

Chapter 2

Linear Algebra

The purpose of this chapter is to briefly review the notations, definitions, and results of linear algebra that are used in the rest of the book. There is no claim of completeness as the reader is assumed to be familiar with the basic concepts of college linear algebra such as vector spaces, linear mappings, matrix decompositions, etc. More systematic exposition and additional material can be found in the books by Demmel [1], Laub [2], or Golub and Van Loan [3].

2.1 Vectors and Matrices

In this book, we work with n -dimensional arithmetic vectors $\mathbf{v} \in \mathbb{R}^n$, where \mathbb{R} denotes the set of real numbers. The only exception is Sect. 2.7, where vectors with complex entries are considered. We denote the i th component of an arithmetic vector $\mathbf{v} \in \mathbb{R}^n$ by $[\mathbf{v}]_i$. Thus $[\mathbf{v}]_i = v_i$ if $\mathbf{v} = [v_i]$ is defined by its components v_i . All the arithmetic vectors are considered by default to be column vectors. The relations between vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ are defined componentwise. Thus $\mathbf{u} \leq \mathbf{v}$ is equivalent to $[\mathbf{u}]_i \leq [\mathbf{v}]_i$, $i = 1, \dots, n$. We sometimes call the elements of \mathbb{R}^n *points* to indicate that the concepts of length and direction are not important.

The vector analog of $0 \in \mathbb{R}$ is the *zero vector* $\mathbf{o}_n \in \mathbb{R}^n$ with all the entries equal to zero. When the dimension can be deduced from the context, possibly using the assumption that all the expressions in our book are well defined, we often drop the subscript and write simply \mathbf{o} .

Given vectors $\mathbf{v}_1, \dots, \mathbf{v}_k \in \mathbb{R}^n$, the set

$$\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\} = \{\mathbf{v} \in \mathbb{R}^n : \mathbf{v} = \alpha_1 \mathbf{v}_1 + \dots + \alpha_k \mathbf{v}_k, \alpha_i \in \mathbb{R}\}$$

is a vector space called the *linear span* of $\mathbf{v}_1, \dots, \mathbf{v}_k$. For example

$$\text{Span}\{\mathbf{s}_1, \dots, \mathbf{s}_n\} = \mathbb{R}^n, \quad [\mathbf{s}_i]_j = \delta_{ij}, \quad i, j = 1, \dots, n,$$

where δ_{ij} denotes the *Kronecker symbol* defined by $\delta_{ij} = 1$ for $i = j$ and $\delta_{ij} = 0$ for $i \neq j$, is the *standard basis* of \mathbb{R}^n .

We sometimes use the componentwise extensions of scalar functions to vectors. Thus, if $\mathbf{v} \in \mathbb{R}^n$, then \mathbf{v}^+ and \mathbf{v}^- are the vectors the i th components of which are $\max\{[\mathbf{v}]_i, 0\}$ and $\min\{[\mathbf{v}]_i, 0\}$, respectively.

If \mathcal{I} is a nonempty subset of $\{1, \dots, n\}$ and $\mathbf{v} \in \mathbb{R}^n$, then we denote by $[\mathbf{v}]_{\mathcal{I}}$ or simply $\mathbf{v}_{\mathcal{I}}$ the subvector of \mathbf{v} with components $[\mathbf{v}]_i$, $i \in \mathcal{I}$. Thus if \mathcal{I} has m elements, then $\mathbf{v}_{\mathcal{I}} \in \mathbb{R}^m$, so we can refer to the components of $\mathbf{v}_{\mathcal{I}}$ either by the *global indices* $i \in \mathcal{I}$ or by the *local indices* $j \in \{1, \dots, m\}$. We usually rely on the reader's judgment to recognize the appropriate type of indexing.

Similarly to the related convention for vectors, the (i, j) th component of a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is denoted by $[\mathbf{A}]_{ij}$, so that $[\mathbf{A}]_{ij} = a_{ij}$ for $\mathbf{A} = [a_{ij}]$ which is defined by its entries a_{ij} . A matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is called an (m, n) -matrix.

The matrix analog of 0 is the *zero matrix* $\mathbf{O}_{mn} \in \mathbb{R}^{m \times n}$ with all the entries equal to zero. When the dimension is clear from the context, we often drop the subscripts and write simply \mathbf{O} .

The matrix counterpart of $1 \in \mathbb{R}$ in $\mathbb{R}^{n \times n}$ is the *identity matrix* $\mathbf{I}_n = [\delta_{ij}]$ of the order n . When the dimension may be deduced from the context, we often drop the subscripts and write simply \mathbf{I} . Thus, we can write

$$\mathbf{A} = \mathbf{I}\mathbf{A} = \mathbf{A}\mathbf{I}$$

for any matrix \mathbf{A} , having in mind that the order of \mathbf{I} on the left may be different from that on the right.

A matrix \mathbf{A} is *positive definite* if $\mathbf{x}^T \mathbf{A} \mathbf{x} > 0$ for any $\mathbf{x} \neq \mathbf{0}$, *positive semidefinite* if $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$ for any \mathbf{x} , and *indefinite* if neither \mathbf{A} nor $-\mathbf{A}$ is positive definite or semidefinite. We are especially interested in *symmetric positive definite (SPD)* or *symmetric positive semidefinite (SPS)* matrices.

If $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathcal{I} \subseteq \{1, \dots, m\}$, and $\mathcal{J} \subseteq \{1, \dots, n\}$, \mathcal{I} and \mathcal{J} nonempty, we denote by $\mathbf{A}_{\mathcal{I}\mathcal{J}}$ the submatrix of \mathbf{A} with the components $[\mathbf{A}]_{ij}$, $i \in \mathcal{I}$, $j \in \mathcal{J}$. The local indexing of the entries of $\mathbf{A}_{\mathcal{I}\mathcal{J}}$ is used whenever it is convenient in a similar way as the local indexing of subvectors which was introduced in Sect. 2.1. The full set of indices may be replaced by $*$ so that $\mathbf{A} = \mathbf{A}_{**}$ and $\mathbf{A}_{\mathcal{I}*}$ denotes the submatrix of \mathbf{A} with the row indices belonging to \mathcal{I} . Occasionally we simplify $\mathbf{A}_{\mathcal{I}} = \mathbf{A}_{\mathcal{I}*}$.

Sometimes it is useful to rearrange the matrix operations into manipulations with submatrices of given matrices called blocks. A *block matrix* $\mathbf{A} \in \mathbb{R}^{m \times n}$ is defined by its blocks $\mathbf{A}_{ij} = \mathbf{A}_{\mathcal{I}_i \mathcal{J}_j}$, where \mathcal{I}_i and \mathcal{J}_j denote nonempty contiguous sets of indices decomposing $\{1, \dots, m\}$ and $\{1, \dots, n\}$, respectively. We can use the block structure to implement matrix operations only when the block structure of the involved matrices matches.

The matrices in our applications are often *sparse* in the sense that they have a small number of nonzero entries distributed in a pattern which can be exploited to the efficient implementation of matrix operations or to the reduction of storage requirements.

2.2 Matrices and Mappings

Each matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ defines the mapping which assigns to each $\mathbf{x} \in \mathbb{R}^n$ the vector $\mathbf{Ax} \in \mathbb{R}^m$. Two important subspaces associated with this mapping are its *range* or *image space* $\text{Im}\mathbf{A}$ and its *kernel* or *null space* $\text{Ker}\mathbf{A}$; they are defined by

$$\text{Im}\mathbf{A} = \{\mathbf{Ax} : \mathbf{x} \in \mathbb{R}^n\} \quad \text{and} \quad \text{Ker}\mathbf{A} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{Ax} = \mathbf{o}\}.$$

The range of \mathbf{A} is the span of its columns. The *rank* and the *defect* of a matrix are defined as the dimension of its image and kernel, respectively.

