H. Numerical linear algebra

V. Blobel

Inst. für Experimentalphysik, Hamburg

Sommersemester 2002

1. Vectors and matrices

The product of the $m \times n$ matrix $\boldsymbol{A}$ and the $n \times k$ matrix $\boldsymbol{B}$ is a $m \times k$ matrix $\boldsymbol{C}$ $\quad (m \cdot n \cdot k$ single products):

$$\boldsymbol{C} = \boldsymbol{A}\,\boldsymbol{B} \qquad\qquad C_{ij} = \sum_{k=1}^{n} A_{ik}\,B_{kj}\ ,$$

in general $\boldsymbol{AB} \neq \boldsymbol{BA}$.

Formulas:

$$\begin{aligned} \boldsymbol{A}\,(\boldsymbol{BC}) &= \boldsymbol{AB} + \boldsymbol{AC} & \qquad \boldsymbol{A}^{-1}\boldsymbol{A} &= \boldsymbol{A}\boldsymbol{A}^{-1} = \mathbf{1} \quad \text{inverse}^{a} \\ (\boldsymbol{A}+\boldsymbol{B})^T &= \boldsymbol{A}^T + \boldsymbol{B}^T & \qquad \left(\boldsymbol{A}^{-1}\right)^T &= \left(\boldsymbol{A}^T\right)^{-1} \\ (\boldsymbol{AB})^T &= \boldsymbol{B}^T\boldsymbol{A}^T & \qquad (\boldsymbol{AB})^{-1} &= \boldsymbol{B}^{-1}\boldsymbol{A}^{-1} \quad \text{note}^{b} \end{aligned}$$

[a] $\boldsymbol{A}$ is square matrix [b] Note: if individual inverses exist.

Library of matrix routines

CERN program library

LAPACK

NAG library

William H. Press, Saul A. Teukolsky, William T. Vetterling, Brian P. Flannery,
Numerical recipes in C (Pascal, Fortran . . .), The Art of Scientific Computing, Cambridge University Press.

Matrix operations

matrix product: direct solution; time $\propto$ product $m \cdot n \cdot k$; order may influence time in multiple products.

matrix inversion: direct solution with time $\propto n^3$, can be reduced for symmetric matrices. Fails for singular matrices, unless . . .
Precision depends on condition number (e.g. large for least squares solution with highly correlated parameters) and may be rather limited.
Iterative solutions exist especially for large matrices.

symmetric matrix diagonalization: iterative methods (no direct solution) with time $\gg$ time for inversion.
Allows to recognize singularities (zero or small eigenvalues).

singular value decomposition: iterative methods (no direct solution) with time $\gg$ time for inversion.
Numerically precise method for solution of system of linear equations.
Allows to recognize singularities (zero or small singular values).

Linear mapping

Set of simultaneous equations

$$\boldsymbol{A}\,\boldsymbol{x} = \boldsymbol{y} \qquad \text{square} \quad n \times n \text{ matrix } \boldsymbol{A}$$

Matrix $\boldsymbol{A}$ defines linear mapping from vector space $\boldsymbol{x}$ to the vector space $\boldsymbol{y}$.

- If $\boldsymbol{A}$ is singular:
  a subspace of $\boldsymbol{x}$, called *nullspace*, is mapped to zero: $\boldsymbol{A}\,\boldsymbol{x} = 0$
- subspace of $\boldsymbol{y}$, called *range*, that can be reached by $\boldsymbol{A}$:
  dimension of range is *rank* of $\boldsymbol{A}$.

The nullspace and the range are determined by a singular value decomposition of the matrix $\boldsymbol{A}$.

2. Symmetric matrices

Many matrices in statistical computation are *symmetric* with $\boldsymbol{A}^T = \boldsymbol{A}$, for example the covariance matrices. They have *real* eigenvalues λ_i.

