonsingular matrix* $\mathbf{A}$ *has an LU factorization if and only if a zero pivot does not emerge during row reduction to upper-triangular form with Type III operations.*

Although it is a bit more theoretical, there is another interesting way to characterize the existence of LU factors. This characterization is given in terms of the ***leading principal submatrices of*** $\mathbf{A}$ that are defined to be those submatrices taken from the upper-left-hand corner of $\mathbf{A}$. That is,

$$\mathbf{A}_1 = \begin{pmatrix} a_{11} \end{pmatrix}, \quad \mathbf{A}_2 = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, \ldots, \mathbf{A}_k = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1k} \\ a_{21} & a_{22} & \cdots & a_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k1} & a_{k2} & \cdots & a_{kk} \end{pmatrix}, \ldots .$$

Existence of LU Factors

Each of the following statements is equivalent to saying that a nonsingular matrix $\mathbf{A}_{n\times n}$ possesses an LU factorization.

- A zero pivot does not emerge during row reduction to upper-triangular form with Type III operations.
- Each leading principal submatrix $\mathbf{A}_k$ is nonsingular. (3.10.12)

Proof. We will prove the statement concerning the leading principal submatrices and leave the proof concerning the nonzero pivots as an exercise. Assume first that $\mathbf{A}$ possesses an LU factorization and partition $\mathbf{A}$ as

$$\mathbf{A} = \mathbf{LU} = \begin{pmatrix} \mathbf{L}_{11} & \mathbf{0} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{pmatrix} \begin{pmatrix} \mathbf{U}_{11} & \mathbf{U}_{12} \\ \mathbf{0} & \mathbf{U}_{22} \end{pmatrix} = \begin{pmatrix} \mathbf{L}_{11}\mathbf{U}_{11} & * \\ * & * \end{pmatrix},$$

where $\mathbf{L}_{11}$ and $\mathbf{U}_{11}$ are each $k \times k$. Thus $\mathbf{A}_k = \mathbf{L}_{11}\mathbf{U}_{11}$ must be nonsingular because $\mathbf{L}_{11}$ and $\mathbf{U}_{11}$ are each nonsingular—they are triangular with nonzero diagonal entries. Conversely, suppose that each leading principal submatrix in $\mathbf{A}$ is nonsingular. Use induction to prove that each $\mathbf{A}_k$ possesses an LU factorization. For $k = 1$, this statement is clearly true because if $\mathbf{A}_1 = (a_{11})$ is nonsingular, then $\mathbf{A}_1 = (1)(a_{11})$ is its LU factorization. Now assume that $\mathbf{A}_k$ has an LU factorization and show that this together with the nonsingularity condition implies $\mathbf{A}_{k+1}$ must also possess an LU factorization. If $\mathbf{A}_k = \mathbf{L}_k\mathbf{U}_k$ is the LU factorization for $\mathbf{A}_k$, then $\mathbf{A}_k^{-1} = \mathbf{U}_k^{-1}\mathbf{L}_k^{-1}$ so that

$$\mathbf{A}_{k+1} = \begin{pmatrix} \mathbf{A}_k & \mathbf{b} \\ \mathbf{c}^T & \alpha_{k+1} \end{pmatrix} = \begin{pmatrix} \mathbf{L}_k & \mathbf{0} \\ \mathbf{c}^T\mathbf{U}_k^{-1} & 1 \end{pmatrix} \begin{pmatrix} \mathbf{U}_k & \mathbf{L}_k^{-1}\mathbf{b} \\ \mathbf{0} & \alpha_{k+1} - \mathbf{c}^T\mathbf{A}_k^{-1}\mathbf{b} \end{pmatrix}, \quad (3.10.13)$$

where $\mathbf{c}^T$ and $\mathbf{b}$ contain the first k components of $\mathbf{A}_{k+1*}$ and $\mathbf{A}_{*k+1}$, respectively. Observe that this is the LU factorization for $\mathbf{A}_{k+1}$ because

$$\mathbf{L}_{k+1} = \begin{pmatrix} \mathbf{L}_k & \mathbf{0} \\ \mathbf{c}^T\mathbf{U}_k^{-1} & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{U}_{k+1} = \begin{pmatrix} \mathbf{U}_k & \mathbf{L}_k^{-1}\mathbf{b} \\ \mathbf{0} & \alpha_{k+1} - \mathbf{c}^T\mathbf{A}_k^{-1}\mathbf{b} \end{pmatrix}$$

are lower- and upper-triangular matrices, respectively, and $\mathbf{L}$ has 1's on its diagonal while the diagonal entries of $\mathbf{U}$ are nonzero. The fact that

$$\alpha_{k+1} - \mathbf{c}^T\mathbf{A}_k^{-1}\mathbf{b} \neq 0$$

follows because $\mathbf{A}_{k+1}$ and $\mathbf{L}_{k+1}$ are each nonsingular, so $\mathbf{U}_{k+1} = \mathbf{L}_{k+1}^{-1}\mathbf{A}_{k+1}$ must also be nonsingular. Therefore, the nonsingularity of the leading principal submatrices implies that each $\mathbf{A}_k$ possesses an LU factorization, and hence $\mathbf{A}_n = \mathbf{A}$ must have an LU factorization. ■

Up to this point we have avoided dealing with row interchanges because if a row interchange is needed to remove a zero pivot, then no LU factorization is possible. However, we know from the discussion in §1.5 that practical computation necessitates row interchanges in the form of partial pivoting. So even if no zero pivots emerge, it is usually the case that we must still somehow account for row interchanges.

