

LU Decomposition: Factoring $A = L \cdot U$

Motivation To solve $Ax = b$, we can row reduce the augmented matrix

$$[a_1 \ a_2 \ \dots \ a_i \ \dots \ a_n \mid b] \sim \dots \sim [c_1 \ c_2 \ \dots \ c_i \ \dots \ c_n \mid d]$$

until $[c_1 \ c_2 \ \dots \ c_i \ \dots \ c_n]$ is in an echelon form (*this is the “forward” part of the row reduction process*). Then we can

- i) use “back substitution” to solve for x , or
- ii) continue on with the second (“backward”) part of the row reduction process until $[c_1 \ c_2 \ \dots \ c_i \ \dots \ c_n]$ is in row reduced echelon form, at which point it's easy to read off the solutions of $Ax = b$.

A computer might take the first approach; the second might be better for hand calculations. But both i) and ii) take about the same number of arithmetic operations (*see the first paragraph about “back substitution” in Sec. 1.2 (p. 19) of the textbook*).

Whether we use i) or ii) to solve, it is sometimes necessary in applications to solve a large system $Ax = b$ many times with the same A but changing b each time – perhaps a system with millions of variables and solving thousands of times! Just for example, read the description of the aircraft design problem at the beginning of Chapter 2. Row reducing $[a_1 \ a_2 \ \dots \ a_i \ \dots \ a_n \mid b]$ each time is too inefficient, even if a computer is doing the work.

One way to avoid these repeated row reductions is to try to factor A , once and for all, in a way that makes it relatively easy to solve $Ax = b$ again and again as b changes. This is one motivation for the “LU decomposition” (also called an “LU factorization”).

The LU Decomposition An LU decomposition of an $m \times n$ matrix A is a factorization

$$A = L \cdot U \text{ where } \begin{cases} U \text{ is an echelon form of } A, \text{ and} \\ L \text{ is a square unit lower triangular matrix} \end{cases}$$

For a square matrix, lower triangular means “all 0's above the diagonal,” and unit lower triangular means only 1's on the diagonal

A couple of observations:

- U has the same $m \times n$ shape as A because U is an echelon form of A .
- If A and U happen to be square, then the echelon form U is automatically upper triangular. But even when U is not square, U has only 0's below the leading entry in each row, so U still “resembles” an upper triangular matrix.
- Since L is square, L must be $m \times m$ for the product LU to make sense. Since L is unit lower triangular, L is always invertible (why? think about the row reduced echelon form for L)

It's not always possible to find an LU-decomposition for A . But when it is possible, the number of steps in the calculation is not much different from the number of steps required to solve

$A\mathbf{x} = \mathbf{b}$ (once!) by row reduction. And if we can factor A in this way, once and for all, then it takes relatively few additional steps to solve $A\mathbf{x} = \mathbf{b}$ for \mathbf{x} each time a new \mathbf{b} is chosen.

Why? If we can write $A\mathbf{x} = LU\mathbf{x} = \mathbf{b}$, then we can substitute $U\mathbf{x} = \mathbf{y}$. Then the original matrix equation $A\mathbf{x} = \mathbf{b}$ is replaced by two new ones.

$$\begin{cases} L\mathbf{y} = \mathbf{b} \\ \mathbf{y} = U\mathbf{x} \end{cases} \quad (**)$$

This is an improvement (as Example 1 shows) because:

- i) Solving $L\mathbf{y} = \mathbf{b}$ for \mathbf{y} is easy since L is lower triangular; and
- ii) After finding \mathbf{y} , then solving $U\mathbf{x} = \mathbf{y}$ is easy because U is already in echelon form
- iii) L and U remain the same even if we keep changing \mathbf{b} .

Example 1 Find an LU decomposition of $A = \begin{bmatrix} 1 & -1 & 2 \\ 2 & 1 & 0 \\ 0 & 4 & 2 \end{bmatrix}$ and use it to solve the matrix

equation $\begin{bmatrix} 1 & -1 & 2 \\ 2 & 1 & 0 \\ 0 & 4 & 2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix}$.