If f is a mapping defined on $\mathcal{D} \subseteq \mathbb{R}^n$ and $\Omega \subseteq \mathcal{D}$, then $f|_{\Omega}$ denotes the *restriction* of f to Ω , that is, the mapping defined on Ω which assigns to each $\mathbf{x} \in \Omega$ the value $f(\mathbf{x})$. If $\mathbf{A} \in \mathbb{R}^{m \times n}$ and V is a subspace of \mathbb{R}^n , we define $\mathbf{A}|V$ as a restriction of the mapping associated with \mathbf{A} to V . The restriction $\mathbf{A}|V$ is said to be positive definite if $\mathbf{x}^T \mathbf{Ax} > 0$ for $\mathbf{x} \in V$, $\mathbf{x} \neq \mathbf{o}$, and positive semidefinite if $\mathbf{x}^T \mathbf{Ax} \geq 0$ for $\mathbf{x} \in V$.

The mapping associated with \mathbf{A} is *injective* if $\mathbf{Ax} = \mathbf{Ay}$ implies $\mathbf{x} = \mathbf{y}$. It is easy to check that the mapping associated with \mathbf{A} is injective if and only if $\text{Ker}\mathbf{A} = \{\mathbf{o}\}$. If $m = n$, then \mathbf{A} is injective if and only if $\text{Im}\mathbf{A} = \mathbb{R}^n$.

A subspace $V \subseteq \mathbb{R}^n$ which satisfies

$$\mathbf{A}V = \{\mathbf{Ax} : \mathbf{x} \in V\} \subseteq V$$

is an *invariant subspace* of \mathbf{A} . Obviously

$$\mathbf{A}(\text{Im}\mathbf{A}) \subseteq \text{Im}\mathbf{A},$$

so that $\text{Im}\mathbf{A}$ is an invariant subspace of \mathbf{A} .

A *projector* is a square matrix \mathbf{P} that satisfies

$$\mathbf{P}^2 = \mathbf{P}.$$

A vector $\mathbf{x} \in \text{Im}\mathbf{P}$ if and only if there is $\mathbf{y} \in \mathbb{R}^n$ such that $\mathbf{x} = \mathbf{Py}$, so that

$$\mathbf{Px} = \mathbf{P}(\mathbf{Py}) = \mathbf{Py} = \mathbf{x}.$$

If \mathbf{P} is a projector, then also $\mathbf{Q} = \mathbf{I} - \mathbf{P}$ and \mathbf{P}^T are projectors as

$$(\mathbf{I} - \mathbf{P})^2 = \mathbf{I} - 2\mathbf{P} + \mathbf{P}^2 = \mathbf{I} - \mathbf{P} \quad \text{and} \quad (\mathbf{P}^T)^2 = (\mathbf{P}^2)^T = \mathbf{P}^T.$$

Since for any $\mathbf{x} \in \mathbb{R}^n$

$$\mathbf{x} = \mathbf{Px} + (\mathbf{I} - \mathbf{P})\mathbf{x},$$

it simply follows that $\text{Im}\mathbf{Q} = \text{Ker}\mathbf{P}$,

$$\mathbb{R}^n = \text{Im}\mathbf{P} + \text{Ker}\mathbf{P}, \quad \text{and} \quad \text{Ker}\mathbf{P} \cap \text{Im}\mathbf{P} = \{\mathbf{o}\}.$$

We say that \mathbf{P} is a projector onto $U = \text{Im}\mathbf{P}$ along $V = \text{Ker}\mathbf{P}$ and \mathbf{Q} is a complementary projector onto V along U . The above relations may be rewritten as

$$\text{Im}\mathbf{P} \oplus \text{Ker}\mathbf{P} = \mathbb{R}^n. \quad (2.1)$$

Let $(\pi(1), \dots, \pi(n))$ be a permutation of numbers $1, \dots, n$. Then, the mapping which assigns to each $\mathbf{v} = [v_i] \in \mathbb{R}^n$ a vector $[v_{\pi(1)}, \dots, v_{\pi(n)}]^T$ is associated with the *permutation matrix*

$$\mathbf{P} = [\mathbf{s}_{\pi(1)}, \dots, \mathbf{s}_{\pi(n)}],$$

where \mathbf{s}_i denotes the i th column of the identity matrix \mathbf{I}_n . If \mathbf{P} is a permutation matrix, then

$$\mathbf{P}\mathbf{P}^T = \mathbf{P}^T\mathbf{P} = \mathbf{I}.$$

Notice that if \mathbf{B} is a matrix obtained from a matrix \mathbf{A} by reordering of the rows of \mathbf{A} , then there is a permutation matrix \mathbf{P} such that $\mathbf{B} = \mathbf{P}\mathbf{A}$. Similarly, if \mathbf{B} is a matrix obtained from \mathbf{A} by reordering of the columns of \mathbf{A} , then there is a permutation matrix \mathbf{P} such that $\mathbf{B} = \mathbf{A}\mathbf{P}$.

2.3 Inverse and Generalized Inverse

If \mathbf{A} is a square full rank matrix, then there is the unique *inverse matrix* \mathbf{A}^{-1} such that

$$\mathbf{A}\mathbf{A}^{-1} = \mathbf{A}^{-1}\mathbf{A} = \mathbf{I}. \quad (2.2)$$

The mapping associated with \mathbf{A}^{-1} is inverse to that associated with \mathbf{A} .

If \mathbf{A}^{-1} exists, we say that \mathbf{A} is *nonsingular*. A square matrix is *singular* if its inverse matrix does not exist. If \mathbf{P} is a permutation matrix, then \mathbf{P} is nonsingular and

$$\mathbf{P}^{-1} = \mathbf{P}^T.$$

If \mathbf{A} is a nonsingular matrix, then $\mathbf{A}^{-1}\mathbf{b}$ is the unique solution of $\mathbf{A}\mathbf{x} = \mathbf{b}$.

If \mathbf{A} is nonsingular, then we can transpose (2.2) to get

$$(\mathbf{A}^{-1})^T \mathbf{A}^T = \mathbf{A}^T (\mathbf{A}^{-1})^T = \mathbf{I},$$

so that

$$(\mathbf{A}^T)^{-1} = (\mathbf{A}^{-1})^T. \quad (2.3)$$

It follows that if \mathbf{A} is symmetric, then \mathbf{A}^{-1} is symmetric.

If $\mathbf{A} \in \mathbb{R}^{n \times n}$ is positive definite, then also \mathbf{A}^{-1} is positive definite, as any vector $\mathbf{x} \neq \mathbf{0}$ can be expressed as $\mathbf{x} = \mathbf{A}\mathbf{y}$, $\mathbf{y} \neq \mathbf{0}$, and

$$\mathbf{x}^T \mathbf{A}^{-1} \mathbf{x} = (\mathbf{A}\mathbf{y})^T \mathbf{A}^{-1} \mathbf{A}\mathbf{y} = \mathbf{y}^T \mathbf{A}^T \mathbf{y} = \mathbf{y}^T \mathbf{A}\mathbf{y} > 0.$$