To understand the effects of row interchanges in the framework of an LU decomposition, let $\mathbf{T}_k = \mathbf{I} - \mathbf{c}_k\mathbf{e}_k^T$ be an elementary lower-triangular matrix as described in (3.10.2), and let $\mathbf{E} = \mathbf{I} - \mathbf{u}\mathbf{u}^T$ with $\mathbf{u} = \mathbf{e}_{k+i} - \mathbf{e}_{k+j}$ be the Type I elementary interchange matrix associated with an interchange of rows $k+i$ and $k+j$. Notice that $\mathbf{e}_k^T\mathbf{E} = \mathbf{e}_k^T$ because $\mathbf{e}_k^T$ has 0's in positions $k+i$ and $k+j$. This together with the fact that $\mathbf{E}^2 = \mathbf{I}$ guarantees

$$\mathbf{E}\mathbf{T}_k\mathbf{E} = \mathbf{E}^2 - \mathbf{E}\mathbf{c}_k\mathbf{e}_k^T\mathbf{E} = \mathbf{I} - \tilde{\mathbf{c}}_k\mathbf{e}_k^T, \quad \text{where} \quad \tilde{\mathbf{c}}_k = \mathbf{E}\mathbf{c}_k.$$

In other words, the matrix

$$\tilde{\mathbf{T}}_k = \mathbf{E}\mathbf{T}_k\mathbf{E} = \mathbf{I} - \tilde{\mathbf{c}}_k\mathbf{e}_k^T \tag{3.10.14}$$

is also an elementary lower-triangular matrix, and $\tilde{\mathbf{T}}_k$ agrees with $\mathbf{T}_k$ in all positions except that the multipliers μ_{k+i} and μ_{k+j} have traded places. As before, assume we are row reducing an $n \times n$ nonsingular matrix $\mathbf{A}$, but suppose that an interchange of rows $k+i$ and $k+j$ is necessary immediately after the k^{th} stage so that the sequence of left-hand multiplications $\mathbf{E}\mathbf{T}_k\mathbf{T}_{k-1}\cdots\mathbf{T}_1$ is applied to $\mathbf{A}$. Since $\mathbf{E}^2 = \mathbf{I}$, we may insert $\mathbf{E}^2$ to the right of each $\mathbf{T}$ to obtain

$$\begin{aligned}
\mathbf{E}\mathbf{T}_k\mathbf{T}_{k-1}\cdots\mathbf{T}_1 &= \mathbf{E}\mathbf{T}_k\mathbf{E}^2\mathbf{T}_{k-1}\mathbf{E}^2\cdots\mathbf{E}^2\mathbf{T}_1\mathbf{E}^2 \\
&= (\mathbf{E}\mathbf{T}_k\mathbf{E})(\mathbf{E}\mathbf{T}_{k-1}\mathbf{E})\cdots(\mathbf{E}\mathbf{T}_1\mathbf{E})\,\mathbf{E} \\
&= \tilde{\mathbf{T}}_k\tilde{\mathbf{T}}_{k-1}\cdots\tilde{\mathbf{T}}_1\mathbf{E}.
\end{aligned}$$

In such a manner, the necessary interchange matrices $\mathbf{E}$ can be "factored" to the far-right-hand side, and the matrices $\tilde{\mathbf{T}}$ retain the desirable feature of being elementary lower-triangular matrices. Furthermore, (3.10.14) implies that $\tilde{\mathbf{T}}_k\tilde{\mathbf{T}}_{k-1}\cdots\tilde{\mathbf{T}}_1$ differs from $\mathbf{T}_k\mathbf{T}_{k-1}\cdots\mathbf{T}_1$ only in the sense that the multipliers in rows $k+i$ and $k+j$ have traded places. Therefore, row interchanges in Gaussian elimination can be accounted for by writing $\tilde{\mathbf{T}}_{n-1}\cdots\tilde{\mathbf{T}}_2\tilde{\mathbf{T}}_1\mathbf{P}\mathbf{A} = \mathbf{U}$, where $\mathbf{P}$ is the product of all elementary interchange matrices used during the reduction and where the $\tilde{\mathbf{T}}_k$'s are elementary lower-triangular matrices in which the multipliers have been permuted according to the row interchanges that were implemented. Since all of the $\tilde{\mathbf{T}}_k$'s are elementary lower-triangular matrices, we may proceed along the same lines discussed in (3.10.4)—(3.10.6) to obtain

$$\mathbf{P}\mathbf{A} = \mathbf{L}\mathbf{U}, \quad \text{where} \quad \mathbf{L} = \tilde{\mathbf{T}}_1^{-1}\tilde{\mathbf{T}}_2^{-1}\cdots\tilde{\mathbf{T}}_{n-1}^{-1}. \tag{3.10.15}$$