For a general $n \times m$ matrix $\boldsymbol{A}$ the product $\boldsymbol{W} = \boldsymbol{A}^T\boldsymbol{A}$ is a symmetric $n \times n$ matrix. The quadratic form

$$Q(\boldsymbol{x}) = \boldsymbol{x}^T\boldsymbol{W}\,\boldsymbol{x} = \boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} = (\boldsymbol{A}\boldsymbol{x})^T\,(\boldsymbol{A}\boldsymbol{x})$$

is a non-negative scalar. The matrix $\boldsymbol{W}$ is called *positive definite*, if for the quadratic form $Q(\boldsymbol{x}) > 0$ for all (non-vanishing) vectors $\boldsymbol{x}$ (all eigenvalues $\lambda_i > 0$.

The $(n^2+n)/2$ different elements W_{ij} of a symmetric matrix can be mapped to a one-dimensional array with index

$$q = \begin{cases} i + j(j-1)/2 & \text{for} \quad i \leq j \\ j + i(i-1)/2 & \text{for} \quad i \geq j \end{cases} \tag{1}$$

Eigenvalues and condition number

Symmetric matrices have *real* eigenvalues λ_i.

Formulae:

$$\det \boldsymbol{A} = |\boldsymbol{A}| = \prod_{i=1}^{n} \lambda_i \qquad \text{trace}\, \boldsymbol{A} = \sum_{i=1}^{n} \lambda_i$$

$$\text{Condition number} \qquad \kappa = \sqrt{\frac{\lambda_{\text{max}}}{\lambda_{\text{min}}}}$$

A matrix is called *ill-conditioned* if its condition number is large ($1/\kappa < 10{-}7$ for single precision arithmetic); it is singular if $\kappa = \infty$.

3. Inversion by partitioning

A square $n \times n$ matrix $\boldsymbol{A}$ can be partitioned in submatrices:

$$\boldsymbol{A} = \left(\begin{array}{c|c} \boldsymbol{A}_{11} & \boldsymbol{A}_{12} \\ \hline \boldsymbol{A}_{21} & \boldsymbol{A}_{22} \end{array}\right),$$

where submatrix $\boldsymbol{A}_{11}$ is a $p \times p$ square matrix and submatrix $\boldsymbol{A}_{22}$ is a $q \times q$ square matrix with $p + q = n$ (values $p = 1$ or $q = 1$ are possible).

The inverse matrix $\boldsymbol{B} = \boldsymbol{A}^{-1}$ can be partitioned in the same way:

$$\boldsymbol{B} = \left(\begin{array}{c|c} \boldsymbol{B}_{11} & \boldsymbol{B}_{12} \\ \hline \boldsymbol{B}_{21} & \boldsymbol{B}_{22} \end{array}\right)$$

Matrix product

Matrix product $\boldsymbol{AB}$ (mit $\boldsymbol{AB} = \boldsymbol{BA} = \boldsymbol{I}$) is given by the following matrix equations;

$$\begin{aligned} \boldsymbol{A}_{11}\,\boldsymbol{B}_{11} + \boldsymbol{A}_{12}\,\boldsymbol{B}_{21} &= \mathbf{1} \\ \boldsymbol{A}_{11}\,\boldsymbol{B}_{12} + \boldsymbol{A}_{12}\,\boldsymbol{B}_{22} &= \mathbf{0} \\ \boldsymbol{A}_{21}\,\boldsymbol{B}_{11} + \boldsymbol{A}_{22}\,\boldsymbol{B}_{21} &= \mathbf{0} \\ \boldsymbol{A}_{21}\,\boldsymbol{B}_{12} + \boldsymbol{A}_{22}\,\boldsymbol{B}_{22} &= \mathbf{1} \end{aligned}$$

In the case of *symmetric* matrices $\boldsymbol{A}$ und $\boldsymbol{B}$ the calculation can be further simplified because of

$$\boldsymbol{A}_{12} = (\boldsymbol{A}_{21})^T \qquad \text{und} \qquad \boldsymbol{B}_{12} = (\boldsymbol{B}_{21})^T \ .$$

