

Programming, numerics and optimization

Lecture B-3:

Linear systems I: Direct and iterative methods

Łukasz Jankowski

ljank@ippt.pan.pl

Institute of Fundamental Technological Research

Room 4.32, Phone +22.8261281 ext. 428

April 13, 2021¹

¹Current version is available at <http://info.ippt.pan.pl/~ljank>.

Methods of solution

The iterative methods

- Compute an *approximate* solution in an iterative way.
- In general, the number of iterations to obtain a good approximation is unknown in advance.
- The convergence properties of the methods are often hard to analyze, especially in finite precision arithmetics. In some cases, finite-precision iterations may not converge at all.
- Many of the methods require access only to a matrix–vector multiplication procedure (\mathbf{Ax} , sometimes also $\mathbf{A}^T \mathbf{x}$), so that \mathbf{A} may be given only implicitly (e.g. \mathbf{Ax} can be the result of an experiment or a simulation).
- They are most useful for solving large sparse or structured systems, for which any factorization (necessary in direct methods) would take too much time, destroy the sparsity or the matrix structure, or be too inaccurate.

Outline

2 Existence and uniqueness of solution

Square diagonal matrices

$$\mathbf{D} = \text{diag}(d_1, d_2, \dots, d_n) = \begin{bmatrix} d_1 & & & 0 \\ & d_2 & & \\ & & \ddots & \\ 0 & & & d_n \end{bmatrix}$$

- Elements of a **diagonal matrix** are all zero except the diagonal.
- If $d_i \neq 0$ for all i , then \mathbf{D} is full-rank and bijective. The system $\mathbf{D}\mathbf{x} = \mathbf{y}$ is then uniquely solvable for all \mathbf{y} .
- If $d_i = 0$ for some i , then \mathbf{D} is singular and neither surjective nor bijective. Depending on \mathbf{y} , the system $\mathbf{D}\mathbf{x} = \mathbf{y}$ has either infinitely many solutions or no solutions at all.

Square diagonal matrices

The linear system $\mathbf{D}\mathbf{x} = \mathbf{y}$, where \mathbf{D} is an $n \times n$ diagonal matrix, is a system of n decoupled equations:

$$d_i x_i = y_i, \quad i = 1, 2, \dots, n,$$

which can be solved in $O(n)$ operations:

- If all $d_i \neq 0$, then the *unique solution* is given by $x_i = y_i/d_i$.
- Otherwise there is i such that $d_i = 0$. The matrix \mathbf{D} (and thus also \mathbf{A}) is singular and neither surjective nor injective.

Existence of solution depends on \mathbf{y} :

- If $y_i = 0$ for all i such that $d_i = 0$, then there are *infinitely many solutions*, since equation $0x_i = 0$ is satisfied by any x_i .
- If there exists i such that $d_i = 0$ and $y_i \neq 0$, then there are *no solutions*, since no x_i can satisfy $0x_i = y_i$, where $y_i \neq 0$.

The case of a non-square diagonal matrix can be treated in a similar way.

Diagonal matrices — examples

Square diagonal matrix

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ y \\ 1 \end{bmatrix}$$

If $y = 0$, then the equation has infinitely many solutions:

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 1 \\ 1/2 \\ 0 \\ 1/4 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ c \\ 0 \end{bmatrix}, \quad c \in \mathbb{R},$$

where $[1 \ 1/2 \ 0 \ 1/4]^T$ is a particular solution and $[0 \ 0 \ c \ 0]^T$ belongs to the null space of the system matrix. Otherwise (if $y \neq 0$), the equation has no solutions.

Diagonal matrices — examples

Non-square diagonal matrix

The equation

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

has infinitely many solutions,

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} 1 \\ 1/2 \\ 1/3 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ c \end{bmatrix}, \quad c \in \mathbb{R},$$

where $[1 \ 1/2 \ 1/3 \ 0]^T$ is a particular solution and $[0 \ 0 \ 0 \ c]^T$ belongs to the null space of the system matrix.