When row interchanges are allowed, zero pivots can always be avoided when the original matrix $\mathbf{A}$ is nonsingular. Consequently, we may conclude that *for every nonsingular matrix* $\mathbf{A}$, *there exists a* ***permutation matrix*** $\mathbf{P}$ *(a product of elementary interchange matrices) such that* $\mathbf{PA}$ *has an LU factorization*. Furthermore, because of the observation in (3.10.14) concerning how the multipliers in $\mathbf{T}_k$ and $\tilde{\mathbf{T}}_k$ trade places when a row interchange occurs, and because

$$\tilde{\mathbf{T}}_k^{-1} = \left(\mathbf{I} - \tilde{\mathbf{c}}_k\mathbf{e}_k^T\right)^{-1} = \mathbf{I} + \tilde{\mathbf{c}}_k\mathbf{e}_k^T,$$

it is not difficult to see that the same line of reasoning used to arrive at (3.10.6) can be applied to conclude that the multipliers in the matrix $\mathbf{L}$ in (3.10.15) are permuted according to the row interchanges that are executed. More specifically, *if rows* k *and* $k+i$ *are interchanged to create the* k^{th} *pivot, then the multipliers*

$$\begin{pmatrix} \ell_{k1} & \ell_{k2} & \cdots & \ell_{k,k-1} \end{pmatrix} \quad \textit{and} \quad \begin{pmatrix} \ell_{k+i,1} & \ell_{k+i,2} & \cdots & \ell_{k+i,k-1} \end{pmatrix}$$

trade places in the formation of $\mathbf{L}$.

This means that we can proceed just as in the case when no interchanges are used and successively overwrite the array originally containing $\mathbf{A}$ with each multiplier replacing the position it annihilates. Whenever a row interchange occurs, the corresponding multipliers will be correctly interchanged as well. The permutation matrix $\mathbf{P}$ is simply the cumulative record of the various interchanges used, and the information in $\mathbf{P}$ is easily accounted for by a simple technique that is illustrated in the following example.

Example 3.10.4

Problem: Use partial pivoting on the matrix

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & -3 & 4 \\ 4 & 8 & 12 & -8 \\ 2 & 3 & 2 & 1 \\ -3 & -1 & 1 & -4 \end{pmatrix}$$

and determine the LU decomposition $\mathbf{PA} = \mathbf{LU}$, where $\mathbf{P}$ is the associated permutation matrix.

Solution: As explained earlier, the strategy is to successively overwrite the array $\mathbf{A}$ with components from $\mathbf{L}$ and $\mathbf{U}$. For the sake of clarity, the multipliers ℓ_{ij} are shown in boldface type. Adjoin a "permutation counter column" $\mathbf{p}$ that is initially set to the natural order 1,2,3,4. Permuting components of $\mathbf{p}$ as the various row interchanges are executed will accumulate the desired permutation. The matrix $\mathbf{P}$ is obtained by executing the final permutation residing in $\mathbf{p}$ to the rows of an appropriate size identity matrix:

$$[\mathbf{A}|\mathbf{p}] = \left(\begin{array}{rrrr|r} 1 & 2 & -3 & 4 & 1 \\ 4 & 8 & 12 & -8 & 2 \\ 2 & 3 & 2 & 1 & 3 \\ -3 & -1 & 1 & -4 & 4 \end{array}\right) \longrightarrow \left(\begin{array}{rrrr|r} 4 & 8 & 12 & -8 & 2 \\ 1 & 2 & -3 & 4 & 1 \\ 2 & 3 & 2 & 1 & 3 \\ -3 & -1 & 1 & -4 & 4 \end{array}\right)$$

$$\longrightarrow \left(\begin{array}{rrrr|r} 4 & 8 & 12 & -8 & 2 \\ \mathbf{1/4} & 0 & -6 & 6 & 1 \\ \mathbf{1/2} & -1 & -4 & 5 & 3 \\ \mathbf{-3/4} & 5 & 10 & -10 & 4 \end{array}\right) \longrightarrow \left(\begin{array}{rrrr|r} 4 & 8 & 12 & -8 & 2 \\ \mathbf{-3/4} & 5 & 10 & -10 & 4 \\ \mathbf{1/2} & -1 & -4 & 5 & 3 \\ \mathbf{1/4} & 0 & -6 & 6 & 1 \end{array}\right)$$

$$\longrightarrow \left(\begin{array}{rrrr|r} 4 & 8 & 12 & -8 & 2 \\ \mathbf{-3/4} & 5 & 10 & -10 & 4 \\ \mathbf{1/2} & \mathbf{-1/5} & -2 & 3 & 3 \\ \mathbf{1/4} & \mathbf{0} & -6 & 6 & 1 \end{array}\right) \longrightarrow \left(\begin{array}{rrrr|r} 4 & 8 & 12 & -8 & 2 \\ \mathbf{-3/4} & 5 & 10 & -10 & 4 \\ \mathbf{1/4} & \mathbf{0} & -6 & 6 & 1 \\ \mathbf{1/2} & \mathbf{-1/5} & -2 & 3 & 3 \end{array}\right)$$

$$\longrightarrow \left(\begin{array}{rrrr|r} 4 & 8 & 12 & -8 & 2 \\ \mathbf{-3/4} & 5 & 10 & -10 & 4 \\ \mathbf{1/4} & \mathbf{0} & -6 & 6 & 1 \\ \mathbf{1/2} & \mathbf{-1/5} & \mathbf{1/3} & 1 & 3 \end{array}\right).$$

Therefore,

$$\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -3/4 & 1 & 0 & 0 \\ 1/4 & 0 & 1 & 0 \\ 1/2 & -1/5 & 1/3 & 1 \end{pmatrix}, \quad \mathbf{U} = \begin{pmatrix} 4 & 8 & 12 & -8 \\ 0 & 5 & 10 & -10 \\ 0 & 0 & -6 & 6 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{P} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

It is easy to combine the advantages of partial pivoting with the LU decomposition in order to solve a nonsingular system $\mathbf{Ax} = \mathbf{b}$. Because permutation matrices are nonsingular, the system $\mathbf{Ax} = \mathbf{b}$ is equivalent to

$$\mathbf{PAx} = \mathbf{Pb},$$

and hence we can employ the LU solution techniques discussed earlier to solve this permuted system. That is, if we have already performed the factorization $\mathbf{PA} = \mathbf{LU}$ —as illustrated in Example 3.10.4—then we can solve $\mathbf{Ly} = \mathbf{Pb}$ for $\mathbf{y}$ by forward substitution, and then solve $\mathbf{Ux} = \mathbf{y}$ by back substitution.

