Section 3.5 LU Decomposition (Factorization)

Suppose we need to solve several linear systems, all with the same coefficient matrix, but each with a different right-hand side vector. If all of the right-hand side vectors are known from the outset, we can place the coefficient matrix and all of the vectors into a large augmented matrix. Gaussian elimination with back substitution applied to this large augmented matrix would then produce a simultaneous solution to all of the systems.

But what if the right-hand side vectors are not all known from the outset? For example, the solution vector for one system may be the right-hand side vector for the next system. Several methods developed later in the text will work in precisely this manner. Although it is the elements in the coefficient matrix which dictate the operations to perform during Gaussian elimination, these operations are also carried out on the right-hand side vector. As a result, each time we change the right-hand side vector, exactly the same sequence of operations has to be repeated on the new augmented matrix. That's $O(n^3)$ operations repeated again and again. From an efficiency standpoint, it would be better to have a solution scheme which treats the coefficient matrix and the right-hand side vector separately, thereby reducing the effort which must be expended when the right-hand side vector is changed. The objective of this section is to develop such a scheme.

In matrix analysis as implemented in modern software the idea of factoring a matrix into a product of matrices of special form is particularly important. The strategy used is to factor the matrix and then use the factors to efficiently and quickly solve the problem.

Case of a General Nonsingular Linear System **Ax** = **b**

As we saw with GEM, triangular matrices are important. The notion of using row operations to transform the augmented matrix [**A** | **b**] to upper triangular form [**U** | **c**] and then use back substitution provides a reliable technique, especially when combined with a pivoting strategy.

Related Idea: If **L** is a lower triangular matrix, then linear system **Lx** = **b** can solved by **forward substitution**. Assuming that no diagonal entry l_{ii} is zero we can proceed as follows.

$$x_1 = \frac{b_1}{l_{11}}$$

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

$$\vdots$$

$$x_n = \frac{b_n - \sum_{j=1}^{n-1} a_{nj}x_j}{l_{nn}}$$

For many nonsingular linear systems **Ax** = **b** it can be shown that the coefficient matrix can be factored as a product of a lower triangular matrix and an upper triangular matrix. That is, **A** = **LU**, and we say we have an **LU**-factorization or **LU**-decomposition of **A**. If any row interchanges are required to perform the factorization or partial pivoting is incorporated, then the equivalent linear system must be expressed as **LUx** = **Pb**, where **P** is a permutation matrix that embodies the row interchanges that were used. If no row interchanges are used then the equivalent system is **LUx** = **b**. In either case, the equivalent system is easily solved as we now show.

> The solution of a linear system **LUx** = **c** is done as follows:
> 1. Name **Ux** to be **z**.
> 2. Solve **Lz** = **c** by forward substitution. We now have vector **z**.
> 3. Solve system **Ux** = **z** by back substitution.

We say that **LUx** = **c** is solved by a forward substitution followed by a back substitution.

Example 1. Given that **Ax** = **b** has $\mathbf{A} = \begin{bmatrix} 2 & 0 & 1 \\ 4 & -3 & 1 \\ 0 & 3 & -1 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 3 \\ -4 \\ 8 \end{bmatrix}$ with LU-factorization of **A** as

$\mathbf{L} = \begin{bmatrix} 1 & 0 & 0 \\ 2 & -1 & 0 \\ 0 & 1 & 2 \end{bmatrix}$ and $\mathbf{U} = \begin{bmatrix} 2 & 0 & 1 \\ 0 & 3 & 1 \\ 0 & 0 & -1 \end{bmatrix}$. Use the LU-factorization to solve the system. We have **LUx** = **b**,

so let **z** = **Ux** and solve **Lz** = **b** by forward substitution; we get $\mathbf{z} = \begin{bmatrix} 3 \\ 10 \\ -1 \end{bmatrix}$. (Verify.) Then we solve

Ux = **z** by back substitution; we get $\mathbf{x} = \begin{bmatrix} 1 \\ 3 \\ 1 \end{bmatrix}$. (Verify.)

+++

> Comment: There can be more than one LU-factorization for a matrix **A**.

Constructing an LU-factorization

We develop an LU-factorization procedure that utilizes row operations in the same way as we applied them to GEM. Initially we assume that no row interchanges are required to get a nonzero pivot and later generalize the procedure to incorporate row interchanges so that partial pivoting can be used.