First, reduce A to an echelon form U . (It turns out that we need to be careful about what EROs we use for the row reduction. For now, just follow along; the discussion that comes later will explain what caution is necessary.) Each elementary row operation we use corresponds to left multiplication by an elementary matrix E_i , and we record those matrices at each step.

$$A = \begin{bmatrix} 1 & -1 & 2 \\ 2 & 1 & 0 \\ 0 & 4 & 2 \end{bmatrix} \sim \begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 4 & 2 \end{bmatrix} \sim \begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix} = U \text{ (an echelon form)}$$

$$E_1 = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad E_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -\frac{4}{3} & 1 \end{bmatrix}$$

Expressed in terms of multiplication by elementary matrices, $E_2E_1A = U$, so

$$A = (E_2E_1)^{-1}U = \begin{matrix} E_1^{-1} & E_2^{-1} & (E_1^{-1}E_2^{-1}) & U \\ \downarrow & \downarrow & \downarrow & \downarrow \end{matrix}$$

$$\begin{bmatrix} 1 & -1 & 2 \\ 2 & 1 & 0 \\ 0 & 4 & 2 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \frac{4}{3} & 1 \end{bmatrix} U = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & \frac{4}{3} & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix}$$

In this example, $E_1^{-1}E_2^{-1}$ is unit lower triangular, so we let $L = E_1^{-1}E_2^{-1}$, so $A = LU$ is the decomposition we want. (Were we just lucky? No, $E_1^{-1}E_2^{-1}$ turned out to be unit lower triangular because we were careful in how we did the row reduction: see the discussion below.)

To solve
$$\begin{bmatrix} 1 & -1 & 2 \\ 2 & 1 & 0 \\ 0 & 4 & 2 \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & \frac{4}{3} & 1 \end{bmatrix} \begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix} \mathbf{x} = \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix}.$$

$$\begin{array}{ccccc} \uparrow & & \uparrow & & \uparrow \\ A\mathbf{x} & = & L & \cdot & U & \mathbf{x} = \mathbf{b} \end{array}$$

Let $U\mathbf{x} = \mathbf{y}$ and substitute. Then we get two equations

$$(**) \quad \begin{cases} L\mathbf{y} = \mathbf{b} \\ \mathbf{y} = U\mathbf{x} \end{cases} \quad \text{that is,} \quad \begin{cases} L\mathbf{y} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & \frac{4}{3} & 1 \end{bmatrix} \mathbf{y} = \begin{bmatrix} 2 \\ 1 \\ -1 \end{bmatrix} \\ \mathbf{y} = U\mathbf{x} = \begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix} \mathbf{x} \end{cases}$$

First solve the top equation for \mathbf{y} . To do this, we could reduce $[L \ \mathbf{b}]$ to row reduced echelon form, or we can use “forward substitution” – just like back substitution but starting with the top equation. Either way we need only a few steps to get \mathbf{y} because L is lower triangular. Using forward substitution gives:

$$y_1 = 2;$$

$$y_2 = 1 - 2y_1 = -3; \text{ and then}$$

$$y_3 = -1 - \frac{4}{3}y_2 = 3, \quad \text{so } \mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} 2 \\ -3 \\ 3 \end{bmatrix}$$

Then the second equation becomes $\begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix} \mathbf{x} = \begin{bmatrix} 2 \\ -3 \\ 3 \end{bmatrix}$, which we can quickly

solve with back substitution or further row reduction. Using back substitution gives:

$$x_3 = 3\left(\frac{3}{\frac{22}{3}}\right) = \frac{9}{22}$$

$$3x_2 = -3 + 4(x_3) = -\frac{15}{11}, \text{ so } x_2 = -\frac{5}{11}, \text{ and}$$

$$x_1 = 2 + x_2 - 2(x_3) = \frac{8}{11}. \quad \text{Therefore } \mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} \frac{8}{11} \\ -\frac{5}{11} \\ \frac{9}{22} \end{bmatrix}$$

