

Chapter 1

Linear and Matrix Algebra

This chapter summarizes some important results of linear and matrix algebra that are instrumental in deriving many statistical results in subsequent chapters. Our emphasis is given to special matrices and their properties. Although the coverage of these mathematical topics is rather brief, it is self-contained. Readers may also consult other linear and matrix algebra textbooks for more detailed discussions; see e.g., Anton (1981), Basilevsky (1983), Graybill (1969), and Noble and Daniel (1977).

1.1 Basic Notations

A *matrix* is an array of numbers. In what follows, a matrix is denoted by an upper-case alphabet in boldface (e.g., \mathbf{A}), and its (i, j) th element (the element at the i th row and j th column) is denoted by the corresponding lower-case alphabet with subscripts ij (e.g., a_{ij}). Specifically, a $m \times n$ matrix \mathbf{A} contains m rows and n columns and can be expressed as

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{bmatrix}.$$

An $n \times 1$ ($1 \times n$) matrix is an n -dimensional column (row) *vector*. Every vector will be denoted by a lower-case alphabet in boldface (e.g., \mathbf{z}), and its i th element is denoted by the corresponding lower-case alphabet with subscript i (e.g., z_i). An 1×1 matrix is just a *scalar*. For a matrix \mathbf{A} , its i th column is denoted as \mathbf{a}_i .

A matrix is *square* if its number of rows equals the number of columns. A matrix is said to be *diagonal* if its off-diagonal elements (i.e., a_{ij} , $i \neq j$) are all zeros and at least one of its diagonal elements is non-zero, i.e., $a_{ii} \neq 0$ for some $i = 1, \dots, n$. A diagonal matrix

whose diagonal elements are all ones is an *identity* matrix, denoted as \mathbf{I} ; we also write the $n \times n$ identity matrix as \mathbf{I}_n . A matrix \mathbf{A} is said to be *lower (upper) triangular* if $a_{ij} = 0$ for $i < (>) j$. We let $\mathbf{0}$ denote the matrix whose elements are all zeros.

For a vector-valued function $\mathbf{f}: \mathbb{R}^m \rightarrow \mathbb{R}^n$, $\nabla_{\boldsymbol{\theta}} \mathbf{f}(\boldsymbol{\theta})$ is the $m \times n$ matrix of the first-order derivatives of \mathbf{f} with respect to the elements of $\boldsymbol{\theta}$:

$$\nabla_{\boldsymbol{\theta}} \mathbf{f}(\boldsymbol{\theta}) = \begin{bmatrix} \frac{\partial f_1(\boldsymbol{\theta})}{\partial \theta_1} & \frac{\partial f_2(\boldsymbol{\theta})}{\partial \theta_1} & \cdots & \frac{\partial f_n(\boldsymbol{\theta})}{\partial \theta_1} \\ \frac{\partial f_1(\boldsymbol{\theta})}{\partial \theta_2} & \frac{\partial f_2(\boldsymbol{\theta})}{\partial \theta_2} & \cdots & \frac{\partial f_n(\boldsymbol{\theta})}{\partial \theta_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial f_1(\boldsymbol{\theta})}{\partial \theta_m} & \frac{\partial f_2(\boldsymbol{\theta})}{\partial \theta_m} & \cdots & \frac{\partial f_n(\boldsymbol{\theta})}{\partial \theta_m} \end{bmatrix}.$$

When $n = 1$, $\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta})$ is the (column) *gradient* vector of $f(\boldsymbol{\theta})$. The $m \times m$ *Hessian* matrix of the second-order derivatives of the real-valued function $f(\boldsymbol{\theta})$ is

$$\nabla_{\boldsymbol{\theta}}^2 f(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}}(\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta})) = \begin{bmatrix} \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_1 \partial \theta_1} & \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_1 \partial \theta_2} & \cdots & \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_1 \partial \theta_m} \\ \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_2 \partial \theta_1} & \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_2 \partial \theta_2} & \cdots & \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_2 \partial \theta_m} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_m \partial \theta_1} & \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_m \partial \theta_2} & \cdots & \frac{\partial^2 f(\boldsymbol{\theta})}{\partial \theta_m \partial \theta_m} \end{bmatrix}.$$

1.2 Matrix Operations

Two matrices are said to be of the same size if they have the same number of rows and same number of columns. Matrix equality is defined for two matrices of the same size. Given two $m \times n$ matrices \mathbf{A} and \mathbf{B} , $\mathbf{A} = \mathbf{B}$ if $a_{ij} = b_{ij}$ for every i, j . The *transpose* of an $m \times n$ matrix \mathbf{A} , denoted as \mathbf{A}' , is the $n \times m$ matrix whose (i, j) th element is the (j, i) th element of \mathbf{A} . The transpose of a column vector is a row vector; the transpose of a scalar is just the scalar itself. A matrix \mathbf{A} is said to be *symmetric* if $\mathbf{A} = \mathbf{A}'$, i.e., $a_{ij} = a_{ji}$ for all i, j . Clearly, a diagonal matrix is symmetric, but a triangular matrix is not.

Matrix addition is also defined for two matrices of the same size. Given two $m \times n$ matrices \mathbf{A} and \mathbf{B} , their sum, $\mathbf{C} = \mathbf{A} + \mathbf{B}$, is the $m \times n$ matrix with the (i, j) th element $c_{ij} = a_{ij} + b_{ij}$. Note that matrix addition, if defined, is commutative:

$$\mathbf{A} + \mathbf{B} = \mathbf{B} + \mathbf{A},$$

and associative:

$$\mathbf{A} + (\mathbf{B} + \mathbf{C}) = (\mathbf{A} + \mathbf{B}) + \mathbf{C}.$$

Also, $\mathbf{A} + \mathbf{0} = \mathbf{A}$.