> Case of NO row interchanges:
>
> We use row operations to transform just the coefficient matrix **A** to upper triangular form **U**. As we proceed **we construct the lower triangular matrix using the negatives of the multipliers k** of the row operations k*Row(i) + Row(j) ==> Row(j). **The negative of the multiplier is stored in row-column position in L that was zeroed out in A by the row operation.**

We illustrate this process in Example 2.

Example 2. Compute an LU-factorization for $\mathbf{A}=\begin{bmatrix}2&3&-4\\-4&-10&9\\6&13&-10\end{bmatrix}$. Initialize matrix $\mathbf{L}=\begin{bmatrix}1&0&0\\&1&0\\&&1\end{bmatrix}$, leaving the entries in the (strict) lower triangular part unassigned. We will be inserting the negatives of the multipliers k of the row operations k*Row(i) + Row(j) ==> Row(j) which are used to zero out below the pivots. As we operate on matrix **A** it will be transformed into upper triangular matrix **U**.

- Use $\begin{cases}2\mathbf{Row}_1+\mathbf{Row}_2\to\mathbf{Row}_2\\-3\mathbf{Row}_1+\mathbf{Row}_3\to\mathbf{Row}_3\end{cases}$ on **A** to get $\begin{bmatrix}2&3&-4\\0&-4&1\\0&4&2\end{bmatrix}$ and store the negatives of the multipliers into **L**; so $\mathbf{L}=\begin{bmatrix}1&0&0\\\mathbf{-2}&1&0\\\mathbf{3}&&1\end{bmatrix}$.
- Use $1\mathbf{Row}_2+\mathbf{Row}_3\to\mathbf{Row}_3$ on $\begin{bmatrix}2&3&-4\\0&-4&1\\0&4&2\end{bmatrix}$ to get $\mathbf{U}=\begin{bmatrix}2&3&-4\\0&-4&1\\0&0&3\end{bmatrix}$ and store the negative of the multiplier into **L**; so $\mathbf{L}=\begin{bmatrix}1&0&0\\\mathbf{-2}&1&0\\\mathbf{3}&\mathbf{-1}&1\end{bmatrix}$.

Thus $\mathbf{A}=\begin{bmatrix}2&3&-4\\-4&-10&9\\6&13&-10\end{bmatrix}=\mathbf{LU}=\begin{bmatrix}1&0&0\\-2&1&0\\3&-1&1\end{bmatrix}\begin{bmatrix}2&3&-4\\0&-4&1\\0&0&3\end{bmatrix}$. (Verify.)

++

Example 3. Use LU-factorization to solve the linear system

$$\mathbf{Ax}=\mathbf{b} \text{ where } \mathbf{A}=\begin{bmatrix}2&3&-4\\-4&-10&9\\6&13&-10\end{bmatrix} \text{ and } \mathbf{b}=\begin{bmatrix}-14\\46\\-54\end{bmatrix}.$$

From Example 2 we have $\mathbf{Ax}==\mathbf{LUx}=\begin{bmatrix}1&0&0\\-2&1&0\\3&-1&1\end{bmatrix}\begin{bmatrix}2&3&-4\\0&-4&1\\0&0&3\end{bmatrix}=\mathbf{b}$. Let $\mathbf{z}=\mathbf{Ux}$, then solve $\mathbf{Lz}=\mathbf{b}$ using forward substitution to get $\mathbf{z}=\begin{bmatrix}-14\\18\\6\end{bmatrix}$. Now solve $\mathbf{Ux}=\mathbf{z}$ to get $\mathbf{x}=\begin{bmatrix}3\\-4\\2\end{bmatrix}$. (Verify.)

++

Case of INCORPORATING row interchanges:

As in the case of NO row interchanges we build L and U using row operations. However, we must employ an indirect addressing scheme for the rows since we physically do want to make row interchanges. To implement this we use the pivot vector idea. **In this case L and U are truly not triangular in general, rather their rows can be interchanged to get triangular form.**

Summary of the Pivot Vector Concept

The **pivot vector** concept is just a book keeping process. At the start of the reduction process for an n × n linear system we initialize the pivot vector **p** as

$$\mathbf{p} = \begin{bmatrix} 1 \\ 2 \\ \vdots \\ \mathbf{n} \end{bmatrix}$$

We denote the entries of **p** by p_i. If at some point we interchange row 2 and row 5 we interchange the contents of p_2 and p_5. **We use the contents of the pivot vector as an indirect addressing scheme for row numbers.**

We illustrate the incorporation of a pivot vector into LU-factorization in Example 4 which uses the matrix of Example 2 this time using partial pivoting to select the pivots.

