Appendix B

Mathematical Refresher

This appendix presents mathematical concepts we use in developing our main arguments in the text of this book. This appendix can be read in the order in which it appears, or as a reference. Items are ordered so that simpler concepts appear earlier and, except where noted, each concept introduced does not depend on anything appearing after it.

B.1 Real Analysis

A *vector space* (section B.1.1) is a set over which one can define a meaningful notion of a "sum" between two elements, and "multiplication of an element by a scalar." For our applications, we impose additional structure on a vector space, and thus use *normed spaces* (section B.1.3), which require also the notion of "length," and *scalar product and Euclidean spaces* (section B.1.4), which add the notion of "projection." We introduce the notion of a vector space here only for its role as a basic building block for these other spaces, and we do not use it by itself. We will see that the structure associated with a scalar product or Euclidean space is "more restrictive" than the one associated with a normed space, in the sense that with each Euclidean space we can associate in a natural way a normed space, but the opposite is in general not true.

Another useful construction is the *metric space* (section B.1.2), which is a set over which we can define a notion of "distance" between two elements. One does not need the notion of sum between two elements in order to define the distance, and so a metric space is not necessarily a vector space. However, many vector spaces can be endowed with the structure of a metric space. In particular, with every scalar product, Euclidean, and normed spaces, we can always associate a metric space in a natural way.

B.1.1 Vector Space

Let X be a set where each element is defined as a vector, and where the following two operations are defined:

1. **Addition:** To every pair of vectors $x, y \in X$ corresponds an element $x + y \in X$ such that the commutative and associative properties hold

$$x + y = y + x \qquad \text{and} \qquad x + (y + z) = (x + y) + z \ , \ \ \forall z \in X.$$

In addition, X contains a unique vector, denoted by 0 (called the *zero vector*), such that $x + 0 = x$, $\forall x \in X$. Moreover, to each $x \in X$ corresponds a unique vector $-x$ such that $x + (-x) = 0$.

2. **Multiplication by a Scalar:** To every $a \in \mathbb{R}$ and $x \in X$ corresponds a vector ax such that

$$1x = x \quad \text{and} \quad a(bx) = (ab)x\ , \quad \forall b \in \mathbb{R}.$$

In addition, the following distributive properties hold:

$$a(x+y) = ax + ay \quad \text{and} \quad (a+b)x = ax + bx\ , \quad \forall a, b \in \mathbb{R} \quad \forall x, y \in X.$$

In this book if we refer to a set X as "the space X," we always mean "the vector space X."

Example 1 The set $\mathbb{R}^n$, $n = 1, 2, \ldots$ is a vector space with respect to the familiar operations of addition between two vectors and multiplication by a scalar. ☒

Example 2 The set $C[a, b]$ of continuous functions over the interval $[a, b]$ is a vector space. Any continuous function on $[a, b]$ is a vector belonging to this vector space. The operations of addition and scalar multiplication correspond to the usual operations of addition between two functions and multiplication of a function by a scalar. ☒

Example 3 The two-dimensional unit sphere $S^2 \equiv \{(x_1, x_2, x_3) \in \mathbb{R}^3 \mid x_1^2 + x_2^2 + x_3^2 = 1\}$ is *not* a vector space with respect to the usual operations of addition between vectors and multiplication of a vector by a scalar. For example, the sum of two vectors on the sphere is a vector which lies off of it. ☒

B.1.2 Metric Space

A metric space is a pair (X, d), where X is a set and $d : X \times X \to [0, +\infty)$, is a function, called *distance* or *metric* with the following properties:

1. **Positiveness:** $d(x, y) \geq 0$ and $d(x, y) = 0$ if and only if $x = y$.
2. **Symmetry:** $d(x, y) = d(y, x)$.
3. **Triangle Inequality:** $d(x, y) + d(y, z) \geq d(x, z)$.

A *semidistance* is a distance except that $d(x, y) = 0$ does not imply that $x = y$.

Example 1 (R^n, d), where $d(x, y) = \sqrt{(x-y)'(x-y)}$, is a metric space. This distance is known as the *Euclidean distance*. ☒

Example 2 (R, d), where $d(x, y) = \frac{|x-y|}{1+|x-y|}$, is a metric space. ☒

Example 3 $(C[0, 1], d)$, where $d(f, g) = \max_{x\in[0,1]} |f(x) - g(x)|$, $f, g \in C[0, 1]$, is a metric space. ☒

The set X of a metric space (X, d) does not have to be a vector space, as we see in the following:

Example 4 (S^n, d), where S^n is the n-dimensional sphere and $d(x, y)$ is the geodesic distance between two points (i.e., the distance measured along the shortest path), is a metric space. This metric space is not a vector space because there is no notion of a zero vector or addition on a sphere. ☒

B.1.3 Normed Space

A vector space X is called a *normed space* if it is equipped with a *norm* $\|\cdot\|$. The norm is a function $\|\cdot\| : X \to [0, +\infty)$ with the following properties:

1. $\|x\| \geq 0$ and $\|x\| = 0$ if and only if $x = 0$.
2. $\|ax\| = |a|\|x\|$ for all $a \in \mathbb{R}$.
3. $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in X$ (triangle inequality).

If condition (1) is replaced with the weaker condition that $\|x\| \geq 0$, then $\|x\|$ is called a *seminorm*. The only difference between a seminorm and a norm is that for a semi-norm it is possible that $\|x\| = 0$ without x being the zero vector.

A normed space is often denoted by the pair $(X, \|\cdot\|)$, because the same vector space X can be endowed with different norms. Every normed space $(X, \|\cdot\|)$ can be made into a metric space (X, d) by defining the distance $d(x, y) \equiv \|x - y\|$. We often refer to this distance as the distance *induced* by the norm $\|\cdot\|$.

Example 1 Denote a normed space as $(R^n, \|\cdot\|_A)$, where A is a strictly positive definite symmetric matrix (see B.2.3, page 234), and we have defined:

$$\|x\|_A \equiv \left(x'Ax\right)^{\frac{1}{2}},$$

which is known as the *Mahalanobis norm*. When $A = I$, this norm is called the *Euclidean norm* and is simply denoted by $\|x\|$. ☒

Example 2 The pair $(R^n, \|\cdot\|_p)$, where $\|x\|_p = (\sum_{i=1}^n x_i^p)^{1/p}$ and $p \geq 1$, is a normed space. This normed space is often denoted by $L_p(R^n)$, and the norm $\|\cdot\|_p$ is often referred to as the L_p norm. ☒

Example 3 The set $(C[a, b], \|\cdot\|_{L_2})$, with the norm defined as

$$\|x\|_{L_2} \equiv \left(\int_a^b dx\ f(x)^2\right)^{\frac{1}{2}},$$

is a normed space, which is usually denoted by $L_2[a, b]$. ☒

Example 4 The set $(C[a, b], \|\cdot\|_{W_2^1})$, with the seminorm defined as

$$\|x\|_{W_2^1} \equiv \left(\int_a^b dx\ \left(\frac{df(x)}{dx}\right)^2\right)^{\frac{1}{2}},$$

is a seminormed space, which is usually denoted by $W_2^1[a, b]$. In fact, $\|x\|_{W_2^1}$ is a seminorm, rather than a norm, because there are functions that are not zero (like $f(x) = a$, for all $a \in \mathbb{R}$) but whose norm is zero. ☒

B.1.4 Scalar Product Space

A vector space (or normed space) is called a *scalar product space* (or sometimes "inner product space") if, with each ordered pair of vectors x and y, we can associate a positive number (x, y), called the *scalar, or inner, product* of x and y, such that the following properties hold:

1. $(x, y) = (y, x)$, for all $x, y \in X$,
2. $(x + y, z) = (x, z) + (y, z)$, for all $x, y, z \in X$,
3. $(ax, y) = a(x, y)$ for all $x, y \in X, a \in \mathbb{R}$.
4. $(x, x) \geq 0$ and $(x, x) = 0$ only if $x = 0$.