It should be evident that the permutation matrix $\mathbf{P}$ is not really needed. All that is necessary is knowledge of the LU factors along with the final permutation contained in the permutation counter column $\mathbf{p}$ illustrated in Example 3.10.4. The column $\tilde{\mathbf{b}} = \mathbf{Pb}$ is simply a rearrangement of the components of $\mathbf{b}$ according to the final permutation shown in $\mathbf{p}$. In other words, the strategy is to first permute $\mathbf{b}$ into $\tilde{\mathbf{b}}$ according to the permutation $\mathbf{p}$, and then solve $\mathbf{Ly} = \tilde{\mathbf{b}}$ followed by $\mathbf{Ux} = \mathbf{y}$.

Example 3.10.5

Problem: Use the LU decomposition obtained with partial pivoting to solve the system $\mathbf{Ax} = \mathbf{b}$, where

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & -3 & 4 \\ 4 & 8 & 12 & -8 \\ 2 & 3 & 2 & 1 \\ -3 & -1 & 1 & -4 \end{pmatrix} \quad \text{and} \quad \mathbf{b} = \begin{pmatrix} 3 \\ 60 \\ 1 \\ 5 \end{pmatrix}.$$

Solution: The LU decomposition with partial pivoting was computed in Example 3.10.4. Permute the components in $\mathbf{b}$ according to the permutation $\mathbf{p} = (2 \quad 4 \quad 1 \quad 3)$, and call the result $\tilde{\mathbf{b}}$. Now solve $\mathbf{Ly} = \tilde{\mathbf{b}}$ by applying forward substitution:

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ -3/4 & 1 & 0 & 0 \\ 1/4 & 0 & 1 & 0 \\ 1/2 & -1/5 & 1/3 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} = \begin{pmatrix} 60 \\ 5 \\ 3 \\ 1 \end{pmatrix} \implies \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} = \begin{pmatrix} 60 \\ 50 \\ -12 \\ -15 \end{pmatrix}.$$

Then solve $\mathbf{Ux} = \mathbf{y}$ by applying back substitution:

$$\begin{pmatrix} 4 & 8 & 12 & -8 \\ 0 & 5 & 10 & -10 \\ 0 & 0 & -6 & 6 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{pmatrix} = \begin{pmatrix} 60 \\ 50 \\ -12 \\ -15 \end{pmatrix} \implies \mathbf{x} = \begin{pmatrix} 12 \\ 6 \\ -13 \\ -15 \end{pmatrix}.$$

LU Factorization with Row Interchanges

- For each nonsingular matrix $\mathbf{A}$, there exists a permutation matrix $\mathbf{P}$ such that $\mathbf{PA}$ possesses an LU factorization $\mathbf{PA} = \mathbf{LU}$.

- To compute $\mathbf{L}$, $\mathbf{U}$, and $\mathbf{P}$, successively overwrite the array originally containing $\mathbf{A}$. Replace each entry being annihilated with the multiplier used to execute the annihilation. Whenever row interchanges such as those used in partial pivoting are implemented, the multipliers in the array will automatically be interchanged in the correct manner.

- Although the entire permutation matrix $\mathbf{P}$ is rarely called for, it can be constructed by permuting the rows of the identity matrix $\mathbf{I}$ according to the various interchanges used. These interchanges can be accumulated in a "permutation counter column" $\mathbf{p}$ that is initially in natural order $(1, 2, \ldots, n)$—see Example 3.10.4.

- To solve a nonsingular linear system $\mathbf{Ax} = \mathbf{b}$ using the LU decomposition with partial pivoting, permute the components in $\mathbf{b}$ to construct $\tilde{\mathbf{b}}$ according to the sequence of interchanges used—i.e., according to $\mathbf{p}$—and then solve $\mathbf{Ly} = \tilde{\mathbf{b}}$ by forward substitution followed by the solution of $\mathbf{Ux} = \mathbf{y}$ using back substitution.

Example 3.10.6

The LDU factorization. There's some asymmetry in an LU factorization because the lower factor has 1's on its diagonal while the upper factor has a nonunit diagonal. This is easily remedied by factoring the diagonal entries out of the upper factor as shown below:

$$\begin{pmatrix} u_{11} & u_{12} & \cdots & u_{1n} \\ 0 & u_{22} & \cdots & u_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & u_{nn} \end{pmatrix} = \begin{pmatrix} u_{11} & 0 & \cdots & 0 \\ 0 & u_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & u_{nn} \end{pmatrix} \begin{pmatrix} 1 & u_{12}/u_{11} & \cdots & u_{1n}/u_{11} \\ 0 & 1 & \cdots & u_{2n}/u_{22} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}.$$

Setting $\mathbf{D} = \operatorname{diag}(u_{11}, u_{22}, \ldots, u_{nn})$ (the diagonal matrix of pivots) and redefining $\mathbf{U}$ to be the rightmost upper-triangular matrix shown above allows any LU factorization to be written as $\mathbf{A} = \mathbf{LDU}$, where $\mathbf{L}$ and $\mathbf{U}$ are lower- and upper-triangular matrices with 1's on both of their diagonals. This is called the ***LDU factorization*** of $\mathbf{A}$. It is uniquely determined, and when $\mathbf{A}$ is symmetric, the LDU factorization is $\mathbf{A} = \mathbf{LDL}^T$ (Exercise 3.10.9).