The scalar multiplication of the scalar c and matrix \mathbf{A} is the matrix $c\mathbf{A}$ whose (i, j) th element is ca_{ij} . Clearly, $c\mathbf{A} = \mathbf{A}c$, and $-\mathbf{A} = -1 \times \mathbf{A}$. Thus, $\mathbf{A} + (-\mathbf{A}) = \mathbf{A} - \mathbf{A} = \mathbf{0}$. Given two matrices \mathbf{A} and \mathbf{B} , the matrix multiplication \mathbf{AB} is defined only when the number of columns of \mathbf{A} is the same as the number of rows of \mathbf{B} . Specifically, when \mathbf{A} is $m \times n$ and \mathbf{B} is $n \times p$, their product, $\mathbf{C} = \mathbf{AB}$, is the $m \times p$ matrix whose (i, j) th element is

$$c_{ij} = \sum_{k=1}^n a_{ik}b_{kj}.$$

Matrix multiplication is not commutative, i.e., $\mathbf{AB} \neq \mathbf{BA}$; in fact, when \mathbf{AB} is defined, \mathbf{BA} need not be defined. On the other hand, matrix multiplication is associative:

$$\mathbf{A}(\mathbf{BC}) = (\mathbf{AB})\mathbf{C},$$

and distributive with respect to matrix addition:

$$\mathbf{A}(\mathbf{B} + \mathbf{C}) = \mathbf{AB} + \mathbf{AC}.$$

It is easy to verify that $(\mathbf{AB})' = \mathbf{B}'\mathbf{A}'$. For an $m \times n$ matrix \mathbf{A} , $\mathbf{I}_m\mathbf{A} = \mathbf{A}\mathbf{I}_n = \mathbf{A}$.

The *inner product* of two d -dimensional vectors \mathbf{y} and \mathbf{z} is the scalar

$$\mathbf{y}'\mathbf{z} = \sum_{i=1}^d y_i z_i.$$

If \mathbf{y} is m -dimensional and \mathbf{z} is n -dimensional, their *outer product* is the matrix $\mathbf{y}\mathbf{z}'$ whose (i, j) th element is $y_i z_j$. In particular,

$$\mathbf{z}'\mathbf{z} = \sum_{i=1}^d z_i^2,$$

which is non-negative and induces the standard *Euclidean norm* of \mathbf{z} as $\|\mathbf{z}\| = (\mathbf{z}'\mathbf{z})^{1/2}$. The vector with Euclidean norm zero must be a zero vector; the vector with Euclidean norm one is referred to as a *unit vector*. For example,

$$(1 \ 0 \ 0), \quad \left(0 \ \frac{1}{2} \ \frac{\sqrt{3}}{2}\right), \quad \left(\frac{1}{\sqrt{2}} \ \frac{1}{\sqrt{3}} \ \frac{1}{\sqrt{6}}\right),$$

are all unit vectors. A vector whose i th element is one and the remaining elements are all zero is called the i th Cartesian unit vector.

Let θ denote the angle between \mathbf{y} and \mathbf{z} . By the law of cosine,

$$\|\mathbf{y} - \mathbf{z}\|^2 = \|\mathbf{y}\|^2 + \|\mathbf{z}\|^2 - 2\|\mathbf{y}\|\|\mathbf{z}\|\cos\theta,$$

where the left-hand side is $\|\mathbf{y}\|^2 + \|\mathbf{z}\|^2 - 2\mathbf{y}'\mathbf{z}$. Thus, the inner product of \mathbf{y} and \mathbf{z} can be expressed as

$$\mathbf{y}'\mathbf{z} = \|\mathbf{y}\|\|\mathbf{z}\|\cos\theta.$$

When $\theta = \pi/2$, $\cos\theta = 0$ so that $\mathbf{y}'\mathbf{z} = 0$. In this case, we say that \mathbf{y} and \mathbf{z} are *orthogonal* to each other. A square matrix \mathbf{A} is said to be *orthogonal* if $\mathbf{A}'\mathbf{A} = \mathbf{A}\mathbf{A}' = \mathbf{I}$. Hence, each column (row) vector of an orthogonal matrix is a unit vector and orthogonal to all remaining column (row) vectors. When $\mathbf{y} = c\mathbf{z}$ for some $c \neq 0$, $\theta = 0$ or π , and \mathbf{y} and \mathbf{z} are said to be *linearly dependent*.

As $-1 \leq \cos\theta \leq 1$, we immediately obtain the so-called *Cauchy-Schwarz inequality*.

Lemma 1.1 (Cauchy-Schwarz) For two d -dimensional vectors \mathbf{y} and \mathbf{z} ,

$$|\mathbf{y}'\mathbf{z}| \leq \|\mathbf{y}\|\|\mathbf{z}\|,$$

where the equality holds when \mathbf{y} and \mathbf{z} are linearly dependent.

It follows from the Cauchy-Schwarz inequality that

$$\begin{aligned} \|\mathbf{y} + \mathbf{z}\|^2 &= \|\mathbf{y}\|^2 + \|\mathbf{z}\|^2 + 2\mathbf{y}'\mathbf{z} \\ &\leq \|\mathbf{y}\|^2 + \|\mathbf{z}\|^2 + 2\|\mathbf{y}\|\|\mathbf{z}\| \\ &= (\|\mathbf{y}\| + \|\mathbf{z}\|)^2. \end{aligned}$$

This leads to the following *triangle inequality*.

Lemma 1.2 For two d -dimensional vectors \mathbf{y} and \mathbf{z} ,

$$\|\mathbf{y} + \mathbf{z}\| \leq \|\mathbf{y}\| + \|\mathbf{z}\|,$$

where the equality holds when $\mathbf{y} = c\mathbf{z}$ for some $c > 0$.