Example 1, continued: same coefficient matrix A but different b

Now that we have factored $A = LU$ we can solve another equation $A\mathbf{x} = \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$ easily:

See above: here only \mathbf{b} has changed !

$$\left\{ \begin{array}{l} \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & \frac{4}{3} & 1 \end{bmatrix} \mathbf{y} = \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix} \\ \begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix} \mathbf{x} = \mathbf{y} \end{array} \right.$$

By forward substitution $y_1 = 1$; $y_2 = 2 - 2(1) = 0$; $y_3 = -\frac{4}{3}(0) = 0$: so $\mathbf{y} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$.

Then the second equation becomes $\begin{bmatrix} 1 & -1 & 2 \\ 0 & 3 & -4 \\ 0 & 0 & \frac{22}{3} \end{bmatrix} \mathbf{x} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ and, by back substitution,

$x_3 = 0$; $3x_2 = 0 + 4(0)$ so $x_2 = 0$; and $x_1 = 1 + 1(0) - 2(0) = 1$: so $\mathbf{x} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$.

When does the method in Example 1 work? We can always perform the steps illustrated in Example 1 to get a factorization $A = Z * U$, where U is an echelon form of A and where Z is a product of elementary matrices (*in Example 1, $Z = E_1^{-1}E_2^{-1}$*). But Z won't always turn out to be lower triangular (so that Z might not deserve the name “ L ”). So when does the method work? (*Compare the comments below to Example 1.*)

Suppose (*as in Example 1*) that the only type of ERO used to row reduce A is “add a multiple of a row to a lower row.” Let these EROs correspond to the elementary matrices E_1, \dots, E_p . Then $E_p \cdot \dots \cdot E_1 A = U$ and therefore $A = E_1^{-1} \cdot \dots \cdot E_p^{-1} U$. Each matrix E_i and its inverse E_i^{-1} will be unit lower triangular (*why?*). Then $L = E_1^{-1} \cdot \dots \cdot E_p^{-1}$ is also lower triangular with 1's on its diagonal. (*Convince yourself that a product of unit lower triangular matrices is a unit lower triangular matrix.*)

On the other hand:

i) if a row rescaling ERO had been used, then the corresponding elementary matrix would not have only 1's on the diagonal – and therefore the product $E_1^{-1} \cdot \dots \cdot E_p^{-1}$ might not be unit lower triangular.

ii) if a row interchange (“swap”) ERO had been used, then the corresponding elementary matrix would not be lower triangular – and therefore the product $E_1^{-1} \cdot \dots \cdot E_p^{-1}$ might not be lower triangular.

To summarize: the method illustrated in Example 1 always produces an LU decomposition for A if

the only EROs used to row reduce A to U are of the form
“add a multiple of a row to a lower row.” (*)

In that case: let the EROs used correspond to the elementary matrices E_1, \dots, E_p . Then $E_p \cdot \dots \cdot E_1 A = U$ is in echelon form and

$$A = LU, \text{ where } L = (E_p \cdot \dots \cdot E_1)^{-1} = E_1^{-1} \cdot \dots \cdot E_p^{-1} \quad (**)$$

is unit lower triangular

The restriction in (*) is not as severe as it sounds. For example,

- When you do row reduction following the standard procedure outlined in the text, you never add a multiple of a row to a higher row. You do it to create 0's below a pivot. So adding multiples of rows only to lower rows is no special restriction: it's really just “standard procedure.”
- “No row rescaling” is at most an inconvenience. Rescaling can be helpful (when doing the row reduction arithmetic by hand), but rescaling is never necessary to reduce A to echelon form. (Of course, row rescaling might be needed to continue from echelon form to row reduced echelon form — to do that, we may need rescaling to convert leading entries into 1's.)

Side comment: If we did allow row rescaling, it would just mean that we might end up with a lower triangular “ L ” with some diagonal entries $\neq 1$. From an equation-solving point of view, this would not be a disaster. However, the text follows the rule that L should have only 1's on its diagonal, so we'll stick to that convention.