Example 3.10.7

The Cholesky Factorization.[22] A symmetric matrix $\mathbf{A}$ possessing an LU factorization in which each pivot is positive is said to be ***positive definite***.

Problem: Prove that $\mathbf{A}$ is positive definite if and only if $\mathbf{A}$ can be uniquely factored as $\mathbf{A} = \mathbf{R}^T\mathbf{R}$, where $\mathbf{R}$ is an upper-triangular matrix with positive diagonal entries. This is known as the *Cholesky factorization* of $\mathbf{A}$, and $\mathbf{R}$ is called the *Cholesky factor* of $\mathbf{A}$.

Solution: If $\mathbf{A}$ is positive definite, then, as pointed out in Example 3.10.6, it has an LDU factorization $\mathbf{A} = \mathbf{LDL}^T$ in which $\mathbf{D} = \operatorname{diag}(p_1, p_2, \ldots, p_n)$ is the diagonal matrix containing the pivots $p_i > 0$. Setting $\mathbf{R} = \mathbf{D}^{1/2}\mathbf{L}^T$ where $\mathbf{D}^{1/2} = \operatorname{diag}\left(\sqrt{p_1}, \sqrt{p_2}, \ldots, \sqrt{p_n}\right)$ yields the desired factorization because $\mathbf{A} = \mathbf{LD}^{1/2}\mathbf{D}^{1/2}\mathbf{L}^T = \mathbf{R}^T\mathbf{R}$, and $\mathbf{R}$ is upper triangular with positive diagonal

---

[22] This is named in honor of the French military officer Major André-Louis Cholesky (1875–1918). Although originally assigned to an artillery branch, Cholesky later became attached to the Geodesic Section of the Geographic Service in France where he became noticed for his extraordinary intelligence and his facility for mathematics. From 1905 to 1909 Cholesky was involved with the problem of adjusting the triangularization grid for France. This was a huge computational task, and there were arguments as to what computational techniques should be employed. It was during this period that Cholesky invented the ingenious procedure for solving a positive definite system of equations that is the basis for the matrix factorization that now bears his name. Unfortunately, Cholesky's mathematical talents were never allowed to flower. In 1914 war broke out, and Cholesky was again placed in an artillery group—but this time as the commander. On August 31, 1918, Major Cholesky was killed in battle. Cholesky never had time to publish his clever computational methods—they were carried forward by word-of-mouth. Issues surrounding the Cholesky factorization have been independently rediscovered several times by people who were unaware of Cholesky, and, in some circles, the Cholesky factorization is known as the *square root method.*

entries. Conversely, if $\mathbf{A} = \mathbf{R}\mathbf{R}^T$, where $\mathbf{R}$ is lower triangular with a positive diagonal, then factoring the diagonal entries out of $\mathbf{R}$ as illustrated in Example 3.10.6 produces $\mathbf{R} = \mathbf{L}\mathbf{D}$, where $\mathbf{L}$ is lower triangular with a unit diagonal and $\mathbf{D}$ is the diagonal matrix whose diagonal entries are the r_{ii}'s. Consequently, $\mathbf{A} = \mathbf{L}\mathbf{D}^2\mathbf{L}^T$ is the LDU factorization for $\mathbf{A}$, and thus the pivots must be positive because they are the diagonal entries in $\mathbf{D}^2$. We have now proven that $\mathbf{A}$ is positive definite if and only if it has a Cholesky factorization. To see why such a factorization is unique, suppose $\mathbf{A} = \mathbf{R}_1\mathbf{R}_1^T = \mathbf{R}_2\mathbf{R}_2^T$, and factor out the diagonal entries as illustrated in Example 3.10.6 to write $\mathbf{R}_1 = \mathbf{L}_1\mathbf{D}_1$ and $\mathbf{R}_2 = \mathbf{L}_2\mathbf{D}_2$, where each $\mathbf{R}_i$ is lower triangular with a unit diagonal and $\mathbf{D}_i$ contains the diagonal of $\mathbf{R}_i$ so that $\mathbf{A} = \mathbf{L}_1\mathbf{D}_1^2\mathbf{L}_1^T = \mathbf{L}_2\mathbf{D}_2^2\mathbf{L}_2^T$. The uniqueness of the LDU factors insures that $\mathbf{L}_1 = \mathbf{L}_2$ and $\mathbf{D}_1 = \mathbf{D}_2$, so $\mathbf{R}_1 = \mathbf{R}_2$. **Note:** More is said about the Cholesky factorization and positive definite matrices on pp. 313, 345, and 559.

Exercises for section 3.10

3.10.1. Let $\mathbf{A} = \begin{pmatrix} 1 & 4 & 5 \\ 4 & 18 & 26 \\ 3 & 16 & 30 \end{pmatrix}$.

(a) Determine the LU factors of $\mathbf{A}$.