If \mathbf{A} and \mathbf{B} are nonsingular matrices, then it is easy to check that also \mathbf{AB} is nonsingular and

$$(\mathbf{AB})^{-1} = \mathbf{B}^{-1}\mathbf{A}^{-1}.$$

If

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_{\mathcal{I}\mathcal{I}} & \mathbf{H}_{\mathcal{I}\mathcal{J}} \\ \mathbf{H}_{\mathcal{J}\mathcal{I}} & \mathbf{H}_{\mathcal{J}\mathcal{J}} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{C} \end{bmatrix}$$

is an SPD block matrix, then we can directly evaluate

$$\mathbf{H}^{-1} = \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{C} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{B}^T\mathbf{S}^{-1}\mathbf{B}\mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{B}^T\mathbf{S}^{-1} \\ -\mathbf{S}^{-1}\mathbf{B}\mathbf{A}^{-1} & \mathbf{S}^{-1} \end{bmatrix}, \quad (2.4)$$

where $\mathbf{S} = \mathbf{C} - \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T$ denotes the *Schur complement* of \mathbf{H} with respect to \mathbf{A} . Thus

$$[\mathbf{A}^{-1}]_{\mathcal{J}\mathcal{J}} = \mathbf{S}^{-1}. \quad (2.5)$$

If $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \text{Im}\mathbf{A}$, then we can express a solution of the system of linear equations $\mathbf{Ax} = \mathbf{b}$ by means of a *left generalized inverse matrix* $\mathbf{A}^+ \in \mathbb{R}^{n \times m}$ which satisfies $\mathbf{AA}^+\mathbf{A} = \mathbf{A}$. Indeed, if $\mathbf{b} \in \text{Im}\mathbf{A}$, then there is \mathbf{y} such that $\mathbf{b} = \mathbf{Ay}$ and $\bar{\mathbf{x}} = \mathbf{A}^+\mathbf{b}$ satisfies

$$\mathbf{A}\bar{\mathbf{x}} = \mathbf{AA}^+\mathbf{b} = \mathbf{AA}^+\mathbf{Ay} = \mathbf{Ay} = \mathbf{b}.$$

Thus \mathbf{A}^+ acts on the range of \mathbf{A} like the inverse matrix. If \mathbf{A} is a nonsingular square matrix, then obviously

$$\mathbf{A}^+ = \mathbf{A}^{-1}.$$

Moreover, if $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{S} \in \mathbb{R}^{n \times p}$ are such that $\mathbf{AS} = \mathbf{O}$, then $(\mathbf{A}^+) + \mathbf{SS}^T$ is also a left generalized inverse as

$$\mathbf{A}(\mathbf{A}^+ + \mathbf{SS}^T)\mathbf{A} = \mathbf{AA}^+\mathbf{A} + \mathbf{ASS}^T\mathbf{A} = \mathbf{A}.$$

If \mathbf{A} is a symmetric singular matrix, then there is a permutation matrix \mathbf{P} such that

$$\mathbf{A} = \mathbf{P}^T \begin{bmatrix} \mathbf{B} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{CB}^{-1}\mathbf{C}^T \end{bmatrix} \mathbf{P},$$

where \mathbf{B} is a nonsingular matrix the dimension of which is equal to the rank of \mathbf{A} . It may be verified directly that the matrix

$$\mathbf{A}^\# = \mathbf{P}^T \begin{bmatrix} \mathbf{B}^{-1} & \mathbf{O}^T \\ \mathbf{O} & \mathbf{O} \end{bmatrix} \mathbf{P} \quad (2.6)$$

is a left generalized inverse of \mathbf{A} . If \mathbf{A} is SPS, then $\mathbf{A}^\#$ is also SPS. Notice that if $\mathbf{AS} = \mathbf{O}$, then $\mathbf{A}^+ = \mathbf{A}^\# + \mathbf{SS}^T$ is also an SPS generalized inverse.

2.4 Direct Methods for Solving Linear Equations

The inverse matrix is a useful tool for theoretical developments, but not for computations. It is often much more efficient to implement the multiplication of a vector by the inverse matrix by solving the related system of linear equations. We recall here briefly the *direct methods*, which reduce solving of the original system of linear equations to solving of a system or systems of equations with triangular matrices.

A matrix $\mathbf{L} = [l_{ij}]$ is *lower triangular* if $l_{ij} = 0$ for $i < j$. It is easy to solve a system $\mathbf{L}\mathbf{x} = \mathbf{b}$ with the nonsingular lower triangular matrix $\mathbf{L} \in \mathbb{R}^n$. As there is only one unknown in the first equation, we can find it and then substitute it into the remaining equations to obtain a system with the same structure, but with only $n - 1$ remaining unknowns. Repeating the procedure, we can find all the components of \mathbf{x} .

A similar procedure, but starting from the last equation, can be applied to a system with the nonsingular *upper triangular matrix* $\mathbf{U} = [u_{ij}]$ with $u_{ij} = 0$ for $i > j$.

The solution costs of a system with triangular matrices is proportional to the number of its nonzero entries. In particular, the solution of a system of linear equations with a *diagonal matrix* $\mathbf{D} = [d_{ij}]$, $d_{ij} = 0$ for $i \neq j$, reduces to the solution of a sequence of linear equations with one unknown.

If we are to solve the system of linear equations with a nonsingular matrix, we can use systematically *equivalent transformations* that do not change the solution in order to modify the original system to that with an upper triangular matrix. It is well known that the solutions of a system of linear equations are the same as the solutions of a system of linear equations obtained from the original system by interchanging two equations, replacing an equation by its nonzero multiple, or adding a multiple of one equation to another equation. The *Gauss elimination* for the solution of a system of linear equations with a nonsingular matrix thus consists of two steps: the *forward reduction*, which exploits equivalent transformations to reduce the original system to the system with an upper triangular matrix, and the *backward substitution*, which solves the resulting system with the upper triangular matrix.

Alternatively, we can use suitable matrix factorizations. For example, it is well known that any SPD matrix \mathbf{A} can be decomposed into the product

$$\mathbf{A} = \mathbf{L}\mathbf{L}^T, \quad (2.7)$$

where \mathbf{L} is a nonsingular lower triangular matrix with positive diagonal entries. Having the decomposition, we can evaluate $\mathbf{z} = \mathbf{A}^{-1}\mathbf{x}$ by solving the systems

$$\mathbf{L}\mathbf{y} = \mathbf{x} \quad \text{and} \quad \mathbf{L}^T\mathbf{z} = \mathbf{y}.$$

The factorization-based solvers may be especially useful when we are to solve several systems of equations with the same coefficients but different right-hand sides.

The method of evaluation of the factor \mathbf{L} is known as the *Cholesky factorization*. The Cholesky factor \mathbf{L} can be computed column by column. Suppose that

$$\mathbf{A} = \begin{bmatrix} a_{11} & \mathbf{a}_1^T \\ \mathbf{a}_1 & \mathbf{A}_{22} \end{bmatrix} \quad \text{and} \quad \mathbf{L} = \begin{bmatrix} l_{11} & \mathbf{o} \\ \mathbf{l}_1 & \mathbf{L}_{22} \end{bmatrix}.$$

Substituting for \mathbf{A} and \mathbf{L} into (2.7) and comparing the corresponding terms immediately reveals that

$$l_{11} = \sqrt{a_{11}}, \quad \mathbf{l}_1 = l_{11}^{-1} \mathbf{a}_1, \quad \mathbf{L}_{22} \mathbf{L}_{22}^T = \mathbf{A}_{22} - \mathbf{l}_1 \mathbf{l}_1^T. \quad (2.8)$$

This gives us the first column of \mathbf{L} , and the remaining factor \mathbf{L}_{22} is simply the Cholesky factor of the Schur complement $\mathbf{A}_{22} - \mathbf{l}_1 \mathbf{l}_1^T$ which is known to be positive definite, so we can find its first column by the above procedure. The algorithm can be implemented to exploit a sparsity pattern of \mathbf{A} , e.g., when $\mathbf{A} = [a_{ij}] \in \mathbb{R}^{n \times n}$ is a *band matrix* with $a_{ij} = 0$ for $|i - j| > b$, $b \ll n$.