Example 4. Compute an LU-factorization for $\mathbf{A} = \begin{bmatrix} 2 & 3 & -4 \\ -4 & -10 & 9 \\ 6 & 13 & -10 \end{bmatrix}$ using partial pivoting. Initialize the pivot vector as $\mathbf{p} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$. To use partial pivoting we look in $\text{col}_1(\mathbf{A}) = \begin{bmatrix} 2 \\ -4 \\ 6 \end{bmatrix}$ for the 1st pivot which is the (3,1)-entry = 6. We interchange the 1st and 3rd entries (p_1 and p_3) of the pivot vector to get $\mathbf{p} = \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}$; thus the first pivot row is the row whose number appears in the 1st entry of vector **p**. We begin forming the rows of **L** by inserting a 1 into the $(p_1,1)$-entry, and zeros in the other entries of row p_1. Thus we have $\mathbf{L} = \begin{bmatrix} & & \\ & & \\ 1 & 0 & 0 \end{bmatrix}$. The value of p_1 is 3 so we use the third row to generate row operations to zero all the entries in column 1 of **A** except the pivot row. Here that means we want to zero-out entries $(p_2,1)$ and $(p_3,1)$ of **A**.

- Use $\begin{cases} \dfrac{-\mathbf{a}_{\mathbf{p}_2 1}}{\mathbf{a}_{\mathbf{p}_1 1}}\mathbf{Row}_{\mathbf{p}_1} + \mathbf{Row}_{\mathbf{p}_2} \to \mathbf{Row}_{\mathbf{p}_2} \text{ so the multiplier is } \dfrac{-(-4)}{6} = \dfrac{2}{3} \\ \dfrac{-\mathbf{a}_{\mathbf{p}_3 1}}{\mathbf{a}_{\mathbf{p}_1 1}}\mathbf{Row}_{\mathbf{p}_1} + \mathbf{Row}_{\mathbf{p}_3} \to \mathbf{Row}_{\mathbf{p}_3} \text{ so the multiplier is } \dfrac{-(2)}{6} = \dfrac{-1}{3} \end{cases}$

  on **A** to get $\mathbf{U} = \begin{bmatrix} 0 & -4/3 & -2/3 \\ 0 & -4/3 & 7/3 \\ 6 & 13 & -10 \end{bmatrix}$ and <u>store the negatives of the multipliers into</u> **L** in the first column at the rows numbered p_2 and p_3 respectively; so $\mathbf{L} = \begin{bmatrix} 1/3 & & \\ -2/3 & & \\ 1 & 0 & 0 \end{bmatrix}$.

To determine the second pivot we inspect the entries in column 2 of **U**, except for the entry previous pivot row. So here we are looking at $\begin{bmatrix} \dfrac{-4}{3} \\ \dfrac{-4}{3} \end{bmatrix}$. In the pivot vector the candidates for the second pivot are in rows whose numbers appear in the second and third entries; that is, in rows $p_2 = 2$ and $p_3 = 1$ of column 2 of **U**. Since the magnitude of this pair of entries is the same we choose the first one encountered from the contents of pivot vector entries. Thus the second pivot is the $(p_2,2)$-entry of **U**. So the pivot vector contents are <u>not</u> rearranged. We insert a 1 into the $(p_2,2)$-entry of **L** and <u>zeros</u> into any other entries of row p_2 that are blank. We have $\mathbf{L} = \begin{bmatrix} 1/3 & & \\ -2/3 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$.

- Use $\dfrac{-\mathbf{a}_{\mathbf{p}_3 2}}{\mathbf{a}_{\mathbf{p}_2 2}}\mathbf{Row}_{\mathbf{p}_2} + \mathbf{Row}_{\mathbf{p}_3} \to \mathbf{Row}_{\mathbf{p}_3} \left(\text{the multiplier is } \dfrac{\text{-(-4/3)}}{\text{-4/3}} = -1 \right)$ on $\begin{bmatrix} 0 & -4/3 & -2/3 \\ 0 & -4/3 & 7/3 \\ 6 & 13 & -10 \end{bmatrix}$

  to get $\mathbf{U} = \begin{bmatrix} 0 & 0 & -3 \\ 0 & -4/3 & 7/3 \\ 6 & 13 & -10 \end{bmatrix}$ and <u>store the negative of the multiplier into the $(p_3,2)$-entry of</u> **L**;

  so $\mathbf{L} = \begin{bmatrix} 1/3 & 1 & \\ -2/3 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$.