But (*) really does impose some restriction: for example, $A = \begin{bmatrix} 0 & 2 & 0 \\ 2 & 0 & 0 \\ 0 & 0 & 2 \end{bmatrix}$ simply cannot be

reduced to echelon form without swapping rows 1 and 2. Here the method in Example 1 simply will not work to produce an LU decomposition. However (*see the last section of these notes*), there is a work-around that is “almost as good” as an LU decomposition in such situations.

Finding L more efficiently

If A is large, then row reducing A to echelon form may create very many elementary matrices E_i in the formula $(**)$ — so that formula involves too much work to find L . Fortunately, there is a technique, illustrated in Example 2, to write down L , entry by entry, as A is being row reduced. *(The text gives a slightly different presentation, and an explanation of why it works. It's really nothing more than careful bookkeeping. Since we're just surveying this technique, we will only write down "the method.")*

Assume A is $m \times n$ and that it can be row-reduced to U following the rule $(*)$. Then L will be a square $m \times m$ matrix, and you can write down L using these steps:

- Put 1's on the diagonal of L and 0's everywhere above the diagonal
- Whenever a row operation “add c times row i to a lower row j ” is used, then enter $-c$ in the (j, i) position of L (careful: in the (j, i) position, not (i, j))
- When A has been reduced to echelon form, fill in any remaining empty positions in L with 0's.

If we apply these steps to the computations in Example 1, here's what they look like:

- A was 3×3 , so start with $L = \begin{bmatrix} 1 & 0 & 0 \\ * & 1 & 0 \\ * & * & 1 \end{bmatrix}$
- The row operations used in Example 1 were:

add -2 (row 1) to row 2, so use “2” as the $(2, 1)$ entry in L : that is, $L_{21} = 2$
add $-\frac{4}{3}$ (row 2) to row 3, so use “ $\frac{4}{3}$ ” as the $(3, 2)$ entry in L : that is, $L_{32} = \frac{4}{3}$

Now we have $L = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ * & \frac{4}{3} & 1 \end{bmatrix}$

- Fill in the rest of L with 0's, to get

$$L = \begin{bmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & \frac{4}{3} & 1 \end{bmatrix}, \text{ the same as we got by multiplying elementary matrices}$$

Example 2 In this example, A is not square and is a little larger. But, except for those differences, everything is quite similar to Example 1.

If possible, find an LU factorization of $A = \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ -2 & 0 & -1 & -2 & 8 & 3 \\ -1 & 0 & -1 & -1 & 3 & 1 \end{bmatrix}$ and use it to solve

$$A\mathbf{x} = \mathbf{b} = \begin{bmatrix} -1 \\ 2 \\ 3 \end{bmatrix}$$

$$\begin{aligned} A &= \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ -2 & 0 & -1 & -2 & 8 & 3 \\ -1 & 0 & -1 & -1 & 3 & 1 \end{bmatrix} \sim \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 5 & -2 & 12 & 5 \\ -1 & 0 & -1 & -1 & 3 & 1 \end{bmatrix} \\ &\sim \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 5 & -2 & 12 & 5 \\ 0 & 0 & 2 & -1 & 5 & 2 \end{bmatrix} \sim \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 5 & -2 & 12 & 5 \\ 0 & 0 & 0 & -\frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix} = U \end{aligned}$$

U is in echelon form, and no rescaling or row swaps were used in the row reduction.