When \mathbf{y} and \mathbf{z} are orthogonal,

$$\|\mathbf{y} + \mathbf{z}\|^2 = \|\mathbf{y}\|^2 + \|\mathbf{z}\|^2,$$

which is the celebrated *Pythagoras theorem*.

A special type of matrix multiplication, known as the *Kronecker product*, is defined for matrices without size restrictions. Specifically, the Kronecker product of two matrices \mathbf{A} ($m \times n$) and \mathbf{B} ($p \times q$) is the $mp \times nq$ matrix:

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \cdots & a_{1n}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \cdots & a_{2n}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1}\mathbf{B} & a_{m2}\mathbf{B} & \cdots & a_{mn}\mathbf{B} \end{bmatrix}.$$

The Kronecker product is not commutative:

$$\mathbf{A} \otimes \mathbf{B} \neq \mathbf{B} \otimes \mathbf{A},$$

but it is associative:

$$(\mathbf{A} \otimes \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes (\mathbf{B} \otimes \mathbf{C}).$$

It also obeys the distributive law:

$$\mathbf{A} \otimes (\mathbf{B} + \mathbf{C}) = \mathbf{A} \otimes \mathbf{B} + \mathbf{A} \otimes \mathbf{C}.$$

It can be verified that

$$(\mathbf{A} \otimes \mathbf{B})' = \mathbf{A}' \otimes \mathbf{B}'.$$

Consider now differentiation involving vectors and matrices. Let \mathbf{a} and $\boldsymbol{\theta}$ be two d -dimensional vectors. We have

$$\nabla_{\boldsymbol{\theta}} (\mathbf{a}' \boldsymbol{\theta}) = \mathbf{a}.$$

For a symmetric matrix \mathbf{A} ,

$$\nabla_{\boldsymbol{\theta}} (\boldsymbol{\theta}' \mathbf{A} \boldsymbol{\theta}) = 2\mathbf{A}\boldsymbol{\theta}, \quad \nabla_{\boldsymbol{\theta}}^2 (\boldsymbol{\theta}' \mathbf{A} \boldsymbol{\theta}) = 2\mathbf{A}.$$

1.3 Matrix Determinant and Trace

Given a square matrix \mathbf{A} , let \mathbf{A}_{ij} denote the sub-matrix obtained from \mathbf{A} by deleting its i th row and j th column. The *determinant* of \mathbf{A} is

$$\det(\mathbf{A}) = \sum_{i=1}^m (-1)^{i+j} a_{ij} \det(\mathbf{A}_{ij}),$$

for any $j = 1, \dots, n$, where $(-1)^{i+j} \det(\mathbf{A}_{ij})$ is called the *cofactor* of a_{ij} . This definition is based on the cofactor expansion along the j th column. Equivalently, the determinant can also be defined using the cofactor expansion along the i th row:

$$\det(\mathbf{A}) = \sum_{j=1}^n (-1)^{i+j} a_{ij} \det(\mathbf{A}_{ij}),$$

for any $i = 1, \dots, m$. The determinant of a scalar is the scalar itself; the determinant of a 2×2 matrix \mathbf{A} is simply $a_{11}a_{22} - a_{12}a_{21}$. A square matrix with non-zero determinant is said to be *nonsingular*; otherwise, it is *singular*.

Clearly, $\det(\mathbf{A}) = \det(\mathbf{A}')$. From the definition of determinant, it is straightforward to see that for a scalar c and an $n \times n$ matrix \mathbf{A} ,

$$\det(c\mathbf{A}) = c^n \det(\mathbf{A}),$$

and that for a square matrix with a column (or row) of zeros, its determinant must be zero. Also, the determinant of a diagonal or triangular matrix is simply the product of all the diagonal elements. It can also be shown that the determinant of the product of two square matrices of the same size is the product of their determinants:

$$\det(\mathbf{AB}) = \det(\mathbf{A}) \det(\mathbf{B}) = \det(\mathbf{BA}).$$

Also, for an $m \times m$ matrix \mathbf{A} and a $p \times p$ matrix \mathbf{B} ,

$$\det(\mathbf{A} \otimes \mathbf{B}) = \det(\mathbf{A})^m \det(\mathbf{B})^p.$$

If \mathbf{A} is an orthogonal matrix, we know $\mathbf{AA}' = \mathbf{I}$ so that

$$\det(\mathbf{I}) = \det(\mathbf{AA}') = [\det(\mathbf{A})]^2.$$

As the determinant of the identity matrix is one, the determinant of an orthogonal matrix must be either 1 or -1 .

The *trace* of a square matrix is the sum of its diagonal elements; i.e., $\text{trace}(\mathbf{A}) = \sum_i a_{ii}$. For example, $\text{trace}(\mathbf{I}_n) = n$. Clearly, $\text{trace}(\mathbf{A}) = \text{trace}(\mathbf{A}')$. The trace function has the linear property:

$$\text{trace}(c\mathbf{A} + d\mathbf{B}) = c \text{trace}(\mathbf{A}) + d \text{trace}(\mathbf{B}),$$

where c and d are scalars. It can also be shown that

$$\text{trace}(\mathbf{AB}) = \text{trace}(\mathbf{BA}),$$

provided that both \mathbf{AB} and \mathbf{BA} are defined. For two square matrices \mathbf{A} and \mathbf{B} ,

$$\text{trace}(\mathbf{A} \otimes \mathbf{B}) = \text{trace}(\mathbf{A}) \text{trace}(\mathbf{B}).$$

1.4 Matrix Inverse

A nonsingular matrix \mathbf{A} possesses a unique *inverse* \mathbf{A}^{-1} in the sense that $\mathbf{AA}^{-1} = \mathbf{A}^{-1}\mathbf{A} = \mathbf{I}$. A singular matrix cannot be inverted, however. Thus, saying that a matrix is invertible is equivalent to saying that it is nonsingular.

