1 Systems of linear equations

1.1 Introduction

A *system of linear equations* (or *linear system*) is a collection of linear equations involving the same set of unknown variables. A general system of n linear equations with m unknowns can be written as

$$
\begin{cases}
A_{11}x_1 & + & A_{12}x_2 & + & \cdots & + & A_{1m}x_m & = & b_1 \\
A_{21}x_1 & + & A_{22}x_2 & + & \cdots & + & A_{2m}x_m & = & b_2 \\
\vdots & & \vdots & & & & \vdots & & \vdots \\
A_{n1}x_1 & + & A_{n2}x_2 & + & \cdots & + & A_{nm}x_m & = & b_n
\end{cases}
, \tag{1}
$$

where $x_1, x_2, \ldots, x_m$ are the unknown variables, $A_{11}, A_{12}, \ldots, A_{nm}$ are the (constant) coefficients, and $b_1, b_2, \ldots, b_n$ are the (constant) right-hand side terms.

The system can be equivalently written in the matrix form,

$$
\mathbf{A}\mathbf{x} = \mathbf{b} , \tag{2}
$$

where $\mathbf{A} \doteq \{A_{ij}\}$ is the $n \times m$ matrix of the coefficients, $\mathbf{x} \doteq \{x_j\}$ is the size-m column-vector of the unknown variables, and $\mathbf{b} \doteq \{b_i\}$ is the size-n column-vector of right-hand side terms.

A solution to a linear system is a set of values for the variables $\mathbf{x}$ which satisfies all equations.

Systems of linear equations occur quite regularly in applied mathematics. Therefore computational algorithms for finding solutions of linear systems are an important part of numerical methods. A system of non-linear equations can often be approximated by a linear system – a helpful technique (called *linearization*) in creating a mathematical model of an otherwise a more complex system.

If $m = n$ the matrix $\mathbf{A}$ is called *square*. A square system has a unique solution if $\mathbf{A}$ is invertible.

1.2 Triangular systems

An efficient algorithm to numerically solve a square system of linear equations is to transform the original system into an equivalent *triangular system*,

$$
\mathbf{T}\mathbf{y} = \mathbf{c} , \tag{3}
$$

where $\mathbf{T}$ is a *triangular matrix* – a special kind of square matrix where the matrix elements either below (upper triangular) or above (lower triangular) the main diagonal are zero.

Indeed, an upper triangular system $\mathbf{Uy} = \mathbf{c}$ can be easily solved by *back-substitution*,

$$y_i = \frac{1}{U_{ii}} \left(c_i - \sum_{k=i+1}^{n} U_{ik} y_k \right), \ i = n, n-1, \ldots, 1 \ , \tag{4}$$

where one first computes $y_n = b_n/U_{nn}$, then substitutes *back* into the previous equation to solve for y_{n-1}, and repeats through y_1.

Here is a C-function implementing the in-place[1] back-substitution[2]:

```
#include<gsl/gsl_vector.h>
#include<gsl/gsl_matrix.h>
void backsub(gsl_matrix *U, gsl_vector *c){
 for(int i=c->size-1; i>=0; i--){
  double s=gsl_vector_get(c,i);
  for(int k=i+1;k<c->size;k++)
   s-=gsl_matrix_get(U,i,k)*gsl_vector_get(c,k);
  gsl_vector_set(c,i,s/gsl_matrix_get(U,i,i)); }}
```

For a lower triangular system $\mathbf{Ly} = \mathbf{c}$ the equivalent procedure is called *forward-substitution*,

$$y_i = \frac{1}{L_{ii}} \left(c_i - \sum_{k=1}^{i-1} L_{ik} y_k \right), \ i = 1, 2, \ldots, n \ . \tag{5}$$

1.3 Reduction to triangular form

Popular algorithms for reducing a square system of linear equations to a triangular form are *LU-decomposition* and *QR-decomposition*.