Diagonal matrices — examples

Non-square diagonal matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \\ y \end{bmatrix}$$

If $y = 0$, the equation has a unique solution

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 1/2 \\ 1/3 \end{bmatrix}.$$

Otherwise, the equation has no solutions.

Unitary matrices

A real square matrix \mathbf{Q} is called **unitary**, if $\mathbf{Q}^{-1} = \mathbf{Q}^T$.

If a matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$ is unitary, then

- \mathbf{Q} has orthonormal rows and columns, i.e.

$$\mathbf{Q}^T \mathbf{Q} = \mathbf{Q} \mathbf{Q}^T = \mathbf{I},$$

where \mathbf{I} is the n by n diagonal matrix.

- Columns $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n$ (and rows) of \mathbf{Q} form an orthonormal basis in \mathbf{R}^n , that is

$$\mathbf{q}_i^T \mathbf{q}_j = \delta_{ij},$$

where δ_{ij} is Kronecker's delta.

- Unitary matrices are thus always full-rank.

Permutation matrices

A square $n \times n$ matrix $\mathbf{\Pi}$ is said to be a **permutation matrix**, if it is obtained from an $n \times n$ identity matrix by permuting its rows.

- Every row and every column of a permutation matrix has exactly single 1 and everywhere else 0s.
- There are $n!$ different permutations of an n -element sequence. So, there are exactly $n!$ permutation matrices of the dimensions $n \times n$.
- A permutation matrix $\mathbf{\Pi}$ satisfies $\mathbf{\Pi}^T \mathbf{\Pi} = \mathbf{I}$, therefore it is a special case of a unitary matrix.
- When applied to an $n \times n$ matrix \mathbf{A} :
 - $\mathbf{\Pi A}$ is the matrix \mathbf{A} with permuted rows.
 - $\mathbf{A \Pi}$ is the matrix \mathbf{A} with permuted columns.

Upper triangular systems — back-substitution

Similarly, an $n \times n$ upper triangular system $\mathbf{U}\mathbf{x} = \mathbf{y}$,

$$\begin{bmatrix} u_{11} & u_{12} & \cdots & u_{1n} \\ & u_{22} & \cdots & u_{2n} \\ & & \ddots & \vdots \\ \mathbf{0} & & & u_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix},$$

can be solved with $O(n^2)$ operations by [back-substitution](#).

Back substitution

$$x_n = \frac{y_n}{u_{11}},$$

$$x_i = \frac{y_i - \sum_{j=i+1}^n u_{ij}x_j}{u_{ij}}.$$

```
for (int i=n-1; i>=0; --i) {
    sum = 0;
    for (int j=i+1; j<n; ++j)
        sum += U(i, j)*x(j);
    x(i) = (b(i)-sum)/U(i, i);
}
```

Factorizations and decompositions

Solution of the system $\mathbf{Ax} = \mathbf{y}$ by direct methods requires a factorization of \mathbf{A} into a product of two or more special matrices that make the system easier to solve, e.g.

$$\mathbf{A} = \mathbf{Q}_1\mathbf{Q}_2 \quad \text{or} \quad \mathbf{A} = \mathbf{Q}_1\mathbf{Q}_2\mathbf{Q}_3.$$

Direct methods can be broadly classified into two groups

- **Decomposition methods** use a factorization with unitary matrices and (usually) a diagonal matrix, which directly provide important information about the matrix and the related mapping (dimensionality and basis of the null-space, eigen- or singular values, etc.). The factorization (decomposition) often amounts to solving the related eigenproblem.
- The other direct methods are known under the general name of **factorization methods**.

Factorizations

The probably most commonly used factorizations are

- **LU factorization:** $\mathbf{A} = \mathbf{LU}$, where \mathbf{L} and \mathbf{U} are respectively lower and upper triangular matrices. It exists for any square nonsingular matrix³.
- **QR factorization:** $\mathbf{A} = \mathbf{QR}$, where \mathbf{Q} is a unitary (or orthogonal) matrix and \mathbf{R} is an upper triangular matrix. Similarly, there exist QL, RQ and LQ factorizations.
- **LDL factorization:** $\mathbf{A} = \mathbf{LDL}^T$, where \mathbf{L} is a lower triangular matrix and \mathbf{D} is a diagonal matrix with positive elements. LDL factorization exists for symmetric positive definite matrices (for other matrices it may not exist).
- **Cholesky factorization:** $\mathbf{A} = \mathbf{LL}^T$, where \mathbf{L} is a lower triangular matrix. Cholesky factorization exists only for symmetric positive definite matrices.