If property (4) is replaced by $(x, x) \geq 0$, the result is a semiscalar, or semi-inner, product space. Because we can define different scalar products on the same vector space X, it is convenient to think of a scalar product space as a pair $(X, (\cdot, \cdot))$. To every scalar product space $(X, (\cdot, \cdot))$, we can associate in a natural way a normed space $(X, \|\cdot\|)$ by defining the norm $\|x\| = \sqrt{(x, x)}$. We refer to this norm as the norm *induced* by the scalar product $(\cdot, \cdot)$. Therefore, any scalar product space can be a normed space. The opposite of this proposition is not true in general, but it is true if the norm has the following property:

$$\|x + y\|^2 + \|x - y\|^2 = 2(\|x\|^2 + \|y\|^2), \quad \forall x, y \in X.$$

When this is the case then it is possible to show that one can define a meaningful scalar product by setting $(x, y) \equiv \frac{1}{4}(\|x + y\|^2 - \|x - y\|^2)$.

Example 1 Let A be any strictly positive definite $n \times n$ symmetric matrix. The set $\mathbb{R}^n$ with the inner product

$$(x, y) \equiv x'Ay, \quad \forall x, y \in \mathbb{R}^n$$

is a scalar product space. The norm induced by this scalar product is $\|x\|^2 = x'Ax$ (see "Normed Spaces," page 221). A *Euclidean space* is a scalar product space where $A = I$, and the norm induced by the scalar product in the Euclidean space is the Euclidean norm. ⊠

Example 2 Let $C[a, b]$ be the set of continuous functions on the interval $[a, b]$. We can make this set an inner product space by defining the following inner product between any two functions $f, g \in C[a, b]$:

$$(f, g) \equiv \int_a^b dx f(x)g(x). \qquad \boxtimes$$

Example 3 Let $C^1[a, b]$ be the set of continuous and differentiable functions on the interval $[a, b]$. We can make this set an inner product space by defining the following inner product between any two functions $f, g \in C[a, b]$:

$$(f, g) \equiv \int_a^b dx f(x)g(x) + \int_a^b dx \frac{df(x)}{dx}\frac{dg(x)}{dx}. \qquad \boxtimes$$

Example 4 Let $C^1[a, b]$ be the set of continuous and differentiable functions on the interval $[a, b]$. We can make this set a semi-inner product space by defining the following semi-inner product between any two functions $f, g \in C[a, b]$:

$$(f, g) \equiv \int_a^b dx \frac{df(x)}{dx} \frac{dg(x)}{dx}. \quad \boxtimes$$

This semi-inner product naturally defines a semi norm $\|f\| \equiv \sqrt{(f, f)}$, which coincides with the semi norm defined in example 4 under the "Normed Space" heading (B.1.3, page 221).

Example 5 Let $M^{p,q}$ be the set of $p \times q$ matrices. This is a vector space that becomes an inner product space if endowed with the *Frobenius inner product* between any two of its elements A and B, which is defined as follows:

$$(A, B) \equiv \text{tr}(AB').$$

The norm associated to this inner product is called the *Frobenius norm*. Given a generic matrix A, its Frobenius norm can be easily computed as follows:

$$\|A\|^2 = \text{tr}(AA') = \sum_{i=1}^{p} \sum_{j=1}^{q} A_{ij}^2.$$

Therefore, the square of the Frobenius norm of a matrix is the sum of the squares of its elements. The Frobenius norm can be used to define the distance between two matrices A and B by setting $d(A, B) \equiv \|A - B\|$. If the matrix B is an approximation of the matrix A, the Frobenius distance between A and B is a natural measure of the approximation error, because it coincides with the error given by a least-squares criterion. $\boxtimes$

B.1.5 Functions, Mappings, and Operators

Let X and Y be two sets. A rule that associates each element $x \in X$ with a unique element $y \in Y$ is called a *mapping* from X into Y and is written as $y = f(x)$. The element y is called the *image* of x under the mapping f. The set X is the *domain* of the map f. The set of elements $y \in Y$ such that $y = f(x)$ for some $x \in X$ is called the *range* of f and is often denoted by $f(X)$. By definition, $f(X) \subset Y$ (which, in words, means that the set $f(X)$ is a subset of the set Y).

When $f(X) = Y$, we say that f maps X *onto* Y. If f maps X onto Y and, to every element $y \in Y$ we can associate a unique element $x \in X$ such that $f(x) = y$, then we say that f is *invertible*. In this case, we denote the element $x \in X$, which corresponds to $y \in Y$ by $x = f^{-1}(y)$, and the mapping f^{-1} from Y to X is called the *inverse* of f.

Other words instead of "mapping" are sometimes used, depending on the properties of X and/or Y. For example, when Y is the set of real numbers we also refer to f as a *real function* on X. Mappings are also called *operators*, although this word is usually reserved for cases in which neither X nor Y is the set of real numbers.

B.1.6 Functional

A *functional* is a real-valued function defined on a vector space X. When $X \subset \mathbb{R}^d$, this coincides with the definition of a real function on X.

Example 1 Let X be a normed space. The norm $\|x\|$ is a functional over X. ⊠

Example 2 Let $C[a, b]$ be the set of continuous functions on the interval $[a, b]$, and let $x_0 \in [a, b]$. Then, for $f \in C[a, b]$, we can define the functional $F[f] \equiv f(x_0)$. ⊠

Example 3 Let $C^1[a, b]$ be the set of functions whose first derivative is continuous on $[a, b]$. Then, for $f \in C^1[a, b]$, we can define the functional

$$F[f] \equiv \int_a^b dx \left(\frac{df(x)}{dx}\right)^2$$

as one simple measure of the smoothness of the function f. ⊠

B.1.7 Span

Let $x_1, \ldots, x_l \in \mathbb{R}^d$ be l vectors in $\mathbb{R}^d$. The span of these vectors is a linear space defined by all the possible linear combinations of these vectors:

$$X \equiv \left\{ x \in \mathbb{R}^d \mid x = \sum_{i=1}^{l} c_i x_i \ , \ c_i \in \mathbb{R} \ , i = 1, \ldots, l \right\}.$$

We also say that the vectors $x_1, \ldots, x_l \in \mathbb{R}^d$ span the space X, or that X is spanned by the vectors $x_1, \ldots, x_l \in \mathbb{R}^d$.

B.1.8 Basis and Dimension

A basis for a vector space X is a set of vectors $x_1, \ldots, x_d$ that are linearly independent and that span X. If the vectors $x_1, \ldots, x_d$ form a basis for X then every vector $x \in X$ can be uniquely written as

$$x = c_1 x_1 + c_2 x_2 + \cdots + c_d x_d, \qquad c_i \in \mathbb{R}, i = 1, \ldots, d.$$

If a vector space X has a basis with d elements, then the number d is known as the *dimension* of the vector space and denoted by $\dim(X)$. We then say that the vector space X is d-dimensional, or has dimension d, and write $\dim(X) = d$.

A vector space may have many different bases, but they must all have the same dimension. The number of elements of a basis can be infinite: in this case, we refer to the vector space as an infinite dimensional vector space.

Example The vectors $x_1 = (1, 0, 0)$, $x_2 = (0, 1, 0)$, and $x_3 = (0, 0, 1)$ form a basis for $\mathbb{R}^3$. The vectors $y_1 = (2, 1, 5)$, $y_2 = (3, 3, 2)$, and $y_3 = (0, -1, 6)$ also form a basis for $\mathbb{R}^3$.

The vectors $y_1 = (2, 1, 5)$, $y_2 = (3, 3, 2)$, and $y_3 = (4, 2, 10)$ do not form a basis for $\mathbb{R}^3$, because they are not linearly independent ($y_3 = 2y_1$) and therefore do not span $\mathbb{R}^3$. ☒

B.1.9 Orthonormality

Let $x_1, \ldots, x_l \in \mathbb{R}^d$ be l vectors in $\mathbb{R}^d$. We say that these vectors are orthonormal if they are mutually orthogonal and of length 1:

$$x_i' x_j = \delta_{ij}, \quad i, j = 1, \ldots, p,$$

where, as always, $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. If a set of orthonormal vectors forms a basis for a vector space X, we refer to them as an orthonormal basis.