1.3.1 QR-decomposition

QR-decomposition is a factorization of a matrix into a product of an orthogonal matrix $\mathbf{Q}$, such that $\mathbf{Q}^\mathsf{T}\mathbf{Q} = 1$, where $^\mathsf{T}$ denotes transposition, and a right triangular matrix $\mathbf{R}$,

$$\mathbf{A} = \mathbf{QR} \ . \tag{6}$$

---

[1] here *in-place* means the right-hand side $\mathbf{c}$ is replaced by the solution $\mathbf{y}$.

[2] the functions `gsl_vector_get`, `gsl_vector_set`, and `gsl_matrix_get` are assumed to implement fetching and setting the vector- and matrix-elements.

QR-decomposition can be used to convert a linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ into the triangular form (by multiplying with $\mathbf{Q}^\mathsf{T}$ from the left),

$$\mathbf{R}\mathbf{x} = \mathbf{Q}^\mathsf{T}\mathbf{b} \, , \tag{7}$$

which can be solved directly by back-substitution.

QR-decomposition can also be performed on non-square matrices with few long columns. Generally speaking a rectangular $n \times m$ matrix $\mathbf{A}$ can be represented as a product, $\mathbf{A} = \mathbf{Q}\mathbf{R}$, of an orthogonal $n \times m$ matrix $\mathbf{Q}$, $\mathbf{Q}^\mathsf{T}\mathbf{Q} = 1$, and a right-triangular $m \times m$ matrix $\mathbf{R}$.

QR-decomposition of a matrix can be computed using several methods, such as Gram-Schmidt orthogonalization, Householder transformation [**?**], or Givens rotation [**?**].

Gram-Schmidt orthogonalization   *Gram-Schmidt orthogonalization* is an algorithm for orthogonalization of a set of vectors in a given inner product space. It takes a linearly independent set of vectors $\mathbf{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_m\}$ and generates an orthogonal set $\mathbf{Q} = \{\mathbf{q}_1, \dots, \mathbf{q}_m\}$ which spans the same subspace as $\mathbf{A}$. The algorithm is given as

```
for  i = 1 to  m :
    q_i ← a_i/‖a_i‖
    for  j = i + 1 to  m :    a_j ← a_j − ⟨q_i|a_j⟩ q_i
```

where $\langle \mathbf{a} | \mathbf{b} \rangle$ is the inner product of two vectors, and $\|\mathbf{a}\| \doteq \sqrt{\langle \mathbf{a} | \mathbf{a} \rangle}$ is the vector's norm. This variant of the algorithm, where all remaining vectors $\mathbf{a}_j$ are made orthogonal to $\mathbf{q}_i$ as soon as the latter is calculated, is considered to be numerically stable and is referred to as *stabilized* or *modified*.

Stabilized Gram-Schmidt orthogonalization can be used to compute QR-decomposition of a matrix $\mathbf{A}$ by orthogonalization of its column-vectors $\mathbf{a}_i$ with the inner product

$$\langle \mathbf{a} | \mathbf{b} \rangle = \mathbf{a}^\mathsf{T}\mathbf{b} \equiv \sum_{k=1}^{n} (\mathbf{a})_k (\mathbf{b})_k \, , \tag{8}$$

where n is the length of column-vectors $\mathbf{a}$ and $\mathbf{b}$, and $(\mathbf{a})_k$ is the kth element of the column-vector,

```
for  i = 1 to  m :
    R_ii = √(a_iᵀa_i) ;   q_i = a_i/R_ii
    for  j = i + 1 to  m :
        R_ij = q_iᵀa_j ;   a_j = a_j − q_iR_ij .
```

3

After orthogonalization the matrices $\mathbf{Q} = \{\mathbf{q}_1 \ldots \mathbf{q}_m\}$ and $\mathbf{R}$ are the sought orthogonal and right-triangular factors of matrix $\mathbf{A}$.

The factorization is unique under requirement that the diagonal elements of R are positive. For a $n \times m$ matrix the complexity of the algorithm is $O(m^2 n)$.