(b) Use the LU factors to solve $\mathbf{A}\mathbf{x}_1 = \mathbf{b}_1$ as well as $\mathbf{A}\mathbf{x}_2 = \mathbf{b}_2$, where

$$\mathbf{b}_1 = \begin{pmatrix} 6 \\ 0 \\ -6 \end{pmatrix} \quad \text{and} \quad \mathbf{b}_2 = \begin{pmatrix} 6 \\ 6 \\ 12 \end{pmatrix}.$$

(c) Use the LU factors to determine $\mathbf{A}^{-1}$.

3.10.2. Let $\mathbf{A}$ and $\mathbf{b}$ be the matrices

$$\mathbf{A} = \begin{pmatrix} 1 & 2 & 4 & 17 \\ 3 & 6 & -12 & 3 \\ 2 & 3 & -3 & 2 \\ 0 & 2 & -2 & 6 \end{pmatrix} \quad \text{and} \quad \mathbf{b} = \begin{pmatrix} 17 \\ 3 \\ 3 \\ 4 \end{pmatrix}.$$

(a) Explain why $\mathbf{A}$ does not have an LU factorization.

(b) Use partial pivoting and find the permutation matrix $\mathbf{P}$ as well as the LU factors such that $\mathbf{P}\mathbf{A} = \mathbf{L}\mathbf{U}$.

(c) Use the information in $\mathbf{P}$, $\mathbf{L}$, and $\mathbf{U}$ to solve $\mathbf{A}\mathbf{x} = \mathbf{b}$.

3.10.3. Determine all values of ξ for which $\mathbf{A} = \begin{pmatrix} \xi & 2 & 0 \\ 1 & \xi & 1 \\ 0 & 1 & \xi \end{pmatrix}$ fails to have an LU factorization.

3.10.4. If $\mathbf{A}$ is a nonsingular matrix that possesses an LU factorization, prove that the pivot that emerges after $(k+1)$ stages of standard Gaussian elimination using only Type III operations is given by

$$p_{k+1} = a_{k+1,k+1} - \mathbf{c}^T \mathbf{A}_k^{-1} \mathbf{b},$$

where $\mathbf{A}_k$ and

$$\mathbf{A}_{k+1} = \begin{pmatrix} \mathbf{A}_k & \mathbf{b} \\ \mathbf{c}^T & a_{k+1,k+1} \end{pmatrix}$$

are the leading principal submatrices of orders k and $k+1$, respectively. Use this to deduce that all pivots must be nonzero when an LU factorization for $\mathbf{A}$ exists.

3.10.5. If $\mathbf{A}$ is a matrix that contains only integer entries and all of its pivots are 1, explain why $\mathbf{A}^{-1}$ must also be an integer matrix. **Note:** This fact can be used to construct random integer matrices that possess integer inverses by randomly generating integer matrices $\mathbf{L}$ and $\mathbf{U}$ with unit diagonals and then constructing the product $\mathbf{A} = \mathbf{LU}$.

3.10.6. Consider the tridiagonal matrix $\mathbf{T} = \begin{pmatrix} \beta_1 & \gamma_1 & 0 & 0 \\ \alpha_1 & \beta_2 & \gamma_2 & 0 \\ 0 & \alpha_2 & \beta_3 & \gamma_3 \\ 0 & 0 & \alpha_3 & \beta_4 \end{pmatrix}$.

(a) Assuming that $\mathbf{T}$ possesses an LU factorization, verify that it is given by

$$\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ \alpha_1/\pi_1 & 1 & 0 & 0 \\ 0 & \alpha_2/\pi_2 & 1 & 0 \\ 0 & 0 & \alpha_3/\pi_3 & 1 \end{pmatrix}, \quad \mathbf{U} = \begin{pmatrix} \pi_1 & \gamma_1 & 0 & 0 \\ 0 & \pi_2 & \gamma_2 & 0 \\ 0 & 0 & \pi_3 & \gamma_3 \\ 0 & 0 & 0 & \pi_4 \end{pmatrix},$$

where the π_i 's are generated by the recursion formula

$$\pi_1 = \beta_1 \quad \text{and} \quad \pi_{i+1} = \beta_{i+1} - \frac{\alpha_i \gamma_i}{\pi_i}.$$

Note: This holds for tridiagonal matrices of arbitrary size thereby making the LU factors of these matrices very easy to compute.

(b) Apply the recursion formula given above to obtain the LU factorization of

$$\mathbf{T} = \begin{pmatrix} 2 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 1 \end{pmatrix}.$$

3.10.7. $\mathbf{A}_{n\times n}$ is called a ***band matrix*** if $a_{ij} = 0$ whenever $|i - j| > w$ for some positive integer w, called the ***bandwidth.*** In other words, the nonzero entries of $\mathbf{A}$ are constrained to be in a band of w diagonal lines above and below the main diagonal. For example, tridiagonal matrices have bandwidth one, and diagonal matrices have bandwidth zero. If $\mathbf{A}$ is a nonsingular matrix with bandwidth w, and if $\mathbf{A}$ has an LU factorization $\mathbf{A} = \mathbf{LU}$, then $\mathbf{L}$ inherits the lower band structure of $\mathbf{A}$, and $\mathbf{U}$ inherits the upper band structure in the sense that $\mathbf{L}$ has "lower bandwidth" w, and $\mathbf{U}$ has "upper bandwidth" w. Illustrate why this is true by using a generic 5×5 matrix with a bandwidth of $w = 2$.