If $\mathbf{A} \in \mathbb{R}^{n \times n}$ is only positive semidefinite, it can happen that $a_{11} = 0$. Then

$$0 \leq \mathbf{x}^T \mathbf{A} \mathbf{x} = \mathbf{y}^T \mathbf{A}_{22} \mathbf{y} + 2x_1 \mathbf{a}_1^T \mathbf{y}$$

for any vector $\mathbf{x} = [x_1, \mathbf{y}^T]^T$. The inequality implies that $\mathbf{a}_1 = \mathbf{o}$, as otherwise we could take $\mathbf{y} = -\mathbf{a}_1$ and large x_1 to get

$$\mathbf{y}^T \mathbf{A}_{22} \mathbf{y} + 2x_1 \mathbf{a}_1^T \mathbf{y} = \mathbf{a}_1^T \mathbf{A}_{22} \mathbf{a}_1 - 2x_1 \|\mathbf{a}_1\|^2 < 0.$$

Thus for \mathbf{A} symmetric positive semidefinite and $a_{11} = 0$, (2.8) reduces to

$$l_{11} = 0, \quad \mathbf{l}_1 = \mathbf{o}, \quad \mathbf{L}_{22} \mathbf{L}_{22}^T = \mathbf{A}_{22}. \quad (2.9)$$

This simple modification assumes exact arithmetics. In the computer arithmetics, the decision whether a_{11} is to be treated as zero depends on some small $\varepsilon > 0$. Alternatively, it is possible to exploit some additional information. For example, any orthonormal basis of the kernel of a matrix can be used to identify the zero rows (and columns) of a Cholesky factor by means of the following lemma.

Lemma 2.1 *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ denote an SPS matrix the kernel of which is spanned by the full column rank matrix $\mathbf{R} \in \mathbb{R}^{n \times d}$ with orthonormal columns. Let*

$$\mathcal{J} = \{i_1, \dots, i_d\}, \quad 1 \leq i_1 < i_2 < \dots < i_d \leq n,$$

denote a set of indices, and let $\mathcal{J} = \mathcal{N} - \mathcal{J}$, $\mathcal{N} = \{1, 2, \dots, n\}$.

Then

$$\lambda_{\min}(\mathbf{A}_{\mathcal{J}\mathcal{J}}) \geq \bar{\lambda}_{\min}(\mathbf{A}) \sigma_{\min}^4(\mathbf{R}_{\mathcal{J}*}), \quad (2.10)$$

where $\bar{\lambda}_{\min}(\mathbf{A})$ and $\sigma_{\min}(\mathbf{R}_{\mathcal{J}})$ denote the least nonzero eigenvalue of \mathbf{A} and the least singular value of $\mathbf{R}_{\mathcal{J}*}$, respectively.*

Proof See Dostál et al. [4]. □

If we can identify a nonsingular block $\mathbf{A}_{\mathcal{J}\mathcal{J}}$ of \mathbf{A} , then we can reduce the problems related to the manipulation with singular matrices to those with smaller matrices using the decomposition

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{\mathcal{J}\mathcal{J}} & \mathbf{A}_{\mathcal{J}\mathcal{S}} \\ \mathbf{A}_{\mathcal{S}\mathcal{J}} & \mathbf{A}_{\mathcal{S}\mathcal{S}} \end{bmatrix} = \begin{bmatrix} \mathbf{L}_{\mathcal{J}\mathcal{J}} & \mathbf{O} \\ \mathbf{L}_{\mathcal{S}\mathcal{J}} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_{\mathcal{J}\mathcal{J}}^T & \mathbf{L}_{\mathcal{S}\mathcal{J}}^T \\ \mathbf{O} & \mathbf{S} \end{bmatrix}, \quad (2.11)$$

where $\mathbf{L}_{\mathcal{J}\mathcal{J}} \in \mathbb{R}^{r \times r}$ is a lower factor of the Cholesky decomposition of $\mathbf{A}_{\mathcal{J}\mathcal{J}}$, $\mathbf{L}_{\mathcal{S}\mathcal{J}} \in \mathbb{R}^{s \times r}$, $r = n - s$, $\mathbf{L}_{\mathcal{S}\mathcal{J}} = \mathbf{A}_{\mathcal{S}\mathcal{J}} \mathbf{L}_{\mathcal{J}\mathcal{J}}^{-T}$, and $\mathbf{S} \in \mathbb{R}^{s \times s}$ is the Schur complement matrix of the block $\mathbf{A}_{\mathcal{J}\mathcal{J}}$ defined by

$$\mathbf{S} = \mathbf{A}_{\mathcal{S}\mathcal{S}} - \mathbf{A}_{\mathcal{S}\mathcal{J}} \mathbf{A}_{\mathcal{J}\mathcal{J}}^{-1} \mathbf{A}_{\mathcal{J}\mathcal{S}}.$$

The decomposition (2.11) is a useful tool for the effective construction of a generalized inverse or for the effective evaluation of the multiplication of a vector by a generalized inverse.

Lemma 2.2 Let \mathbf{A} denote an SPS block matrix as in (2.11), let $\mathbf{e} \in \text{Ker} \mathbf{A}$, and let \mathbf{S} denote the Schur complement matrix of the block $\mathbf{A}_{\mathcal{J}\mathcal{J}}$. Then $\mathbf{e}_{\mathcal{J}} \in \text{Ker} \mathbf{S}$ and if \mathbf{S}^+ denote any generalized inverse of \mathbf{S} , then the matrix \mathbf{A}^+ defined by

$$\mathbf{A}^+ = \begin{bmatrix} \mathbf{L}_{\mathcal{J}\mathcal{J}}^{-T} & -\mathbf{L}_{\mathcal{J}\mathcal{J}}^{-T} \mathbf{L}_{\mathcal{S}\mathcal{J}}^T \mathbf{S}^+ \\ \mathbf{O} & \mathbf{S}^+ \end{bmatrix} \begin{bmatrix} \mathbf{L}_{\mathcal{J}\mathcal{J}}^{-1} & \mathbf{O} \\ -\mathbf{L}_{\mathcal{S}\mathcal{J}} \mathbf{L}_{\mathcal{J}\mathcal{J}}^{-1} & \mathbf{I} \end{bmatrix} \quad (2.12)$$

is a generalized inverse of \mathbf{A} which satisfies

$$[\mathbf{A}^+]_{\mathcal{J}\mathcal{J}} = \mathbf{S}^+. \quad (2.13)$$

Proof If $\mathbf{A}\mathbf{e} = \mathbf{o}$, then

$$\mathbf{A}_{\mathcal{J}\mathcal{J}} \mathbf{e}_{\mathcal{J}} + \mathbf{A}_{\mathcal{J}\mathcal{S}} \mathbf{e}_{\mathcal{S}} = \mathbf{o}, \quad \mathbf{A}_{\mathcal{S}\mathcal{J}} \mathbf{e}_{\mathcal{J}} + \mathbf{A}_{\mathcal{S}\mathcal{S}} \mathbf{e}_{\mathcal{S}} = \mathbf{o},$$

and

$$\begin{aligned} \mathbf{S} \mathbf{e}_{\mathcal{J}} &= (\mathbf{A}_{\mathcal{S}\mathcal{S}} - \mathbf{A}_{\mathcal{S}\mathcal{J}} \mathbf{A}_{\mathcal{J}\mathcal{J}}^{-1} \mathbf{A}_{\mathcal{J}\mathcal{S}}) \mathbf{e}_{\mathcal{J}} \\ &= \mathbf{A}_{\mathcal{S}\mathcal{S}} \mathbf{e}_{\mathcal{J}} - \mathbf{A}_{\mathcal{S}\mathcal{J}} \mathbf{A}_{\mathcal{J}\mathcal{J}}^{-1} (-\mathbf{A}_{\mathcal{J}\mathcal{S}} \mathbf{e}_{\mathcal{J}}) = \mathbf{o}, \end{aligned}$$

i.e., $\mathbf{e}_{\mathcal{J}} \in \text{Ker} \mathbf{S}$. The rest can be verified directly. \square

Notice that

$$\mathbf{S}^+ = [\mathbf{A}^+]_{\mathcal{J}\mathcal{J}}. \quad (2.14)$$

If \mathbf{A} is a full rank matrix, then we get

$$\mathbf{S}^{-1} = [\mathbf{A}^{-1}]_{\mathcal{J}\mathcal{J}}, \quad (2.15)$$

which agrees with (2.5).

2.5 Norms

General concepts of size and distance in a vector space are expressed by norms. A *norm* on \mathbb{R}^n is a function which assigns to each $\mathbf{x} \in \mathbb{R}^n$ a number $\|\mathbf{x}\| \in \mathbb{R}$ in such a way that for any vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and any scalar $\alpha \in \mathbb{R}$, the following three conditions are satisfied:

- (i) $\|\mathbf{x}\| \geq 0$, and $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{o}$.
- (ii) $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$.
- (iii) $\|\alpha\mathbf{x}\| = |\alpha| \|\mathbf{x}\|$.

