1

Mathematical preliminaries

This chapter is intended to serve as a review of mathematical concepts to be used throughout this book, and also as a reference to be consulted as subsequent chapters are studied, if the need should arise. The first section focuses on linear algebra, and the second on analysis and related topics. Unlike the other chapters in this book, the present chapter does not include proofs, and is not intended to serve as a primary source for the material it reviews—a collection of references provided at the end of the chapter may be consulted by readers interested in a proper development of this material.

1.1 Linear algebra

The theory of quantum information relies heavily on linear algebra in finite-dimensional spaces. The subsections that follow present an overview of the aspects of this subject that are most relevant within the theory of quantum information. It is assumed that the reader is already familiar with the most basic notions of linear algebra, including those of linear dependence and independence, subspaces, spanning sets, bases, and dimension.

1.1.1 Complex Euclidean spaces

The notion of a complex Euclidean space is used throughout this book. One associates a complex Euclidean space with every discrete and finite system; and fundamental notions such as states and measurements of systems are represented in linear-algebraic terms that refer to these spaces.

Definition of complex Euclidean spaces

An *alphabet* is a finite and nonempty set, whose elements may be considered to be *symbols*. Alphabets will generally be denoted by capital Greek letters,

including Σ, Γ, and Λ, while lower case Roman letters near the beginning of the alphabet, including a, b, c, and d, will be used to denote symbols in alphabets. Examples of alphabets include the *binary alphabet* $\{0,1\}$, the n-fold Cartesian product $\{0,1\}^n$ of the binary alphabet with itself, and the alphabet $\{1,\ldots,n\}$, for n being a fixed positive integer.

For any alphabet Σ, one denotes by $\mathbb{C}^\Sigma$ the set of all functions from Σ to the complex numbers $\mathbb{C}$. The set $\mathbb{C}^\Sigma$ forms a vector space of dimension $|\Sigma|$ over the complex numbers when addition and scalar multiplication are defined in the following standard way:

1. Addition: for vectors $u, v \in \mathbb{C}^\Sigma$, the vector $u + v \in \mathbb{C}^\Sigma$ is defined by the equation $(u+v)(a) = u(a) + v(a)$ for all $a \in \Sigma$.
2. Scalar multiplication: for a vector $u \in \mathbb{C}^\Sigma$ and a scalar $\alpha \in \mathbb{C}$, the vector $\alpha u \in \mathbb{C}^\Sigma$ is defined by the equation $(\alpha u)(a) = \alpha u(a)$ for all $a \in \Sigma$.

A vector space defined in this way will be called a *complex Euclidean space*.[1] The value $u(a)$ is referred to as the *entry* of u indexed by a, for each $u \in \mathbb{C}^\Sigma$ and $a \in \Sigma$. The vector whose entries are all zero is simply denoted 0.

Complex Euclidean spaces will be denoted by scripted capital letters near the end of the alphabet, such as $\mathcal{W}$, $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$. Subsets of these spaces will also be denoted by scripted letters, and when possible this book will follow a convention to use letters such as $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$ near the beginning of the alphabet when these subsets are not necessarily vector spaces. Vectors will be denoted by lowercase Roman letters, again near the end of the alphabet, such as u, v, w, x, y, and z.

When n is a positive integer, one typically writes $\mathbb{C}^n$ rather than $\mathbb{C}^{\{1,\ldots,n\}}$, and it is also typical that one views a vector $u \in \mathbb{C}^n$ as an n-tuple of the form $u = (\alpha_1, \ldots, \alpha_n)$, or as a column vector of the form

$$u = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{pmatrix}, \tag{1.1}$$

for complex numbers $\alpha_1, \ldots, \alpha_n$.

For an arbitrary alphabet Σ, the complex Euclidean space $\mathbb{C}^\Sigma$ may be viewed as being equivalent to $\mathbb{C}^n$ for $n = |\Sigma|$; one simply fixes a bijection

$$f : \{1,\ldots,n\} \to \Sigma \tag{1.2}$$

and associates each vector $u \in \mathbb{C}^\Sigma$ with the vector in $\mathbb{C}^n$ whose k-th entry

[1] Many quantum information theorists prefer to use the term *Hilbert space*. The term *complex Euclidean space* will be preferred in this book, however, as the term *Hilbert space* refers to a more general notion that allows the possibility of infinite index sets.

is $u(f(k))$, for each $k \in \{1, \ldots, n\}$. This may be done implicitly when there is a natural or obviously preferred choice for the bijection f. For example, the elements of the alphabet $\Sigma = \{0,1\}^2$ are naturally ordered 00, 01, 10, 11. Each vector $u \in \mathbb{C}^\Sigma$ may therefore be associated with the 4-tuple

$$(u(00), u(01), u(10), u(11)), \tag{1.3}$$

or with the column vector

$$\begin{pmatrix} u(00) \\ u(01) \\ u(10) \\ u(11) \end{pmatrix}, \tag{1.4}$$

when it is convenient to do this. While little or no generality would be lost in restricting one's attention to complex Euclidean spaces of the form $\mathbb{C}^n$ for this reason, it is both natural and convenient within computational and information-theoretic settings to allow complex Euclidean spaces to be indexed by arbitrary alphabets.

Inner products and norms of vectors

The *inner product* $\langle u, v \rangle$ of two vectors $u, v \in \mathbb{C}^\Sigma$ is defined as

$$\langle u, v \rangle = \sum_{a \in \Sigma} \overline{u(a)}\, v(a). \tag{1.5}$$

It may be verified that the inner product satisfies the following properties:

1. Linearity in the second argument:

   $$\langle u, \alpha v + \beta w \rangle = \alpha \langle u, v \rangle + \beta \langle u, w \rangle \tag{1.6}$$

   for all $u, v, w \in \mathbb{C}^\Sigma$ and $\alpha, \beta \in \mathbb{C}$.
2. Conjugate symmetry:

   $$\langle u, v \rangle = \overline{\langle v, u \rangle} \tag{1.7}$$

   for all $u, v \in \mathbb{C}^\Sigma$.
3. Positive definiteness:

   $$\langle u, u \rangle \geq 0 \tag{1.8}$$

   for all $u \in \mathbb{C}^\Sigma$, with equality if and only if $u = 0$.

It is typical that any function satisfying these three properties is referred to as an inner product, but this is the only inner product for vectors in complex Euclidean spaces that is considered in this book.

The *Euclidean norm* of a vector $u \in \mathbb{C}^{\Sigma}$ is defined as

$$\|u\| = \sqrt{\langle u, u \rangle} = \sqrt{\sum_{a \in \Sigma} |u(a)|^2}. \tag{1.9}$$

The Euclidean norm possesses the following properties, which define the more general notion of a norm:

1. Positive definiteness: $\|u\| \geq 0$ for all $u \in \mathbb{C}^{\Sigma}$, with $\|u\| = 0$ if and only if $u = 0$.
2. Positive scalability: $\|\alpha u\| = |\alpha| \|u\|$ for all $u \in \mathbb{C}^{\Sigma}$ and $\alpha \in \mathbb{C}$.
3. The triangle inequality: $\|u + v\| \leq \|u\| + \|v\|$ for all $u, v \in \mathbb{C}^{\Sigma}$.

