Lecture 5 - Triangular Factorizations & Operation Counts

LU Factorization

We have seen that the process of GE essentially factors a matrix A into LU. Now we want to see how this factorization allows us to solve linear systems and why in many cases it is the preferred algorithm compared with GE. Remember on paper, these methods are the same but computationally they can be different.

First, suppose we want to solve $A\vec{x} = \vec{b}$ and we are given the factorization $A = LU$. It turns out that the system $LU\vec{x} = \vec{b}$ is "easy" to solve because we do a

$$\textit{forward solve } L\vec{y} = \vec{b} \text{ and then } \textit{back solve } \ U\vec{x} = \vec{y}.$$

We have seen that we can easily implement the equations for the back solve and for homework you will write out the equations for the forward solve.

Example If

$$A = \begin{pmatrix} 2 & -1 & 2 \\ 4 & 1 & 9 \\ 8 & 5 & 24 \end{pmatrix} = LU = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 4 & 3 & 1 \end{pmatrix} \begin{pmatrix} 2 & -1 & 2 \\ 0 & 3 & 5 \\ 0 & 0 & 1 \end{pmatrix}$$

solve the linear system $A\vec{x} = \vec{b}$ where $\vec{b} = (0, -5, -16)^T$.

We first solve $L\vec{y} = \vec{b}$ to get $y_1 = 0$, $2y_1 + y_2 = -5$ implies $y_2 = -5$ and $4y_1 + 3y_2 + y_3 = -16$ implies $y_3 = -1$. Now we solve $U\vec{x} = \vec{y} = (0, -5, -1)^T$. Back solving yields $x_3 = -1$, $3x_2 + 5x_3 = -5$ implies $x_2 = 0$ and finally $2x_1 - x_2 + 2x_3 = 0$ implies $x_1 = 1$ giving the solution $(1, 0, -1)^T$.

If GE and LU factorization are equivalent on paper, why would one be computationally advantageous in some settings? Recall that when we solve $A\vec{x} = \vec{b}$ by GE we must also multiply the right hand side by the Gauss transformation matrices. Often in applications, we have to solve many linear systems where the coefficient matrix is the same but the right hand side vector changes. If we have all of the right hand side vectors at one time, then we can treat them as a rectangular matrix and multiply this by the Gauss transformation matrices. However, in many instances we solve a single linear system and use its solution to compute a new right hand side, i.e., we don't have all the right hand sides at once. When we perform an LU factorization then we overwrite the factors onto A and if the right hand side changes, we simply do another forward and back solve to find the solution.

One can easily derive the equations for an LU factorization by writing $A = LU$ and equating entries. Consider the matrix equation $A = LU$ written as

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ a_{31} & a_{32} & a_{33} & \cdots & a_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ a_{n1} & a_{n2} & a_{n3} & \cdots & a_{nn} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ \ell_{21} & 1 & 0 & \cdots & 0 \\ \ell_{31} & \ell_{32} & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \ell_{n1} & \ell_{n2} & \ell_{n3} & \cdots & 1 \end{pmatrix} \begin{pmatrix} u_{11} & u_{12} & u_{13} & \cdots & u_{1n} \\ 0 & u_{22} & u_{23} & \cdots & u_{2n} \\ 0 & 0 & u_{33} & \cdots & u_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & u_{nn} \end{pmatrix}$$

Now equating the $(1,1)$ entry gives

$$a_{11} = 1 \cdot u_{11} \Rightarrow u_{11} = a_{11}$$

In fact, if we equate each entry of the first row of A, i.e., a_{1j} we get

$$u_{1j} = a_{1j} \quad \text{for } j = 1, \ldots, n.$$

Now we move to the second row and look at the (2,1) entry to get $a_{21} = \ell_{21} \cdot u_{11}$ implies $\ell_{21} = a_{21}/u_{11}$. Now we can determine the remaining terms in the first column of L by

$$\ell_{i1} = a_{i1}/u_{11} \quad \text{for } i = 2, \ldots, n.$$

We now find the second row of U. Equating the (2,2) entry gives $a_{22} = \ell_{21}u_{12} + u_{22}$ implies $u_{22} = a_{22} - \ell_{21}u_{12}$. In general

$$u_{2j} = a_{2j} - \ell_{21}u_{1j} \quad \text{for } j = 2, \ldots, n.$$

We now obtain formulas for the second column of L. Equating the (3,2) entries gives

$$\ell_{31}u_{12} + \ell_{32}u_{22} = a_{32} \Rightarrow \ell_{32} = \frac{a_{32} - \ell_{31}u_{12}}{u_{22}}$$

and equating $(i,2)$ entries for $i = 3, 4, \ldots, n$ gives

$$\ell_{i2} = \frac{a_{i2} - \ell_{i1}u_{12}}{u_{22}} \quad i = 3, 4, \ldots, n$$

Continuing in this manner, we get the following algorithm.