It is easy to check that the functions

$$\|\mathbf{x}\|_2 = \sqrt{\mathbf{x}_1^2 + \cdots + \mathbf{x}_n^2} \quad \text{and} \quad \|\mathbf{x}\|_\infty = \max\{|\mathbf{x}_1|, \dots, |\mathbf{x}_n|\}$$

are norms. They are called ℓ_2 (Euclidean) and ℓ_∞ norms, respectively.

Given a norm defined on the domain and the range of a matrix \mathbf{A} , we can define the *induced norm* $\|\mathbf{A}\|$ of \mathbf{A} by

$$\|\mathbf{A}\| = \sup_{\|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\| = \sup_{\mathbf{x} \neq \mathbf{o}} \frac{\|\mathbf{A}\mathbf{x}\|}{\|\mathbf{x}\|}.$$

If $\mathbf{B} \neq \mathbf{O}$, then

$$\|\mathbf{AB}\| = \sup_{\mathbf{x} \neq \mathbf{o}} \frac{\|\mathbf{AB}\mathbf{x}\|}{\|\mathbf{x}\|} = \sup_{\mathbf{B}\mathbf{x} \neq \mathbf{o}} \frac{\|\mathbf{AB}\mathbf{x}\|}{\|\mathbf{B}\mathbf{x}\|} \frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|} \leq \sup_{\substack{\mathbf{y} \in \text{Im } \mathbf{B}, \\ \mathbf{y} \neq \mathbf{o}}} \frac{\|\mathbf{A}\mathbf{y}\|}{\|\mathbf{y}\|} \sup_{\mathbf{x} \neq \mathbf{o}} \frac{\|\mathbf{B}\mathbf{x}\|}{\|\mathbf{x}\|}.$$

It follows easily that the induced norm is *submultiplicative*, i.e.,

$$\|\mathbf{AB}\| \leq \|\mathbf{A}\| \|\mathbf{B}\| \leq \|\mathbf{A}\| \|\mathbf{B}\|. \quad (2.16)$$

If $\mathbf{A} = [a_{ij}] \in \mathbb{R}^{m \times n}$ and $\mathbf{x} = [x_i] \in \mathbb{R}^n$, then

$$\|\mathbf{Ax}\|_\infty = \max_{i=1,\dots,m} \left| \sum_{j=1}^n a_{ij}x_j \right| \leq \max_{i=1,\dots,m} \sum_{j=1}^n |a_{ij}| |x_j| \leq \|\mathbf{x}\|_\infty \max_{i=1,\dots,m} \sum_{j=1}^n |a_{ij}|,$$

that is, $\|\mathbf{A}\|_\infty \leq \max_{i=1,\dots,m} \sum_{j=1}^n |a_{ij}|$. Since the last inequality turns into the equality for a vector \mathbf{x} with suitably chosen entries $x_i \in \{1, -1\}$, we have

$$\|\mathbf{A}\|_\infty = \max_{i=1,\dots,m} \sum_{j=1}^n |a_{ij}|. \quad (2.17)$$

2.6 Scalar Products

General concepts of length and angle in a vector space are introduced by means of a *scalar product*; it is the mapping which assigns to each couple $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ a number $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}$ in such a way that for any vectors $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbb{R}^n$ and any scalar $\alpha \in \mathbb{R}$, the following four conditions are satisfied:

- (i) $(\mathbf{x}, \mathbf{y} + \mathbf{z}) = (\mathbf{x}, \mathbf{y}) + (\mathbf{x}, \mathbf{z})$.
- (ii) $(\alpha \mathbf{x}, \mathbf{y}) = \alpha(\mathbf{x}, \mathbf{y})$.
- (iii) $(\mathbf{x}, \mathbf{y}) = (\mathbf{y}, \mathbf{x})$.
- (iv) $(\mathbf{x}, \mathbf{x}) > 0$ for $\mathbf{x} \neq \mathbf{o}$.

The scalar product is an SPD form, see also Chap. 4.

We often use the *Euclidean scalar product* or the *Euclidean inner product* which assigns to each couple of vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ a number defined by

$$(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y}.$$

In more complicated expressions, we often denote the Euclidean scalar product in \mathbb{R}^3 by dot, so that

$$\mathbf{x} \cdot \mathbf{y} = \mathbf{x}^T \mathbf{y}.$$

If \mathbf{A} is an SPD matrix, then we can define the more general *A-scalar product* on \mathbb{R}^n by

$$(\mathbf{x}, \mathbf{y})_{\mathbf{A}} = \mathbf{x}^T \mathbf{A} \mathbf{y}.$$

We denote for any $\mathbf{x} \in \mathbb{R}^n$ its *Euclidean norm* and *A-norm* by

$$\|\mathbf{x}\| = (\mathbf{x}, \mathbf{x})^{1/2}, \quad \|\mathbf{x}\|_{\mathbf{A}} = (\mathbf{x}, \mathbf{x})_{\mathbf{A}}^{1/2}.$$

It is easy to see that any norm induced by a scalar product satisfies the properties (i) and (iii) of the norm. The property (ii) follows from the *Cauchy–Schwarz inequality*

$$(\mathbf{x}, \mathbf{y})^2 \leq \|\mathbf{x}\|^2 \|\mathbf{y}\|^2, \quad (2.18)$$

which is valid for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and any scalar product. The bound is tight in the sense that the inequality becomes the equality when \mathbf{x}, \mathbf{y} are dependent.

A pair of vectors \mathbf{x} and \mathbf{y} is *orthogonal* (with respect to a given scalar product) if

$$(\mathbf{x}, \mathbf{y}) = 0.$$

The vectors \mathbf{x} and \mathbf{y} that are orthogonal in \mathbf{A} -scalar product are called *\mathbf{A} -conjugate* or briefly *conjugate*.

Two sets of vectors \mathcal{E} and \mathcal{F} are *orthogonal* (also stated “ \mathcal{E} orthogonal to \mathcal{F} ”) if any $\mathbf{x} \in \mathcal{E}$ is orthogonal to any $\mathbf{y} \in \mathcal{F}$. The set \mathcal{E}^\perp of all the vectors of \mathbb{R}^n that are orthogonal to $\mathcal{E} \subseteq \mathbb{R}^n$ is a vector space called an *orthogonal complement* of \mathcal{E} . If $\mathcal{E} \subseteq \mathbb{R}^n$, then

$$\mathbb{R}^n = \text{Span } \mathcal{E} \oplus \mathcal{E}^\perp.$$

A set of vectors \mathcal{E} is *orthogonal* if its elements are pairwise orthogonal, i.e., any $\mathbf{x} \in \mathcal{E}$ is orthogonal to any $\mathbf{y} \in \mathcal{E}$, $\mathbf{y} \neq \mathbf{x}$. A set of vectors \mathcal{E} is *orthonormal* if it is orthogonal and $(\mathbf{x}, \mathbf{x}) = 1$ for any $\mathbf{x} \in \mathcal{E}$.

