Tutorials and worked examples for simulation, curve fitting, statistical analysis, and plotting.
http://www.simfit.org.uk

The LU factorization of a matrix, the matrix norms, and condition numbers of a matrix are of considerable value in solving non-singular linear systems.

From the main SIMFIT menu choose [Statistics], then [Numerical analysis], and select the option to calculate the LU decomposition of a matrix, noting that the following arbitrary 5 by 5 matrix A is contained in the default test file `matrix.tf1`.

Matrix A

1.20	4.50	6.10	7.20	8.00
3.00	5.60	3.70	9.10	12.5
17.1	23.4	5.50	9.20	3.30
7.15	5.87	9.94	8.82	10.8
12.4	4.30	7.70	8.95	1.60

Proceeding to the analysis of this matrix yields the following results.

LU factorization of matrix A

Matrix 1-norm = 43.67, Condition number = 3.6940420E+01
Matrix ∞-norm = 58.50, Condition number = 2.6184088E+01

Lower triangular/trapezoidal L where $A = PLU$

1				
7.2514620E-01	1			
7.0175439E-02	-2.2559202E-01	1		
1.7543860E-01	-1.1798920E-01	4.8433085E-01	1	
4.1812865E-01	3.0897383E-01	9.9116389E-01	-6.3185748E-01	1

Upper triangular/trapezoidal U where $A = PLU$

1.7100000E+01	2.3400000E+01	5.5000000E+00	9.2000000E+00	3.3000000E+00
	-1.2668421E+01	3.7116959E+00	2.2786550E+00	-7.9298246E-01
		6.5513641E+00	7.0684324E+00	7.5895305E+00
			4.3313617E+00	8.1516455E+00
				7.2933958E+00

This calculation produces the L and U factors after the standard pivoting operations, familiar in Gaussian elimination, so that the matrix equation is not simply $A = LU$ but

$$A = PLU$$

where P is a permutation matrix. To appreciate this fact, note that multiplying L and U together does not give matrix A directly, but yields the following matrix instead with rows of A permuted.

Matrix LU

17.1	23.4	5.50	9.20	3.30
12.4	4.30	7.70	8.95	1.60
1.20	4.50	6.10	7.20	8.00
3.00	5.60	3.70	9.10	12.5
7.15	5.87	9.94	8.82	10.8

So it is convenient to consider the sequence of pivots applied to matrix A which are contained in the vector

$$IPIV = (3, 5, 3, 5, 5).$$

This is the sequence of row exchanges applied to matrix A in order to rearrange it as the calculation proceeds. That is, row i of a m by n matrix is interchanged with row $IPIV(i)$ for $i = 1, 2, \ldots, \min(m, n)$. Clearly matrix A was transformed into LU by these pivots, which can therefore be used to calculate the permutation matrix P needed to represent A directly in terms of L and U if required.

As the LU representation is of interest in the solution of linear equations, this procedure also calculates the matrix norms and condition numbers needed to assess the sensitivity of the solutions to perturbations when the matrix is square. Given a vector norm $\|.\|$, a matrix A, and the set of vectors x where $\|x\| = 1$, the matrix norm subordinate to the vector norm is

$$\|A\| = \max_{\|x\|=1} \|Ax\|.$$

For a m by n matrix A, the three most important norms are

$$\begin{aligned}
\|A\|_1 &= \max_{1 \le j \le n} \left(\sum_{i=1}^{m} |a_{ij}|\right) \\
\|A\|_2 &= (\lambda_{\max}|A^T A|)^{\frac{1}{2}} \\
\|A\|_\infty &= \max_{1 \le i \le m} \left(\sum_{j=1}^{n} |a_{ij}|\right),
\end{aligned}$$

so that the 1-norm is the maximum absolute column sum, the 2-norm is the square root of the largest eigenvalue of $A^T A$, and the infinity norm is the maximum absolute row sum. The condition numbers estimated are

$$\begin{aligned}
\kappa_1(A) &= \|A\|_1 \|A^{-1}\|_1 \\
\kappa_\infty(A) &= \|A\|_\infty \|A^{-1}\|_\infty \\
&= \kappa_1(A^T)
\end{aligned}$$

which satisfy $\kappa_1 \ge 1$, and $\kappa_\infty \ge 1$ and they are included in the tabulated output unless A is in singular, when they are infinite. For a perturbation δb to the right hand side of a linear system with $m = n$ we have

$$\begin{aligned}
Ax &= b \\
A(x + \delta x) &= b + \delta b \\
\frac{\|\delta x\|}{\|x\|} &\le \kappa(A) \frac{\|\delta b\|}{\|b\|},
\end{aligned}$$

while a perturbation δA to the matrix A leads to

$$\begin{aligned}
(A + \delta A)(x + \delta x) &= b \\
\frac{\|\delta x\|}{\|x + \delta x\|} &\le \kappa(A) \frac{\|\delta A\|}{\|A\|},
\end{aligned}$$

and, for complete generality,

$$\begin{aligned}
(A + \delta A)(x + \delta x) &= b + \delta b \\
\frac{\|\delta x\|}{\|x\|} &\le \frac{\kappa(A)}{1 - \kappa(A)\|\delta A\|/\|A\|} \left(\frac{\|\delta A\|}{\|A\|} + \frac{\|\delta b\|}{\|b\|}\right)
\end{aligned}$$

provided $\kappa(A)\|\delta A\|/\|A\| < 1$. These inequalities estimate bounds for the relative error in computed solutions of linear equations, so that a small condition number indicates a well-conditioned problem, a large condition number indicates an ill-conditioned problem, while an infinite condition number indicates a singular matrix and no solution. To a rough approximation; if the condition number is 10^k and computation involves n-digit precision, then the computed solution will have about $(n - k)$-digit precision.