Example 1 The vectors $x_1 = (1, 0, 0)$ and $x_2 = (0, 1, 0)$ are orthonormal. However, they do not form an orthonormal basis for $\mathbb{R}^3$, because they do not span $\mathbb{R}^3$ and therefore they are not a basis. ☒

Example 2 The vectors $x_1 = \frac{1}{\sqrt{2}}(1, 1)$ and $x_2 = \frac{1}{\sqrt{2}}(1, -1)$ are orthonormal and form an orthonormal basis for $\mathbb{R}^2$. ☒

B.1.10 Subspace

A set $Y \subset X$ (which should be read as "Y, which is a subset of the set X") is a *subspace* of the vector space X if Y itself is a vector space, with respect to the same operations.

Example 1 The set $\mathbb{R}^2$ is a subspace of $\mathbb{R}^4$, because $\mathbb{R}^2 \subset \mathbb{R}^4$ and both $\mathbb{R}^2$ and $\mathbb{R}^4$ are vector spaces. ☒

Example 2 Let $a \in \mathbb{R}^3$ be a fixed column vector, and let V be the set $V \equiv \{x \in \mathbb{R}^3 \mid a' x = 0\}$. For any two elements $x, y \in V$, we have $x + y \in V$, and, as such, the set V is a subspace of $\mathbb{R}^3$. It is easy to see that V is a two-dimensional plane going through the origin. ☒

Example 3 Let $a \in \mathbb{R}^3$ be a fixed column vector and let V be the set $V \equiv \{x \in \mathbb{R}^3 \mid a' x = 1\}$. For any two elements $x, y \in V$ we have $x + y \notin V$. As such, the set V is *not* a subspace of $\mathbb{R}^3$. It is easy to see that V is a two-dimensional plane that does *not* go through the origin. ☒

Example 4 Let $C[a, b]$ be the set of continuous functions on the interval $[a, b]$. Polynomials of degree n are continuous functions, and the sum of two polynomials of degree n is also a polynomial of degree n. The set Π_n of polynomials of degree n, $n > 0$, is a subspace of $C[a, b]$. ☒

Example 5 Let $M^{p,q}$ be the set of $p \times q$ matrices, and let $M_r^{p,q}$, with $r \leq \min(p, q)$, be the set of elements of $M^{p,q}$ with rank r. While $M^{p,q}$ is a vector space (of dimension pq), the subset $M_r^{p,q}$ of $M^{p,q}$ is *not* a subspace of $M^{p,q}$, because the sum of two matrices of rank r does not necessarily have rank r. ☒

B.1.11 Orthogonal Complement

Subspaces of $\mathbb{R}^d$ (and of some infinite-dimensional spaces) enjoy some particular properties of great usefulness in linear algebra. Let Y be an n dimensional subspace of $\mathbb{R}^d$, with $n < d$, and let us endow $\mathbb{R}^d$ with the Euclidean scalar product. Then we can define the *orthogonal complement* of Y in $\mathbb{R}^d$, and denote it by $Y_\perp$, as the set of all vectors in $\mathbb{R}^d$ that are orthogonal to every element $y \in Y$:

$$Y_\perp = \{x \in \mathbb{R}^d \mid x'y = 0 \ \ \forall y \in Y\}.$$

The set $Y_\perp$ is a subspace of $\mathbb{R}^d$ and has dimension $r = d - n$. An important result is the following: if X is an inner product space and $Y \subset X$ is a subspace of X, then every element x of X has a unique representation as the sum of an element $x_\circ$ of Y and an element $x_\perp$ of $Y_\perp$. In other words, the following representation is unique:

$$\forall x \in X: \quad x = x_\circ + x_\perp, \quad x_\circ \in Y, \ \ x_\perp \in Y_\perp.$$

The preceding statement is often summarized by writing: $X = Y \oplus Y_\perp$, where the symbol $\oplus$ is defined in section B.1.12. The vectors $x_\circ$ and $x_\perp$ are called the *projections* of x onto Y and $Y_\perp$, respectively. Given a vector $x \in X$ and a subspace Y, we can always find the projection of x onto Y using the projection operator defined in section B.1.13 (page 227).

B.1.12 Direct Sum

Let X be a vector space and let $Y, Z \subset X$ be subspaces of X. We say that X is the *direct sum* of Y and Z and write $X = Y \oplus Z$ if every $x \in X$ can be written in a unique way as

$$x = y + z, \ \ y \in Y, \ \ z \in Z.$$

If $X = Y \oplus Z$, then we say that Y and Z are *complementary* subspaces. It is important to note that if Y and Z are complementary subspaces, then $Y \cap Z = 0$. The notion of a direct sum applies to generic vector spaces, and no inner product structure is required. When X is an inner product space, the subspaces Y and Z are orthogonal complements (see section B.1.11). ☒

Example 1 Let $X = \mathbb{R}^3$, $Y = \{x \in X \mid x = (a, b, 0)\ , a, b \in \mathbb{R}\}$, and $Z = \{x \in X \mid x = (c, c, c)\ , c \in \mathbb{R}\}$. The subspace Y is the two-dimensional (x_1, x_2) plane, and the subspace Z is the diagonal of the positive orthant (i.e., the line at a 45° angle with all the coordinate axis). Then Y and Z are complementary subspaces, and $X = Y \oplus Z$. This is equivalent to saying that that every vector in $\mathbb{R}^3$ can be uniquely written as the sum of a vector in the two-dimensional plane and a vector "slanted" at 45° with respect to all the axis. ☒

Example 2 Let $X = \mathbb{R}^3$, $Y = \{x \in X \mid x = (a, b, 0)\ , a, b \in \mathbb{R}\}$, and $Z = \{x \in X \mid x = (0, c, d)\ , c, d \in \mathbb{R}\}$. The subspace Y is the two-dimensional (x_1, x_2) plane, and the subspace Z is the two-dimensional $(x_2 - x_3)$ plane. Although it is true that every $x \in X$ can be written as the sum of an element of Y and an element of Z, this representation is clearly not unique, and therefore X is not the direct sum of Y and Z, and Y and Z are *not* complementary subspaces (in fact, $Y \cap Z = \{x \in X \mid x = (0, b, 0)\ , b \in \mathbb{R}\} \neq \{0\}$). ☒

B.1.13 Projection Operators

Definition and Properties Let X be a vector space, and let $P: X \to X$ be a linear operator that maps X into itself. The operator P is called a *projection operator*, or a *projector* if $P(Px) = Px$, $\forall x \in X$ (in short: $P^2 = P$). In all the cases in which we use projectors in this book, the vector space X is $\mathbb{R}^d$, and therefore we can think of a projector P simply as an $d \times d$ matrix such that $P^2 = P$. In addition, where projectors are involved, we always assume that $\mathbb{R}^d$ is endowed with the usual Euclidean inner product. Projectors are very useful whenever we are given a vector $x \in X$ and are interested in picking the part of X that lies in a subspace Y (i.e., which is "explained" by Y). In order to see the connection between projectors and subspaces, remember that if Y is a subspace of an inner product vector space X, then $X = Y \oplus Y_\perp$, that is, any vector $x \in X$ can be uniquely written as $x = x_\circ + x_\perp$, with $x_\circ \in Y$ and $x_\perp \in Y_\perp$. This means that there is a well-defined map $P_\circ$ that associates with the vector $x \in X$, the vector $x_\circ \in Y$, so that $x_\circ = P_\circ x$. This map is clearly linear, and because $x_\circ \in Y$, it satisfies the identity $P_\circ(P_\circ x) = P_\circ x_\circ = x_\circ$. Therefore $P_\circ^2 = P_\circ$, and $P_\circ$ is a projector, which is often referred to as the projector of X onto Y. We also say that $x_\circ$ is the projection of x onto Y.