³Sometimes a pre-multiplication by a permutation matrix $\mathbf{\Pi}$ is necessary, so that $\mathbf{\Pi A} = \mathbf{LU}$.

Decompositions — eigen decomposition

Eigenvalues and eigenvectors

A fundamental notion in linear algebra is that of an eigenvalue and the corresponding eigenvectors of a square matrix.

Let \mathbf{A} be a square $n \times n$ matrix. A number λ is called an **eigenvalue** of \mathbf{A} and a vector \mathbf{v} is called a corresponding **eigenvector** if and only if $\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$.

As the above condition yields $(\mathbf{A} - \lambda\mathbf{I})\mathbf{v} = \mathbf{0}$, the eigenvalues of \mathbf{A} are the roots of its characteristic polynomial,

$$f_{\mathbf{A}}(\lambda) = \det(\mathbf{A} - \lambda\mathbf{I}),$$

which always has n complex roots. Thus, every square $n \times n$ matrix has always n complex eigenvalues (some or all of which can be real). Every eigenvalue has a **multiplicity**, which is defined as the multiplicity of the corresponding root of $f_{\mathbf{A}}(\lambda)$.

Decompositions — eigen decomposition

Symmetric (Hermitian) matrices

The probably most known decomposition is the [eigen decomposition](#) of a symmetric (Hermitian) matrix:

Every square $n \times n$ symmetric (Hermitian) matrix \mathbf{A} has n linearly independent eigenvectors and can be expressed as

$$\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1},$$

where \mathbf{D} is a diagonal matrix with the eigenvalues of \mathbf{A} on the diagonal, $\mathbf{D} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$, and \mathbf{P} is the matrix with the corresponding eigenvectors as columns. The eigenvalues are all real. The eigenvectors are orthogonal, and if they are scaled to be orthonormal, then $\mathbf{P}^{-1} = \mathbf{P}^T$ and

$$\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^T.$$

Decompositions — eigen decomposition

If \mathbf{A} is non-symmetric (non-Hermitian), the existence of the eigen decomposition depends on the number of the eigenvectors:

If a square $n \times n$ matrix \mathbf{A} has n linearly independent eigenvectors, then it can be expressed as

$$\mathbf{A} = \mathbf{P}\mathbf{D}\mathbf{P}^{-1},$$

where \mathbf{P} collects all the eigenvectors as columns, and \mathbf{D} is a diagonal matrix with the corresponding eigenvalues on the diagonal, $\mathbf{D} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$.

If \mathbf{A}

has an eigen decomposition, it is called *diagonalizable*⁴. A non-symmetric (non-Hermitian) matrix may have complex eigenvalues.

⁴It is diagonal in the coordinates defined by the columns of \mathbf{P} .

Decompositions — eigen decomposition — examples

Symmetric matrix

The matrix

$$\mathbf{A} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \quad f_{\mathbf{A}}(\lambda) = \lambda^2 - 2\lambda,$$

is symmetric, thus it has two linearly independent orthogonal eigenvectors with two corresponding real eigenvalues,

$$\begin{aligned} \mathbf{v}_1 &= [-1 \ 1]^T & \mathbf{v}_2 &= [1 \ 1]^T, \\ \lambda_1 &= 2 & \lambda_2 &= 0. \end{aligned}$$

One of the eigenvalues is zero, so \mathbf{A} is singular.

$$\mathbf{A} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} -0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}$$