The *Cauchy–Schwarz inequality* states that

$$|\langle u, v \rangle| \leq \|u\| \, \|v\| \tag{1.10}$$

for all $u, v \in \mathbb{C}^{\Sigma}$, with equality if and only if u and v are linearly dependent. The collection of all unit vectors in a complex Euclidean space $\mathcal{X}$ is called the *unit sphere* in that space, and is denoted

$$\mathcal{S}(\mathcal{X}) = \{u \in \mathcal{X} \,:\, \|u\| = 1\}. \tag{1.11}$$

The Euclidean norm represents the case $p = 2$ of the class of *p-norms*, defined for each $u \in \mathbb{C}^{\Sigma}$ as

$$\|u\|_p = \left(\sum_{a \in \Sigma} |u(a)|^p \right)^{\frac{1}{p}} \tag{1.12}$$

for $p < \infty$, and

$$\|u\|_{\infty} = \max\{|u(a)| \,:\, a \in \Sigma\}. \tag{1.13}$$

The above three norm properties (positive definiteness, positive scalability, and the triangle inequality) hold for $\|\cdot\|$ replaced by $\|\cdot\|_p$ for any choice of $p \in [1, \infty]$.

Orthogonality and orthonormality

Two vectors $u, v \in \mathbb{C}^{\Sigma}$ are said to be *orthogonal* if $\langle u, v \rangle = 0$. The notation $u \perp v$ is also used to indicate that u and v are orthogonal. More generally, for any set $\mathcal{A} \subseteq \mathbb{C}^{\Sigma}$, the notation $u \perp \mathcal{A}$ indicates that $\langle u, v \rangle = 0$ for all vectors $v \in \mathcal{A}$.

A collection of vectors

$$\{u_a \,:\, a \in \Gamma\} \subset \mathbb{C}^{\Sigma}, \tag{1.14}$$

indexed by an alphabet Γ, is said to be an *orthogonal set* if it holds that

$\langle u_a, u_b \rangle = 0$ for all choices of $a, b \in \Gamma$ with $a \neq b$. A collection of nonzero orthogonal vectors is necessarily linearly independent.

An orthogonal set of *unit* vectors is called an *orthonormal set*, and when such a set forms a basis it is called an *orthonormal basis*. It holds that an orthonormal set of the form (1.14) is an orthonormal basis of $\mathbb{C}^{\Sigma}$ if and only if $|\Gamma| = |\Sigma|$. The *standard basis* of $\mathbb{C}^{\Sigma}$ is the orthonormal basis given by $\{e_a : a \in \Sigma\}$, where

$$e_a(b) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{if } a \neq b \end{cases} \tag{1.15}$$

for all $a, b \in \Sigma$.

Direct sums of complex Euclidean spaces

The *direct sum* of n complex Euclidean spaces $\mathcal{X}_1 = \mathbb{C}^{\Sigma_1}, \ldots, \mathcal{X}_n = \mathbb{C}^{\Sigma_n}$ is the complex Euclidean space

$$\mathcal{X}_1 \oplus \cdots \oplus \mathcal{X}_n = \mathbb{C}^{\Sigma_1 \sqcup \cdots \sqcup \Sigma_n}, \tag{1.16}$$

where $\Sigma_1 \sqcup \cdots \sqcup \Sigma_n$ denotes the *disjoint union* of the alphabets $\Sigma_1, \ldots, \Sigma_n$, defined as

$$\Sigma_1 \sqcup \cdots \sqcup \Sigma_n = \bigcup_{k \in \{1,\ldots,n\}} \{(k, a) : a \in \Sigma_k\}. \tag{1.17}$$

For vectors $u_1 \in \mathcal{X}_1, \ldots, u_n \in \mathcal{X}_n$, the notation $u_1 \oplus \cdots \oplus u_n \in \mathcal{X}_1 \oplus \cdots \oplus \mathcal{X}_n$ refers to the vector for which

$$(u_1 \oplus \cdots \oplus u_n)(k, a) = u_k(a), \tag{1.18}$$

for each $k \in \{1, \ldots, n\}$ and $a \in \Sigma_k$. If each u_k is viewed as a column vector of dimension $|\Sigma_k|$, the vector $u_1 \oplus \cdots \oplus u_n$ may be viewed as a column vector

$$\begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix} \tag{1.19}$$

having dimension $|\Sigma_1| + \cdots + |\Sigma_n|$.

Every element of the space $\mathcal{X}_1 \oplus \cdots \oplus \mathcal{X}_n$ can be written as $u_1 \oplus \cdots \oplus u_n$ for a unique choice of vectors $u_1, \ldots, u_n$. The following identities hold for

every choice of $u_1, v_1 \in \mathcal{X}_1, \ldots, u_n, v_n \in \mathcal{X}_n$, and $\alpha \in \mathbb{C}$:

$$u_1 \oplus \cdots \oplus u_n + v_1 \oplus \cdots \oplus v_n = (u_1 + v_1) \oplus \cdots \oplus (u_n + v_n), \tag{1.20}$$

$$\alpha(u_1 \oplus \cdots \oplus u_n) = (\alpha u_1) \oplus \cdots \oplus (\alpha u_n), \tag{1.21}$$

$$\langle u_1 \oplus \cdots \oplus u_n, v_1 \oplus \cdots \oplus v_n \rangle = \langle u_1, v_1 \rangle + \cdots + \langle u_n, v_n \rangle. \tag{1.22}$$

Tensor products of complex Euclidean spaces

The *tensor product* of n complex Euclidean spaces $\mathcal{X}_1 = \mathbb{C}^{\Sigma_1}, \ldots, \mathcal{X}_n = \mathbb{C}^{\Sigma_n}$ is the complex Euclidean space

$$\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n = \mathbb{C}^{\Sigma_1 \times \cdots \times \Sigma_n}. \tag{1.23}$$

For vectors $u_1 \in \mathcal{X}_1, \ldots, u_n \in \mathcal{X}_n$, the notation $u_1 \otimes \cdots \otimes u_n \in \mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n$ refers to the vector for which

$$(u_1 \otimes \cdots \otimes u_n)(a_1, \ldots, a_n) = u_1(a_1) \cdots u_n(a_n). \tag{1.24}$$

Vectors of the form $u_1 \otimes \cdots \otimes u_n$ are called *elementary tensors*. They span the space $\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n$, but not every element of $\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n$ is an elementary tensor.