Let A be a given $n \times n$ matrix. Then if no pivoting is needed, the LU factorization of A into a unit lower triangular matrix L with entries ℓ_{ij} and an upper triangular matrix U with entries u_{ij} is given by the following algorithm.

Set $u_{1j} = a_{1j}$ for $j = 1, \ldots, n$

For $k = 1, 2, 3 \ldots, n-1$

for $i = k+1, \ldots, n$

$$\ell_{i,k} = \frac{a_{i,k} - \displaystyle\sum_{m=1}^{k-1} \ell_{im}u_{m,k}}{u_{k,k}}$$

for $j = k+1, \ldots, n$

$$u_{k+1,j} = a_{k+1,j} - \sum_{m=1}^{k} \ell_{k+1,m}u_{m,j}$$

Note that this algorithm clearly demonstrates that you can NOT find all of L and then all of U or vice versa. One must determine a row of U, then a column of L, then a row of U, etc.

Example Perform an LU decomposition of

$$A = \begin{pmatrix} 2 & -1 & 2 \\ 4 & 1 & 9 \\ 8 & 5 & 24 \end{pmatrix}$$

The result is given in a previous example and can be found directly by equating elements of A with the corresponding element of LU.

- Does LU factorization work for all systems that have a unique solution?

Example Consider $A\vec{x} = \vec{b}$ where

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

which has the unique solution $\vec{x} = (1,1)^T$. Can you find an LU factorization of A?

Just like in GE the (1,1) entry is a zero pivot and so we can't find u_{11}.

Theorem *Let A be an $n \times n$ matrix. Then there exists a permutation matrix P such that*

$$PA = LU$$

where L is unit lower triangular and U is upper triangular.

Example For the matrix above find the permutation matrix P which makes PA have an LU decomposition and then find the decomposition.

We want to interchange the first and second rows so we need a permutation matrix with the first two rows of the identity interchanged.

$$PA = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

Variants of LU Factorization

There are several variants of LU factorization.

1. $A = LU$ where L is lower triangular and U is unit upper triangular. This is explored in the homework.

2. $A = LDU$ where L is unit lower triangular, U is unit upper triangular and D is diagonal.

Example If

$$A = \begin{pmatrix} 2 & -1 & 2 \\ 4 & 1 & 9 \\ 8 & 5 & 24 \end{pmatrix} = L\tilde{U} = \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 4 & 3 & 1 \end{pmatrix} \begin{pmatrix} 2 & -1 & 2 \\ 0 & 3 & 5 \\ 0 & 0 & 1 \end{pmatrix}$$

perform an LDU decomposition.

All we need to do here is write our $\tilde{U}$ as DU where U is unit upper triangular. We have

$$\begin{pmatrix} 2 & -1 & 2 \\ 0 & 3 & 5 \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & \frac{-1}{2} & 1 \\ 0 & 1 & \frac{5}{3} \\ 0 & 0 & 1 \end{pmatrix}$$

Definition *An $n \times n$ matrix is positive definite provided*

$$\vec{x}^T A \vec{x} > 0 \quad \textit{for all } \vec{x} \neq 0$$

3. If A is symmetric and positive definite then $A = LL^T$ where L is lower triangular. This is known as Cholesky decomposition. If the diagonal entries of L are chosen to be positive, then the decomposition is unique.