Given an invertible matrix \mathbf{A} , its inverse can be calculated as

$$\mathbf{A}^{-1} = \frac{1}{\det(\mathbf{A})} \mathbf{F}',$$

where \mathbf{F} is the matrix of cofactors, i.e., the (i, j) th element of \mathbf{F} is the cofactor $(-1)^{i+j} \det(\mathbf{A}_{ij})$. The matrix \mathbf{F}' is known as the *adjoint* of \mathbf{A} . For example, when \mathbf{A} is 2×2 ,

$$\mathbf{A}^{-1} = \frac{1}{a_{11}a_{22} - a_{12}a_{21}} \begin{bmatrix} a_{22} & -a_{12} \\ -a_{21} & a_{11} \end{bmatrix}.$$

Matrix inversion and transposition can be interchanged, i.e., $(\mathbf{A}')^{-1} = (\mathbf{A}^{-1})'$. For two nonsingular matrices \mathbf{A} and \mathbf{B} of the same size, we have $\mathbf{A}\mathbf{B}\mathbf{B}^{-1}\mathbf{A}^{-1} = \mathbf{I}$, so that

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

Some special matrices can be easily inverted. For example, for a diagonal matrix \mathbf{A} , \mathbf{A}^{-1} is also diagonal with the diagonal elements a_{ii}^{-1} ; for an orthogonal matrix \mathbf{A} , $\mathbf{A}^{-1} = \mathbf{A}'$.

A formula for computing the inverse of a partitioned matrix is

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}\mathbf{F}^{-1}\mathbf{C}\mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{B}\mathbf{F}^{-1} \\ -\mathbf{F}^{-1}\mathbf{C}\mathbf{A}^{-1} & \mathbf{F}^{-1} \end{bmatrix},$$

where $\mathbf{F} = \mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$, or equivalently,

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{G}^{-1} & -\mathbf{G}^{-1}\mathbf{B}\mathbf{D}^{-1} \\ -\mathbf{D}^{-1}\mathbf{C}\mathbf{G}^{-1} & \mathbf{D}^{-1} + \mathbf{D}^{-1}\mathbf{C}\mathbf{G}^{-1}\mathbf{B}\mathbf{D}^{-1} \end{bmatrix},$$

where $\mathbf{G} = \mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C}$, provided that the matrix inverses in the expressions above are well defined. In particular, if this matrix is *block diagonal* so that the off-diagonal blocks are zero matrices, we have

$$\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{D} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^{-1} \end{bmatrix},$$

provided that \mathbf{A} and \mathbf{D} are invertible.

1.5 Matrix Rank

The vectors $\mathbf{z}_1, \dots, \mathbf{z}_n$ are said to be *linearly independent* if the only solution to

$$c_1\mathbf{z}_1 + c_2\mathbf{z}_2 + \dots + c_n\mathbf{z}_n = \mathbf{0}$$

is the trivial solution: $c_1 = \dots = c_n = 0$; otherwise, they are *linearly dependent*. When two (three) vectors are linearly dependent, they are on the same line (plane).

The *column (row) rank* of a matrix \mathbf{A} is the maximum number of linearly independent column (row) vectors of \mathbf{A} . When the column (row) rank equals the number of column (row) vectors, this matrix is said to be of full column (row) rank. The space *spanned* by the vectors $\mathbf{z}_1, \dots, \mathbf{z}_n$ is the collection of all linear combinations of these vectors, denoted as $\text{span}(\mathbf{z}_1, \dots, \mathbf{z}_n)$. The space spanned by the column vectors of \mathbf{A} is therefore $\text{span}(\mathbf{A})$, which is also known as the *column space* of \mathbf{A} . A vector \mathbf{z} is in $\text{span}(\mathbf{A})$ if it can be expressed as $\mathbf{A}\mathbf{c}$ for some vector $\mathbf{c} \neq \mathbf{0}$. Similarly, the space spanned by the row vectors of \mathbf{A} is $\text{span}(\mathbf{A}')$ and known as the *row space* of \mathbf{A} . The column (row) rank of \mathbf{A} is the *dimension* of the column (row) space of \mathbf{A} .

Given an $n \times k$ matrix \mathbf{A} with $k \leq n$, suppose that \mathbf{A} has row rank $r \leq n$ and column rank $c \leq k$. Without loss of generality, assume that the first r row vectors are linear independent. Hence, each row vector \mathbf{a}_i can be expressed as

$$\mathbf{a}_i = q_{i1}\mathbf{a}_1 + q_{i2}\mathbf{a}_2 + \cdots + q_{ir}\mathbf{a}_r, \quad i = 1, \dots, n,$$

with the j^{th} element

$$a_{ij} = q_{i1}a_{1j} + q_{i2}a_{2j} + \cdots + q_{ir}a_{rj}, \quad i = 1, \dots, n, \quad j = 1, \dots, k.$$

Fixing j , we immediately see that every column vector of \mathbf{A} can be written as a linear combination of the vectors $\mathbf{q}_1, \dots, \mathbf{q}_r$. As such, the column rank of \mathbf{A} must be less than or equal to r . Similarly, the column rank of \mathbf{A}' , which is also the row rank of \mathbf{A} , must be less than or equal to c . This proves the following result.

Lemma 1.3 *The column rank and row rank of a matrix are equal.*

By Lemma 1.3, we can then define the *rank* of \mathbf{A} as the maximum number of linearly independent column (or row) vectors of \mathbf{A} , denoted as $\text{rank}(\mathbf{A})$. Clearly, $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}')$. An $n \times n$ matrix \mathbf{A} is said to be of *full rank* if $\text{rank}(\mathbf{A}) = n$.