Householder transformation   A square matrix $\mathbf{H}$ of the form

$$\mathbf{H} = 1 - \frac{2}{\mathbf{u}^\mathsf{T}\mathbf{u}}\mathbf{u}\mathbf{u}^\mathsf{T} \tag{9}$$

is called *Householder matrix*, where the vector $\mathbf{u}$ is called a *Householder vector*. Householder matrices are symmetric and orthogonal,

$$\mathbf{H}^\mathsf{T} = \mathbf{H}\,,\ \mathbf{H}^\mathsf{T}\mathbf{H} = 1\,. \tag{10}$$

The transformation induced by the Householder matrix on a given vector $\mathbf{a}$,

$$\mathbf{a} \rightarrow \mathbf{H}\mathbf{a}\,, \tag{11}$$

is called a *Householder transformation* or *Householder reflection*. The transformation changes the sign of the affected vector's component in the $\mathbf{u}$ direction, or, in other words, makes a reflection of the vector about the hyperplane perpendicular to $\mathbf{u}$, hence the name.

Householder transformation can be used to zero selected components of a given vector $\mathbf{a}$. For example, one can zero all components but the first one, such that

$$\mathbf{H}\mathbf{a} = \gamma\mathbf{e}_1\,, \tag{12}$$

where γ is a number and $\mathbf{e}_1$ is the unit vector in the first direction. The factor γ can be easily calculated,

$$\|\mathbf{a}\|^2 \doteq \mathbf{a}^\mathsf{T}\mathbf{a} = \mathbf{a}^\mathsf{T}\mathbf{H}^\mathsf{T}\mathbf{H}\mathbf{a} = (\gamma\mathbf{e}_1)^\mathsf{T}(\gamma\mathbf{e}_1) = \gamma^2\,, \tag{13}$$

$$\Rightarrow\ \gamma = \pm\|a\|\,. \tag{14}$$

To find the Householder vector, we notice that

$$\mathbf{a} = \mathbf{H}^\mathsf{T}\mathbf{H}\mathbf{a} = \mathbf{H}^\mathsf{T}\gamma\mathbf{e}_1 = \gamma\mathbf{e}_1 - \frac{2(\mathbf{u})_1}{\mathbf{u}^\mathsf{T}\mathbf{u}}\mathbf{u}\,, \tag{15}$$

$$\Rightarrow\ \frac{2(\mathbf{u})_1}{\mathbf{u}^\mathsf{T}\mathbf{u}}\mathbf{u} = \gamma\mathbf{e}^1 - \mathbf{a}\,, \tag{16}$$

where $(\mathbf{u})_1$ is the first component of the vector $\mathbf{u}$. One usually chooses $(\mathbf{u})_1 = 1$ (for the sake of the possibility to store the other components of the Householder vector in the zeroed elements of the vector $\mathbf{a}$) and stores the factor

$$\frac{2}{\mathbf{u}^\mathsf{T}\mathbf{u}} \equiv \tau \tag{17}$$

separately. With this convention one readily finds τ from the first component of equation (16),

$$\tau = \gamma - (\mathbf{a})_1 \,. \tag{18}$$

where $(\mathbf{a})_1$ is the first element of the vector $\mathbf{a}$. For the sake of numerical stability the sign of γ has to be chosen opposite to the sign of $(\mathbf{a})_1$,

$$\gamma = -\mathrm{sign}\left((\mathbf{a})_1\right)\|\mathbf{a}\| \,. \tag{19}$$

Finally, the Householder reflection, which zeroes all component of a vector $\mathbf{a}$ but the first, is given as

$$\mathbf{H} = 1 - \tau\mathbf{u}\mathbf{u}^\mathsf{T}, \ \ \tau = -\mathrm{sign}((\mathbf{a})_1)\|\mathbf{a}\| - (\mathbf{a})_1, \ (\mathbf{u})_1 = 1, \ (\mathbf{u})_{i>1} = -\frac{1}{\tau}(\mathbf{a})_i. \tag{20}$$

Now, a QR-decomposition of an $n \times n$ matrix $\mathbf{A}$ by Householder transformations can be performed in the following way:

1. Build the size-n Householder vector $\mathbf{u}_1$ which zeroes the sub-diagonal elements of the first column of matrix $\mathbf{A}$, such that

$$\mathbf{H}_1\mathbf{A} = \begin{bmatrix} \star & \star & \cdots & \star \\ \hline 0 & & & \\ \vdots & & \mathbf{A}_1 & \\ 0 & & & \end{bmatrix}, \tag{21}$$

where $\mathbf{H}_1 = 1 - \tau_1\mathbf{u}_1\mathbf{u}_1^\mathsf{T}$ and where $\star$ denotes (generally) non-zero matrix elements. In practice one does not build the matrix $\mathbf{H}_1$ explicitly, but rather calculates the matrix $\mathbf{H}_1\mathbf{A}$ in-place, consecutively applying the Householder reflection to columns the matrix $\mathbf{A}$, thus avoiding computationally expensive matrix-matrix operations. The zeroed sub-diagonal elements of the first column of the matrix $\mathbf{A}$ can be used to store the elements of the Householder vector $\mathbf{u}_1$ while the factor τ_1 has to be stored separately in a special array. This is the storage scheme used by LAPACK and GSL.

2. Similarly, build the size-$(n-1)$ Householder vector $\mathbf{u}_2$ which zeroes the sub-diagonal elements of the first column of matrix $\mathbf{A}_1$ from eq. (21). With the transformation matrix $\mathbf{H}_2$ defined as

$$
\mathbf{H}_2 = \left[\begin{array}{c|ccc} 1 & 0 & \cdots & 0 \\ \hline 0 & & & \\ \vdots & & 1 - \tau_2 \mathbf{u}_2 \mathbf{u}_2^\mathsf{T} & \\ 0 & & & \end{array}\right] .
\tag{22}
$$

the two transformations together zero the sub-diagonal elements of the two first columns of matrix $\mathbf{A}$,

$$
\mathbf{H}_2 \mathbf{H}_1 \mathbf{A} = \left[\begin{array}{cc|cccc} \star & \star & \star & \cdots & \star \\ 0 & \star & \star & \cdots & \star \\ 0 & 0 & & & \\ \vdots & \vdots & & \mathbf{A}_3 & \\ 0 & 0 & & & \end{array}\right] ,
\tag{23}
$$

3. Repeating the process zero the sub-diagonal elements of the remaining columns. For column k the corresponding Householder matrix is

$$
\mathbf{H}_k = \left[\begin{array}{c|c} \mathbf{I}_{k-1} & 0 \\ \hline 0 & 1 - \tau_k \mathbf{u}_k \mathbf{u}_k^\mathsf{T} \end{array}\right] ,
\tag{24}
$$

where $\mathbf{I}_{k-1}$ is an identity matrix of size $k-1$, $\mathbf{u}_k$ is the size-(n-k+1) Householder vector that zeroes the sub-diagonal elements of matrix $\mathbf{A}_{k-1}$ from the previous step. The corresponding transformation step is

$$
\mathbf{H}_k \ldots \mathbf{H}_2 \mathbf{H}_1 \mathbf{A} = \left[\begin{array}{c|c} \mathbf{R}_k & \star \\ \hline 0 & \mathbf{A}_k \end{array}\right] ,
\tag{25}
$$

where $\mathbf{R}_k$ is a size-k right-triangular matrix.

After $n-1$ steps the matrix $\mathbf{A}$ will be transformed into a right triangular matrix,

$$
\mathbf{H}_{n-1} \cdots \mathbf{H}_2 \mathbf{H}_1 \mathbf{A} = \mathbf{R} .
\tag{26}
$$

4. Finally, introducing an orthogonal matrix $\mathbf{Q} = \mathbf{H}_1^\mathsf{T}\mathbf{H}_2^\mathsf{T}\ldots\mathbf{H}_{n-1}^\mathsf{T}$ and multiplying eq. (26) by $\mathbf{Q}$ from the left, we get the sought QR-decomposition,

$$\mathbf{A} = \mathbf{QR} \ . \tag{27}$$

In practice one does not explicitly builds the $\mathbf{Q}$ matrix but rather applies the successive Householder reflections stored during the decomposition.