Solution

Submatrices of the inverse matrix are

$$\begin{aligned} \boldsymbol{B}_{22} &= \left(\boldsymbol{A}_{22} - \boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\right)^{-1} \\ \boldsymbol{B}_{21} &= -\boldsymbol{B}_{22}\boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1} \\ \boldsymbol{B}_{12} &= -\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\boldsymbol{B}_{22} \\ \boldsymbol{B}_{11} &= \boldsymbol{A}_{11}^{-1} - \boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\boldsymbol{B}_{21} \end{aligned}$$

Here the inverse $p \times p$ matrix $\boldsymbol{A}_{11}^{-1}$ is used, while the original matrix $\boldsymbol{A}_{11}$ is not explicitly required.

One matrix inversion is required for the $q \times q$ matrix $\boldsymbol{B}_{22}$.

Result: $$\boldsymbol{B} = \begin{pmatrix} \boldsymbol{A}_{11}^{-1} - \boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\boldsymbol{B}_{21} & -\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\boldsymbol{B}_{22} \\ -\boldsymbol{B}_{22}\boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1} & \left(\boldsymbol{A}_{22} - \boldsymbol{A}_{21}\boldsymbol{A}_{11}^{-1}\boldsymbol{A}_{12}\right)^{-1} \end{pmatrix}$$

Solution with exchanged elements

By exchange of indices 1 and 2 a second set of equations is obtained:

$$\begin{aligned}
\boldsymbol{B}_{11} &= \left(\boldsymbol{A}_{11} - \boldsymbol{A}_{12}\boldsymbol{A}_{22}^{-1}\boldsymbol{A}_{21}\right)^{-1} \\
\boldsymbol{B}_{12} &= -\boldsymbol{B}_{11}\boldsymbol{A}_{12}\boldsymbol{A}_{22}^{-1} \\
\boldsymbol{B}_{21} &= -\boldsymbol{A}_{22}^{-1}\boldsymbol{A}_{21}\boldsymbol{B}_{11} \\
\boldsymbol{B}_{22} &= \boldsymbol{A}_{22}^{-1} - \boldsymbol{A}_{22}^{-1}\boldsymbol{A}_{21}\boldsymbol{B}_{12}
\end{aligned}$$

Result: $$\boldsymbol{B} = \begin{pmatrix} \left(\boldsymbol{A}_{11} - \boldsymbol{A}_{12}\boldsymbol{A}_{22}^{-1}\boldsymbol{A}_{21}\right)^{-1} & -\boldsymbol{B}_{11}\boldsymbol{A}_{12}\boldsymbol{A}_{22}^{-1} \\ -\boldsymbol{A}_{22}^{-1}\boldsymbol{A}_{21}\boldsymbol{B}_{11} & \boldsymbol{A}_{22}^{-1} - \boldsymbol{A}_{22}^{-1}\boldsymbol{A}_{21}\boldsymbol{B}_{12} \end{pmatrix}$$

Inversion by bordering

Partitioning in the case $q = 1$ is called bordering.

The $n \times n$ matrix $\boldsymbol{A}_n$ has the form

$$\boldsymbol{A}_n = \left(\begin{array}{cccc|c} a_{11} & a_{12} & \cdots & a_{1,n-1} & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2,n-1} & a_{2n} \\ \cdots & & & & \cdots \\ a_{n-1,1} & a_{n-1,2} & \cdots & a_{n-1,n-1} & a_{n-1,n} \\ \hline a_{n1} & a_{n2} & \cdots & a_{n,n-1} & a_{nn} \end{array}\right) = \left(\begin{array}{c|c} \boldsymbol{A}_{n-1} & \boldsymbol{u} \\ \hline \boldsymbol{v}^T & a_{nn} \end{array}\right)$$

with two $(n-1)$ vectors $\boldsymbol{u}$ und $\boldsymbol{v}$ and a scalar a_{nn}.