For an $n \times k$ matrix \mathbf{A} , its *left inverse* is a $k \times n$ matrix \mathbf{A}_L^{-1} such that $\mathbf{A}_L^{-1}\mathbf{A} = \mathbf{I}_k$. Similarly, a *right inverse* of \mathbf{A} is a $k \times n$ matrix \mathbf{A}_R^{-1} such that $\mathbf{A}\mathbf{A}_R^{-1} = \mathbf{I}_n$. The left and right inverses are not unique, however. It can be shown that a matrix possesses a left (right) inverse if and only if it has full column (row) rank. Thus, for a square matrix with full rank, it has both inverses, which are just the unique matrix inverse. Thus, a nonsingular (invertible) matrix must be of full rank and vice versa.

It can be shown that for two $n \times k$ matrices \mathbf{A} and \mathbf{B} ,

$$\text{rank}(\mathbf{A} + \mathbf{B}) \leq \text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B}).$$

If \mathbf{A} is $n \times k$ and \mathbf{B} is $k \times m$,

$$\text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B}) - k \leq \text{rank}(\mathbf{AB}) \leq \min[\text{rank}(\mathbf{A}), \text{rank}(\mathbf{B})].$$

For the Kronecker product, we have

$$\text{rank}(\mathbf{A} \otimes \mathbf{B}) = \text{rank}(\mathbf{A}) \text{rank}(\mathbf{B}).$$

If \mathbf{A} is a nonsingular matrix, we have from the inequality above that

$$\text{rank}(\mathbf{AB}) \leq \text{rank}(\mathbf{B}) = \text{rank}(\mathbf{A}^{-1}\mathbf{AB}) \leq \text{rank}(\mathbf{AB});$$

i.e., $\text{rank}(\mathbf{AB}) = \text{rank}(\mathbf{B})$. This also implies that for a nonsingular matrix \mathbf{C} ,

$$\text{rank}(\mathbf{BC}) = \text{rank}(\mathbf{C}'\mathbf{B}') = \text{rank}(\mathbf{B}') = \text{rank}(\mathbf{B}).$$

Thus, the rank of a matrix is preserved under nonsingular transformations.

Lemma 1.4 *Let \mathbf{A} ($n \times n$) and \mathbf{C} ($k \times k$) be nonsingular matrices. Then for any $n \times k$ matrix \mathbf{B} ,*

$$\text{rank}(\mathbf{B}) = \text{rank}(\mathbf{AB}) = \text{rank}(\mathbf{BC}).$$

1.6 Eigenvalue and Eigenvector

Given a square matrix \mathbf{A} , if $\mathbf{Ac} = \lambda\mathbf{c}$ for some scalar λ and non-zero vector \mathbf{c} , then \mathbf{c} is an *eigenvector* of \mathbf{A} corresponding to the *eigenvalue* λ . The system $(\mathbf{A} - \lambda\mathbf{I})\mathbf{c} = \mathbf{0}$ has a non-trivial solution if and only if

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

This is known as the *characteristic equation* of \mathbf{A} , from which we can solve for the eigenvalues of \mathbf{A} . Hence, eigenvalues (eigenvectors) are also referred to as *characteristic roots* (*characteristic vectors*). Note that the eigenvalues and eigenvectors of a real-valued matrix need not be real-valued.

When \mathbf{A} is $n \times n$, the characteristic equation is an n th-order polynomial in λ and has at most n distinct solutions. These solutions (eigenvalues) are usually complex-valued. If some eigenvalues take the same value, there may exist several eigenvectors corresponding to the same eigenvalue. Given an eigenvalue λ , let $\mathbf{c}_1, \dots, \mathbf{c}_k$ be associated eigenvectors. Then,

$$\mathbf{A}(a_1\mathbf{c}_1 + a_2\mathbf{c}_2 + \dots + a_k\mathbf{c}_k) = \lambda(a_1\mathbf{c}_1 + a_2\mathbf{c}_2 + \dots + a_k\mathbf{c}_k),$$

so that any linear combination of these eigenvectors is again an eigenvector corresponding to λ . That is, these eigenvectors are closed under scalar multiplication and vector addition and form the *eigenspace* corresponding to λ . As such, for a common eigenvalue, we are mainly concerned with those eigenvectors that are linearly independent.

If \mathbf{A} ($n \times n$) possesses n distinct eigenvalues, each eigenvalue must correspond to one eigenvector, unique up to scalar multiplications. It is therefore typical to normalize eigenvectors such that they have Euclidean norm one. It can also be shown that if the eigenvalues of a matrix are all distinct, their associated eigenvectors must be linearly independent. Let \mathbf{C} denote the matrix of these eigenvectors and $\mathbf{\Lambda}$ denote the diagonal matrix with diagonal elements being the eigenvalues of \mathbf{A} . We can write $\mathbf{AC} = \mathbf{C}\mathbf{\Lambda}$. As \mathbf{C} is nonsingular, we have

$$\mathbf{C}^{-1}\mathbf{AC} = \mathbf{\Lambda}, \quad \text{or} \quad \mathbf{A} = \mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1}.$$

In this case, \mathbf{A} is said to be *similar* to $\mathbf{\Lambda}$.

When \mathbf{A} has n distinct eigenvalues $\lambda_1, \dots, \lambda_n$, it is readily seen that

$$\det(\mathbf{A}) = \det(\mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1}) = \det(\mathbf{\Lambda}) \det(\mathbf{C}) \det(\mathbf{C}^{-1}) = \det(\mathbf{\Lambda}),$$

and

$$\text{trace}(\mathbf{A}) = \text{trace}(\mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1}) = \text{trace}(\mathbf{C}^{-1}\mathbf{C}\mathbf{\Lambda}) = \text{trace}(\mathbf{\Lambda}).$$

This yields the following result.