The following identities hold for all vectors $u_1, v_1 \in \mathcal{X}_1, \ldots, u_n, v_n \in \mathcal{X}_n$, scalars $\alpha, \beta \in \mathbb{C}$, and indices $k \in \{1, \ldots, n\}$:

$$\begin{aligned} &u_1 \otimes \cdots \otimes u_{k-1} \otimes (\alpha u_k + \beta v_k) \otimes u_{k+1} \otimes \cdots \otimes u_n \\ &\qquad = \alpha \left(u_1 \otimes \cdots \otimes u_{k-1} \otimes u_k \otimes u_{k+1} \otimes \cdots \otimes u_n\right) \\ &\qquad + \beta \left(u_1 \otimes \cdots \otimes u_{k-1} \otimes v_k \otimes u_{k+1} \otimes \cdots \otimes u_n\right), \end{aligned} \tag{1.25}$$

$$\langle u_1 \otimes \cdots \otimes u_n, v_1 \otimes \cdots \otimes v_n \rangle = \langle u_1, v_1 \rangle \cdots \langle u_n, v_n \rangle. \tag{1.26}$$

Tensor products are often defined in a way that is more abstract (and more generally applicable) than the definition above, which is sometimes known more specifically as the *Kronecker product*. The following proposition is a reflection of the more abstract definition.

Proposition 1.1 *Let $\mathcal{X}_1, \ldots, \mathcal{X}_n$ and $\mathcal{Y}$ be complex Euclidean spaces and let*

$$\phi : \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \to \mathcal{Y} \tag{1.27}$$

be a multilinear function, meaning a function for which the mapping

$$u_k \mapsto \phi(u_1, \ldots, u_n) \tag{1.28}$$

is linear for all $k \in \{1, \ldots, n\}$ *and every fixed choice of vectors* $u_1, \ldots, u_{k-1}$, $u_{k+1}, \ldots, u_n$. *There exists a unique linear mapping*

$$A : \mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n \to \mathcal{Y} \tag{1.29}$$

such that

$$\phi(u_1, \ldots, u_n) = A(u_1 \otimes \cdots \otimes u_n) \tag{1.30}$$

for all choices of $u_1 \in \mathcal{X}_1, \ldots, u_n \in \mathcal{X}_n$.

If $\mathcal{X}$ is a complex Euclidean space, $u \in \mathcal{X}$ is a vector, and n is a positive integer, then the notations $\mathcal{X}^{\otimes n}$ and $u^{\otimes n}$ refer to the n-fold tensor product of either $\mathcal{X}$ or u with itself. It is often convenient to make the identification

$$\mathcal{X}^{\otimes n} = \mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n, \tag{1.31}$$

under the assumption that $\mathcal{X}_1, \ldots, \mathcal{X}_n$ and $\mathcal{X}$ all refer to the same complex Euclidean space; this allows one to refer to the different tensor factors in $\mathcal{X}^{\otimes n}$ individually, and to express $\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n$ more concisely.

Remark A rigid interpretation of the definitions above suggests that tensor products of complex Euclidean spaces (or of vectors in complex Euclidean spaces) are not associative, insofar as Cartesian products are not associative. For instance, given alphabets Σ, Γ, and Λ, the alphabet $(\Sigma \times \Gamma) \times \Lambda$ contains elements of the form $((a, b), c)$, the alphabet $\Sigma \times (\Gamma \times \Lambda)$ contains elements of the form $(a, (b, c))$, and the alphabet $\Sigma \times \Gamma \times \Lambda$ contains elements of the form (a, b, c), for $a \in \Sigma$, $b \in \Gamma$, and $c \in \Lambda$. For $\mathcal{X} = \mathbb{C}^{\Sigma}$, $\mathcal{Y} = \mathbb{C}^{\Gamma}$, and $\mathcal{Z} = \mathbb{C}^{\Lambda}$, one may therefore view the complex Euclidean spaces $(\mathcal{X} \otimes \mathcal{Y}) \otimes \mathcal{Z}$, $\mathcal{X} \otimes (\mathcal{Y} \otimes \mathcal{Z})$, and $\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z}$ as being different.

However, the alphabets $(\Sigma \times \Gamma) \times \Lambda$, $\Sigma \times (\Gamma \times \Lambda)$, and $\Sigma \times \Gamma \times \Lambda$ can of course be viewed as equivalent by simply removing parentheses. For this reason, there is a natural equivalence between the complex Euclidean spaces $(\mathcal{X} \otimes \mathcal{Y}) \otimes \mathcal{Z}$, $\mathcal{X} \otimes (\mathcal{Y} \otimes \mathcal{Z})$, and $\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z}$. Whenever it is convenient, identifications of this sort are made implicitly throughout this book. For example, given vectors $u \in \mathcal{X} \otimes \mathcal{Y}$ and $v \in \mathcal{Z}$, the vector $u \otimes v$ may be treated as an element of $\mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z}$ rather than $(\mathcal{X} \otimes \mathcal{Y}) \otimes \mathcal{Z}$.

Although such instances are much less common in this book, a similar convention applies to direct sums of complex Euclidean spaces.

Real Euclidean spaces

Real Euclidean spaces are defined in a similar way to complex Euclidean spaces, except that the field of complex numbers $\mathbb{C}$ is replaced by the field of real numbers $\mathbb{R}$ in each of the definitions and concepts in which it arises.

Naturally, complex conjugation acts trivially in the real case, and therefore may be omitted.

Complex Euclidean spaces will play a more prominent role than real ones in this book. Real Euclidean spaces will, nevertheless, be important in those settings that make use of concepts from the theory of convexity. The space of Hermitian operators acting on a given complex Euclidean space is an important example of a real vector space that can be identified with a real Euclidean space, as is discussed in the subsection following this one.

1.1.2 Linear operators

Given complex Euclidean spaces $\mathcal{X}$ and $\mathcal{Y}$, one writes $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ to refer to the collection of all linear mappings of the form

$$A : \mathcal{X} \to \mathcal{Y}. \tag{1.32}$$

Such mappings will be referred to as *linear operators*, or simply *operators*, from $\mathcal{X}$ to $\mathcal{Y}$ in this book. Parentheses are omitted when expressing the action of linear operators on vectors when no confusion arises in doing so. For instance, one writes Au rather than $A(u)$ to denote the vector resulting from the application of an operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ to a vector $u \in \mathcal{X}$.

The set $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ forms a complex vector space when addition and scalar multiplication are defined as follows:

1. Addition: for operators $A, B \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, the operator $A + B \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ is defined by the equation
$$(A + B)u = Au + Bu \tag{1.33}$$
for all $u \in \mathcal{X}$.
2. Scalar multiplication: for an operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ and a scalar $\alpha \in \mathbb{C}$, the operator $\alpha A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ is defined by the equation
$$(\alpha A)u = \alpha Au \tag{1.34}$$
for all $u \in \mathcal{X}$.