A square matrix \mathbf{U} is *orthogonal* if $\mathbf{U}^T \mathbf{U} = \mathbf{I}$, that is, $\mathbf{U}^{-1} = \mathbf{U}^T$. Multiplication by an orthogonal matrix \mathbf{U} preserves both the angles between any two vectors and the Euclidean norm of any vector as

$$(\mathbf{U}\mathbf{x})^T \mathbf{U}\mathbf{y} = \mathbf{x}^T \mathbf{U}^T \mathbf{U}\mathbf{y} = \mathbf{x}^T \mathbf{y}.$$

A matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ is an *orthogonal projector* if \mathbf{P} is a projector, i.e., $\mathbf{P}^2 = \mathbf{P}$, and $\text{Im } \mathbf{P}$ is orthogonal to $\text{Ker } \mathbf{P}$. The latter condition can be rewritten equivalently as

$$\mathbf{P}^T (\mathbf{I} - \mathbf{P}) = \mathbf{O}.$$

It simply follows that

$$\mathbf{P}^T = \mathbf{P}^T \mathbf{P} = \mathbf{P},$$

so that orthogonal projectors are symmetric matrices and symmetric projectors are orthogonal projectors. If \mathbf{P} is an orthogonal projector, then $\mathbf{I} - \mathbf{P}$ is also an orthogonal projector as

$$(\mathbf{I} - \mathbf{P})^2 = \mathbf{I} - 2\mathbf{P} + \mathbf{P}^2 = \mathbf{I} - \mathbf{P} \quad \text{and} \quad (\mathbf{I} - \mathbf{P})^T \mathbf{P} = (\mathbf{I} - \mathbf{P})\mathbf{P} = \mathbf{O}.$$

If $U \subseteq \mathbb{R}^n$ is the subspace spanned by the columns of a full column rank matrix $\mathbf{U} \in \mathbb{R}^{m \times n}$, then

$$\mathbf{P} = \mathbf{U}(\mathbf{U}^T \mathbf{U})^{-1} \mathbf{U}^T$$

is an orthogonal projector as

$$P^2 = U(U^T U)^{-1} U^T U (U^T U)^{-1} U^T = P \quad \text{and} \quad P^T = P.$$

Since any vector $\mathbf{x} \in U$ may be written in the form $\mathbf{x} = U\mathbf{y}$ and

$$P\mathbf{x} = U(U^T U)^{-1} U^T U\mathbf{y} = U\mathbf{y} = \mathbf{x},$$

it follows that

$$U = \text{Im}P.$$

Observe that $U^T U$ is nonsingular; since $U^T U\mathbf{x} = \mathbf{0}$ implies

$$\|U\mathbf{x}\|^2 = \mathbf{x}^T (U^T U\mathbf{x}) = 0,$$

it follows that $\mathbf{x} = \mathbf{0}$ by the assumption on the full column rank of U .

2.7 Eigenvalues and Eigenvectors

Let $\mathbf{A} \in \mathbb{C}^{n \times n}$ denote a square matrix with complex entries. If a vector $\mathbf{e} \in \mathbb{C}^n$ and a scalar $\lambda \in \mathbb{C}$ satisfy

$$\mathbf{A}\mathbf{e} = \lambda\mathbf{e}, \tag{2.19}$$

then \mathbf{e} is said to be an *eigenvector* of \mathbf{A} associated with an *eigenvalue* λ . A vector \mathbf{e} is an eigenvector of \mathbf{A} if and only if $\text{Span}\{\mathbf{e}\}$ is an invariant subspace of \mathbf{A} ; the restriction $\mathbf{A}|_{\text{Span}\{\mathbf{e}\}}$ reduces to the multiplication by λ . If $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ are eigenvectors of a symmetric matrix \mathbf{A} , then it is easy to check that $\text{Span}\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ and $\text{Span}\{\mathbf{e}_1, \dots, \mathbf{e}_k\}^\perp$ are invariant subspaces.

The set of all eigenvalues of \mathbf{A} is called the *spectrum* of \mathbf{A} ; we denote it by $\sigma(\mathbf{A})$. Obviously, $\lambda \in \sigma(\mathbf{A})$ if and only if $\mathbf{A} - \lambda\mathbf{I}$ is singular, and $0 \in \sigma(\mathbf{A})$ if and only if \mathbf{A} is singular. If $\lambda \neq 0$, $\lambda \in \sigma(\mathbf{A})$, then we can multiply (2.19) by $\lambda^{-1}\mathbf{A}^{-1}$ to get $\mathbf{A}^{-1}\mathbf{e} = \lambda^{-1}\mathbf{e}$, so we can write

$$\sigma(\mathbf{A}^{-1}) = \sigma^{-1}(\mathbf{A}).$$

If $U \subseteq \mathbb{C}^n$ is an invariant subspace of $\mathbf{A} \in \mathbb{C}^{n \times n}$, then we denote by $\sigma(\mathbf{A}|_U)$ the eigenvalues of \mathbf{A} that correspond to the eigenvectors belonging to U .

Since it is well known that a matrix is singular if and only if its determinant is equal to zero, it follows that the eigenvalues of \mathbf{A} are the roots of the *characteristic equation*

$$\det(\mathbf{A} - \lambda\mathbf{I}) = 0. \tag{2.20}$$

The *characteristic polynomial* $p_{\mathbf{A}}(\lambda) = \det(\mathbf{A} - \lambda\mathbf{I})$ is of the degree n . Thus there are at most n distinct eigenvalues and $\sigma(\mathbf{A})$ is not the empty set.

Even though it is in general difficult to evaluate the eigenvalues of a given matrix \mathbf{A} , it is still possible to get nontrivial information about $\sigma(\mathbf{A})$ without heavy computations. Useful information about the location of eigenvalues can be obtained by *Gershgorin's theorem*, which guarantees that every eigenvalue of $\mathbf{A} = [a_{ij}] \in \mathbb{C}^{n \times n}$ is located in at least one of the n circular disks in the complex plane with the centers a_{ii} and radii $r_i = \sum_{j \neq i} |a_{ij}|$.

The *eigenvalues of a real symmetric matrix are real*. Since it is easy to check whether a matrix is symmetric, this gives us useful information about the location of eigenvalues.

Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ denote a real symmetric matrix, let $\mathcal{J} = \{1, \dots, n-1\}$, and let $\mathbf{A}^1 = \mathbf{A}_{\mathcal{J}, \mathcal{J}}$. Let $\lambda_1 \geq \dots \geq \lambda_n$ and $\lambda_1^1 \geq \dots \geq \lambda_{n-1}^1$ denote the eigenvalues of \mathbf{A} and \mathbf{A}^1 , respectively. Then by the *Cauchy interlacing theorem*

$$\lambda_1 \geq \lambda_1^1 \geq \lambda_2 \geq \lambda_2^1 \geq \dots \geq \lambda_{n-1}^1 \geq \lambda_n. \quad (2.21)$$

2.8 Matrix Decompositions

If $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a symmetric matrix, then it is possible to find n orthonormal eigenvectors $\mathbf{e}_1, \dots, \mathbf{e}_n$ that form the basis of \mathbb{R}^n . Moreover, the corresponding eigenvalues are real. Denoting by $\mathbf{U} = [\mathbf{e}_1, \dots, \mathbf{e}_n] \in \mathbb{R}^{n \times n}$ an orthogonal matrix the columns of which are the eigenvectors, we may write the *spectral decomposition* of \mathbf{A} as

$$\mathbf{A} = \mathbf{U} \mathbf{D} \mathbf{U}^T, \quad (2.22)$$

where $\mathbf{D} = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^{n \times n}$ is the diagonal matrix the diagonal entries of which are the eigenvalues corresponding to the eigenvectors $\mathbf{e}_1, \dots, \mathbf{e}_n$. Reordering the columns of \mathbf{U} , we can achieve that $\lambda_1 \geq \dots \geq \lambda_n$.

The spectral decomposition reveals close relations between the properties of a symmetric matrix and its eigenvalues. Thus, a symmetric matrix is SPD if and only if all its eigenvalues are positive, and it is SPS if and only if they are nonnegative. The rank of a symmetric matrix is equal to the number of nonzero entries of \mathbf{D} .