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ a_{31} & a_{32} & a_{33} & \cdots & a_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ a_{n1} & a_{n2} & a_{n3} & \cdots & a_{nn} \end{pmatrix} = \begin{pmatrix} \ell_{11} & 0 & 0 & \cdots & 0 \\ \ell_{21} & \ell_{22} & 0 & \cdots & 0 \\ \ell_{31} & \ell_{32} & \ell_{33} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \ell_{n1} & \ell_{n2} & \ell_{n3} & \cdots & \ell_{nn} \end{pmatrix} \begin{pmatrix} \ell_{11} & \ell_{21} & \ell_{31} & \cdots & \ell_{n1} \\ 0 & \ell_{22} & \ell_{32} & \cdots & \ell_{n2} \\ 0 & 0 & \ell_{33} & \cdots & \ell_{n3} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \ell_{nn} \end{pmatrix}$$

Equating the (1,1) entry gives

$$\ell_{11} = \sqrt{a_{11}}$$

Clearly, a_{11} must be ≥ 0 which is guaranteed by the fact that A is positive definite (just choose $\vec{x} = (1, 0, \dots, 0)^T$). Next we see that

$$\ell_{11}\ell_{i1} = a_{1i} = a_{i1} \Rightarrow \ell_{i1} = \frac{a_{i1}}{\ell_{11}}, \quad i = 2, 3, \dots, n$$

Then to find the next diagonal entry we have

$$\ell_{21}^2 + \ell_{22}^2 = a_{22} \Rightarrow \ell_{22} = \left(a_{22} - \ell_{21}^2\right)^{1/2}$$

and the remaining terms in the second row found from

$$\ell_{i2} = \frac{a_{i2} - \ell_{i1}\ell_{21}}{\ell_{22}} \quad i = 3, 4, \dots, n$$

Continuing in this manner and similar to obtaining the equations for the LU factorization we have the following algorithm.

Let A be a symmetric, positive definite matrix. Then the Cholesky factorization $A = LL^T$ is given by the following algorithm

For $i = 1, 2, 3, \ldots, n$

$$\ell_{ii} = \Big(a_{ii} - \sum_{j=1}^{i-1} \ell_{ij}^2\Big)^{1/2}$$

for $k = i+1, \ldots, n$

$$\ell_{ki} = \ell_{ik} = \frac{1}{\ell_{ii}}\Big[a_{ki} - \sum_{j=1}^{i-1} \ell_{kj}\ell_{ij}\Big]$$

Operation Count

One way to compare the work required to solve a linear system by different methods is to determine the number of operations required to find the solution. We first look at the number of operations required to multiply a vector by a matrix, then to perform a back solve and finally to perform an LU decomposition. For homework you will be asked to do an operation count for the decomposition of a tridiagonal matrix.

Multiplication of an n-vector by an $n \times n$ matrix. Suppose we want to perform the multiplication

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ a_{31} & a_{32} & a_{33} & \cdots & a_{3n} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ a_{n1} & a_{n2} & a_{n3} & \cdots & a_{nn} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n \end{pmatrix}$$

We know that the result is a vector $\vec{b}$ so if we can determine the number of operations required to compute one component of $\vec{b}$ then we can simply multiply this result by n. To compute b_1 we have to perform

$$b_1 = \sum_{j=1}^{n} a_{1j} x_j$$

Consequently we need to multiply each a_{1j} times x_j and because there are n terms we have n multiplications. We also need to sum up the n terms so we have $n-1$ additions. So for all n components we have $n(n)$ multiplications and $n(n-1)$ additions. We say that it requires order n^2 multiplications and a like number of additions which means a constant times n^2. We denote this as $\mathcal{O}(n^2)$ and say "order n squared". Note that we have not included the $-n$ terms in the additions because this is negligible for large n.

Back solve. Recall the equations for performing a backsolve.

Given an $n \times n$ upper triangular matrix U with entries u_{ij} and an n-vector $\vec{b}$ with components b_i then the solution of $U\vec{x} = \vec{b}$ is given by the following algorithm.