Matrices and their correspondence with operators

A *matrix* over the complex numbers is a mapping of the form

$$M : \Gamma \times \Sigma \to \mathbb{C} \tag{1.35}$$

for alphabets Σ and Γ. For $a \in \Gamma$ and $b \in \Sigma$ the value $M(a, b)$ is called the (a, b) *entry* of M, and the elements a and b are referred to as *indices* in this

context: a is the *row index* and b is the *column index* of the entry $M(a, b)$. Addition and scalar multiplication of matrices are defined in a similar way to vectors in complex Euclidean spaces:

1. Addition: for matrices $M : \Gamma \times \Sigma \to \mathbb{C}$ and $N : \Gamma \times \Sigma \to \mathbb{C}$, the matrix $M + N$ is defined as
$$(M + N)(a, b) = M(a, b) + N(a, b) \tag{1.36}$$
for all $a \in \Gamma$ and $b \in \Sigma$.
2. Scalar multiplication: for a matrix $M : \Gamma \times \Sigma \to \mathbb{C}$ and a scalar $\alpha \in \mathbb{C}$, the matrix αM is defined as
$$(\alpha M)(a, b) = \alpha M(a, b) \tag{1.37}$$
for all $a \in \Gamma$ and $b \in \Sigma$.

In addition, one defines matrix multiplication as follows:

3. Matrix multiplication: for matrices $M : \Gamma \times \Lambda \to \mathbb{C}$ and $N : \Lambda \times \Sigma \to \mathbb{C}$, the matrix $MN : \Gamma \times \Sigma \to \mathbb{C}$ is defined as
$$(MN)(a, b) = \sum_{c \in \Lambda} M(a, c)N(c, b) \tag{1.38}$$
for all $a \in \Gamma$ and $b \in \Sigma$.

For any choice of complex Euclidean spaces $\mathcal{X} = \mathbb{C}^{\Sigma}$ and $\mathcal{Y} = \mathbb{C}^{\Gamma}$, there is a bijective linear correspondence between the set of operators $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ and the collection of all matrices taking the form $M : \Gamma \times \Sigma \to \mathbb{C}$ that is obtained as follows. With each operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, one associates the matrix M defined as
$$M(a, b) = \langle e_a, Ae_b \rangle \tag{1.39}$$
for $a \in \Gamma$ and $b \in \Sigma$. The operator A is uniquely determined by M, and may be recovered from M by the equation
$$(Au)(a) = \sum_{b \in \Sigma} M(a, b)u(b) \tag{1.40}$$
for all $a \in \Gamma$. With respect to this correspondence, matrix multiplication is equivalent to operator composition.

Hereafter in this book, linear operators will be associated with matrices implicitly, without the introduction of names that distinguish matrices from the operators with which they are associated. With this in mind, the notation
$$A(a, b) = \langle e_a, Ae_b \rangle \tag{1.41}$$

is introduced for each $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, $a \in \Gamma$, and $b \in \Sigma$ (where it is to be assumed that $\mathcal{X} = \mathbb{C}^{\Sigma}$ and $\mathcal{Y} = \mathbb{C}^{\Gamma}$, as above).

The standard basis of a space of operators

For every choice of complex Euclidean spaces $\mathcal{X} = \mathbb{C}^{\Sigma}$ and $\mathcal{Y} = \mathbb{C}^{\Gamma}$, and each choice of symbols $a \in \Gamma$ and $b \in \Sigma$, the operator $E_{a,b} \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ is defined as

$$E_{a,b}\, u = u(b) e_a \tag{1.42}$$

for every $u \in \mathcal{X}$. Equivalently, $E_{a,b}$ is defined by the equation

$$E_{a,b}(c,d) = \begin{cases} 1 & \text{if } (c,d) = (a,b) \\ 0 & \text{otherwise} \end{cases} \tag{1.43}$$

holding for all $c \in \Gamma$ and $d \in \Sigma$. The collection

$$\{E_{a,b} \,:\, a \in \Gamma,\ b \in \Sigma\} \tag{1.44}$$

forms a basis of $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ known as the *standard basis* of this space. The number of elements in this basis is, of course, consistent with the fact that the dimension of $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ is given by $\dim(\mathrm{L}(\mathcal{X}, \mathcal{Y})) = \dim(\mathcal{X}) \dim(\mathcal{Y})$.

The entry-wise conjugate, transpose, and adjoint

For every operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, for complex Euclidean spaces $\mathcal{X} = \mathbb{C}^{\Sigma}$ and $\mathcal{Y} = \mathbb{C}^{\Gamma}$, one defines three additional operators,

$$\overline{A} \in \mathrm{L}(\mathcal{X}, \mathcal{Y}) \quad \text{and} \quad A^{\mathsf{T}}, A^{*} \in \mathrm{L}(\mathcal{Y}, \mathcal{X}), \tag{1.45}$$

as follows:

1. The operator $\overline{A} \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ is the operator whose matrix representation has entries that are complex conjugates to the matrix representation of A:

   $$\overline{A}(a,b) = \overline{A(a,b)} \tag{1.46}$$

   for all $a \in \Gamma$ and $b \in \Sigma$.
2. The operator $A^{\mathsf{T}} \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$ is the operator whose matrix representation is obtained by *transposing* the matrix representation of A:

   $$A^{\mathsf{T}}(b,a) = A(a,b) \tag{1.47}$$

   for all $a \in \Gamma$ and $b \in \Sigma$.

3. The operator $A^* \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$ is the uniquely determined operator that satisfies the equation

$$\langle v, Au \rangle = \langle A^* v, u \rangle \tag{1.48}$$

for all $u \in \mathcal{X}$ and $v \in \mathcal{Y}$. It may be obtained by performing both of the operations described in items 1 and 2:

$$A^* = \overline{A^{\mathsf{T}}}. \tag{1.49}$$

The operators $\overline{A}$, A^{T}, and A^* are called the *entry-wise conjugate*, *transpose*, and *adjoint* operators to A, respectively.

The mappings $A \mapsto \overline{A}$ and $A \mapsto A^*$ are conjugate linear and $A \mapsto A^{\mathsf{T}}$ is linear:

$$\begin{aligned}
\overline{\alpha A + \beta B} &= \overline{\alpha}\,\overline{A} + \overline{\beta}\,\overline{B}, \\
(\alpha A + \beta B)^* &= \overline{\alpha} A^* + \overline{\beta} B^*, \\
(\alpha A + \beta B)^{\mathsf{T}} &= \alpha A^{\mathsf{T}} + \beta B^{\mathsf{T}},
\end{aligned}$$

for all $A, B \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ and $\alpha, \beta \in \mathbb{C}$. These mappings are bijections, each being its own inverse.

Each vector $u \in \mathcal{X}$ in a complex Euclidean space $\mathcal{X}$ may be identified with the linear operator in $\mathrm{L}(\mathbb{C}, \mathcal{X})$ defined as $\alpha \mapsto \alpha u$ for all $\alpha \in \mathbb{C}$. Through this identification, the linear mappings $\overline{u} \in \mathrm{L}(\mathbb{C}, \mathcal{X})$ and $u^{\mathsf{T}}, u^* \in \mathrm{L}(\mathcal{X}, \mathbb{C})$ are defined as above. As an element of $\mathcal{X}$, the vector $\overline{u}$ is simply the entry-wise complex conjugate of u, i.e., if $\mathcal{X} = \mathbb{C}^\Sigma$ then

$$\overline{u}(a) = \overline{u(a)} \tag{1.50}$$

for every $a \in \Sigma$. For each vector $u \in \mathcal{X}$ the mapping $u^* \in \mathrm{L}(\mathcal{X}, \mathbb{C})$ satisfies $u^* v = \langle u, v \rangle$ for all $v \in \mathcal{X}$.