Decompositions — eigen decomposition — examples

Non-symmetric diagonalizable matrix

The non-symmetric matrix

$$\mathbf{A} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix}, \quad f_{\mathbf{A}}(\lambda) = \lambda^2 - 2\lambda + 2,$$

is diagonalizable, since it has two linearly independent eigenvectors,

$$\mathbf{v}_1 = [-i \ 1]^T \quad \mathbf{v}_2 = [i \ 1]^T.$$

The two corresponding eigenvalues are complex:

$$\lambda_1 = 1 + i \quad \lambda_2 = 1 - i.$$

$$\mathbf{A} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} = \begin{bmatrix} -i & i \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1+i & 0 \\ 0 & 1-i \end{bmatrix} \begin{bmatrix} 0.5i & 0.5 \\ -0.5i & 0.5 \end{bmatrix}$$

Decompositions — eigen decomposition — examples

Non-symmetric non-diagonalizable matrix

The non-symmetric matrix

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad f_{\mathbf{A}}(\lambda) = \lambda^2 - 2\lambda + 1,$$

is non-diagonalizable, since it has only one linearly independent eigenvector,

$$\mathbf{v}_1 = [0 \ 1]^T,$$

even if its single eigenvalue has the multiplicity of two:

$$\lambda_1 = \lambda_2 = 1.$$

The matrix \mathbf{A} has no eigen decomposition and is non-singular.

Decompositions — singular value decomposition

The probably most important decomposition is the [singular value decomposition \(SVD\)](#):

Every rectangular matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ can be expressed as

$$\mathbf{A} = \mathbf{U} \begin{bmatrix} \boldsymbol{\Sigma}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{V}^T = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T,$$

where \mathbf{U} and \mathbf{V} are unitary matrices, and $\boldsymbol{\Sigma}$ is an $n \times m$ diagonal matrix with nonnegative diagonal elements ordered in a nonincreasing way.

Decompositions — singular value decomposition


Every rectangular matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ can be expressed as

$$\mathbf{A} = \mathbf{U} \begin{bmatrix} \boldsymbol{\Sigma}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{V}^T = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T,$$

where \mathbf{U} and \mathbf{V} are unitary, and $\boldsymbol{\Sigma}$ is a diagonal matrix.

full-rank $\mathbf{A} \in \mathbb{R}^{n \times m}$ (rank $\mathbf{A} = n < m$)

$$\mathbf{A} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T =$$



Decompositions — singular value decomposition

Every rectangular matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ can be expressed as

$$\mathbf{A} = \mathbf{U} \begin{bmatrix} \boldsymbol{\Sigma}_0 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{V}^T = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T,$$

where \mathbf{U} and \mathbf{V} are unitary, and $\boldsymbol{\Sigma}$ is a diagonal matrix.

- Diagonal elements σ_i of $\boldsymbol{\Sigma}$ are called the **singular values** of \mathbf{A} .
- The number r of positive singular values equals to $\text{rank } \mathbf{A}$.
The matrix $\boldsymbol{\Sigma}_0$ is thus $\text{rank } \mathbf{A} \times \text{rank } \mathbf{A}$.
- The SVD is unique up to the ordering of the singular vectors (columns of \mathbf{U} and \mathbf{V}) corresponding to equal singular values.
- The columns of \mathbf{V} , which correspond to the vanishing singular values, form a basis for the null-space of \mathbf{A} .
- $\mathbf{A}^T \mathbf{A} = \mathbf{V} \boldsymbol{\Sigma}^T \mathbf{U}^T \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T = \mathbf{V} \boldsymbol{\Sigma}^T \boldsymbol{\Sigma} \mathbf{V}^T$ is the eigen decomposition of $\mathbf{A}^T \mathbf{A}$. Therefore, $\sigma_i^2(\mathbf{A}) = \lambda_i(\mathbf{A}^T \mathbf{A})$.
- The singular values provide full information about conditioning of \mathbf{A} (see Lecture B-4).