To find L : start with $L = \begin{bmatrix} 1 & 0 & 0 \\ * & 1 & 0 \\ * & * & 1 \end{bmatrix}$. The row operations used were

$$\left\{ \begin{array}{l} \text{add } 2*(\text{row } 1) \text{ to } (\text{row } 2) \\ \text{add } 1*(\text{row } 1) \text{ to } (\text{row } 3) \\ \text{add } -\frac{2}{5}(\text{row } 2) \text{ to } (\text{row } 3) \end{array} \right. \quad \text{and these give us that } L = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -1 & \frac{2}{5} & 1 \end{bmatrix}. \text{ So we have}$$

$$A = \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ -2 & 0 & -1 & -2 & 8 & 3 \\ -1 & 0 & -1 & -1 & 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -1 & \frac{2}{5} & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 5 & -2 & 12 & 5 \\ 0 & 0 & 0 & -\frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix}$$

We can solve $A\mathbf{x} = \mathbf{b} = \begin{bmatrix} -1 \\ 2 \\ 3 \end{bmatrix}$ by writing

$$\left\{ \begin{array}{l} L\mathbf{y} = \mathbf{b}, \text{ that is } \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -1 & \frac{2}{5} & 1 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \mathbf{b} = \begin{bmatrix} -1 \\ 2 \\ 3 \end{bmatrix} \\ \mathbf{y} = U\mathbf{x}, \text{ that is } \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 5 & -2 & 12 & 5 \\ 0 & 0 & 0 & -\frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} \end{array} \right.$$

The first equation gives $\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix} = \begin{bmatrix} -1 \\ 2 + 2(y_1) \\ 3 + y_1 - \frac{2}{5}y_2 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix}$, and then the second

equation becomes $\begin{bmatrix} \mathbf{1} & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & \mathbf{5} & -2 & 12 & 5 \\ 0 & 0 & 0 & -\frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \\ 2 \end{bmatrix}$. So the solution is

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} 11 + 4x_5 + 2x_6 \\ x_2 \\ -2x_5 - x_6 - 4 \\ x_5 - 10 \\ x_5 \\ x_6 \end{bmatrix} = \begin{bmatrix} 11 \\ 0 \\ -4 \\ -10 \\ 0 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + x_5 \begin{bmatrix} 4 \\ 0 \\ -2 \\ 1 \\ 1 \\ 0 \end{bmatrix} + x_6 \begin{bmatrix} 2 \\ 0 \\ -1 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

Slight Variations on the LU Decomposition

LDU decompositions If we accept a slightly more complicated factorization and only a tiny bit more work, we can also arrange that U has only 1's in the pivot positions. Since L is unit lower triangular, this introduces more symmetry in appearance between L and U . (For example: if A is square and happens to be invertible, then U will be unit upper triangular – why?) But frankly, this nicer appearance really doesn't contribute much to solving the system of equations.

Starting with $A = LU$, just factor from the rows of U the numbers that are in pivot positions. Make these numbers the entries in the diagonal matrix D . If a row of U has no pivot position, then “factor out” a 0 and put it on the diagonal of D . Then $A = LDU$, which is cleverly called an LDU decomposition of A . Here's an illustration using the matrices from Example 2.

$$\begin{aligned} A &= \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ -2 & 0 & -1 & -2 & 8 & 3 \\ -1 & 0 & -1 & -1 & 3 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -1 & \frac{2}{5} & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 5 & -2 & 12 & 5 \\ 0 & 0 & 0 & -\frac{1}{5} & \frac{1}{5} & 0 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & 0 \\ -2 & 1 & 0 \\ -1 & \frac{2}{5} & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & -\frac{1}{5} \end{bmatrix} \begin{bmatrix} 1 & 0 & 3 & 0 & 2 & 1 \\ 0 & 0 & 1 & -\frac{2}{5} & \frac{12}{5} & 1 \\ 0 & 0 & 0 & 1 & -1 & 0 \end{bmatrix} \\ &= \begin{matrix} L & D & U \end{matrix} \end{aligned}$$

↑
this U is a new matrix, so we shouldn't really give it the same name as the original matrix U . But we will do so just because it is the standard name for the last matrix in an “ LDU decomposition.” Here “original U ” = D (“new matrix U ”)

This always works because when computing DU , each diagonal entry in D just multiplies the corresponding row in U (check!).