Since the matrix was 3 by 3, the 3rd pivot is in the $(p_3,3)$-entry of **U** so we insert a 1 into that entry of **L** giving $\mathbf{L} = \begin{bmatrix} 1/3 & 1 & 1 \\ -2/3 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$. So we have constructed both L and U at this point:

$$\mathbf{L} = \begin{bmatrix} 1/3 & 1 & 1 \\ -2/3 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \qquad \mathbf{U} = \begin{bmatrix} 0 & 0 & -3 \\ 0 & -4/3 & 7/3 \\ 6 & 13 & -10 \end{bmatrix}$$

The **L** and **U** in this case are sometimes called pseudo triangular, meaning that if we rearranged their rows according to the "directions" in the pivot vector we would have corresponding triangular matrices. The pivot vector $\mathbf{p} = \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix}$ "says" move the 3rd row to the 1st row position, move the 2nd row to the 2nd row position, and move the 1st row to the 3rd row position. In terms of a permutation matrix **P**, interchange the rows of the identity according to the contents of the pivot vector so $\mathbf{P} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$, then **P** times our pseudo triangular matrices **L** or **U** above will yield the corresponding triangular matrices.

++

Example 5. Suppose we had the linear system **Ax** = **b**, where $\mathbf{A} = \begin{bmatrix} 2 & 3 & -4 \\ -4 & -10 & 9 \\ 6 & 13 & -10 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 11 \\ -26 \\ 37 \end{bmatrix}$.

To solve this system by LU-factorization that incorporates partial pivoting as in Example 4, we proceed as follows. Using the results of Example 4 we have the pseudo triangular matrices $\mathbf{L} = \begin{bmatrix} 1/3 & 1 & 1 \\ -2/3 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$ and $\mathbf{U} = \begin{bmatrix} 0 & 0 & -3 \\ 0 & -4/3 & 7/3 \\ 6 & 13 & -10 \end{bmatrix}$. In order to solve **Ax** = **b** using a forward substitution followed by a back substitution, one way to proceed is use the permutation matrix to get (true) triangular matrices **L** and **U**;

$$\mathbf{PL} = \begin{bmatrix} 1 & 0 & 0 \\ -2/3 & 1 & 0 \\ 1/3 & 1 & 1 \end{bmatrix}, \ \mathbf{PU} = \begin{bmatrix} 6 & 13 & -10 \\ 0 & -4/3 & 7/3 \\ 0 & 0 & -3 \end{bmatrix}.$$

Unfortunately, the solution of **Ax** = **b** is **not** obtained from

$$\begin{bmatrix} 1 & 0 & 0 \\ -2/3 & 1 & 0 \\ 1/3 & 1 & 1 \end{bmatrix} \begin{bmatrix} 6 & 13 & -10 \\ 0 & -4/3 & 7/3 \\ 0 & 0 & 5/3 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} = \mathbf{b} = \begin{bmatrix} 11 \\ -26 \\ 37 \end{bmatrix}.$$

However, the solution of **Ax** = **b** is obtained from

$$\begin{bmatrix} 1 & 0 & 0 \\ -2/3 & 1 & 0 \\ 1/3 & 1 & 1 \end{bmatrix} \begin{bmatrix} 6 & 13 & -10 \\ 0 & -4/3 & 7/3 \\ 0 & 0 & -3 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} = \mathbf{Pb} = \mathbf{P} \begin{bmatrix} 11 \\ -26 \\ 37 \end{bmatrix}.$$

Verify that the solution is $\mathbf{x} = \begin{bmatrix} 4 \\ 1 \\ 0 \end{bmatrix}$.

++

Comment:

It is not necessary to have a separate matrix L. The contents of **L** can be stored in the entries of **A** that are zeroed out as a result of the row operations k*Row(i) + Row(j) ==> Row(j). In this regard we "remember" that **L** is to have 1's in diagonal entries if no interchanges are made and correspondingly when a strategy such as partial pivoting is used, the pivot vector contains information on where the 1's should appear. Hence there is storage economization with this device.