Kernel, image, and rank

The *kernel* of an operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ is the subspace of $\mathcal{X}$ defined as

$$\ker(A) = \{u \in \mathcal{X} \,:\, Au = 0\}, \tag{1.51}$$

while the *image* of A is the subspace of $\mathcal{Y}$ defined as

$$\mathrm{im}(A) = \{Au \,:\, u \in \mathcal{X}\}. \tag{1.52}$$

For every operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, one has that

$$\ker(A) = \ker(A^* A) \quad \text{and} \quad \mathrm{im}(A) = \mathrm{im}(AA^*), \tag{1.53}$$

as well as the equation

$$\dim(\ker(A)) + \dim(\mathrm{im}(A)) = \dim(\mathcal{X}). \tag{1.54}$$

The *rank* of an operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, denoted $\mathrm{rank}(A)$, is the dimension of the image of A:

$$\mathrm{rank}(A) = \dim(\mathrm{im}(A)). \tag{1.55}$$

By (1.53) and (1.54), one may conclude that

$$\mathrm{rank}(A) = \mathrm{rank}(AA^*) = \mathrm{rank}(A^*A) \tag{1.56}$$

for every $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$.

For any choice of vectors $u \in \mathcal{X}$ and $v \in \mathcal{Y}$, the operator $vu^* \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ satisfies

$$(vu^*)w = v(u^*w) = \langle u, w \rangle v \tag{1.57}$$

for all $w \in \mathcal{X}$. Assuming that u and v are nonzero, the operator vu^* has rank equal to one, and every rank one operator in $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ can be expressed in this form for vectors u and v that are unique up to scalar multiples.

Operators involving direct sums of complex Euclidean spaces

Suppose that

$$\mathcal{X}_1 = \mathbb{C}^{\Sigma_1}, \ldots, \mathcal{X}_n = \mathbb{C}^{\Sigma_n} \quad \text{and} \quad \mathcal{Y}_1 = \mathbb{C}^{\Gamma_1}, \ldots, \mathcal{Y}_m = \mathbb{C}^{\Gamma_m} \tag{1.58}$$

are complex Euclidean spaces, for alphabets $\Sigma_1, \ldots, \Sigma_n$ and $\Gamma_1, \ldots, \Gamma_m$. For a given operator

$$A \in \mathrm{L}(\mathcal{X}_1 \oplus \cdots \oplus \mathcal{X}_n, \mathcal{Y}_1 \oplus \cdots \oplus \mathcal{Y}_m), \tag{1.59}$$

there exists a unique collection of operators

$$\bigl\{A_{j,k} \in \mathrm{L}(\mathcal{X}_k, \mathcal{Y}_j) \,:\, 1 \leq j \leq m,\, 1 \leq k \leq n\bigr\} \tag{1.60}$$

for which the equation

$$A_{j,k}(a, b) = A\bigl((j, a), (k, b)\bigr) \tag{1.61}$$

holds for all $j \in \{1, \ldots, m\}$, $k \in \{1, \ldots, n\}$, $a \in \Gamma_j$, and $b \in \Sigma_k$. For all vectors $u_1 \in \mathcal{X}_1, \ldots, u_n \in \mathcal{X}_n$, one has that

$$A(u_1 \oplus \cdots \oplus u_n) = v_1 \oplus \cdots \oplus v_m \tag{1.62}$$

for $v_1 \in \mathcal{Y}_1, \ldots, v_m \in \mathcal{Y}_m$ being defined as

$$v_j = \sum_{k=1}^{n} A_{j,k} u_k \tag{1.63}$$

for each $j \in \{1, \ldots, m\}$. Conversely, for any collection of operators of the form (1.60), there is a unique operator A of the form (1.59) that obeys the equations (1.62) and (1.63) for all vectors $u_1 \in \mathcal{X}_1, \ldots, u_n \in \mathcal{X}_n$.

There is therefore a bijective correspondence between operators of the form (1.59) and collections of operators of the form (1.60). With respect to the matrix representations of these operators, this correspondence may be expressed succinctly as

$$A = \begin{pmatrix} A_{1,1} & \cdots & A_{1,n} \\ \vdots & \ddots & \vdots \\ A_{m,1} & \cdots & A_{m,n} \end{pmatrix}. \tag{1.64}$$

One interprets the right-hand side of (1.64) as the specification of the operator having the form (1.59) that is defined by the collection (1.60) in this way.

Tensor products of operators

Suppose that

$$\mathcal{X}_1 = \mathbb{C}^{\Sigma_1}, \ldots, \mathcal{X}_n = \mathbb{C}^{\Sigma_n} \quad \text{and} \quad \mathcal{Y}_1 = \mathbb{C}^{\Gamma_1}, \ldots, \mathcal{Y}_n = \mathbb{C}^{\Gamma_n} \tag{1.65}$$

are complex Euclidean spaces, for alphabets $\Sigma_1, \ldots, \Sigma_n$ and $\Gamma_1, \ldots, \Gamma_n$. For any choice of operators

$$A_1 \in \mathrm{L}(\mathcal{X}_1, \mathcal{Y}_1), \ \ldots, \ A_n \in \mathrm{L}(\mathcal{X}_n, \mathcal{Y}_n), \tag{1.66}$$

one defines the tensor product

$$A_1 \otimes \cdots \otimes A_n \in \mathrm{L}(\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n, \mathcal{Y}_1 \otimes \cdots \otimes \mathcal{Y}_n) \tag{1.67}$$

of these operators to be the unique operator that satisfies the equation

$$(A_1 \otimes \cdots \otimes A_n)(u_1 \otimes \cdots \otimes u_n) = (A_1 u_1) \otimes \cdots \otimes (A_n u_n) \tag{1.68}$$

for all choices of $u_1 \in \mathcal{X}_1, \ldots, u_n \in \mathcal{X}_n$. This operator may equivalently be defined in terms of its matrix representation as

$$\begin{aligned} (A_1 \otimes \cdots \otimes A_n)((a_1, \ldots, a_n), (b_1, \ldots, b_n)) \\ = A_1(a_1, b_1) \cdots A_n(a_n, b_n) \end{aligned} \tag{1.69}$$

for all $a_1 \in \Gamma_1, \ldots, a_n \in \Gamma_n$ and $b_1 \in \Sigma_1, \ldots, b_n \in \Sigma_n$.