Set $x_n = \dfrac{b_n}{u_{nn}}$

For $i = n-1, n-2, \ldots, 1$

$$x_i = \frac{b_i - \sum_{j=i+1}^{n} u_{i,j} x_j}{u_{ii}}$$

Now computing the number of operations in this case is a bit more complicated than a matrix vector multiplication. For x_n we require one division; we will count multiplications and divisions the same. For x_{n-1} we have one multiplication, one division and one addition. For x_{n-2} we have two multiplications, one division and two additions. We have

Component	multiplications	divisions	additions
x_n	0	1	0
x_{n-1}	1	1	1
x_{n-2}	2	1	2
x_{n-3}	3	1	3
$\vdots$			
x_1	n	1	n-1

So counting multiplications and divisions as the same we have

$$(n) + (1 + 2 + 3 + \cdots + n) = n + \sum_{i=1}^{n} i \quad \text{multiplications/divisions}$$

and

$$\sum_{i=1}^{n-1} i \quad \text{additions}$$

Now we would like to have the result in terms of $\mathcal{O}(n^r)$ for some r. If you recall from calculus

$$\sum_{i=1}^{p} i = \frac{p(p+1)}{2} \qquad \sum_{i=1}^{p} i^2 = \frac{p(p+1)(2p-1)}{6}$$

Using this first expression we obtain

$$n + \frac{n^2+n}{2} = \mathcal{O}(n^2) \quad \text{multiplications/divisions}$$

and

$$\frac{(n-1)^2 + (n-1)}{2} = \mathcal{O}(n^2) \quad \text{additions}$$

So we say that performing a back solve requires $\mathcal{O}(n^2)$ operations.

LU factorization We now want to demonstrate that an LU factorization requires $\mathcal{O}(n^3)$ operations so that the bulk of the work is required for the LU factorization. Note that when n is small there is not that much difference in $\mathcal{O}(n^2)$ and $\mathcal{O}(n^3)$ but when n is large, it is hugely different. It is important to know the power of n in each of the major operations. The same is true of GE; to transform a system to an equivalent upper triangular system requires $\mathcal{O}(n^3)$. Recall that back solving (and consequently forward solving) requires $\mathcal{O}(n^2)$ operations so to solve $A\vec{x} = \vec{b}$ by LU or GE requires $\mathcal{O}(n^3)$. If we have a symmetric, positive definite system, then the operation count should be approximately half what it is for the full matrix but that still means it's $\mathcal{O}(n^3)$.

Lets look at the first few terms until we see a pattern.

Component	multiplications/div	additions/sub
u_{1j}	0	0
$\ell_{i1}, i = 2, \ldots, n$	(n-1)(1)	0
$u_{2j}, j = 2, \ldots, n$	(n-1)(1)	(n-1)(1)
$\ell_{i2}, i = 3, \ldots, n$	(n-2)(2)	(n-1)(1)
$u_{3j}, j = 3, \ldots, n$	(n-2)(2)	(n-2)(2)

If we consider L we have

$$(n-1)(1) + (n-2)(2) + (n-3)(3) + \cdots + (1)(n-1) = \sum_{i=1}^{n-1} i(n-i) = n\sum_{i=1}^{n-1} i - \sum_{i=1}^{n-1} i^2$$

multiplications/divisions and using our formulas from calculus

$$n\sum_{i=1}^{n-1} i - \sum_{i=1}^{n-1} i^2 = n\Big(\frac{n(n-1)}{2}\Big) - \Big(\frac{(n-1)n(2n-3)}{6}\Big) \approx \frac{n^3}{2} - \frac{n^3}{3} = \mathcal{O}(n^3)$$

We should get the same operation count for U so that the LU factorization requires $\mathcal{O}(n^3)$ operations.

Norms

We said that one of our goals was to determine whether a matrix was well-conditioned; that is, if we perturb the data in a linear system by a small amount then we expect the solution to be changed by a small amount. We have seen that the linear system can be written as $A\vec{x} = \vec{b}$ so what we are really saying is if we perturb the vector $\vec{b}$ or the matrix A by a small amount then the vector solution $\vec{x}$ should change by a small amount. But what do we mean formally by changing a vector or a matrix by a small amount? The concept of norm will be useful in this case and in many more settings.

The Euclidean length of a vector is actually a norm. We call that this is found by

$$\sqrt{\sum_{i=1}^{n} x_i^2}$$

We want to generalize this concept to include other measures of a norm. We can view the Euclidean length as a map (or function) whose domain is R^n and whose range is all scalars i.e., $f : \mathrm{R}^n \rightarrow \mathrm{R}^1$. What properties does this Euclidean length have? We know that the length is always ≥ 0 and only $=0$ if the vector is identically zero. We know what multiplication of a vector by a scalar k does to the length; i.e., it changes by the length by $|k|$. Also, from the triangle inequality we know that the length of the sum of two vectors is always $\leq$ the sum of the two lengths. We combine these properties into a formal definition of a vector.