Decompositions — singular value decomposition

Examples

Non-symmetric non-diagonalizable square matrix

The non-symmetric and non-diagonalizable matrix

$$\mathbf{A} = \begin{bmatrix} 2 & 0 \\ 3 & 2 \end{bmatrix}$$

has the following singular value decomposition:

$$\mathbf{A} = \begin{bmatrix} \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} 4 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \end{bmatrix}^T$$

Decompositions — singular value decomposition

Examples

Non-square matrix

The non-square matrix

$$\mathbf{A} = \begin{bmatrix} 2 & 0 & 0 \\ 3 & 2 & 0 \end{bmatrix}$$

has the following singular value decomposition:

$$\mathbf{A} = \begin{bmatrix} \frac{1}{\sqrt{5}} & -\frac{2}{\sqrt{5}} \\ \frac{2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{bmatrix} \begin{bmatrix} 4 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \frac{2}{\sqrt{5}} & -\frac{1}{\sqrt{5}} & 0 \\ \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} & 0 \\ 0 & 0 & 1 \end{bmatrix}^T$$

Decompositions — rank-revealing QR factorization

Another decomposition is called the **rank-revealing QR factorization**:

Let the rank of $\mathbf{A} \in \mathbb{R}^{n \times m}$ be $r = \text{rank } \mathbf{A} \leq \min(n, m)$.
There exists a rank-revealing QR factorization,

$$\mathbf{A}\mathbf{\Pi} = \begin{bmatrix} \mathbf{Q}_1 & \mathbf{Q}_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_1 & \mathbf{R}_2 \\ \mathbf{0} & \mathbf{0} \end{bmatrix},$$

where \mathbf{R}_1 is an $r \times r$ upper triangular matrix with positive diagonal elements, \mathbf{R}_2 is an $r \times (m - r)$ matrix and $\mathbf{\Pi}$ is a permutation matrix such that the first r columns of $\mathbf{A}\mathbf{\Pi}$ are linearly independent. The matrix $[\mathbf{Q}_1 \mathbf{Q}_2]$ is unitary. The

matrices \mathbf{Q}_1 and \mathbf{Q}_2 are (column) orthonormal:

- The r columns of \mathbf{Q}_1 form a basis for $\text{range } \mathbf{A}\mathbf{\Pi}$.
- The $n - r$ columns of \mathbf{Q}_2 form a basis for $\ker(\mathbf{A}\mathbf{\Pi})^\top$.

Gaussian elimination

Gaussian elimination is a method of

- solving a full-rank $\mathbf{Ax} = \mathbf{y}$ by
- performing the LU factorization of \mathbf{A} .

Gaussian elimination uses two elementary operations

- 1 adding a multiple of the i th row to the j th row and
- 2 interchanging two rows/equations (or columns/unknowns), called *pivoting*

to eliminate the unknowns x_i in order to obtain an equivalent upper triangular system $\mathbf{Ux} = \hat{\mathbf{y}}$, which can be solved by back-substitution.

Gaussian elimination

$$\mathbf{Ax} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = \mathbf{y}$$

First, x_1 is eliminated from the other equations by subtracting the multiple $l_{i1} = a_{i1}/a_{11}$ of the first equation from the i th equation. This way a reduced system of $n - 1$ equations with $n - 1$ unknowns x_2, \dots, x_n is obtained,

$$\mathbf{A}^{(2)}\mathbf{x} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ 0 & a_{22}^{(2)} & \cdots & a_{2n}^{(2)} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & a_{n2}^{(2)} & \cdots & a_{nn}^{(2)} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2^{(2)} \\ \vdots \\ y_n^{(2)} \end{bmatrix} = \mathbf{y}^{(2)},$$

where $a_{ij}^{(2)} = a_{ij} - l_{i1}a_{j1}$ and $y_i^{(2)} = b_i - l_{i1}b_1$.

Gaussian elimination

The procedure is repeated $n - 1$ times: x_k , $k = 2, \dots, n - 1$ is eliminated from the rest $i = k + 1, \dots, n$ equations using the multiplier $l_{ik} = a_{ik}^{(k)} / a_{kk}^{(k)}$. This yields an upper triangular system, which can be solved by back-substitution:

$$\mathbf{A}^{(n)} \mathbf{x} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ & a_{22}^{(2)} & \cdots & a_{2n}^{(2)} \\ & & \ddots & \vdots \\ \mathbf{0} & & & a_{nn}^{(n)} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2^{(2)} \\ \vdots \\ y_n^{(n)} \end{bmatrix} = \mathbf{y}^{(n)},$$

where $a_{ij}^{(k+1)} = a_{ij}^{(k)} - l_{ik} a_{ik}^{(k)}$ and $y_i^{(k+1)} = b_i^{(k)} - l_{ik} b_k^{(k)}$.