Example 1 Let $X = \mathbb{R}^3$ and let Y be the (x_1, x_2) plane, that is, $Y = \{x \in X \mid x = (a, b, 0)\ , a, b \in \mathbb{R}\}$. Because Y is a subspace of X, $X = Y \oplus Y_\perp$. The orthogonal complement $Y_\perp$ of Y is the vertical axis: $Y_\perp = \{x \in X \mid x = (0, 0, c)\ , c \in \mathbb{R}\}$. In words: any vector $x = (a, b, c)$ can be written as $x = (a, b, 0) + (0, 0, c)$, and the vectors $(a, b, 0)$ and $(0, 0, c)$ are orthogonal. Therefore, the projection $x_\circ$ of $x = (a, b, c)$ onto Y is $x_\circ = (a, b, 0)$, and it is easily verified that the projector $P_\circ$ of X onto Y has the form:

$$P_\circ = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}. \qquad \boxtimes$$

If Y is a subspace of X, with its associated projector $P_\circ$, its orthogonal complement is also a subspace, and it has a corresponding projector that we define as $P_\perp$, with the property that if $x = x_\circ + x_\perp$, then $x_\perp = P_\perp x$. Therefore, the two projectors $P_\circ$ and $P_\perp$ are related:

$$P_\circ + P_\perp = I.$$

Writing this relationship as $P_\perp = I - P_\circ$ makes clear that $P_\perp$ should be interpreted as the "residual making matrix" with respect to Y. In fact $x_\circ$ is the component of x "explained" by Y, and the residual is $x_\perp = x - x_\circ = P_\perp x$. Because the residual $x_\perp$ is orthogonal to $x_\circ$, it also follows that the two projectors $P_\circ$ and $P_\perp$ have the property that $P_\circ P_\perp = P_\perp P_\circ = 0$. We summarize the properties of these projectors as

$$P_\circ P_\circ = P_\circ, \quad P_\perp P_\perp = P_\perp, \quad P_\circ + P_\perp = I, \quad P_\circ P_\perp = P_\perp P_\circ = 0.$$

Constructing Projectors Assume that $X = \mathbb{R}^d$ and Y is an n-dimensional subspace of X. Let $\{u_1, \ldots, u_n \in X\}$ be a basis for Y, that is, a set of n vectors that span Y (they do not have to be orthonormal). This means that any vector in Y can be written as a linear combination of the vectors $u_1, \ldots, u_n$. Our goal is to find a unique decomposition of an arbitrary vector x as $x_\circ + x_\perp$, with $x_\circ \in Y$ and $x_\perp \in Y_\perp$. Because $x_\circ \in Y$, then it can be

written as $x_\circ = \sum_{i=1}^{n} a_i u_i$ for some coefficients a_i. Therefore, our goal is to find, for every $x \in X$, coefficients a_i such that the following two conditions are satisfied:

$$x = \sum_{i=1}^{n} a_i u_i + x_\perp, \quad x'_\perp \sum_{i=1}^{n} a_i u_i = 0, \tag{B.1}$$

where the last condition ensures that $x_\perp$ and $x_\circ$ are orthogonal. If we define an n-dimensional vector $a = (a_1, \ldots, a_n)$ and the $n \times d$ matrix U with the vectors $u_1, \ldots, u_n$ on its columns, the preceding expression can be rewritten as

$$x = Ua + x_\perp, \quad x'_\perp Ua = 0. \tag{B.2}$$

Substituting $x_\perp = x - Ua$ in the preceding orthogonality condition, we conclude that the vector of coefficients a must satisfy

$$(x - Ua)'Ua = 0, \quad \Rightarrow a'(U'x - U'Ua) = 0.$$

Because we know that the decomposition of equation B.1 is unique, the solution of the preceding equation is obtained by setting $U'x - U'Ua$ to 0, so that we obtain $a = (U'U)^{-1}U'x$. Therefore, $x_\circ = Ua = U(U'U)^{-1}U'x$. This implies that the projector $P_\circ$ of X on the subspace Y is the following matrix:

$$P_\circ = U(U'U)^{-1}U'. \tag{B.3}$$

Therefore, all we need in order to construct the projector $P_\circ$ onto an arbitrary subspace Y of X is a basis for the subspace Y. If the basis for Y is orthonormal then $U'U = I$, and the formula above simplifies to

$$P_\circ = UU'. \tag{B.4}$$

The derivation of $P_\circ$ has an obvious interpretation in linear regression theory. In fact, equation B.2 can be seen as a linear specification for the vector of observations x, in which a is the vector of unknown regression coefficients, U is the matrix of covariates, and $x_\perp$ is a residual disturbance. The condition $x'_\perp Ua = 0$ expresses the well-known fact that the residuals and the fitted values (Ua) are mutually orthogonal, and the residual making matrix $I - U(U'U)^{-1}U'$ is immediately identified with $P_\perp$. Equation B.3 is often called the "hat matrix."

The connection with linear regression helps to explain an important property of projection operators, and one of the reasons for which they are so useful in practice: the projection of x on the subspace Y is the vector of Y, which has minimum distance from x. This follows from the observation that the vector of coefficients a, which we have derived previously, is also the vector that minimizes the least-squares error $\|x - Ua\|^2$, which is exactly the Euclidean distance between x and a generic element Ua of Y. In other words, if we have a vector space X and a subspace Y, and we want to approximate a vector $x \in X$ with an element of Y, the solution of this problem is simply $P_\circ x$, where $P_\circ$ is the projector of X on Y and can be computed using equation B.3.

Example 2 Let $X = \mathbb{R}^3$, let $w \in X$ be a given vector of norm 1, and let $Y = \{x \in X \mid w'x = 0\}$. Y is a two-dimensional subspace, and more precisely a two-dimensional "slanted" plane going through the origin. Because Y is constructed as the set of points that are orthogonal to the vector w, the orthogonal complement $Y_\perp$ of Y is simply the set of vectors that are multiples of w, that is, a line through the origin. Let x be a generic point in X: we wish to find the closest point to x on the plane Y. From what we have seen previously, this point is simply $P_\circ x$, where $P_\circ$ is the projector of X onto Y. In order to find $P_\circ$, we need a basis for Y, which is not readily available. However, we have a basis for $Y_\perp$, which is given by w. Therefore, we can find the projector $P_\perp$ of X on $Y_\perp$ and obtain $P_\circ$ as $I - P_\perp$. Applying formula B.4, we have $P_\perp = ww'$, and therefore:

$$P_\circ = I - ww'. \qquad \boxtimes$$

B.2 Linear Algebra

B.2.1 Range, Null Space, Rank, and Nullity

In order to understand the properties of a matrix, it is important to understand the effect it has when it acts on a vector. Crucial to this understanding are two subspaces associated with a matrix: its range and its null space, which we now describe. Let A be a $q \times d$ matrix, and let x be a d-dimensional column vector. The first question we ask is, What happens (i.e., what kind of vectors do we obtain) when we operate with A on x? To answer this question, we need to study the *range* of A, that is, the set of all the vectors $y \in \mathbb{R}^q$ that can be written as $y = Ax$ for some $x \in \mathbb{R}^d$. Formally, we have

$$\text{range}(A) \equiv \{y \in \mathbb{R}^q \mid y = Ax, \text{ for some } x \in \mathbb{R}^d\}. \tag{B.5}$$

Because the expression Ax can be read as "a linear combination of the columns of A with coefficients given by the components of x," we can also define the range of A as the vector space spanned by the columns of A. For this reason, the range of A is often referred to as the *column space* of A. Denoting by $a_1, \ldots a_d \in R^q$ the columns of A, this definition is formalized as follows:

$$\text{range}(A) \equiv \{y \in \mathbb{R}^q \mid y = \sum_{i=1}^{d} x_i a_i, \text{ for some } x_i \in \mathbb{R},\ i = 1, \ldots, d\}. \tag{B.6}$$

Let us assume that $q \geq d$: this definition makes clear that in this case the range of A is a vector space whose dimensionality r is at most d. The reason for which we say "at most" d, rather than equal to d, is that the columns of A may not be linearly independent, and therefore they may not span $\mathbb{R}^d$. Let us assume instead that $q < d$: then, since the d vectors a_i are q-dimensional, they can span at most a q-dimensional space. Therefore, by defining the *rank* of A as the dimensionality r of range(A), and denoting it by rank(A), we conclude that

$$\text{rank}(A) \leq \min(q, d).$$

When $\text{rank}(A) = \min(q, d)$, we say that the matrix A has *full rank*; otherwise, we say that it is *rank deficient*.