3.10.8. (a) Construct an example of a nonsingular symmetric matrix that fails to possess an LU (or LDU) factorization.

(b) Construct an example of a nonsingular symmetric matrix that has an LU factorization but is not positive definite.

3.10.9. (a) Determine the LDU factors for $\mathbf{A} = \begin{pmatrix} 1 & 4 & 5 \\ 4 & 18 & 26 \\ 3 & 16 & 30 \end{pmatrix}$ (this is the same matrix used in Exercise 3.10.1).

(b) Prove that if a matrix has an LDU factorization, then the LDU factors are uniquely determined.

(c) If $\mathbf{A}$ is symmetric and possesses an LDU factorization, explain why it must be given by $\mathbf{A} = \mathbf{LDL}^T$.

3.10.10. Explain why $\mathbf{A} = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 8 & 12 \\ 3 & 12 & 27 \end{pmatrix}$ is positive definite, and then find the Cholesky factor $\mathbf{R}$.

3.9.9. If $\mathbf{A} = \mathbf{u}\mathbf{v}^T$, where $\mathbf{u}_{m\times 1}$ and $\mathbf{v}_{n\times 1}$ are nonzero columns, then

$$\mathbf{u} \overset{\text{row}}{\sim} \mathbf{e}_1 \quad \text{and} \quad \mathbf{v}^T \overset{\text{col}}{\sim} \mathbf{e}_1^T \implies \mathbf{A} = \mathbf{u}\mathbf{v}^T \sim \mathbf{e}_1\mathbf{e}_1^T = \mathbf{N}_1 \implies \mathit{rank}\,(\mathbf{A}) = 1.$$

Conversely, if $\mathit{rank}\,(\mathbf{A}) = 1$, then the existence of $\mathbf{u}$ and $\mathbf{v}$ follows from Exercise 3.9.8. If you do not wish to rely on Exercise 3.9.8, write $\mathbf{PAQ} = \mathbf{N}_1 = \mathbf{e}_1\mathbf{e}_1^T$, where $\mathbf{e}_1$ is $m \times 1$ and $\mathbf{e}_1^T$ is $1 \times n$ so that

$$\mathbf{A} = \mathbf{P}^{-1}\mathbf{e}_1\mathbf{e}_1^T\mathbf{Q}^{-1} = \left(\mathbf{P}^{-1}\right)_{*1}\left(\mathbf{Q}^{-1}\right)_{1*} = \mathbf{u}\mathbf{v}^T.$$

3.9.10. Use Exercise 3.9.9 and write

$$\mathbf{A} = \mathbf{u}\mathbf{v}^T \implies \mathbf{A}^2 = \left(\mathbf{u}\mathbf{v}^T\right)\left(\mathbf{u}\mathbf{v}^T\right) = \mathbf{u}\left(\mathbf{v}^T\mathbf{u}\right)\mathbf{v}^T = \tau\mathbf{u}\mathbf{v}^T = \tau\mathbf{A},$$

where $\tau = \mathbf{v}^T\mathbf{u}$. Recall from Example 3.6.5 that $\mathit{trace}\,(\mathbf{AB}) = \mathit{trace}\,(\mathbf{BA})$, and write

$$\tau = \mathit{trace}(\tau) = \mathit{trace}\left(\mathbf{v}^T\mathbf{u}\right) = \mathit{trace}\left(\mathbf{u}\mathbf{v}^T\right) = \mathit{trace}\,(\mathbf{A}).$$

Solutions for exercises in section 3. 10

3.10.1. (a) $\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ 4 & 1 & 0 \\ 3 & 2 & 1 \end{pmatrix}$ and $\mathbf{U} = \begin{pmatrix} 1 & 4 & 5 \\ 0 & 2 & 6 \\ 0 & 0 & 3 \end{pmatrix}$ (b) $\mathbf{x}_1 = \begin{pmatrix} 110 \\ -36 \\ 8 \end{pmatrix}$ and $\mathbf{x}_2 = \begin{pmatrix} 112 \\ -39 \\ 10 \end{pmatrix}$

(c) $\mathbf{A}^{-1} = \frac{1}{6}\begin{pmatrix} 124 & -40 & 14 \\ -42 & 15 & -6 \\ 10 & -4 & 2 \end{pmatrix}$

3.10.2. (a) The second pivot is zero. (b) $\mathbf{P}$ is the permutation matrix associated with the permutation $\mathbf{p} = (\,2 \quad 4 \quad 1 \quad 3\,)$. $\mathbf{P}$ is constructed by permuting the rows of $\mathbf{I}$ in this manner.

$$\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1/3 & 0 & 1 & 0 \\ 2/3 & -1/2 & 1/2 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{U} = \begin{pmatrix} 3 & 6 & -12 & 3 \\ 0 & 2 & -2 & 6 \\ 0 & 0 & 8 & 16 \\ 0 & 0 & 0 & -5 \end{pmatrix}$$