Gaussian elimination with pivoting

Gaussian elimination uses in the k th step the multiplier

$l_{ik} = a_{ik}^{(k)} / a_{kk}^{(k)}$ to eliminate the unknown x_k . This works seamlessly, unless the diagonal element $a_{kk}^{(k)} \approx 0$.

- If $a_{kk}^{(k)} = 0$, elimination is not possible.
- If $a_{kk}^{(k)}$ is very small, elimination can be numerically unstable.

The second elementary operation (row pivoting) can be then used. The k th row is interchanged with one of the next rows, so that the diagonal element is maximized. Some procedures pivot also columns (interchange the unknowns). Before performing the Gaussian elimination with pivoting, the matrix should be rescaled

- The maximum magnitude of the elements in each row to one.
- The sum of the magnitudes of the row elements to one.

There is no need for pivoting, if \mathbf{A} is symmetric positive definite.

Gaussian elimination with pivoting — example

Example (Dahlquist and Björck)

$$\mathbf{Ax} = \begin{bmatrix} 10^{-6} & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \mathbf{y}.$$

\mathbf{A} is nonsingular and well-conditioned. The exact solution is

$$\mathbf{x} = \frac{1}{1 - 10^{-6}} \begin{bmatrix} -1 \\ 1 \end{bmatrix} \approx \begin{bmatrix} -1 \\ 1 \end{bmatrix}.$$

However, the solution computed by Gaussian elimination *without pivoting* in the precision of 6 decimal digits yields

$$\mathbf{x} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

Gaussian elimination — numerical costs

After all steps an upper triangular system is produced,

$$\mathbf{A}^{(n)} \mathbf{x} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ & a_{22}^{(2)} & \cdots & a_{2n}^{(2)} \\ & & \ddots & \vdots \\ \mathbf{0} & & & a_{nn}^{(n)} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2^{(2)} \\ \vdots \\ y_n^{(n)} \end{bmatrix} = \mathbf{y}^{(n)}.$$

The number of operations in Gaussian elimination is $\sim n^3/3$. This is substantially more than $\sim n^2/2$ necessary to solve the resulting upper triangular system.

If the multipliers l_{ik} are stored (together with the information about row interchange and scaling, if necessary), then $\mathbf{y}^{(n)}$ for different right-hand side vectors \mathbf{y} can be computed at later times at the cost of $\sim n^2/2$ only. Hence the total cost of each subsequent computation would be $\sim n^2$ only.

Stationary methods for $\mathbf{Ax} = \mathbf{y}$

- Split the system matrix \mathbf{A} into

$$\mathbf{A} = \mathbf{M} - \mathbf{N},$$

such that \mathbf{M} is a nonsingular special matrix (diagonal, lower triangular, etc.) Direct access to \mathbf{A} is thus required.

- Compute iteratively the solution using an easily solvable⁵

$$\mathbf{M}\mathbf{x}_{k+1} = \mathbf{N}\mathbf{x}_k + \mathbf{y}.$$

- Convergence properties are usually well-analyzed, e.g.

If $\mathbf{A} \in \mathbb{R}^{n \times n}$ is nonsingular and the spectral radius⁶ of $\mathbf{M}^{-1}\mathbf{N}$ satisfies $\rho(\mathbf{M}^{-1}\mathbf{N}) < 1$, then

- \mathbf{x}_k converge to $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y}$ for any starting \mathbf{x}_0 ,
- the error $\|\mathbf{x} - \mathbf{x}_k\|$ tends to zero like $\rho(\mathbf{M}^{-1}\mathbf{N})^k$.

⁵The subscript k in \mathbf{x}_k denotes the iteration number, and not the k th component x_k of the (k th iterate) vector \mathbf{x}_k .

⁶Magnitude of the maximum-magnitude eigenvalue.