To summarize: the matrix A takes $\mathbb{R}^d$ and maps into a subspace of $\mathbb{R}^q$ whose dimensionality is rank(A) and is at most d. A fundamental result of linear algebra, which we do not prove here but that we will use later, is the following:

$$\text{rank}(A) = \text{rank}(A'). \tag{B.7}$$

We now present four examples that exhaust all the possibilities for the values of the rank of a matrix.

Example 1: $q \geq d$, full rank Consider the following matrix A:

$$A \equiv \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{pmatrix}. \tag{B.8}$$

The rank of A is 2, because the two columns of A are linearly independent. Because $2 = \min(3, 2)$, the matrix has full rank. The range of A is a two-dimensional subspace of $\mathbb{R}^3$, that is, a slanted plane going through the origin. Every point on the slanted plane is the image of at least one point $x \in \mathbb{R}^2$ (see section B.1.5 , page 223, for the definition of image). Vectors in $\mathbb{R}^3$ that are not in the range of A, that is, do not lie on the slanted plane, are not the image of any point in $\mathbb{R}^2$. ⊠

Example 2: $q \geq d$, rank deficient Consider the following matrix A:

$$A \equiv \begin{pmatrix} 1 & 2 \\ 1 & 2 \\ 1 & 2 \end{pmatrix}. \tag{B.9}$$

The rank of A is 1, because the second column is a multiple of the first. Because $1 < \min(3, 2)$, the matrix is rank deficient. The range of A is a one-dimensional subspace of $\mathbb{R}^3$, that is, the set of vectors x that are multiples of the vector (1, 1, 1). This subspace is therefore the diagonal of the positive orthant. Every point on the diagonal is the image of at least a point $x \in \mathbb{R}^2$. Vectors in $\mathbb{R}^3$ that are not in the range of A, that is, do not lie on the diagonal of the positive orthant, are not the image of any point in $\mathbb{R}^2$. ⊠

Example 3: $q < d$, full rank Consider the following matrix A:

$$A \equiv \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}. \tag{B.10}$$

The rank of A is 2, because the third column of A is the sum of the first 2. Because the columns are two-dimensional vectors the rank is as high as it can be and the matrix has full rank ($2 = \min(2, 3)$). The range of A is the entire space $\mathbb{R}^2$: every two-dimensional vector y is the image of at least one point x under A. ⊠

Example 4: $q < d$, rank deficient Consider the following matrix A:

$$A \equiv \begin{pmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \end{pmatrix}. \tag{B.11}$$

The rank of A is 1, because the second and third columns of A are both multiples of the first. Because $1 < \min(2, 3)$ the matrix is rank deficient. The range of A is not the entire space $\mathbb{R}^2$, but a one-dimensional subspace, which is easily identified with the diagonal of the positive quadrant: only points on the diagonal are images of points in $\mathbb{R}^3$ under A. ⊠

The range of a matrix is important because it allows us identify those vectors for which the matrix equation $Ax = y$ has at least one solution. Clearly, if y is not in the range of A, then there is no solution, and if y is in the range of A, there is at least one solution. The reason for which there may be more than one solution is that there may be vectors $x_\circ$ such that $Ax_\circ = 0$. In this case, if we have $Ax = y$, then we also have $A(x + x_\circ) = y$. The set of vectors $x \in \mathbb{R}^d$ such that $Ax = 0$ is called the *null space* of the matrix A, and it is denoted by $\mathfrak{N}(A)$ (or just $\mathfrak{N}$ if no confusion arises). Formally, we have

$$\mathfrak{N}(A) \equiv \{x \in \mathbb{R}^d \mid Ax = 0\}. \tag{B.12}$$

By rewriting the condition $Ax = 0$ as $\sum_{i=1}^{d} x_i a_i = 0$, we see that if the columns of A are linearly independent, then $\mathfrak{N}(A) = \{0\}$, and we say that the null space is trivial. In fact, if the columns of A are linearly independent, the only numbers x_i such that $\sum_{i=1}^{d} x_i a_i = 0$ are zeros (if the x_i were not 0, one could express one of the a_i as a linear combination of the others). Therefore, if A has full rank, its null space is trivial, and if a solution to $Ax = y$ exists, it is unique. When A is rank deficient, we can expect a nontrivial null space: in this case the equation $Ax = y$ has an infinite set of solutions, which differ by an element of the null space of A (if x_1 and x_2 are solutions, then $A(x_1 - x_2) = 0$ and $x_1 - x_2 \in \mathfrak{N}(A)$).

We will also need to know "how big" is the set of solutions making up the null space. To do this, we note that $Ax = 0$ implies that x is orthogonal to every row of A, or every column of A', or any linear combination of columns of A'. This is equivalent to saying that x is orthogonal to the span of the columns of A', which in turn is the same as saying that x is orthogonal to the range(A') (because range(A') is the span of the columns of A'). This sequence of reasoning leads us to the key result that *the null space of a matrix is the orthogonal complement of the range of its transpose*:

$$\mathfrak{N}(A) = \text{range}(A')_\perp. \tag{B.13}$$

This result is important because it allows us to compute the dimensionality of the null space of A, which is called the *nullity* of A and denoted by $\text{nullity}(A) \equiv \dim(\mathfrak{N})$. In fact, because the range of A' is a subspace of $\mathbb{R}^d$ of dimension $\text{rank}(A')$, we know that its orthogonal complement must have dimension $\dim(\text{range}(A')_\perp) = d - \text{rank}(A)$. Therefore, we conclude with the fundamental decomposition:

$$\text{nullity}(A) = d - \text{rank}(A), \tag{B.14}$$

where throughout d is the number of columns of A. As anticipated previously, then, the nullity of a matrix is zero (and the null space is trivial) only if the matrix has full rank.

The results of this section can be summarized as follows: the range of A allows us to characterize the vectors y for which the linear equation $Ax = y$ has a solution. The null space of A allows us to characterize whether this solution is unique, and in the case it is not, the whole set of solutions.