Givens rotations   A Givens rotation is a transformation in the form

$$\mathbf{A} \rightarrow \mathbf{G}(p, q, \theta)\mathbf{A} \ , \tag{28}$$

where $\mathbf{A}$ is the object to be transformed—matrix of vector—and $\mathbf{G}(p, q, \theta)$ is the Givens rotation matrix (also known as Jacobi rotation matrix): an orthogonal matrix in the form

$$G(p, q, \theta) = \begin{bmatrix} 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots & & \vdots \\ 0 & \cdots & \cos\theta & \cdots & \sin\theta & \cdots & 0 \\ \vdots & & \vdots & \ddots & \vdots & & \vdots \\ 0 & \cdots & -\sin\theta & \cdots & \cos\theta & \cdots & 0 \\ \vdots & & \vdots & & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 0 & \cdots & 1 \end{bmatrix} \begin{matrix} \\ \\ \leftarrow \text{row } p \\ \\ \leftarrow \text{row } q \\ \\ \\ \end{matrix} \ . \tag{29}$$

When a Givens rotation matrix $\mathbf{G}(p, q, \theta)$ multiplies a vector $\mathbf{x}$, only elements x_p and x_q are affected. Considering only these two affected elements, the Givens rotation is given explicitly as

$$\begin{bmatrix} x_p' \\ x_q' \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x_p \\ x_q \end{bmatrix} = \begin{bmatrix} x_p\cos\theta + x_q\sin\theta \\ -x_p\sin\theta + x_q\cos\theta \end{bmatrix} \ . \tag{30}$$

Apparently the rotation can zero the element x_q', if the angle θ is chosen as

$$\tan\theta = \frac{x_q}{x_p} \quad \Rightarrow \quad \theta = \mathrm{atan2}(x_q, x_p) \ . \tag{31}$$

A sequence of Givens rotations,

$$\mathbf{G} = \prod_{n\geq q>p=1}^{m} \mathbf{G}(p, q, \theta_{qp}) \ , \tag{32}$$

7

(where $n \times m$ is the dimension of the matrix $\mathbf{A}$) can zero all elements of a matrix below the main diagonal if the angles θ_{qp} are chosen to zero the elements with indices q, p of the partially transformed matrix just before applying the matrix $\mathbf{G}(p, q, \theta_{qp})$. The resulting matrix is obviously the $\mathbf{R}$-matrix of the sought QR-decomposition of the matrix $\mathbf{A}$ where $\mathbf{G} = \mathbf{Q}^\mathsf{T}$.

In practice one does not explicitly builds the $\mathbf{G}$ matrix but rather stores the θ angles in the places of the corresponding zeroed elements of the original matrix:

```
#include<gsl/gsl_matrix.h>
#include<math.h>
void givens_qr(gsl_matrix* A){ /* A <- Q,R */
 for(int q=0;q<A->size2;q++)for(int p=q+1;p<A->size1;p++){
  double theta=atan2(gsl_matrix_get(A,p,q),gsl_matrix_get(A,q,q));
  for(int k=q;k<A->size2;k++){
   double xq=gsl_matrix_get(A,q,k), xp=gsl_matrix_get(A,p,k);
   gsl_matrix_set(A,q,k, xq*cos(theta)+xp*sin(theta));
   gsl_matrix_set(A,p,k,-xq*sin(theta)+xp*cos(theta)); }
  gsl_matrix_set(A,p,q,theta); } }
```

When solving the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ one transforms it into the equivalent triangular system $\mathbf{R}\mathbf{x} = \mathbf{G}\mathbf{b}$ where one calculates $\mathbf{G}\mathbf{b}$ by successively applying the individual Givens rotations with the stored θ-angles:

```
#include<gsl/gsl_vector.h>
#include<gsl/gsl_matrix.h>
#include<math.h>
void givens_qr_QTvec(gsl_matrix* QR, gsl_vector* v){/* v <- Q^Tv */
 for(int q=0; q<QR->size2; q++)for(int p=q+1; p<QR->size1; p++){
  double theta = gsl_matrix_get(QR,p,q);
  double vq=gsl_vector_get(v,q), vp=gsl_vector_get(v,p);
  gsl_vector_set(v,q, vq*cos(theta)+vp*sin(theta));
  gsl_vector_set(v,p,-vq*sin(theta)+vp*cos(theta)); } }
```