Inversion by bordering contnd.

Inversion is simple if the matrix $\boldsymbol{A}_{n-1}^{-1}$ is already known or easily calculable. One has to calculate

$$\alpha_n = a_{nn} - \boldsymbol{v}^T \boldsymbol{A}_{n-1}^{-1} \boldsymbol{u} \ ,$$

where α_n corresponds to matrix $\boldsymbol{B}_{22}^{-1}$. The other submatrices are obtained by matrix products:

$$\boldsymbol{A}_n^{-1} = \left(\begin{array}{c|c} \boldsymbol{A}_{n-1}^{-1} + \dfrac{\boldsymbol{A}_{n-1}^{-1} \boldsymbol{u}\boldsymbol{v}^T \boldsymbol{A}_{n-1}^{-1}}{\alpha_n} & -\dfrac{\boldsymbol{A}_{n-1}^{-1} \boldsymbol{u}}{\alpha_n} \\ \hline -\dfrac{\boldsymbol{v}^T \boldsymbol{A}_{n-1}^{-1}}{\alpha_n} & \dfrac{1}{\alpha_n} \end{array} \right)$$

The calculation is simplified in the case of a symmetric matrix (with $\boldsymbol{u} = \boldsymbol{v}$).

4. Diagonalization of symmetric matrices

The eigenvalues of real symmetric matrices $\boldsymbol{V}$ are real.

Eigenvalues and eigenvectors can be determined by a transformation with orthogonal matrices.

A matrix $\boldsymbol{U}$ is called *orthogonal*, if

$$\boldsymbol{U}^T \boldsymbol{U} = \boldsymbol{D} \ , \qquad \boldsymbol{D} = \text{diagonal matrix},$$

(coulumn vectors of $\boldsymbol{U}$ orthogonal).

The matrix $\boldsymbol{U}$ is called *orthonormal*, if the diagonal matrix $\boldsymbol{D}$ is a unit matrix with $\boldsymbol{U}^T \boldsymbol{U} = \boldsymbol{U} \boldsymbol{U}^T = \boldsymbol{I}$. If the columns vectors of a square matrix are orthogonal, then the column vectors are orthogonal and normalized too.

Diagonalization

For any real symmetric matrix $\boldsymbol{V}$ a orthonormal matrix $\boldsymbol{U}$ can be constructed, which transforms the matrix $\boldsymbol{V}$ to a diagonal matrix:

$$\boldsymbol{U}^T \boldsymbol{V} \boldsymbol{U} = \boldsymbol{D} \quad \text{und} \quad \boldsymbol{V} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^T$$

with square matrices $\boldsymbol{V}$, $\boldsymbol{U}$ and $\boldsymbol{D}$.

$$\boxed{\boldsymbol{V}} = \boxed{\boldsymbol{U}} \times \boxed{\boldsymbol{D}} \times \boxed{\boldsymbol{U}^T}$$

The diagonal elements $D_{11}, D_{22} \ldots D_{nn}$ of $\boldsymbol{D}$ are the eigenvalues of $\boldsymbol{V}$, the column vectors of the orthogonal matrix $\boldsymbol{U}$ are the corresponding eigenvectors.

Quadratic forms

The quadratic form $Q = \boldsymbol{x}^T \boldsymbol{V} \boldsymbol{x}$ of the transformed vector $\boldsymbol{y} = \boldsymbol{U}^T \boldsymbol{x}$ can be written in the form

$$Q = \boldsymbol{x}^T \boldsymbol{V} \boldsymbol{x} = \boldsymbol{x}^T \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^T \boldsymbol{x} = \boldsymbol{y}^T \boldsymbol{D} \boldsymbol{y} = \sum_{i=1}^{n} D_{ii} y_i^2 \, ,$$

and for all non-vanishing $\boldsymbol{y}$ the value of the quadratic form $Q > 0$, if for all diagonal elements $D_{ii} > 0$; in this case the matrix $\boldsymbol{V}$ is positive definite.