Example 1 (continued): $q \geq d$, full rank We have considered the matrix A:

$$A \equiv \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{pmatrix}. \tag{B.15}$$

We have already seen that the rank of A is $2 = \min(3, 2)$, so that the matrix has full rank. Therefore its null space is trivial and every point in the range of A (the slanted plane) is the image of only one point in $\mathbb{R}^2$: the map between $\mathbb{R}^2$ and the slanted plane is one-to-one and therefore invertible. In fact, if $y \in \text{range}(A)$, we can solve $Ax = y$ with the usual formula:

$$x = (A'A)^{-1}A'y. \tag{B.16}$$

Notice that this shows that if A has full rank then $A'A$ must be invertible. ☒

Example 2 (continued): $q \geq d$, rank deficient We have considered the following matrix A:

$$A \equiv \begin{pmatrix} 1 & 2 \\ 1 & 2 \\ 1 & 2 \end{pmatrix}. \tag{B.17}$$

We have already seen that A is rank deficient, because its rank is 1. As a consequence, its null space is not trivial: from equation B.14 we have that $\text{nullity}(A) = 2 - \text{rank}(A) = 1$, so $\mathfrak{N}(A)$ is one-dimensional. We now define $\mathfrak{N}(A)$ explicitly. From equation B.13, we know that it is the orthogonal complement of the range of A':

$$A' \equiv \begin{pmatrix} 1 & 1 & 1 \\ 2 & 2 & 2 \end{pmatrix}.$$

This expression confirms the fact that the rank of A and A' are the same, because only 1 of the three column vectors of A' is linearly independent of the others. The range of A' is therefore given by the multiples of the vector $(1, 2)$. Its orthogonal complement, that is, $\mathfrak{N}(A)$, is easily seen to be the set of vectors that are multiples of the vector $(-2, 1)$ (since $(-2, 1)(1, 2)' = 0$). Therefore, every point in $\mathbb{R}^2$ is mapped into a point on the diagonal of the positive quadrant (range(A)), but all the points of the form $x + \alpha(-2, 1)$, for any $\alpha \in \mathbb{R}$, that lie on a straight line through x are mapped into the same point on the diagonal. Therefore, the solution to the linear equation $Ax = y$, when it exists, is not unique, and equation B.16 does not hold anymore. This implies that when A is rank deficient, the matrix $A'A$ must be not invertible. As we will see in section B.2.5, in this case we can still give a meaning to the problem of solving $Ax = y$, but equation B.16 must be replaced by something else. ☒

B.2.2 Eigenvalues and Eigenvectors for Symmetric Matrices

Symmetric matrices and their eigenvectors and eigenvalues play a special role in this book, so we list here some of their properties.

Let A be a $d \times d$ symmetric matrix. If we can find a non-null vector $v \in \mathbb{R}^d$ and a number ℓ such that

$$Av = \ell v,$$

then we say that v is an eigenvector of A with eigenvalue ℓ. Notice that if v is an eigenvector with eigenvalue ℓ, then kv, with $k \in \mathbb{R}$, is also an eigenvector with eigenvalue ℓ. We eliminate this trivial redundancy by using the convention that eigenvectors always have length 1, so that $\|v\| = 1$ unless otherwise specified.

An important property of symmetric $d \times d$ matrices is that they always have exactly d mutually orthogonal eigenvectors $v_1, \ldots v_d$. We denote by $\ell_1, \ldots \ell_d$ the corresponding eigenvalues.

Let R be a $d \times d$ matrix with the eigenvectors of A, $v_1, \ldots v_d$, as its columns. By virtue of the orthogonality of the eigenvectors and the convention that they have length one, it follows that R is an orthogonal matrix, so that $R' = R^{-1}$. Let L be the diagonal matrix with the eigenvalues $\ell_1, \ldots, \ell_d$ on the diagonal. One can prove that the matrix A can always be written as

$$A = RLR'. \tag{B.18}$$

Equation B.18, the *eigenvalue/eigenvector decomposition*, tells us everything we may need to know about the matrix A. If the eigenvalues are all strictly positive then the matrix is invertible, and the inverse is simply:

$$A^{-1} = RL^{-1}R',$$

where L^{-1} is a diagonal matrix with the reciprocals of the eigenvalues $(1/\ell_1, \ldots, 1/\ell_d)$ on the diagonal.

The rank r of A is the number of nonzero eigenvalues, and the nullity n of A is the number of zero eigenvalues. The eigenvectors corresponding to the nonzero eigenvalues, that is, the first r columns of R, span the range of A, while the eigenvectors corresponding to the zero eigenvalues (the last n columns of R) span the null space of A.

When A does not have full rank, the decomposition B.18 can be written in a simplified form, often useful for computational purposes. Let us write the matrix L in block form:

$$L = \begin{pmatrix} \ell & 0 \\ 0 & 0 \end{pmatrix}, \quad \ell = \text{diag}[(\ell_1, \ldots \ell_r)].$$

Let us also write R as $R = (R_\perp R_\circ)$, where $R_\perp$ is a $d \times r$ matrix whose columns are the first r eigenvectors (a basis for $\mathfrak{N}_\perp$), and $R_\circ$ is a $d \times n$ matrix whose columns are the last n eigenvectors (a basis for $\mathfrak{N}$). Then we have the following identity:

$$A = RLR' = (R_\perp R_\circ) \begin{pmatrix} \ell & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} R'_\perp \\ R'_\circ \end{pmatrix} = R_\perp \ell R'_\perp. \tag{B.19}$$

B.2.3 Definiteness

Let A be an $n \times n$ symmetric matrix. If we have

$$x'Ax > 0, \quad \forall x \in \mathbb{R}^n \ , \ x \neq 0,$$

then we say that A is *positive definite*. If this condition is substituted with the weaker condition

$$x'Ax \geq 0, \quad \forall x \in \mathbb{R}^n,$$

then we say that A is *positive semidefinite*. The only difference between positive definite and positive semidefinite matrices is that for a positive semidefinite matrix A the fact that $x'Ax = 0$ does not imply $x = 0$. Similar definitions for negative definite and negative semidefinite matrices are obtained by switching the sign in the preceding inequalities.

If a matrix A is positive definite, then the quantity $x'Ax$ is a norm, while if A is positive semidefinite, then $x'Ax$ is a seminorm.

It can be shown that A is positive definite if and only if its eigenvalues are all strictly positive, where we use the word "strictly" to emphasize the fact that they cannot be equal to 0. Similarly, A is positive semidefinite if and only if its eigenvalues are either positive or equal to zero.

Example 1 Let A be an $n \times d$ matrix of full rank. Then the $d \times d$ matrix $C = A'A$ is positive definite. This is seen by noticing that $x'Cx = x'A'Ax = (Ax)'Ax = \|Ax\|^2 > 0$. The strict inequality is a consequence of the fact that $Ax = 0$ implies $x = 0$, because A has full rank. If A has rank smaller than d, then C is positive semidefinite. ☒

B.2.4 Singular Values Decomposition

For symmetric matrices the eigenvalue-eigenvector decomposition tells us everything we need to know about a matrix. The analog of this decomposition for generic, rectangular matrices is the singular values decomposition (SVD). SVD is one the most useful and common tools of linear algebra, and it has been known for more than a century (Beltrami, 1873; Jordan, 1874; see Stewart, 1992, for its history). Here we give the basics facts, needed in this book, and refer the reader to a linear algebra book (e.g., Strang, 1988) for a full explanation.

Definition

Let A be a $q \times d$ matrix, with $q \geq d$ (if $q < d$ we can look at the transpose of A). It is possible to show that one can always write A as follows:

$$A = UWV', \tag{B.20}$$

where U is an $q \times d$ matrix whose columns are mutually orthonormal ($U'U = I_d$), W is a $d \times d$ diagonal matrix with positive or zero values, and V' is a $d \times d$ orthogonal matrix ($V'V = V'V = I_d$). The diagonal elements of W, denoted by $w_1, \ldots w_d$, are called the

singular values of the matrix A. The decomposition B.20 is known as the *singular values decomposition*, and it is unique, up to a permutation of the columns of U and V and of the corresponding singular values. Here we consider the general case in which n singular values are 0, with $n \geq 0$, and denote by r the number of nonzero singular values, so that $r + n = d$ and

$$w_1 \geq w_2 \geq \dots w_r \geq w_{r+1} = w_{r+2} \dots w_{r+n} = 0.$$

The singular values play here the role played by the eigenvalues in the eigenvalue-eigenvector decomposition of symmetric matrices: the rank of A is equal to r, the number of nonzero singular values, and therefore the nullity of A is n, which is the number of singular values equal to zero. We now list some useful properties of the SVD decomposition. To this end, we define $u_1, \dots u_d$ as the column vectors of U and $v_1, \dots v_d$ as the column vectors of V.