For every choice of complex Euclidean spaces $\mathcal{X}_1, \ldots, \mathcal{X}_n$, $\mathcal{Y}_1, \ldots, \mathcal{Y}_n$, and $\mathcal{Z}_1, \ldots, \mathcal{Z}_n$, operators

$$\begin{aligned} A_1, B_1 \in \mathrm{L}(\mathcal{X}_1, \mathcal{Y}_1), \ \ldots, \ A_n, B_n \in \mathrm{L}(\mathcal{X}_n, \mathcal{Y}_n), \\ C_1 \in \mathrm{L}(\mathcal{Y}_1, \mathcal{Z}_1), \ \ldots, \ C_n \in \mathrm{L}(\mathcal{Y}_n, \mathcal{Z}_n), \end{aligned} \tag{1.70}$$

and scalars $\alpha, \beta \in \mathbb{C}$, the following equations hold:

$$\begin{aligned} &A_1 \otimes \cdots \otimes A_{k-1} \otimes (\alpha A_k + \beta B_k) \otimes A_{k+1} \otimes \cdots \otimes A_n \\ &\quad = \alpha(A_1 \otimes \cdots \otimes A_{k-1} \otimes A_k \otimes A_{k+1} \otimes \cdots \otimes A_n) \\ &\quad + \beta(A_1 \otimes \cdots \otimes A_{k-1} \otimes B_k \otimes A_{k+1} \otimes \cdots \otimes A_n), \end{aligned} \tag{1.71}$$

$$(C_1 \otimes \cdots \otimes C_n)(A_1 \otimes \cdots \otimes A_n) = (C_1 A_1) \otimes \cdots \otimes (C_n A_n), \tag{1.72}$$

$$(A_1 \otimes \cdots \otimes A_n)^{\mathsf{T}} = A_1^{\mathsf{T}} \otimes \cdots \otimes A_n^{\mathsf{T}}, \tag{1.73}$$

$$\overline{A_1 \otimes \cdots \otimes A_n} = \overline{A_1} \otimes \cdots \otimes \overline{A_n}, \tag{1.74}$$

$$(A_1 \otimes \cdots \otimes A_n)^{*} = A_1^{*} \otimes \cdots \otimes A_n^{*}. \tag{1.75}$$

Similar to vectors, for an operator A and a positive integer n, the notation $A^{\otimes n}$ refers to the n-fold tensor product of A with itself.

Square operators

For every complex Euclidean space $\mathcal{X}$, the notation $\mathrm{L}(\mathcal{X})$ is understood to be a shorthand for $\mathrm{L}(\mathcal{X}, \mathcal{X})$. Operators in the space $\mathrm{L}(\mathcal{X})$ will be called *square operators*, due to the fact that their matrix representations are square, with rows and columns indexed by the same set.

The space $\mathrm{L}(\mathcal{X})$ is an *associative algebra*; in addition to being a vector space, the composition of square operators is associative and bilinear:

$$\begin{aligned} (XY)Z &= X(YZ), \\ Z(\alpha X + \beta Y) &= \alpha ZX + \beta ZY, \\ (\alpha X + \beta Y)Z &= \alpha XZ + \beta YZ, \end{aligned} \tag{1.76}$$

for every choice of $X, Y, Z \in \mathrm{L}(\mathcal{X})$ and $\alpha, \beta \in \mathbb{C}$.

The *identity operator* $\mathbb{1} \in \mathrm{L}(\mathcal{X})$ is the operator defined as $\mathbb{1}u = u$ for all $u \in \mathcal{X}$. It may also be defined by its matrix representation as

$$\mathbb{1}(a, b) = \begin{cases} 1 & \text{if } a = b \\ 0 & \text{if } a \neq b \end{cases} \tag{1.77}$$

for all $a, b \in \Sigma$, assuming $\mathcal{X} = \mathbb{C}^{\Sigma}$. One writes $\mathbb{1}_{\mathcal{X}}$ rather than $\mathbb{1}$ when it is helpful to indicate explicitly that this operator acts on $\mathcal{X}$.

For a complex Euclidean space $\mathcal{X}$, an operator $X \in \mathrm{L}(\mathcal{X})$ is *invertible* if there exists an operator $Y \in \mathrm{L}(\mathcal{X})$ such that $YX = \mathbb{1}$. When such an operator Y exists it is necessarily unique and is denoted X^{-1}. When the inverse X^{-1} of X exists, it must also satisfy $XX^{-1} = \mathbb{1}$.

Trace and determinant

The *diagonal* entries of a square operator $X \in \mathrm{L}(\mathcal{X})$, for $\mathcal{X} = \mathbb{C}^{\Sigma}$, are those of the form $X(a, a)$ for $a \in \Sigma$. The *trace* of a square operator $X \in \mathrm{L}(\mathcal{X})$ is defined as the sum of its diagonal entries:

$$\mathrm{Tr}(X) = \sum_{a \in \Sigma} X(a, a). \tag{1.78}$$

Alternatively, the trace is the unique linear function $\mathrm{Tr} : \mathrm{L}(\mathcal{X}) \to \mathbb{C}$ such that, for all vectors $u, v \in \mathcal{X}$, one has

$$\mathrm{Tr}(uv^*) = \langle v, u \rangle. \tag{1.79}$$

For every choice of complex Euclidean spaces $\mathcal{X}$ and $\mathcal{Y}$ and operators $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ and $B \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$, it holds that

$$\mathrm{Tr}(AB) = \mathrm{Tr}(BA). \tag{1.80}$$

This property is known as the *cyclic property* of the trace.

By means of the trace, one defines an inner product on the space $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ as follows:

$$\langle A, B \rangle = \mathrm{Tr}\big(A^* B\big) \tag{1.81}$$

for all $A, B \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$. It may be verified that this inner product satisfies the requisite properties of being an inner product:

1. Linearity in the second argument:

   $$\langle A, \alpha B + \beta C \rangle = \alpha \langle A, B \rangle + \beta \langle A, C \rangle \tag{1.82}$$

   for all $A, B, C \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ and $\alpha, \beta \in \mathbb{C}$.
2. Conjugate symmetry:

   $$\langle A, B \rangle = \overline{\langle B, A \rangle} \tag{1.83}$$

   for all $A, B \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$.
3. Positive definiteness: $\langle A, A \rangle \geq 0$ for all $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, with equality if and only if $A = 0$.

The *determinant* of a square operator $X \in \mathrm{L}(\mathcal{X})$, for $\mathcal{X} = \mathbb{C}^{\Sigma}$, is defined by the equation

$$\mathrm{Det}(X) = \sum_{\pi \in \mathrm{Sym}(\Sigma)} \mathrm{sign}(\pi) \prod_{a \in \Sigma} X(a, \pi(a)). \tag{1.84}$$

Here, the set $\mathrm{Sym}(\Sigma)$ denotes the collection of all permutations $\pi : \Sigma \to \Sigma$,

and $\operatorname{sign}(\pi) \in \{-1, +1\}$ denotes the sign (or parity) of the permutation π. The determinant is multiplicative,

$$\operatorname{Det}(XY) = \operatorname{Det}(X)\operatorname{Det}(Y) \tag{1.85}$$

for all $X, Y \in \mathrm{L}(\mathcal{X})$, and $\operatorname{Det}(X) \neq 0$ if and only if X is invertible.