The triangular system $\mathbf{R}\mathbf{x} = \mathbf{G}\mathbf{b}$ is then solved by the ordinary backsubstitution:

```
#include<gsl/gsl_matrix.h>
#include"givens_qr.h"
void givens_qr_solve(gsl_matrix* QR, gsl_vector* b){
 givens_qr_QTvec(QR,b);
 backsub(QR,b); }
```

If one needs to build the $\mathbf{Q}$-matrix explicitly, one uses

$$Q_{ij} = \mathbf{e}_i^\mathsf{T} \mathbf{Q} \mathbf{e}_j = \mathbf{e}_j^\mathsf{T} \mathbf{Q}^\mathsf{T} \mathbf{e}_i \, , \tag{33}$$

where $\mathbf{e}_i$ is the unit vector in the direction i and where again one can use the succesive rotations to calculate $\mathbf{Q}^\mathsf{T}\mathbf{e}_i$,

```
#include<gsl/gsl_vector.h>
#include<gsl/gsl_matrix.h>
#include"givens_qr.h"
void givens_qr_unpack_Q(gsl_matrix* QR, gsl_matrix* Q){
 gsl_vector* ei = gsl_vector_alloc(QR->size1);
 for(int i=0; i<QR->size1; i++){
  gsl_vector_set_basis(ei,i);
  givens_qr_QTvec(QR, ei);
  for(int j=0; j<QR->size2; j++)
   gsl_matrix_set(Q,i,j,gsl_vector_get(ei,j)); }
 gsl_vector_free(ei); }
```

Since each Givens rotation only affects two rows of the matrix it is possible to apply a set of rotations in parallel. Givens rotations are also more efficient on sparse matrices.

1.3.2 LU-decomposition

LU-decomposition is a factorization of a square matrix $\mathbf{A}$ into a product of a lower triangular matrix $\mathbf{L}$ and an upper triangular matrix $\mathbf{U}$,

$$\mathbf{A} = \mathbf{L}\mathbf{U} \ . \tag{34}$$

The linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ after LU-decomposition of the matrix $\mathbf{A}$ becomes $\mathbf{L}\mathbf{U}\mathbf{x} = \mathbf{b}$ and can be solved by first solving $\mathbf{L}\mathbf{y} = \mathbf{b}$ for $\mathbf{y}$ and then $\mathbf{U}\mathbf{x} = \mathbf{y}$ for $\mathbf{x}$ with two runs of forward and backward substitutions.

If $\mathbf{A}$ is an $n \times n$ matrix, the condition (34) is a set of n^2 equations,

$$\sum_{k=1}^{n} L_{ik}U_{kj} = A_{ij} \Big|_{i,j=1\ldots n} , \tag{35}$$

for $n^2 + n$ unknown elements of the triangular matrices $\mathbf{L}$ and $\mathbf{U}$. The decomposition is thus not unique.

Usually the decomposition is made unique by providing extra n conditions e.g. by the requirement that the elements of the main diagonal of the matrix $\mathbf{L}$ are equal one,

$$L_{ii} = 1 \, , \ i = 1\ldots n \, . \tag{36}$$

The system (35) with the extra conditions (36) can then be easily solved row after row using the *Doolittle's algorithm*,

```
for  i = 1 ... n  :
    L_ii = 1
    for  j = i ... n  :    U_ij = A_ij − ∑_{k<i} L_ik U_kj
    for  j = i + 1 ... n  :    L_ji = 1/U_ii ( A_ji − ∑_{k<j} L_jk U_ki )
```

In a slightly different *Crout's algorithm* it is the matrix **U** that has unit diagonal elements,

```
for  i = 1 ... n  :
    U_ii = 1
    for  j = i ... n  :    L_ji = A_ji − ∑_{k<i} L_jk U_ki
    for  j = i + 1 ... n  :    U_ij = 1/L_ii ( A_ji − ∑_{k<j} L_jk U_ki )
```

Without a proper ordering (permutations) in the matrix, the factorization may fail. For example, it is easy to verify that $A_{11} = L_{11}U_{11}$. If $A_{11} = 0$, then at least one of L_{11} and U_{11} has to be zero, which implies either **L** or **U** is singular, which is impossible if **A** is non-singular. This is however only a procedural problem. It can be removed by simply reordering the rows of **A** so that the first element of the permuted matrix is nonzero (or, even better, the largest in absolute value among all elements of the column below the diagonal). The same problem in subsequent factorization steps can be removed in a similar way. Such algorithm is refered to as *partial pivoting*. It requires an extra integer array to keep track of row permutations.