- Rank$(A) = r$.
- Nullity$(A) = n$.
- The vectors $u_1, \dots u_r$ (first r columns of U) form a basis for the range of A.
- The vectors $v_{r+1}, \dots v_d$ (last n columns of V) form a basis for the null space of A.
- The following relationships hold:

$$Av_i = w_i u_i, \quad A'u_i = w_i v_i, \quad i = 1, \dots d.$$

- If the matrix A is square, then it is invertible if and only if the singular values are all strictly positive. In this case one can easily verify that the inverse of A is given by:

$$A^{-1} = V W^{-1} U'.$$

For Approximation

We now show how to use SVD to approximate a matrix A as a linear combination of "simpler" matrices, and how to bound the corresponding approximation error. If we use this notation, equation B.20 can be rewritten as

$$A = \sum_{i=1}^{r} w_i u_i v_i' \equiv \sum_{i=1}^{r} w_i a_i, \tag{B.21}$$

where we have defined the $q \times d$ matrices $a_i = u_i v_i'$, which are all of rank 1 (the fact that a_i has rank 1 follows from its SVD). Equation B.21 is a powerful result: it says that *any matrix A of rank r can be written as a linear combination of r matrices of rank 1*. Because of the orthonormality properties of U and V, the matrices a_i are mutually orthonormal under the Frobenius inner product (see examples in appendix B.1.4, page 222). In particular, they all have the same "size," where the size is measured by their Frobenius norm, which is equal to 1.

Therefore, if some singular values are much bigger than the others, then the corresponding terms in the expansion B.21 will dominate the others. This suggests that a good approximation to the matrix A can be obtained by retaining, in equation B.21, only the largest singular values. To quantify this observation, define $\tilde{A}_k$ as the approximation of A obtained by retaining only the k ($k < r$) largest singular values in equation B.21:

$$\tilde{A}_k \equiv \sum_{i=1}^{k} w_i a_i.$$

We quantify the approximation error as $\|A - \tilde{A}_k\|^2$, where $\|\cdot\|$ is the Frobenius norm (see appendix B.1.4, page 222). However, this definition of error is not always useful because it depends on the scale of matrix A. Therefore, we introduce a relative measure of error, defined as

$$\Delta E_k \equiv \frac{\|A - \tilde{A}_k\|^2}{\|A\|^2}.$$

Substituting the expansion B.21 in the preceding equation, we obtain

$$\Delta E_k = \frac{\|\sum_{i=k+1}^{d} w_i a_i\|^2}{\|\sum_{i=1}^{d} w_i a_i\|^2} = \frac{\sum_{i=k+1}^{d}\sum_{j=k+1}^{d} w_i w_j (a_i, a_j)}{\sum_{i=1}^{d}\sum_{i=j}^{d} w_i w_j (a_i, a_j)} = \frac{\sum_{i=k+1}^{d} w_i^2}{\sum_{i=1}^{d} w_i^2}, \qquad \text{(B.22)}$$

where we have used the definition of the Frobenius norm $\|\cdot\|$ in terms of Frobenius inner product $(\cdot, \cdot)$ ($\|A\|^2 = (A, A)$) and the orthonormality of the matrices a_i under the Frobenius inner product ($(a_i, a_j) = \delta_{ij}$). Equation B.22 is a very useful result: it allows us to estimate precisely the error we make approximating a matrix A of rank r by a linear combinations of k matrices of rank 1 only in terms of the singular values of A: the faster the decay rate of the singular values, the better the approximation using a small number of terms. The relative error ΔE_k is usually referred to as "the percentage of the variance linearly accounted for by the first k singular values." The term "variance" here refers to the square of the Frobenius norm of A. The reason for this terminology is that in some applications the rows of A are realizations of a random d-dimensional variable with zero mean, and therefore the Frobenius norm of A (squared) is proportional to the variance of this random variable.

B.2.5 Generalized Inverse

Let A be a $d \times d$ matrix. In this section, we consider the problem of finding a solution to the linear system $Ax = y$. We do not consider the more general problem in which A is rectangular to keep notation simple.

When A has full rank the solution to this problem is obviously unique and given by $x = A^{-1}y$. When A is rank deficient, with $\text{rank}(A) = r$, $r < d$, two things happen:

1. The range of A is an r-dimensional subspace of $\mathbb{R}^d$, and therefore the linear system $Ax = y$ has a solution only if $y \in \text{range}(A)$.
2. The null space of A is a subspace of $\mathbb{R}^d$ with dimensionality $n = d - r$ (see equation B.14, page 231). Therefore, when $y \in \text{range}(A)$ and a solution exist, it is not unique. In fact, if x is a solution, then $x + x_\circ$, where $x_\circ \in \mathfrak{N}(A)$, is also a solution.

Here we assume that $y \in \text{range}(A)$ (otherwise an exact solution does not exist), and focus on the problem of having an infinite number of solutions. Even if all these solutions are equivalent, in the sense that they differ from each other for an element of $\mathfrak{N}(A)$, which is "invisible" to A, it is important to have a consistent criterion to pick a particular one, which we can consider as "representative." In order to choose a criterion, we reason as follows.

Because $\mathfrak{N}(A)$ is a subspace of $\mathbb{R}^d$, we write $\mathbb{R}^d = \mathfrak{N}(A) \oplus \mathfrak{N}(A)_\perp$ (see sections B.1.10 and B.1.11 on subspaces and orthogonal complements). Let x be such that $Ax = y$, and let us decompose it as $x = x_\circ + x_\perp$, with $x_\circ \in \mathfrak{N}(A)$ and $x_\perp \in \mathfrak{N}(A)_\perp$. Because x is a solution of $Ax = y$, by adding or subtracting any element of $\mathfrak{N}(A)$, we obtain another solution, and v $x_\circ \in \mathfrak{N}(A)$, therefore $x - x_\circ = x_\perp$ is also a solution. Therefore, there is always a well-defined solution that lies in $\mathfrak{N}(A)_\perp$, that is, a solution whose projection on the null space of A is zero. We take this as the "representative" solution. We will see later how this is equivalent to choosing the solution of minimum norm.

To summarize, we wish to find, among all the possible vectors x such that $Ax = y$, the one such $x_\circ = P_\circ x = 0$. This problem is solved easily by using the SVD of A. First, we notice that the condition $P_\circ x = 0$ can be written as $P_\perp x = x$, because $P_\circ + P_\perp = I$. Therefore, we substitute $P_\perp x$ for x in $Ax = y$ and obtain $AP_\perp x = y$. Now let us use the SVD of A and write $A = UWV'$, where U and V are orthogonal $d \times d$ matrices. The equation $AP_\perp x = y$ becomes

$$UWV'P_\perp x = y, \Rightarrow\ WV'P_\perp x = U'y. \tag{B.23}$$

By the properties of the SVD, the matrix W is diagonal, and because $\text{rank}(a) = r$, $\text{nullity}(A) = n$, and $d = r + n$, it has the following structure:

$$W \equiv \begin{pmatrix} w & 0_{r\times n} \\ 0_{n\times r} & 0_{n\times n} \end{pmatrix},$$

where w is an $r \times r$ diagonal matrix with the first r nonzero singular values on the diagonal. Because our goal is to "isolate" $P_\perp x$ in equation B.23, ideally we would multiply both sides of B.23 by the inverse of W. The inverse of W does not exist, but we can define something that resembles it and see whether that it is enough. We define the matrix W^+ here, which we list together with the one useful identity it satisfies:

$$W^+ \equiv \begin{pmatrix} w^{-1} & 0_{r\times n} \\ 0_{n\times r} & 0_{n\times n} \end{pmatrix}, \quad W^+W = \begin{pmatrix} I_{r\times r} & 0_{r\times n} \\ 0_{n\times r} & 0_{n\times n} \end{pmatrix}.$$

Now we premultiply both sides of equation B.23 by W^+ and obtain

$$\begin{pmatrix} I_{r\times r} & 0_{r\times n} \\ 0_{n\times r} & 0_{n\times n} \end{pmatrix} V'P_\perp x = W^+U'y.$$