Eigenvectors and eigenvalues

If $X \in \mathrm{L}(\mathcal{X})$ is an operator and $u \in \mathcal{X}$ is a nonzero vector for which it holds that

$$Xu = \lambda u \tag{1.86}$$

for some choice of $\lambda \in \mathbb{C}$, then u is said to be an *eigenvector* of X and λ is its corresponding *eigenvalue*.

For every operator $X \in \mathrm{L}(\mathcal{X})$, one has that

$$p_X(\alpha) = \operatorname{Det}(\alpha \mathbb{1}_{\mathcal{X}} - X) \tag{1.87}$$

is a monic polynomial in the variable α having degree $\dim(\mathcal{X})$, known as the *characteristic polynomial* of X. The *spectrum* of X, denoted $\operatorname{spec}(X)$, is the multiset containing the roots of the polynomial p_X, where each root appears a number of times equal to its multiplicity. As p_X is monic, it holds that

$$p_X(\alpha) = \prod_{\lambda \in \operatorname{spec}(X)} (\alpha - \lambda). \tag{1.88}$$

Each element $\lambda \in \operatorname{spec}(X)$ is necessarily an eigenvalue of X, and every eigenvalue of X is contained in $\operatorname{spec}(X)$.

The trace and determinant may be expressed in terms of the spectrum as follows:

$$\operatorname{Tr}(X) = \sum_{\lambda \in \operatorname{spec}(X)} \lambda \quad \text{and} \quad \operatorname{Det}(X) = \prod_{\lambda \in \operatorname{spec}(X)} \lambda \tag{1.89}$$

for every $X \in \mathrm{L}(\mathcal{X})$. The *spectral radius* of an operator $X \in \mathrm{L}(\mathcal{X})$ is the maximum absolute value $|\lambda|$ taken over all eigenvalues λ of X. For every choice of operators $X, Y \in \mathrm{L}(\mathcal{X})$ it holds that

$$\operatorname{spec}(XY) = \operatorname{spec}(YX). \tag{1.90}$$

Lie brackets and commutants

A set $\mathcal{A} \subseteq \mathrm{L}(\mathcal{X})$ is a *subalgebra* of $\mathrm{L}(\mathcal{X})$ if it is closed under addition, scalar multiplication, and operator composition:

$$X + Y \in \mathcal{A}, \quad \alpha X \in \mathcal{A}, \quad \text{and} \quad XY \in \mathcal{A} \tag{1.91}$$

for all $X, Y \in \mathcal{A}$ and $\alpha \in \mathbb{C}$. A subalgebra $\mathcal{A}$ of $\mathrm{L}(\mathcal{X})$ is said to be *self-adjoint* if it holds that $X^* \in \mathcal{A}$ for every $X \in \mathcal{A}$, and is said to be *unital* if it holds that $\mathbb{1} \in \mathcal{A}$.

For any pair of operators $X, Y \in \mathrm{L}(\mathcal{X})$, the *Lie bracket* $[X, Y] \in \mathrm{L}(\mathcal{X})$ is defined as

$$[X, Y] = XY - YX. \tag{1.92}$$

It holds that $[X, Y] = 0$ if and only if X and Y *commute*: $XY = YX$. For any subset of operators $\mathcal{A} \subseteq \mathrm{L}(\mathcal{X})$, one defines the *commutant* of $\mathcal{A}$ as

$$\mathrm{comm}(\mathcal{A}) = \{Y \in \mathrm{L}(\mathcal{X}) \,:\, [X, Y] = 0 \text{ for all } X \in \mathcal{A}\}. \tag{1.93}$$

The commutant of every subset of $\mathrm{L}(\mathcal{X})$ is a unital subalgebra of $\mathrm{L}(\mathcal{X})$.

Important classes of operators

The following classes of operators have particular importance in the theory of quantum information:

1. *Normal operators.* An operator $X \in \mathrm{L}(\mathcal{X})$ is *normal* if it commutes with its adjoint: $[X, X^*] = 0$, or equivalently, $XX^* = X^*X$. The importance of this collection of operators, for the purposes of this book, is mainly derived from two facts: (1) the normal operators are those for which the spectral theorem (discussed later in Section 1.1.3) holds, and (2) most of the special classes of operators that are discussed below are subsets of the normal operators.
2. *Hermitian operators.* An operator $X \in \mathrm{L}(\mathcal{X})$ is *Hermitian* if $X = X^*$. The set of Hermitian operators acting on a complex Euclidean space $\mathcal{X}$ will hereafter be denoted $\mathrm{Herm}(\mathcal{X})$ in this book:

   $$\mathrm{Herm}(\mathcal{X}) = \{X \in \mathrm{L}(\mathcal{X}) \,:\, X = X^*\}. \tag{1.94}$$

   Every Hermitian operator is a normal operator.
3. *Positive semidefinite operators.* An operator $X \in \mathrm{L}(\mathcal{X})$ is *positive semidefinite* if it holds that $X = Y^*Y$ for some operator $Y \in \mathrm{L}(\mathcal{X})$. Positive semidefinite operators will, as a convention, often be denoted by the letters P, Q, and R in this book. The collection of positive semidefinite operators acting on $\mathcal{X}$ is denoted $\mathrm{Pos}(\mathcal{X})$, so that

   $$\mathrm{Pos}(\mathcal{X}) = \{Y^*Y \,:\, Y \in \mathrm{L}(\mathcal{X})\}. \tag{1.95}$$

   Every positive semidefinite operator is Hermitian.