5. Singular value decomposition

Solving the matrix equation

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b} \qquad m \times n \text{ matrix } \boldsymbol{A}$$

for an n vector $\boldsymbol{x}$ of unknowns.

Distinguish cases:

- $n = m$ exact solution usually exists
- $m > n$ overdetermined system – obtain best "approximate" solution
- $m < n$ underdetermined system – insufficient data

Overdetermined system $(m > n)$

In general no solution vector $\boldsymbol{x}$ exist.
Vector of residua $\boldsymbol{r}$, defined by

$$\boldsymbol{r} = \boldsymbol{A}\boldsymbol{x} - \boldsymbol{b},$$

never zero.
What is the "best" solution?
Minimum value of the square of the vector of residua:

$$\begin{aligned} \boldsymbol{r}^T\boldsymbol{r} &= (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b})^T (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) \\ &= \boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} - 2\boldsymbol{x}^T\boldsymbol{A}^T\boldsymbol{b} + \boldsymbol{b}^T\boldsymbol{b} = \text{Minimum}\,. \end{aligned}$$

The minimum is determined, by setting the derivative of the above expression to zero:

$$\boldsymbol{A}^T\boldsymbol{A}\boldsymbol{x} - \boldsymbol{A}^T\boldsymbol{b} = 0 \quad \text{oder} \quad \boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}) = \boldsymbol{A}^T\boldsymbol{r} = 0\,.$$

These equation are called *normal equations*; the vector $\boldsymbol{r}$ of residua is *normal* with respect to the column vectors of matrix $\boldsymbol{A}$.

Solution

In general the system has a unique solution

$$\boldsymbol{x} = \left(\boldsymbol{A}^T\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T\boldsymbol{b}$$

The symmetric matrix $\boldsymbol{W} = \boldsymbol{A}^T\boldsymbol{A}$ and the vector $\boldsymbol{c} = \boldsymbol{A}^T\boldsymbol{b}$ are formed:

$$\boxed{\boldsymbol{W}} = \boxed{\boldsymbol{A}^T} \times \boxed{\boldsymbol{A}} \qquad \boxed{\boldsymbol{c}} = \boxed{\boldsymbol{A}^T} \times \boxed{\boldsymbol{b}}$$

The solution vector $\boldsymbol{x}$ is determined by the solution of the matrix equation $\boldsymbol{W}\boldsymbol{x} = \boldsymbol{c}$:

$$\boxed{\boldsymbol{x}} = \boxed{\boldsymbol{W}^{-1}} \times \boxed{\boldsymbol{c}}$$

A solution in different form ...

The symmetric matrix $\boldsymbol{W}$ can be diagonalized:

$$\boxed{\boldsymbol{W}} = \boxed{\boldsymbol{U}} \times \boxed{\boldsymbol{D}} \times \boxed{\boldsymbol{U}^T}$$

Matrix $\boldsymbol{U}$ is orthonormal and matrix $\boldsymbol{D}$ is diagonal with the eigenvalues λ of matrix $\boldsymbol{W}$ on the diagonal:

$$\boldsymbol{D} = \begin{pmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{pmatrix} \qquad \boldsymbol{D}^{-1} = \begin{pmatrix} 1/\lambda_1 & & & \\ & 1/\lambda_2 & & \\ & & \ddots & \\ & & & 1/\lambda_n \end{pmatrix}$$

$$\boxed{\boldsymbol{W}^{-1}} = \boxed{\boldsymbol{U}} \times \boxed{\boldsymbol{D}^{-1}} \times \boxed{\boldsymbol{U}^T}$$

$$\boldsymbol{x} = \boldsymbol{U}\boldsymbol{D}^{-1}\boldsymbol{U}^T\boldsymbol{c}\,.$$

Case of singular matrix $\boldsymbol{W}$

At least one of the eigenvalues is zero and the solution above is impossible

The columns of matrix $\boldsymbol{U}$ (these are the eigenvectors of matrix $\boldsymbol{W}$), which correspond to the *zero* eigenvalues, form an orthogonal basis of the nullspace of $\boldsymbol{W}$. Any arbitrary vector of the null space can be added to the solution $\boldsymbol{x}$.