Now we remember that from SVD the matrix V is orthogonal and has the form $V = (V_\perp \ \ V_\circ)$, where $V_\circ$ is a basis for $\mathfrak{N}(A)$ and $V_\perp$ is a basis for $\mathfrak{N}(A)_\perp$. Premultiplying both sides of the preceding equation by V we obtain

$$V \begin{pmatrix} I_{r\times r} & 0_{r\times n} \\ 0_{n\times r} & 0_{n\times n} \end{pmatrix} V'P_\perp x = (V_\perp \ \ V_0) \begin{pmatrix} I_{r\times r} & 0_{r\times n} \\ 0_{n\times r} & 0_{n\times n} \end{pmatrix} \begin{pmatrix} V'_\perp \\ V'_\circ \end{pmatrix} P_\perp x.$$

As a final step, we remember that because the columns of $V_\perp$ form an orthonormal basis for $\mathfrak{N}(A)_\perp$, the projector $P_\perp$ on $\mathfrak{N}(A)$ is simply $P_\perp = V_\perp V_\perp'$. Therefore,

$$V_\perp V_\perp' P_\perp x = P_\perp P_\perp x = P_\perp x = V W^+ U' y.$$

Because we started with the assumption $P_\perp x = x$, the solution to our problem is finally:

$$x = V W^+ U' y \equiv A^+ y, \tag{B.24}$$

where the matrix $A^+ = V W^+ U'$ is the so-called *generalized inverse* of A. Notice that in applying this definition the generalized inverse inverse of W is W^+, which justifies our notation. To summarize: among the infinite number of solutions of $Ax = y$, with $y \in$ range(A), the solution computed using the generalized inverse is the one whose projection on the null space of A is zero.

In many books, the generalized inverse is defined in the same way, but it is derived according to a different criterion: among all the solutions of $Ax = y$, the solution computed using the generalized inverse is the one with minimum norm. We now show that these two criteria are equivalent. The set of all solutions can be obtained by adding to a known solution (e.g., $x^* = A^+ y$) arbitrary points in $\mathfrak{N}(A)$, which can always be written as $P_\circ z$, for arbitrary $z \in \mathbb{R}^d$. Therefore, the set of solutions is the set of vectors that can be written as $x = x^* + P_\circ z$, with z varying in $\mathbb{R}^d$. Let us find the vector of this form with minimum norm:

$$\min_z \ \|x^* + P_\circ z\|^2 = \min_z \ [2(P_\circ z)' x^* + (P_\circ z)'(P_\circ z)].$$

We notice that $(P_\circ z)' x^* = z' P_\circ' x^* = z' P_\circ x^* = 0$, where we have used the fact that $P_\circ' = P_\circ$ and $P_\circ x^* = 0$ by definition of x^*. Therefore,

$$\min_z \ \|x^* + P_\circ z\|^2 = \min_z \ (P_\circ z)'(P_\circ z).$$

The minimum is attained for any z such that $P_\circ z = 0$, and its value is $\|x^*\|^2$. Therefore the solution whose projection in the null space of A is 0 and the solution of minimum norm coincide.

B.2.6 Quadratic Form Identity

Let $b_1, \ldots, b_N \in R^d$ be a collection of N vectors in R^d. For any $N \times N$ symmetric matrix s and for any $d \times d$ symmetric matrix Φ, the following identity holds:

$$\frac{1}{2} \sum_{i,j=1}^{N} s_{ij} (b_i - b_j)' \Phi (b_i - b_j) = \sum_{i,j=1}^{N} W_{ij} b_i' \Phi b_j, \tag{B.25}$$

where

$$W \equiv s^+ - s, \quad s^+ \equiv \text{diag}[s_i^+], \quad s_i^+ \equiv \sum_{j=1}^{N} s_{ij}. \tag{B.26}$$

Because the rows of the matrix W sum to 0 ($W\mathbf{1} = 0$), W is not full rank. If the elements of s are all positive, the matrix W is positive semidefinite, but the reverse does not generally hold. The values of the expression in equation B.25 do not depend on the diagonal elements of s. Therefore, for all expressions of the form in equation B.25, we assume $s_{ii} = 0, \forall i$. Under this assumption, given any matrix W such that $W\mathbf{1} = 0$, we can always find a matrix s such that $W \equiv s^+ - s$: it is sufficient to take $s = \text{diag}(W) - W$.

A particular case of equation B.25 that appears often in the book is when $d = 1$ and $\Phi = 1$. We restate it as follows. Let $u \equiv (u_1, \ldots, u_N)'$ be a column vector, then

$$\frac{1}{2} \sum_{i,j=1}^{N} s_{ij}(u_i - u_j)^2 = \sum_{i,j=1}^{N} W_{ij} u_i u_j = u' W u. \tag{B.27}$$

B.3 Probability Densities

B.3.1 The Normal Distribution

Let D be a strictly positive definite $d \times d$ matrix and $\theta > 0$. We say that a d-dimensional random variable x has normal distribution with mean $\bar{x}$ and covariance D^{-1}, and write $x \sim \mathcal{N}(\bar{x}, D^{-1})$, if its probability density is

$$\mathcal{P}(x) = \left(\frac{\theta}{2\pi}\right)^{\frac{d}{2}} \sqrt{\det D}\ \exp\left(-\frac{1}{2}\theta(x - \bar{x})' D (x - \bar{x})\right).$$

Because a density must integrate to 1, we have the multidimensional saussian integral:

$$\int_{\mathbb{R}^d} dx \exp\left(-\frac{1}{2}\theta x' D x\right) = \left(\frac{2\pi}{\theta}\right)^{\frac{d}{2}} \frac{1}{\sqrt{\det D}}. \tag{B.28}$$

B.3.2 The Gamma Distribution

We say that a random variable x with values on the positive real axis has a Gamma density, with $\mathbb{a}, \mathbb{b} > 0$, if its probability density is

$$\mathcal{P}(x) = \frac{1}{\Gamma(\mathbb{a})} \mathbb{b}^{\mathbb{a}} x^{\mathbb{a}-1} e^{-\mathbb{b}x}. \tag{B.29}$$

We also write the preceding equation as

$$x \sim \mathcal{G}(\mathbb{a}, \mathbb{b}). \tag{B.30}$$

The mean and variance of the Gamma($\mathbb{a}, \mathbb{b}$) density are as follows:

$$\mathbb{E}[x] = \frac{\mathbb{a}}{\mathbb{b}}, \quad \mathbb{V}[x] = \frac{\mathbb{a}}{\mathbb{b}^2}. \tag{B.31}$$

B.3.3 The Log-Normal Distribution

We say that a random variable x with values on the positive real axis has a log-normal density with parameters ν and ϱ, and write $x \sim \log \mathcal{N}(\nu, \varrho^2)$, if its probability density is

$$\mathcal{P}(x) = \frac{1}{\sqrt{2\pi}\varrho x} \exp\left[-\frac{1}{2}\left(\frac{\log x - \nu}{\varrho}\right)^2\right]. \tag{B.32}$$

The log-normal density has the property that

$$x \sim \log \mathcal{N}(\nu, \varrho^2) \Longleftrightarrow \log x \sim \mathcal{N}(\nu, \varrho^2).$$

The mean and variance of the log-normal density are as follows:

$$\mathrm{E}[x] = e^{\nu + \frac{\varrho^2}{2}}, \quad \mathrm{V}[x] = e^{2(\nu+\varrho^2)} - e^{2\nu+\varrho^2}. \tag{B.33}$$

It is often useful to be able to express ν and ϱ as functions of the mean and the variance. This is done as:

$$\varrho^2 = \log\left(1 + \frac{\mathrm{V}[x]}{\mathrm{E}[x]^2}\right), \quad \nu = \log\left(\frac{\mathrm{E}[x]^2}{\sqrt{\mathrm{V}[x] + \mathrm{E}[x]^2}}\right). \tag{B.34}$$