(c) $\mathbf{x} = \begin{pmatrix} 2 \\ -1 \\ 0 \\ 1 \end{pmatrix}$

3.10.3. $\xi = 0, \ \pm\sqrt{2}, \ \pm\sqrt{3}$

3.10.4. $\mathbf{A}$ possesses an LU factorization if and only if all leading principal submatrices are nonsingular. The argument associated with equation (3.10.13) proves that

$$\begin{pmatrix} \mathbf{L}_k & \mathbf{0} \\ \mathbf{c}^T\mathbf{U}_k^{-1} & 1 \end{pmatrix}\begin{pmatrix} \mathbf{U}_k & \mathbf{L}_k^{-1}\mathbf{b} \\ \mathbf{0} & a_{k+1,k+1} - \mathbf{c}^T\mathbf{A}_k^{-1}\mathbf{b} \end{pmatrix} = \mathbf{L}_{k+1}\mathbf{U}_{k+1}$$

is the LU factorization for $\mathbf{A}_{k+1}$. The desired conclusion follows from the fact that the $k+1^{th}$ pivot is the $(k+1,k+1)$-entry in $\mathbf{U}_{k+1}$. This pivot must be nonzero because $\mathbf{U}_{k+1}$ is nonsingular.

3.10.5. If $\mathbf{L}$ and $\mathbf{U}$ are both triangular with 1's on the diagonal, then $\mathbf{L}^{-1}$ and $\mathbf{U}^{-1}$ contain only integer entries, and consequently $\mathbf{A}^{-1} = \mathbf{U}^{-1}\mathbf{L}^{-1}$ is an integer matrix.

3.10.6. (b) $\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ -1/2 & 1 & 0 & 0 \\ 0 & -2/3 & 1 & 0 \\ 0 & 0 & -3/4 & 1 \end{pmatrix}$ and $\mathbf{U} = \begin{pmatrix} 2 & -1 & 0 & 0 \\ 0 & 3/2 & -1 & 0 \\ 0 & 0 & 4/3 & -1 \\ 0 & 0 & 0 & 1/4 \end{pmatrix}$

3.10.7. Observe how the LU factors evolve from Gaussian elimination. Following the procedure described in Example 3.10.1 where multipliers ℓ_{ij} are stored in the positions they annihilate (i.e., in the (i,j)-position), and where $\star$'s are put in positions that can be nonzero, the reduction of a 5×5 band matrix with bandwidth $w=2$ proceeds as shown below.

$$\begin{pmatrix} \star & \star & \star & 0 & 0 \\ \star & \star & \star & \star & 0 \\ \star & \star & \star & \star & \star \\ 0 & \star & \star & \star & \star \\ 0 & 0 & \star & \star & \star \end{pmatrix} \longrightarrow \begin{pmatrix} \star & \star & \star & 0 & 0 \\ l_{21} & \star & \star & \star & 0 \\ l_{31} & \star & \star & \star & \star \\ 0 & \star & \star & \star & \star \\ 0 & 0 & \star & \star & \star \end{pmatrix} \longrightarrow \begin{pmatrix} \star & \star & \star & 0 & 0 \\ l_{21} & \star & \star & \star & 0 \\ l_{31} & l_{32} & \star & \star & \star \\ 0 & l_{42} & \star & \star & \star \\ 0 & 0 & \star & \star & \star \end{pmatrix}$$

$$\longrightarrow \begin{pmatrix} \star & \star & \star & 0 & 0 \\ l_{21} & \star & \star & \star & 0 \\ l_{31} & l_{32} & \star & \star & \star \\ 0 & l_{42} & l_{43} & \star & \star \\ 0 & 0 & l_{53} & \star & \star \end{pmatrix} \longrightarrow \begin{pmatrix} \star & \star & \star & 0 & 0 \\ l_{21} & \star & \star & \star & 0 \\ l_{31} & l_{32} & \star & \star & \star \\ 0 & l_{42} & l_{43} & \star & \star \\ 0 & 0 & l_{53} & l_{54} & \star \end{pmatrix}$$

Thus $\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ l_{21} & 1 & 0 & 0 & 0 \\ l_{31} & l_{32} & 1 & 0 & 0 \\ 0 & l_{42} & l_{43} & 1 & 0 \\ 0 & 0 & l_{53} & l_{54} & 1 \end{pmatrix}$ and $\mathbf{U} = \begin{pmatrix} \star & \star & \star & 0 & 0 \\ 0 & \star & \star & \star & 0 \\ 0 & 0 & \star & \star & \star \\ 0 & 0 & 0 & \star & \star \\ 0 & 0 & 0 & 0 & \star \end{pmatrix}$.

3.10.8. (a) $\mathbf{A} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ (b) $\mathbf{A} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$

3.10.9. (a) $\mathbf{L} = \begin{pmatrix} 1 & 0 & 0 \\ 4 & 1 & 0 \\ 3 & 2 & 1 \end{pmatrix}$, $\mathbf{D} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}$, and $\mathbf{U} = \begin{pmatrix} 1 & 4 & 5 \\ 0 & 1 & 3 \\ 0 & 0 & 1 \end{pmatrix}$

(b) Use the same argument given for the uniqueness of the LU factorization with minor modifications.

(c) $\mathbf{A} = \mathbf{A}^T \implies \mathbf{LDU} = \mathbf{U}^T\mathbf{D}^T\mathbf{L}^T = \mathbf{U}^T\mathbf{D}\mathbf{L}^T$. These are each LDU factorizations for $\mathbf{A}$, and consequently the uniqueness of the LDU factorization means that $\mathbf{U} = \mathbf{L}^T$.

3.10.10. $\mathbf{A}$ is symmetric with pivots 1, 4, 9. The Cholesky factor is $\mathbf{R} = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 2 & 0 \\ 3 & 3 & 3 \end{pmatrix}$.