How to avoid this kind of arbitrariness?

Better define a solution with *no* contribution from the null space: define a modified matrix $\boldsymbol{D}^{*-1}$, where the diagonal elements for zero eigenvalues of $\boldsymbol{D}$ are set to zero (thus replacing $1/0 = \infty$ by 0). The modified solution

$$\boldsymbol{x} = \boldsymbol{U}\boldsymbol{D}^{*-1}\boldsymbol{U}^T\boldsymbol{c}$$

contains *no* contribution from the null space.

It can be shown, that this is the solution with smallest length of $|\boldsymbol{x}|$.

Case of almost singular matrix $\boldsymbol{W}$

Sometimes a system of equation is not singular, but ill conditioned: some of the eigenvalues are very small and the inversion is inaccurate.

This property is visible in the condition number of the matrix, the ratio of the largest to the smallest eigenvalue: $\kappa = \lambda_{\text{max}}/\lambda_{\text{min}}$.

In such a case a solution with modified matrix Matrix $\boldsymbol{D}^{*-1}$ is preferred, where the small eigenvalues are replaced by zero.

Singular value decomposition

A $m \times n$ Matrix $\boldsymbol{A}$ can be expressed in the form

$$\boxed{\boldsymbol{A}} = \boxed{\boldsymbol{U}} \times \boxed{\boldsymbol{S}} \times \boxed{\boldsymbol{V}^T}$$

$\boldsymbol{U}$ is a column-orthogonal $m \times n$ matrix: $\boldsymbol{U}^T\boldsymbol{U} = \mathbf{1}$
$\boldsymbol{V}$ is an orthogonal $n \times n$ matrix: $\boldsymbol{V}^T\boldsymbol{V} = \mathbf{1} \quad \boldsymbol{V}\boldsymbol{V}^T = \mathbf{1}$

Matrix S is diagonal and has positive or zero elements, called *singular values*.

Modified solution

Matrix S is diagonal and has positive or zero elements, called *singular values*:

$$\boldsymbol{S} = \begin{pmatrix} \sigma_1 & & & \\ & \sigma_2 & & \\ & & \ddots & \\ & & & \sigma_n \end{pmatrix} \qquad \boldsymbol{S}^{-1} = \begin{pmatrix} 1/\sigma_1 & & & \\ & 1/\sigma_2 & & \\ & & \ddots & \\ & & & 1/\sigma_n \end{pmatrix}$$

Comparison with previous formalism shows:

$$\boldsymbol{W} = \boldsymbol{A}^T\boldsymbol{A} = \boldsymbol{V}\boldsymbol{S}\boldsymbol{S}\boldsymbol{V}^T \qquad \boldsymbol{D} = \boldsymbol{S}\boldsymbol{S}$$

i.e. eigenvalues of $\boldsymbol{D}$ are $\lambda_i = \sigma_i^2$.
The solution vector is the least squares solution with minimum length

$$\boldsymbol{x} = \boldsymbol{V}\boldsymbol{S}^{*-1}\left(\boldsymbol{U}^T\boldsymbol{b}\right) .$$

As above a modified matrix $\boldsymbol{S}^{*-1}$ is defined, where the diagonal elements for small singular values S_{jj} are set to zero.

Comparison

The two methods are equivalent, but there are differences:

- The singular value method is numerically more accurate ($\sigma_i = \sqrt{\lambda_i}$).
- The eigenvalue method works on more compact matrices (small $n \times n$ square matrices instead of large rectangular $m \times n$ matrices).

Both methods avoid those components in the solution vector $\boldsymbol{x}$ from the Null-space, which have no effect on the right hand-side of the initial equations.