Lemma 1.5 *Let \mathbf{A} be an $n \times n$ matrix with distinct eigenvalues $\lambda_1, \dots, \lambda_n$. Then*

$$\det(\mathbf{A}) = \det(\mathbf{\Lambda}) = \prod_{i=1}^n \lambda_i,$$

$$\text{trace}(\mathbf{A}) = \text{trace}(\mathbf{\Lambda}) = \sum_{i=1}^n \lambda_i.$$

When $\mathbf{A} = \mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1}$, we have $\mathbf{A}^{-1} = \mathbf{C}\mathbf{\Lambda}^{-1}\mathbf{C}^{-1}$. This shows that the eigenvectors of \mathbf{A}^{-1} are the same as those of \mathbf{A} , and the corresponding eigenvalues are simply the reciprocals of the eigenvalues of \mathbf{A} . Similarly,

$$\mathbf{A}^2 = (\mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1})(\mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1}) = \mathbf{C}\mathbf{\Lambda}^2\mathbf{C}^{-1},$$

so that the eigenvectors of \mathbf{A}^2 are the same as those of \mathbf{A} , and the corresponding eigenvalues are the squares of the eigenvalues of \mathbf{A} . This result generalizes immediately to \mathbf{A}^k .

1.7 Symmetric Matrix

More can be said about symmetric matrices. Let \mathbf{c}_1 and \mathbf{c}_2 be two eigenvectors of \mathbf{A} corresponding to the distinct eigenvalues λ_1 and λ_2 , respectively. If \mathbf{A} is symmetric, then

$$\mathbf{c}'_2 \mathbf{A} \mathbf{c}_1 = \lambda_1 \mathbf{c}'_2 \mathbf{c}_1 = \lambda_2 \mathbf{c}'_2 \mathbf{c}_1.$$

As $\lambda_1 \neq \lambda_2$, it must be true that $\mathbf{c}'_2 \mathbf{c}_1 = 0$, so that they are orthogonal. Given linearly independent eigenvectors that correspond to a common eigenvalue, they can also be orthogonalized. Thus, a symmetric matrix is *orthogonally diagonalizable*, in the sense that

$$\mathbf{C}' \mathbf{A} \mathbf{C} = \mathbf{\Lambda}, \quad \text{or} \quad \mathbf{A} = \mathbf{C} \mathbf{\Lambda} \mathbf{C}',$$

where $\mathbf{\Lambda}$ is again the diagonal matrix of the eigenvalues of \mathbf{A} , and \mathbf{C} is the orthogonal matrix of associated eigenvectors.

As nonsingular transformations preserve rank (Lemma 1.4), so do orthogonal transformations. We thus have the result below.

Lemma 1.6 *For a symmetric matrix \mathbf{A} , $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{\Lambda})$, the number of non-zero eigenvalues of \mathbf{A} .*

Moreover, when \mathbf{A} is diagonalizable, the assertions of Lemma 1.5 remain valid, whether or not the eigenvalues of \mathbf{A} are distinct.

Lemma 1.7 *Let \mathbf{A} be an $n \times n$ symmetric matrix. Then,*

$$\det(\mathbf{A}) = \det(\mathbf{\Lambda}) = \prod_{i=1}^n \lambda_i,$$

$$\text{trace}(\mathbf{A}) = \text{trace}(\mathbf{\Lambda}) = \sum_{i=1}^n \lambda_i.$$

By Lemma 1.7, a symmetric matrix is nonsingular if its eigenvalues are all non-zero.

A symmetric matrix \mathbf{A} is said to be *positive definite* if $\mathbf{b}' \mathbf{A} \mathbf{b} > 0$ for all vectors $\mathbf{b} \neq \mathbf{0}$; \mathbf{A} is said to be *positive semi-definite* if $\mathbf{b}' \mathbf{A} \mathbf{b} \geq 0$ for all $\mathbf{b} \neq \mathbf{0}$. A positive definite matrix thus must be nonsingular, but a positive semi-definite matrix may be singular. Suppose that \mathbf{A} is a symmetric matrix orthogonally diagonalized as $\mathbf{C}' \mathbf{A} \mathbf{C} = \mathbf{\Lambda}$. If \mathbf{A} is also positive semi-definite, then for any $\mathbf{b} \neq \mathbf{0}$,

$$\mathbf{b}' \mathbf{A} \mathbf{b} = \mathbf{b}' (\mathbf{C}' \mathbf{A} \mathbf{C}) \mathbf{b} = \tilde{\mathbf{b}}' \mathbf{A} \tilde{\mathbf{b}} \geq 0,$$

where $\tilde{\mathbf{b}} = \mathbf{C} \mathbf{b}$. This shows that $\mathbf{\Lambda}$ is also positive semi-definite, and all the diagonal elements of $\mathbf{\Lambda}$ must be non-negative. It can be seen that the converse also holds.

Lemma 1.8 *A symmetric matrix is positive definite (positive semi-definite) if, and only if, its eigenvalues are all positive (non-negative).*

For a symmetric and positive definite matrix \mathbf{A} , $\mathbf{A}^{-1/2}$ is such that $\mathbf{A}^{-1/2'}\mathbf{A}^{-1/2} = \mathbf{A}^{-1}$. In particular, by orthogonal diagonalization,

$$\mathbf{A}^{-1} = \mathbf{C}\mathbf{\Lambda}^{-1}\mathbf{C}' = (\mathbf{C}\mathbf{\Lambda}^{-1/2}\mathbf{C}')(\mathbf{C}\mathbf{\Lambda}^{-1/2}\mathbf{C}'),$$

so that we may choose $\mathbf{A}^{-1/2} = \mathbf{C}\mathbf{\Lambda}^{-1/2}\mathbf{C}'$. The inverse of $\mathbf{A}^{-1/2}$ is $\mathbf{A}^{1/2} = \mathbf{C}\mathbf{\Lambda}^{1/2}\mathbf{C}'$. It follows that $\mathbf{A}^{1/2}\mathbf{A}^{1/2'} = \mathbf{A}$, and $\mathbf{A}^{-1/2}\mathbf{A}\mathbf{A}^{-1/2'} = \mathbf{I}$. Note that $\mathbf{\Lambda}^{-1/2}\mathbf{C}'$ is also a legitimate choice of $\mathbf{A}^{-1/2}$, yet it is not symmetric.