Stationary methods — Jacobi and Gauss-Seidel

Stationary methods for $\mathbf{Ax} = \mathbf{y}$ are defined by

$$\mathbf{M}\mathbf{x}_{k+1} = \mathbf{N}\mathbf{x}_k + \mathbf{y}, \quad \text{where } \mathbf{A} = \mathbf{M} - \mathbf{N}.$$

Split the system matrix \mathbf{A} into its strictly lower triangular \mathbf{L} , diagonal \mathbf{D} and strictly upper triangular \mathbf{U} parts,

$$\mathbf{A} = \mathbf{L} + \mathbf{D} + \mathbf{U}.$$

- The **Jacobi method** is defined by

$$\mathbf{M} = \mathbf{D},$$

$$\mathbf{N} = -\mathbf{L} - \mathbf{U}.$$

- The **Gauss-Seidel method** is defined by

$$\mathbf{M} = \mathbf{D} + \mathbf{L},$$

$$\mathbf{N} = -\mathbf{U}.$$

Both method are convergent if \mathbf{A} is symmetric positive-definite.

Stationary methods — Jacobi and Gauss-Seidel

Therefore, every linear system $\mathbf{Ax} = \mathbf{y}$ can be transformed to the equivalent form

$$\mathbf{D}\mathbf{x} = \mathbf{y} - (\mathbf{L} + \mathbf{U})\mathbf{x}.$$

The Jacobi method takes this formula directly and obtains

$$\mathbf{D}\mathbf{x}_{k+1} = \mathbf{y} - (\mathbf{L} + \mathbf{U})\mathbf{x}_k,$$

where all components of \mathbf{x}_{k+1} are computed using the previous-step iterate \mathbf{x}_k . However, when the i th component of the vector \mathbf{x}_{k+1} is being computed, then all the components preceding it (no. $1, 2, \dots, i - 1$) are already known and can be used instead of these of the previous-step iterate \mathbf{x}_k . This yields the Gauss-Seidel method,

$$\begin{aligned} \mathbf{D}\mathbf{x}_{k+1} &= \mathbf{y} - \mathbf{L}\mathbf{x}_{k+1} - \mathbf{U}\mathbf{x}_k, & \text{that is} \\ (\mathbf{D} + \mathbf{U})\mathbf{x}_{k+1} &= \mathbf{y} - \mathbf{U}\mathbf{x}_k. \end{aligned}$$

Krylov subspace methods — CGLS method

after Björck, Dahlquist

If \mathbf{A} is not symmetric positive-definite, then the CG method can be applied to the *normal equations*,

$$\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{y},$$

which amounts to minimizing

$$\phi(x) = \frac{1}{2} \|\mathbf{A} \mathbf{x} - \mathbf{y}\|^2.$$

This is the [conjugate gradient least-squares \(CGLS\)](#) method. \mathbf{A} is required only in the form of the multiplications $\mathbf{A} \mathbf{x}$ and $\mathbf{A}^T \mathbf{y}$.

$$\mathbf{r} = \mathbf{y} - \mathbf{A} \mathbf{x}$$

$$\mathbf{p} = \mathbf{s}_0 = \mathbf{A}^T \mathbf{r}$$

while $\|\mathbf{r}\| > \epsilon$

$$\mathbf{q} = \mathbf{A} \mathbf{p}$$

$$\alpha = \|\mathbf{s}_0\|^2 / \|\mathbf{q}\|^2$$

$$\mathbf{x} = \mathbf{x} + \alpha \mathbf{p}$$


$$\mathbf{r} = \mathbf{r} - \alpha \mathbf{q}$$

$$\mathbf{s}_1 = \mathbf{A}^T \mathbf{r}$$

$$\beta = \|\mathbf{s}_1\|^2 / \|\mathbf{s}_0\|^2$$

$$\mathbf{p} = \mathbf{s}_1 + \beta \mathbf{p}$$

$$\mathbf{s}_0 = \mathbf{s}_1$$



Outline

5 Homework 7

Homework 7 (25 points)

LU decomposition

Available soon at `http://info.ippt.pan.pl/~ljank`.

E-mail the answer and the source codes to `ljank@ippt.pan.pl`.