If \mathbf{A} is symmetric, then we can use the spectral decomposition (2.22) to check that for any nonzero \mathbf{x}

$$\lambda_1 = \lambda_{\max} \geq \|\mathbf{x}\|^{-2} \mathbf{x}^T \mathbf{A} \mathbf{x} \geq \lambda_{\min} = \lambda_n. \quad (2.23)$$

Thus for any symmetric positive definite matrix \mathbf{A}

$$\|\mathbf{A}\| = \lambda_{\max}, \quad \|\mathbf{A}^{-1}\| = \lambda_{\min}^{-1}, \quad \|\mathbf{x}\|_{\mathbf{A}} \leq \lambda_{\max} \|\mathbf{x}\|, \quad \|\mathbf{x}\|_{\mathbf{A}^{-1}} \leq \lambda_{\min}^{-1} \|\mathbf{x}\|. \quad (2.24)$$

The *spectral condition number* $\kappa(\mathbf{A}) = \|\mathbf{A}\| \|\mathbf{A}^{-1}\|$, which is a measure of departure from the identity, can be expressed for real symmetric matrix by

$$\kappa(\mathbf{A}) = \lambda_{\max}/\lambda_{\min}.$$

If \mathbf{A} is a real symmetric matrix and f is a real function defined on $\sigma(\mathbf{A})$, we can use the spectral decomposition to define the *scalar function* by

$$f(\mathbf{A}) = \mathbf{U}f(\mathbf{D})\mathbf{U}^T,$$

where $f(\mathbf{D}) = \text{diag}(f(\lambda_1), \dots, f(\lambda_n))$. It is easy to check that if a is the identity function on \mathbb{R} defined by $a(x) = x$, then

$$a(\mathbf{A}) = \mathbf{A},$$

and if f and g are real functions defined on $\sigma(\mathbf{A})$, then

$$(f + g)(\mathbf{A}) = f(\mathbf{A}) + g(\mathbf{A}) \quad \text{and} \quad (f \cdot g)(\mathbf{A}) = f(\mathbf{A})g(\mathbf{A}).$$

Moreover, if $f(x) \geq 0$ for $x \in \sigma(\mathbf{A})$, then $f(\mathbf{A})$ is SPS, and if $f(x) > 0$ for $x \in \sigma(\mathbf{A})$, then $f(\mathbf{A})$ is SPD. For example, if \mathbf{A} is SPD, then

$$\mathbf{A} = \mathbf{A}^{1/2}\mathbf{A}^{1/2}.$$

Obviously

$$\sigma(f(\mathbf{A})) = f(\sigma(\mathbf{A})), \quad (2.25)$$

and if \mathbf{e}_i is an eigenvector corresponding to $\lambda_i \in \sigma(\mathbf{A})$, then it is also an eigenvector of $f(\mathbf{A})$ corresponding to $f(\lambda_i)$. It follows easily that for any SPS matrix

$$\text{Im}\mathbf{A} = \text{Im}\mathbf{A}^{1/2} \quad \text{and} \quad \text{Ker}\mathbf{A} = \text{Ker}\mathbf{A}^{1/2}. \quad (2.26)$$

A key to understanding nonsymmetric matrices is the *singular value decomposition* (SVD). If $\mathbf{B} \in \mathbb{R}^{m \times n}$, then SVD of \mathbf{B} is given by

$$\mathbf{B} = \mathbf{U}\mathbf{S}\mathbf{V}^T, \quad (2.27)$$

where $\mathbf{U} \in \mathbb{R}^{m \times m}$ and $\mathbf{V} \in \mathbb{R}^{n \times n}$ are orthogonal, and $\mathbf{S} \in \mathbb{R}^{m \times n}$ is a diagonal matrix with nonnegative diagonal entries $\sigma_1 \geq \dots \geq \sigma_{\min\{m,n\}} = \sigma_{\min}$ called *singular values* of \mathbf{B} . If \mathbf{A} is not a full rank matrix, then it is often more convenient to use the *reduced singular value decomposition* (RSVD)

$$\mathbf{B} = \widehat{\mathbf{U}}\widehat{\mathbf{S}}\widehat{\mathbf{V}}^T, \quad (2.28)$$

where $\widehat{\mathbf{U}} \in \mathbb{R}^{m \times r}$ and $\widehat{\mathbf{V}} \in \mathbb{R}^{n \times r}$ are matrices with orthonormal columns, $\widehat{\mathbf{S}} \in \mathbb{R}^{r \times r}$ is a nonsingular diagonal matrix with positive diagonal entries $\sigma_1 \geq \dots \geq \sigma_r = \bar{\sigma}_{\min}$, and $r \leq \min\{m, n\}$ is the rank of \mathbf{B} . The matrices $\widehat{\mathbf{U}}$ and $\widehat{\mathbf{V}}$ are formed by the first r columns of \mathbf{U} and \mathbf{V} . If $\mathbf{x} \in \mathbb{R}^m$, then

$$\mathbf{B}\mathbf{x} = \widehat{\mathbf{U}}\widehat{\mathbf{S}}\widehat{\mathbf{V}}^T \mathbf{x} = (\widehat{\mathbf{U}}\widehat{\mathbf{S}}\widehat{\mathbf{V}}^T)(\widehat{\mathbf{V}}\widehat{\mathbf{S}}\widehat{\mathbf{U}}^T)(\widehat{\mathbf{U}}\widehat{\mathbf{S}}^{-1}\mathbf{V}^T \mathbf{x}) = \mathbf{B}\mathbf{B}^T \mathbf{y},$$

so that

$$\text{Im } \mathbf{B} = \text{Im } \mathbf{B}\mathbf{B}^T. \quad (2.29)$$

If $\mathbf{B} = \widehat{\mathbf{U}}\widehat{\mathbf{S}}\widehat{\mathbf{V}}^T$ is RSVD, then

$$\text{Im } \mathbf{B} = \text{Im } \widehat{\mathbf{U}}, \quad \text{Ker } \mathbf{B} = (\text{Im } \widehat{\mathbf{V}})^\perp.$$

It follows that

$$\text{Im } \mathbf{B}^T = (\text{Ker } \mathbf{B})^\perp. \quad (2.30)$$

The SVD reveals close relations between the properties of a matrix and its singular values. Thus, the rank of $\mathbf{B} \in \mathbb{R}^{m \times n}$ is equal to the number of its nonzero singular values,

$$\|\mathbf{B}\| = \|\mathbf{B}^T\| = \sigma_1, \quad (2.31)$$

and for any vector $\mathbf{x} \in \mathbb{R}^n$

$$\sigma_{\min} \|\mathbf{x}\| \leq \|\mathbf{B}\mathbf{x}\| \leq \|\mathbf{B}\| \|\mathbf{x}\|. \quad (2.32)$$

Let $\bar{\sigma}_{\min}$ denote the least nonzero singular value of $\mathbf{B} \in \mathbb{R}^{m \times n}$, let $\mathbf{x} \in \text{Im } \mathbf{B}^T$, and consider the RSVD $\mathbf{B} = \widehat{\mathbf{U}}\widehat{\mathbf{S}}\widehat{\mathbf{V}}^T$ with $\widehat{\mathbf{U}} \in \mathbb{R}^{m \times r}$, $\widehat{\mathbf{V}} \in \mathbb{R}^{n \times r}$, and $\widehat{\mathbf{S}} \in \mathbb{R}^{r \times r}$. Then there is $\mathbf{y} \in \mathbb{R}^r$ such that $\mathbf{x} = \widehat{\mathbf{V}}\mathbf{y}$ and

$$\|\mathbf{B}\mathbf{x}\| = \|\widehat{\mathbf{U}}\widehat{\mathbf{S}}\widehat{\mathbf{V}}^T \widehat{\mathbf{V}}\mathbf{y}\| = \|\widehat{\mathbf{U}}\widehat{\mathbf{S}}\mathbf{y}\| = \|\widehat{\mathbf{S}}\mathbf{y}\| \geq \bar{\sigma}_{\min} \|\mathbf{y}\|.$$

Since

$$\|\mathbf{x}\| = \|\widehat{\mathbf{V}}\mathbf{y}\| = \|\mathbf{y}\|,$$

we conclude that

$$\bar{\sigma}_{\min} \|\mathbf{x}\| \leq \|\mathbf{B}\mathbf{x}\| \quad \text{for any } \mathbf{x} \in \text{Im } \mathbf{B}^T, \quad (2.33)$$

or, equivalently,

$$\bar{\sigma}_{\min} \|\mathbf{x}\| \leq \|\mathbf{B}^T \mathbf{x}\| \quad \text{for any } \mathbf{x} \in \text{Im } \mathbf{B}. \quad (2.34)$$