Finally, we know that for two positive real numbers a and b , $a \geq b$ implies $b^{-1} \geq a^{-1}$. This result can be generalized to compare two positive definite matrices, as stated below without proof.

Lemma 1.9 *Given two symmetric and positive definite matrices \mathbf{A} and \mathbf{B} , if $\mathbf{A} - \mathbf{B}$ is positive semi-definite, then so is $\mathbf{B}^{-1} - \mathbf{A}^{-1}$.*

1.8 Orthogonal Projection

A matrix \mathbf{A} is said to be *idempotent* if $\mathbf{A}^2 = \mathbf{A}$. Given a vector \mathbf{y} in the Euclidean space V , a *projection* of \mathbf{y} onto a subspace S of V is a linear transformation of \mathbf{y} to S . The resulting projected vector can be written as $\mathbf{P}\mathbf{y}$, where \mathbf{P} is the associated transformation matrix. Given the projection $\mathbf{P}\mathbf{y}$ in S , further projection to S should have no effect on $\mathbf{P}\mathbf{y}$, i.e.,

$$\mathbf{P}(\mathbf{P}\mathbf{y}) = \mathbf{P}^2\mathbf{y} = \mathbf{P}\mathbf{y}.$$

Thus, a matrix \mathbf{P} is said to be a *projection matrix* if it is *idempotent*.

A projection of \mathbf{y} onto S is orthogonal if the projection $\mathbf{P}\mathbf{y}$ is orthogonal to the difference between \mathbf{y} and $\mathbf{P}\mathbf{y}$. That is,

$$(\mathbf{y} - \mathbf{P}\mathbf{y})'\mathbf{P}\mathbf{y} = \mathbf{y}'(\mathbf{I} - \mathbf{P})'\mathbf{P}\mathbf{y} = \mathbf{0}.$$

As \mathbf{y} is arbitrary, the equality above holds if, and only if, $(\mathbf{I} - \mathbf{P})'\mathbf{P} = \mathbf{0}$. Consequently, $\mathbf{P} = \mathbf{P}'\mathbf{P}$ and $\mathbf{P}' = \mathbf{P}'\mathbf{P}$. This shows that \mathbf{P} must be symmetric. Thus, a matrix is an *orthogonal projection matrix* if, and only if, it is symmetric and idempotent. It can be easily verified that the orthogonal projection $\mathbf{P}\mathbf{y}$ must be unique.

When \mathbf{P} is an orthogonal projection matrix, it is easily seen that $\mathbf{I} - \mathbf{P}$ is idempotent because

$$(\mathbf{I} - \mathbf{P})^2 = \mathbf{I} - 2\mathbf{P} + \mathbf{P}^2 = \mathbf{I} - \mathbf{P}.$$

As $\mathbf{I} - \mathbf{P}$ is also symmetric, it is an orthogonal projection matrix. Since $(\mathbf{I} - \mathbf{P})\mathbf{P} = \mathbf{0}$, the projections $\mathbf{P}\mathbf{y}$ and $(\mathbf{I} - \mathbf{P})\mathbf{y}$ must be orthogonal. This shows that any vector \mathbf{y} can be uniquely decomposed into two orthogonal components:

$$\mathbf{y} = \mathbf{P}\mathbf{y} + (\mathbf{I} - \mathbf{P})\mathbf{y}.$$

Define the *orthogonal complement* of a subspace $S \subseteq V$ as

$$S^\perp = \{\mathbf{v} \in V : \mathbf{v}'\mathbf{s} = 0, \text{ for all } \mathbf{s} \in S\}.$$

If \mathbf{P} is the orthogonal projection matrix that projects vectors onto $S \subseteq V$, we have $\mathbf{P}\mathbf{s} = \mathbf{s}$ for any $\mathbf{s} \in S$. It follows that $(\mathbf{I} - \mathbf{P})\mathbf{y}$ is orthogonal to \mathbf{s} and that $(\mathbf{I} - \mathbf{P})\mathbf{y}$ is the orthogonal projection of \mathbf{y} onto S^\perp .

Intuitively, the orthogonal projection $\mathbf{P}\mathbf{y}$ can be interpreted as the “best approximation” of \mathbf{y} in S , in the sense that $\mathbf{P}\mathbf{y}$ is the closest to \mathbf{y} in terms of the Euclidean norm. To see this, we observe that for any $\mathbf{s} \in S$,

$$\begin{aligned} \|\mathbf{y} - \mathbf{s}\|^2 &= \|\mathbf{y} - \mathbf{P}\mathbf{y} + \mathbf{P}\mathbf{y} - \mathbf{s}\|^2 \\ &= \|\mathbf{y} - \mathbf{P}\mathbf{y}\|^2 + \|\mathbf{P}\mathbf{y} - \mathbf{s}\|^2 + 2(\mathbf{y} - \mathbf{P}\mathbf{y})'(\mathbf{P}\mathbf{y} - \mathbf{s}) \\ &= \|\mathbf{y} - \mathbf{P}\mathbf{y}\|^2 + \|\mathbf{P}\mathbf{y} - \mathbf{s}\|^2. \end{aligned}$$

This establishes the following result.

Lemma 1.10 *Let \mathbf{y} be a vector in V and $\mathbf{P}\mathbf{y}$ its orthogonal projection onto $S \subseteq V$. Then,*

$$\|\mathbf{y} - \mathbf{P}\mathbf{y}\| \leq \|\mathbf{y} - \mathbf{s}\|,$$

for all $\mathbf{s} \in S$.