Related decomposition Sometimes row interchanges are simply unavoidable to reduce A to echelon form. Other times, even when they could be avoided, row interchanges are introduced in computer computations to reduce roundoff errors (*see the Numerical Note about partial pivoting in the text on p. 17*). So what happens if row interchanges are used? We can still get a factorization similar to $A = LU$ and just as useful. Here's how.

- 1) Look at the steps needed to row reduce A to echelon form U and see what row interchanges will be used. Then go back and “repackage” A : do all these row interchanges first. The result is “ A with some rows interchanged.” Look at this in more detail:

Suppose the row interchanges we use are represented by the elementary matrices E_1, \dots, E_k . Then the “repackaged” matrix A with some rows interchanged” is just $E_k \cdot \dots \cdot E_1 A = PA$, where $P = E_k \cdot \dots \cdot E_1$.

Since $P = E_k \cdot \dots \cdot E_1 = E_k \cdot \dots \cdot E_1 I$ we see that P is just the final result of performing the same row interchanges on I – that is, P is just “ I with some rows rearranged.” P is called a permutation matrix; the multiplication PA permutes (“rearranges”) the rows of A in exactly the same way that the rows of I were permuted to create P . *Such a matrix P will always be a square matrix with exactly one 1 in each row and exactly one 1 in each column.*

Every permutation matrix P is invertible (because it is a product of elementary matrices, which are invertible). $P^{-1} = E_1^{-1} \cdot \dots \cdot E_p^{-1}$ is also a permutation matrix, and $P^{-1}PA = A$. This equation says that P^{-1} rearranges the rows of PA to restore the original matrix A . *(If you know the steps to write down P , it's just as easy to write the matrix $P^{-1} = E_1^{-1} \cdot \dots \cdot E_p^{-1}$. Or, perhaps you can convince yourself that $P^{-1} = P^T$.)*

- 2) Having done the need row swaps, we can row reduce PA to echelon form without using any row interchanges, so we can use our previous method to get an LU decomposition of PA :

$$PA = LU \quad (***)$$

- 3) Then $A = (P^{-1}L)U$. $P^{-1}L$ is called a permuted lower triangular matrix, meaning “lower triangular with some rows interchanged.” *MATLAB Help gives $P^{-1}L$ the cute name “psychologically lower triangular matrix,” a term never seen anywhere else to my knowledge.*

This is just as useful as the LU decomposition. For example, we can solve $A\mathbf{x} = (P^{-1}L)U\mathbf{x} = \mathbf{b}$ by first substituting $\mathbf{y} = U\mathbf{x}$. Solving $(P^{-1}L)\mathbf{y} = \mathbf{b}$ is just as easy as when the coefficient matrix is lower triangular.

For example if this equation were

permuted lower triangular

$$\begin{bmatrix} 2 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 2 & 2 & 1 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} \quad \text{we would simply imitate forward substitution,}$$

but solve using the rows in the order that (rearranged) would make a lower triangular matrix. That is,

$$\begin{cases} y_1 = 1, & \text{from the second row, then} \\ y_2 = 0 - 2y_1 = -2, & \text{from the first row, then} \\ y_3 = 0 - y_1 - y_2 = 1, & \text{from the fourth row, then} \\ y_4 = 1 - 2y_1 - 2y_2 - y_3 = 2, & \text{from the third row} \end{cases}$$

(In some books or applications, when you see $A = LU$ it may even be assumed that L is permuted lower triangular rather than lower triangular.)

OR to solve $A\mathbf{x} = (P^{-1}L)U\mathbf{x} = \mathbf{b}$, you can reason as follows:

Notice that the equation $PA\mathbf{x} = P\mathbf{b}$ has exactly the same solutions as $A\mathbf{x} = \mathbf{b}$.

Because: If $PA\mathbf{x} = P\mathbf{b}$ is true, then multiplying both sides by P^{-1} shows that \mathbf{x} is also a solution $A\mathbf{x} = \mathbf{b}$ also, and vice-versa. *If you think in terms of writing out the linear systems of equations: $PA\mathbf{x} = P\mathbf{b}$ is the same system of equations as $A\mathbf{x} = \mathbf{b}$ but with the equations listed in some other interchanged order determined by the permutation matrix P .*

And since we have an LU decomposition for PA , solving $PA\mathbf{x} = LU\mathbf{x} = P\mathbf{b}$ for \mathbf{x} is easy.