The SVD (2.27) can be used to introduce the *Moore–Penrose generalized inverse* of an $m \times n$ matrix \mathbf{B} by

$$\mathbf{B}^\dagger = \mathbf{V}\mathbf{S}^\dagger \mathbf{U}^T,$$

where \mathbf{S}^\dagger is the diagonal matrix with the entries $[\mathbf{S}^\dagger]_{ii} = 0$ if $\sigma_i = 0$ and $[\mathbf{S}^\dagger]_{ii} = \sigma_i^{-1}$ otherwise. It is easy to check that

$$\mathbf{B}\mathbf{B}^\dagger \mathbf{B} = \mathbf{U}\mathbf{S}\mathbf{V}^T \mathbf{V}\mathbf{S}^\dagger \mathbf{U}^T \mathbf{U}\mathbf{S}\mathbf{V}^T = \mathbf{U}\mathbf{S}\mathbf{V}^T = \mathbf{B}, \quad (2.35)$$

so that the Moore–Penrose generalized inverse is a generalized inverse. If \mathbf{B} is a full row rank matrix, then it may be checked directly that

$$\mathbf{B}^\dagger = \mathbf{B}^T (\mathbf{B}\mathbf{B}^T)^{-1}.$$

If \mathbf{B} is a singular matrix and $\mathbf{c} \in \text{Im}\mathbf{B}$, then $\mathbf{x}_{\text{LS}} = \mathbf{B}^\dagger \mathbf{c}$ is a solution of the system of linear equations $\mathbf{B}\mathbf{x} = \mathbf{c}$, i.e.,

$$\mathbf{B}\mathbf{x}_{\text{LS}} = \mathbf{c}.$$

Notice that $\mathbf{x}_{\text{LS}} \in \text{Im}\mathbf{B}^T$, so that if $\bar{\mathbf{x}}$ is any other solution, then $\bar{\mathbf{x}} = \mathbf{x}_{\text{LS}} + \mathbf{d}$, where $\mathbf{d} \in \text{Ker}\mathbf{B}$, $\mathbf{x}_{\text{LS}}^T \mathbf{d} = 0$, and

$$\|\mathbf{x}_{\text{LS}}\|^2 \leq \|\bar{\mathbf{x}}_{\text{LS}}\|^2 + \|\mathbf{d}\|^2 = \|\bar{\mathbf{x}}\|^2. \quad (2.36)$$

The vector \mathbf{x}_{LS} is called the *least square solution* of $\mathbf{B}\mathbf{x} = \mathbf{c}$.

Obviously

$$\|\mathbf{B}^\dagger\| = \bar{\sigma}_{\min}^{-1}, \quad (2.37)$$

where $\bar{\sigma}_{\min}$ denotes the least nonzero singular value of \mathbf{B} , so that

$$\|\mathbf{x}_{\text{LS}}\| = \|\mathbf{B}^\dagger \mathbf{c}\| \leq \bar{\sigma}_{\min}^{-1} \|\mathbf{c}\|. \quad (2.38)$$

It can be verified directly that

$$(\mathbf{B}^\dagger)^T = (\mathbf{B}^T)^\dagger.$$

2.9 Graphs, Walks, and Adjacency Matrices

We shall need some simple results of graph theory and linear algebra. Let us recall that the *vertices* V_i and the *edges* of the *graph of the mesh of the triangulation* $\mathcal{T} = \{\tau_i\}$ of a polyhedral domain Ω are the nodes of the mesh and their adjacent couples $V_i V_j$, respectively. Recall that the edges V_i and V_j are *adjacent* if there is an element $\tau_k \in \mathcal{T}$ such that $V_i V_j$ is the edge of τ_k . The graph is fully described by the *adjacency matrix* \mathbf{D} with the nonzero entries d_{ij} equal to one if the nodes V_i and V_j are adjacent. Since the graph of the mesh is not oriented and does not contain loops, the adjacency matrix is symmetric and its diagonal entries d_{ii} are equal to zero. Let us also recall that the *walk of length k* in the mesh of \mathcal{T} is a sequence of the distinct nodes V_{i_1}, \dots, V_{i_k} such that the edges $V_{i_j} V_{i_{j+1}}$, $j = 1, 2, \dots, k-1$ belong to the graph of the mesh. Thus

$$d_{i_j i_{j+1}} = 1, \quad j = 1, \dots, k-1.$$

The walk $(V_{i_1}, \dots, V_{i_k})$ starts at V_{i_1} and ends at V_{i_k} . Moreover, we call a walk between nodes V_i and V_k an (i, k) -walk. We use the following well-known observation [5].

Lemma 2.3 *Let \mathbf{D} be the adjacency matrix of the mesh of \mathcal{T} and $\mathbf{B} = \mathbf{D}^k$. Then each entry b_{ij} of \mathbf{B} gives the number of distinct (i, j) -walks of length k .*

Proof To see why our lemma holds, we use induction on k . For $k = 1$ the claim follows immediately from the definition of \mathbf{D} . Suppose that for some $k \geq 1$ the entry b_{ij} in $\mathbf{B} = \mathbf{D}^k$ gives the number of distinct (i, j) -walks of length k . For convenience we denote $\mathbf{C} = \mathbf{D}^{k+1}$, so that $\mathbf{C} = \mathbf{B}\mathbf{D}$. The entries of \mathbf{C} are given by

$$c_{ij} = \sum_{\ell=1}^n b_{i\ell} d_{\ell j},$$

where the number $b_{i\ell}$ gives the number of distinct (i, ℓ) -walks of length k and $d_{\ell j} = 0$ or $d_{\ell j} = 1$. If a particular edge $V_\ell V_j$ is not in the mesh (graph) then $d_{\ell j} = 0$ and $b_{i\ell} d_{\ell j} = 0$. Thus, there is no (i, j) -walk of length $k + 1$ with the last-but-one node ℓ . On the other hand, if $d_{\ell j} = 1$ then we can prolong each (i, ℓ) -walk to an (i, j) -walk. Thus c_{ij} gives the number of all distinct (i, j) -walks of length $k + 1$. \square

The following corollary follows easily from Lemma 2.3.

Corollary 2.1 *Let \mathbf{D} denote the adjacency matrix of a given mesh and $\mathbf{e} = [e_i]$, $e_i = 1$, $i = 1, 2, \dots, n$. Then the number $w(i, k)$ of distinct walks of length k starting at node i is given by*

$$w(i, k) = [\mathbf{D}^k \mathbf{e}]_i.$$

If the mesh is approximately regular, we expect that more walks of length k originate from the nodes that are near a center of the mesh than from the vertices that are far from it. It simply follows that the node with the index i which satisfies $w(i, k) \geq w(j, k)$, $j = 1, 2, \dots, n$, for sufficiently large k is in a sense near to the center of the mesh.

References

1. Demmel, J.W.: Applied Numerical Linear Algebra. SIAM, Philadelphia (1997)
2. Laub, A.J.: Matrix Analysis for Scientists and Engineers. SIAM, Philadelphia (2005)
3. Golub, G.H., Van Loan, C.F.: Matrix Computations, 2nd edn. Johns Hopkins University Press, Baltimore (1989)
4. Dostál, Z., Kozubek, T., Markopoulos, A., Menšík, M.: Cholesky factorization of a positive semidefinite matrix with known kernel. Appl. Math. Comput. **217**, 6067–6077 (2011)
5. Diestel, R.: Graph Theory. Springer, Heidelberg (2005)

Scalable Algorithms for Contact Problems

Dostál, Z.; Kozubek, T.; Sadowska, M.; Vondrák, V.

2016, XIX, 340 p. 80 illus., 23 illus. in color., Hardcover

ISBN: 978-1-4939-6832-9