1.3.3   Cholesky decomposition

The Cholesky decomposition of a Hermitian positive-definite matrix **A** is a decomposition in the form

$$A = \mathbf{L}\mathbf{L}^{\dagger} , \tag{37}$$

where **L** is a lower triangular matrix with real and positive diagonal elements, and $\mathbf{L}^{\dagger}$ is the conjugate transpose of **L**.

For real symmetric positive-definite matrices the decomposition reads

$$A = \mathbf{L}\mathbf{L}^{\mathsf{T}} , \tag{38}$$

where **L** is real.

The decomposition can be calculated using the following in-place algorithm,

$$L_{jj} = \sqrt{A_{jj} - \sum_{k=1}^{j-1} L_{jk}^2} , \; L_{ij} = \frac{1}{L_{jj}} \left(A_{ij} - \sum_{k=1}^{j-1} L_{ik}L_{jk} \right)\Bigg|_{i>j} . \tag{39}$$

The expression under the square root is always positive if $\mathbf{A}$ is real and positive-definite.

When applicable, the Cholesky decomposition is about twice as efficient as LU-decomposition for solving systems of linear equations.

1.4 Determinant of a matrix

LU- and QR-decompositions allow $O(n^3)$ calculation of the determinant of a square matrix. Indeed, for the LU-decomposition,

$$\det \mathbf{A} = \det \mathbf{LU} = \det \mathbf{L} \det \mathbf{U} = \det \mathbf{U} = \prod_{i=1}^{n} U_{ii} \ . \tag{40}$$

For the Gram-Schmidt QR-decomposition

$$\det \mathbf{A} = \det \mathbf{QR} = \det \mathbf{Q} \det \mathbf{R} \ . \tag{41}$$

Since $\mathbf{Q}$ is an orthogonal matrix $(\det \mathbf{Q})^2 = 1$,

$$|\det \mathbf{A}| = |\det \mathbf{R}| = \left| \prod_{i=1}^{n} R_{ii} \right| \ . \tag{42}$$

With Gram-Schmidt method one arbitrarily assigns positive sign to diagonal elements of the R-matrix thus removing from the R-matrix the memory of the original sign of the determinant.

However with Givens rotation method the determinant of the individual rotation matrix—and thus the determinant of the total rotation matrix—is equal one, therefore for a square matrix $\mathbf{A}$ the QR-decomposition $\mathbf{A} = \mathbf{GR}$ via Givens rotations allows calculation of the determinant with the correct sign,

$$\det \mathbf{A} = \det \mathbf{R} \equiv \prod_{i=1}^{n} R_{ii} \tag{43}$$

1.5 Matrix inverse

The inverse $\mathbf{A}^{-1}$ of a square $n \times n$ matrix $\mathbf{A}$ can be calculated by solving n linear equations

$$\mathbf{A}\mathbf{x}_i = \mathbf{e}_i \Big|_{i=1,\dots,n} \ , \tag{44}$$

where $\mathbf{e}_i$ is the unit-vector in the i-direction: a column where all elements are equal zero except for the element number i which is equal one. Thus the set of columns $\{\mathbf{e}_i\}_{i=1,\ldots,n}$ form the identity matrix. The matrix made of columns $\mathbf{x}_i$ is apparently the inverse of $\mathbf{A}$.

Here is an implementation of this algorithm using the functions from the Givens rotation chapter,

```
#include<gsl/gsl_vector.h>
#include<gsl/gsl_matrix.h>
#include"givens_qr.h"
void givens_qr_inverse(gsl_matrix* QR, gsl_matrix* B){
  gsl_matrix_set_identity(B);
  for(int i=0; i<QR->size2; i++){
    gsl_vector_view v = gsl_matrix_column(B,i);
    givens_qr_solve(QR,&v.vector); } }
```