Example 3 Let $A = \begin{bmatrix} 1 & 1 & 2 \\ 1 & 1 & 1 \\ 1 & 4 & 4 \\ 0 & 1 & 1 \end{bmatrix}$.

1) On scratch paper, do enough steps row reducing A to echelon form (*no row rescaling allowed, and, following standard procedure, only adding multiples of rows to lower rows*) to see what row interchanges, if any, are necessary. It turns out that only one, interchanging rows 2 and 4 at a certain stage, was necessary.

2) So go back to the start and perform that row interchange first, creating

$$PA = \begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 1 & 4 & 4 \\ 1 & 1 & 1 \end{bmatrix}, \text{ where } P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}$$

3) Row reduce PA to an echelon form U , keeping track of the EROs used:

$$\begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 1 & 4 & 4 \\ 1 & 1 & 1 \end{bmatrix} \xrightarrow{(1)} \begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 0 & 3 & 2 \\ 1 & 1 & 1 \end{bmatrix} \xrightarrow{(2)} \begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 0 & 3 & 2 \\ 0 & 0 & -1 \end{bmatrix} \xrightarrow{(3)} \begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \\ 0 & 0 & -1 \end{bmatrix}$$

$$\xrightarrow{(4)} \begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix} = U$$

The row operations were

- (1) add $-1 \cdot \text{row}(1)$ to row 3
- (2) add $-1 \cdot \text{row}(1)$ to row 4
- (3) add $-3 \cdot \text{row}(2)$ to row 3
- (4) add $-1 \cdot \text{row}(3)$ to row 4

Using the method described earlier for “finding L efficiently” we can write down L step by step as we do the row reduction of PA :

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ * & 1 & 0 & 0 \\ * & * & 1 & 0 \\ * & * & * & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ * & 1 & 0 & 0 \\ 1 & * & 1 & 0 \\ * & * & * & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ * & 1 & 0 & 0 \\ 1 & * & 1 & 0 \\ 1 & * & * & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ * & 1 & 0 & 0 \\ 1 & 3 & 1 & 0 \\ 1 & * & * & 1 \end{bmatrix}$$

$$\rightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ * & 1 & 0 & 0 \\ 1 & 3 & 1 & 0 \\ 1 & * & 1 & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 3 & 1 & 0 \\ 1 & 0 & 1 & 1 \end{bmatrix} = L$$

Then $PA = L U$

$$\begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 1 & 4 & 4 \\ 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 3 & 1 & 0 \\ 1 & 0 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & 2 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix}$$

To solve $A\mathbf{x} = \begin{bmatrix} 2 \\ 1 \\ 3 \\ 7 \end{bmatrix}$, write $PA\mathbf{x} = LU\mathbf{x} = P\mathbf{b} = \begin{bmatrix} 2 \\ 7 \\ 3 \\ 1 \end{bmatrix}$ (rows 2,4 of \mathbf{b} interchanged, as for A earlier). We then solve $LU\mathbf{x} = \begin{bmatrix} 2 \\ 7 \\ 3 \\ 1 \end{bmatrix}$ in the same way as before

For those who use MATLAB

If you create a square matrix A , and then enter the MATLAB command

$$[L, U, P] = \text{lu}(A) \quad (\text{"lu" is lowercase in MATLAB})$$

then MATLAB returns to you

- 1) a square unit lower triangular L
- 2) a square U in echelon form, and
- 3) a permutation matrix P

for which $PA = LU$. *If it's possible simply to factor $A = LU$, then MATLAB just gives $P = I = \text{the identity matrix}$,*

Our earlier work doesn't require A to be square, but the MATLAB command "lu" only works when A is square.