Definition A vector norm, denoted $\|\vec{x}\|$, is a map from R^n to R^1 which has the properties

1. $\|\vec{x}\| \geq 0$ and $= 0$ only if $\vec{x} = \vec{0}$
2. $\|k\vec{x}\| = |k|\|\vec{x}\|$
3. $\|\vec{x}\| + \|\vec{y}\| \leq \|\vec{x}\| + \|\vec{y}\|$

for all $\vec{x}, \vec{y} \in \mathrm{R}^n$.

We can have other ways to measure vectors. All we have to do is find a map which satisfies the above three conditions; however, practically it should be useful. Three of the most useful vector norms are defined below.

Definition Most common vector norms are:

1. Euclidean norm, denoted $\|\vec{x}\|_2$ and defined by $\|\vec{x}\|_2 = \sqrt{\sum_{i=1}^{n} x_i^2}$
2. Max or infinity norm, denoted $\|\vec{x}\|_\infty$ and defined by $\|\vec{x}\|_\infty = \max_{1\leq i\leq n} |x_i|$
3. one-norm, denoted $\|\vec{x}\|_1$ and defined by $\|\vec{x}\|_1 = \sum_{i=1}^{n} |x_i|$

Example Determine $\|\vec{x}\|_1$, $\|\vec{x}\|_2$ and $\|\vec{x}\|_\infty$ for each vector.

$$\vec{x} = \begin{pmatrix} -3 \\ 2 \\ 4 \\ -7 \end{pmatrix}$$

$$\|\vec{x}\|_1 = |-3| + 2 + 4 + |-7| = 16 \qquad \|\vec{x}\|_2 = \sqrt{9+4+16+49} = \sqrt{78}$$
$$\|\vec{x}\|_\infty = \max\{|-3|, 2, 4, |-7|\} = 7$$

Example It is interesting to sketch the unit ball for each norm, i.e., sketch all points in R^2 such that

$$\{(x_1, x_2) \quad \text{such that} \quad \|\vec{x}\|_p = 1\} \quad \text{for } p = 1, 2, \infty$$

Many times we will use a norm to measure the length of an error vector, i.e., we will associate a number with a vector. In the previous example we saw that different norms give us different numbers for the same vector. How different can these numbers be? Each norm actually measures a different attribute of a norm. However, should we be worried that if we can show a particular norm of the error goes to zero, then will the other norms go to zero too? The following definition helps us to quantify this concept.

Definition Let $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ denote any two vector norms Then these norms are norm-equivalent if there exists constants C_1, C_2 greater than zero such that

$$C_1\|\vec{x}\|_\beta \leq \|\vec{x}\|_\alpha \leq C_2\|\vec{x}\|_\beta \quad \text{for all } \vec{x}$$

Note that if this inequality holds, we also have the equivalent statement

$$\frac{1}{C_2}\|\vec{x}\|_\alpha \leq \|\vec{x}\|_\beta \leq \frac{1}{C_1}\|\vec{x}\|_\alpha \quad \text{for all } \vec{x}$$

If two norms are norm-equivalent and we have that $\|\vec{x}\|_\beta \to 0$ then clearly $\|\vec{x}\|_\alpha \to 0$.

We claim that any pair of our three vector norms are norm-equivalent. We show one set here and you are asked to do another set for homework.

Example The norms $\|\vec{x}\|_\infty$ and $\|\vec{x}\|_2$ are equivalent.
We have

$$\|\vec{x}\|_2^2 = \sum_{i=1}^n x_i^2 \geq \max |x_i|^2 = \|\vec{x}\|_\infty^2$$

so $C_1 = 1$.

Also

$$\|\vec{x}\|_2^2 = \sum_{i=1}^n x_i^2 \leq n \max |x_i|^2 = n\|\vec{x}\|_\infty^2$$

so $C_2 = \sqrt{n}$.

$$\|\vec{x}\|_\infty \leq \|\vec{x}\|_2 \leq \sqrt{n}\|\vec{x}\|_\infty \quad \text{for all } \vec{x}$$

Our next goal is to associate a matrix with a number; i.e., we want to define a matrix norm.