Let \mathbf{A} be a symmetric and idempotent matrix and \mathbf{C} be the orthogonal matrix that diagonalizes \mathbf{A} to $\mathbf{\Lambda}$. Then,

$$\mathbf{\Lambda} = \mathbf{C}'\mathbf{A}\mathbf{C} = \mathbf{C}'\mathbf{A}(\mathbf{C}\mathbf{C}')\mathbf{A}\mathbf{C} = \mathbf{\Lambda}^2.$$

This is possible only when the eigenvalues of \mathbf{A} are zero and one. The result below now follows from Lemmas 1.8.

Lemma 1.11 *A symmetric and idempotent matrix is positive semi-definite with the eigenvalues 0 and 1.*

Moreover, $\text{trace}(\mathbf{\Lambda})$ is the number of non-zero eigenvalues of \mathbf{A} and hence $\text{rank}(\mathbf{\Lambda})$. When \mathbf{A} is symmetric, $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{\Lambda})$ by Lemma 1.6, and $\text{trace}(\mathbf{A}) = \text{trace}(\mathbf{\Lambda})$ by Lemma 1.7. Combining these results we have:

Lemma 1.12 *For a symmetric and idempotent matrix \mathbf{A} , $\text{rank}(\mathbf{A}) = \text{trace}(\mathbf{A})$, the number of non-zero eigenvalues of \mathbf{A} .*

Given an $n \times k$ matrix \mathbf{A} , it is easy to see that $\mathbf{A}'\mathbf{A}$ and $\mathbf{A}\mathbf{A}'$ are symmetric and positive semi-definite. Let \mathbf{x} denote a vector orthogonal to the rows of $\mathbf{A}'\mathbf{A}$; i.e., $\mathbf{A}'\mathbf{A}\mathbf{x} = \mathbf{0}$. Hence $\mathbf{x}'\mathbf{A}'\mathbf{A}\mathbf{x} = 0$, so that $\mathbf{A}\mathbf{x}$ must be a zero vector. That is, \mathbf{x} is also orthogonal to the rows of \mathbf{A} . Conversely, $\mathbf{A}\mathbf{x} = \mathbf{0}$ implies $\mathbf{A}'\mathbf{A}\mathbf{x} = \mathbf{0}$. This shows that the orthogonal complement of the row space of \mathbf{A} is the same as the orthogonal complement of the row space of $\mathbf{A}'\mathbf{A}$. Hence, these two row spaces are also the same. Similarly, the column space of \mathbf{A} is the same as the column space of $\mathbf{A}\mathbf{A}'$. It follows from Lemma 1.3 that

$$\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}'\mathbf{A}) = \text{rank}(\mathbf{A}\mathbf{A}').$$

In particular, if \mathbf{A} ($n \times k$) is of full column rank $k < n$, then $\mathbf{A}'\mathbf{A}$ is $k \times k$ and hence of full rank k (nonsingular), but $\mathbf{A}\mathbf{A}'$ is $n \times n$ and hence singular. The result below is now immediate.

Lemma 1.13 *If \mathbf{A} is an $n \times k$ matrix with full column rank $k < n$, then, $\mathbf{A}'\mathbf{A}$ is symmetric and positive definite.*

Given an $n \times k$ matrix \mathbf{A} with full column rank $k < n$, $\mathbf{P} = \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'$ is clearly symmetric and idempotent and hence an orthogonal projection matrix. As

$$\text{trace}(\mathbf{P}) = \text{trace}(\mathbf{A}'\mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}) = \text{trace}(\mathbf{I}_k) = k,$$

we have from Lemmas 1.11 and 1.12 that \mathbf{P} has exactly k eigenvalues equal to 1 and that $\text{rank}(\mathbf{P}) = k$. Similarly, $\text{rank}(\mathbf{I} - \mathbf{P}) = n - k$. Moreover, any vector $\mathbf{y} \in \text{span}(\mathbf{A})$ can be written as $\mathbf{A}\mathbf{b}$ for some non-zero vector \mathbf{b} , and

$$\mathbf{P}\mathbf{y} = \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'(\mathbf{A}\mathbf{b}) = \mathbf{A}\mathbf{b} = \mathbf{y}.$$

This suggests that \mathbf{P} must project vectors onto $\text{span}(\mathbf{A})$. On the other hand, when $\mathbf{y} \in \text{span}(\mathbf{A})^\perp$, \mathbf{y} is orthogonal to the column vectors of \mathbf{A} so that $\mathbf{A}'\mathbf{y} = \mathbf{0}$. It follows that $\mathbf{P}\mathbf{y} = \mathbf{0}$ and $(\mathbf{I} - \mathbf{P})\mathbf{y} = \mathbf{y}$. Thus, $\mathbf{I} - \mathbf{P}$ must project vectors onto $\text{span}(\mathbf{A})^\perp$. These results are summarized below.

Lemma 1.14 *Let \mathbf{A} be an $n \times k$ matrix with full column rank k . Then, $\mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'$ orthogonally projects vectors onto $\text{span}(\mathbf{A})$ and has rank k ; $\mathbf{I}_n - \mathbf{A}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'$ orthogonally projects vectors onto $\text{span}(\mathbf{A})^\perp$ and has rank $n - k$.*

References

- Anton, Howard (1981). *Elementary Linear Algebra*, third edition, New York: Wiley.
- Basilevsky, Alexander (1983). *Applied Matrix Algebra in the Statistical Sciences*, New York: North-Holland.
- Graybill, Franklin A. (1969). *Introduction to Matrices with Applications in Statistics*, second edition, Belmont, CA: Wadsworth.
- Noble, Ben and James W. Daniel. *Applied Linear Algebra*, second edition, Englewood Cliffs, NJ: Prentice-Hall.