4. *Positive definite operators.* A positive semidefinite operator $P \in \mathrm{Pos}(\mathcal{X})$ is said to be *positive definite* if, in addition to being positive semidefinite, it is invertible. The notation
$$\mathrm{Pd}(\mathcal{X}) = \{P \in \mathrm{Pos}(\mathcal{X}) \,:\, \mathrm{Det}(P) \neq 0\} \tag{1.96}$$
will be used to denote the set of such operators for a complex Euclidean space $\mathcal{X}$.
5. *Density operators.* Positive semidefinite operators having trace equal to 1 are called *density operators*. Lowercase Greek letters, such as ρ, ξ, and σ, are conventionally used to denote density operators. The notation
$$\mathrm{D}(\mathcal{X}) = \{\rho \in \mathrm{Pos}(\mathcal{X}) \,:\, \mathrm{Tr}(\rho) = 1\} \tag{1.97}$$
will be used to denote the collection of density operators acting on a complex Euclidean space $\mathcal{X}$.
6. *Projection operators.* A positive semidefinite operator $\Pi \in \mathrm{Pos}(\mathcal{X})$ is said to be a *projection operator*[2] if, in addition to being positive semidefinite, it satisfies the equation $\Pi^2 = \Pi$. Equivalently, a projection operator is a Hermitian operator whose only eigenvalues are 0 and 1. The collection of all projection operators of the form $\Pi \in \mathrm{Pos}(\mathcal{X})$ is denoted $\mathrm{Proj}(\mathcal{X})$. For each subspace $\mathcal{V} \subseteq \mathcal{X}$, there is a uniquely defined projection operator $\Pi \in \mathrm{Proj}(\mathcal{X})$ satisfying $\mathrm{im}(\Pi) = \mathcal{V}$; when it is convenient, the notation $\Pi_{\mathcal{V}}$ is used to refer to this projection operator.
7. *Isometries.* An operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ is an *isometry* if it preserves the Euclidean norm: $\|Au\| = \|u\|$ for all $u \in \mathcal{X}$. This condition is equivalent to $A^*A = \mathbb{1}_{\mathcal{X}}$. The notation
$$\mathrm{U}(\mathcal{X}, \mathcal{Y}) = \{A \in \mathrm{L}(\mathcal{X}, \mathcal{Y}) \,:\, A^*A = \mathbb{1}_{\mathcal{X}}\} \tag{1.98}$$
is used to denote this class of operators. In order for an isometry of the form $A \in \mathrm{U}(\mathcal{X}, \mathcal{Y})$ to exist, it must hold that $\dim(\mathcal{Y}) \geq \dim(\mathcal{X})$. Every isometry preserves not only the Euclidean norm, but inner products as well: $\langle Au, Av \rangle = \langle u, v \rangle$ for all $u, v \in \mathcal{X}$.
8. *Unitary operators.* The set of isometries mapping a complex Euclidean space $\mathcal{X}$ to itself is denoted $\mathrm{U}(\mathcal{X})$, and operators in this set are *unitary operators*. The letters U, V, and W will often be used to refer to unitary operators (and sometimes to isometries more generally) in this book. Every unitary operator $U \in \mathrm{U}(\mathcal{X})$ is necessarily invertible and satisfies the equation $UU^* = U^*U = \mathbb{1}_{\mathcal{X}}$, and is therefore normal.

[2] Sometimes the term *projection operator* refers to an operator $X \in \mathrm{L}(\mathcal{X})$ that satisfies the equation $X^2 = X$, but that might not be Hermitian. This is not the meaning that is associated with this term in this book.

9. *Diagonal operators.* An operator $X \in \mathrm{L}(\mathcal{X})$, for a complex Euclidean space of the form $\mathcal{X} = \mathbb{C}^{\Sigma}$, is a *diagonal operator* if $X(a,b) = 0$ for all $a, b \in \Sigma$ with $a \neq b$. For a given vector $u \in \mathcal{X}$, one writes $\mathrm{Diag}(u) \in \mathrm{L}(\mathcal{X})$ to denote the diagonal operator defined as

$$\mathrm{Diag}(u)(a,b) = \begin{cases} u(a) & \text{if } a = b \\ 0 & \text{if } a \neq b. \end{cases} \tag{1.99}$$

Further remarks on Hermitian and positive semidefinite operators

The sum of two Hermitian operators is Hermitian, as is a real scalar multiple of a Hermitian operator. The inner product of two Hermitian operators is real as well. For every choice of a complex Euclidean space $\mathcal{X}$, the space $\mathrm{Herm}(\mathcal{X})$ therefore forms a vector space over the real numbers on which an inner product is defined.

Indeed, under the assumption that $\mathcal{X} = \mathbb{C}^{\Sigma}$, it holds that the space $\mathrm{Herm}(\mathcal{X})$ and the real Euclidean space $\mathbb{R}^{\Sigma\times\Sigma}$ are *isometrically isomorphic*: there exists a linear bijection

$$\phi : \mathbb{R}^{\Sigma\times\Sigma} \to \mathrm{Herm}(\mathcal{X}) \tag{1.100}$$

with the property that

$$\langle \phi(u), \phi(v) \rangle = \langle u, v \rangle \tag{1.101}$$

for all $u, v \in \mathbb{R}^{\Sigma\times\Sigma}$. The existence of such a linear bijection allows one to directly translate many statements about real Euclidean spaces to the space of Hermitian operators acting on a complex Euclidean space.

One way to define a mapping ϕ as above is as follows. First, assume that a total ordering of Σ has been fixed, and define a collection

$$\{H_{a,b} \,:\, (a,b) \in \Sigma \times \Sigma\} \subset \mathrm{Herm}(\mathcal{X}) \tag{1.102}$$

as

$$H_{a,b} = \begin{cases} E_{a,a} & \text{if } a = b \\ \frac{1}{\sqrt{2}}(E_{a,b} + E_{b,a}) & \text{if } a < b \\ \frac{1}{\sqrt{2}}(iE_{a,b} - iE_{b,a}) & \text{if } a > b \end{cases} \tag{1.103}$$

for each pair $(a,b) \in \Sigma \times \Sigma$. It holds that (1.102) is an orthonormal set (with respect to the usual inner product defined on $\mathrm{L}(\mathcal{X})$), and moreover every element of $\mathrm{Herm}(\mathcal{X})$ can be expressed uniquely as a real linear combination of the operators in this set. The mapping ϕ defined by the equation

$$\phi\big(e_{(a,b)}\big) = H_{a,b}, \tag{1.104}$$

and extended to all of $\mathbb{R}^{\Sigma\times\Sigma}$ by linearity, satisfies the requirement (1.101).

The eigenvalues of a Hermitian operator are necessarily real numbers, and can therefore be ordered from largest to smallest. For every complex Euclidean space $\mathcal{X}$ and every Hermitian operator $H \in \mathrm{Herm}(\mathcal{X})$, the vector

$$\lambda(H) = (\lambda_1(H),\, \lambda_2(H),\, \ldots,\, \lambda_n(H)) \in \mathbb{R}^n \tag{1.105}$$

is defined so that

$$\mathrm{spec}(H) = \{\lambda_1(H),\, \lambda_2(H),\, \ldots,\, \lambda_n(H)\} \tag{1.106}$$

and

$$\lambda_1(H) \geq \lambda_2(H) \geq \cdots \geq \lambda_n(H). \tag{1.107}$$

The notation $\lambda_k(H)$ may also be used in isolation to refer to the k-th largest eigenvalue of a Hermitian operator H.

The eigenvalues of Hermitian operators can be characterized by a theorem known as the *Courant–Fischer theorem*, which is as follows.

Theorem 1.2 (Courant–Fischer theorem) *Let $\mathcal{X}$ be a complex Euclidean space of dimension n and let $H \in \mathrm{Herm}(\mathcal{X})$ be a Hermitian operator. For every $k \in \{1, \ldots, n\}$ it holds that*

$$\begin{aligned} \lambda_k(H) &= \max_{u_1,\ldots,u_{n-k}\in\mathcal{S}(\mathcal{X})} \; \min_{\substack{v\in\mathcal{S}(\mathcal{X}) \\ v\perp\{u_1,\ldots,u_{n-k}\}}} v^*Hv \\ &= \min_{u_1,\ldots,u_{k-1}\in\mathcal{S}(\mathcal{X})} \; \max_{\substack{v\in\mathcal{S}(\mathcal{X}) \\ v\perp\{u_1,\ldots,u_{k-1}\}}} v^*Hv \end{aligned} \tag{1.108}$$

(It is to be interpreted that the maximum or minimum is omitted if it is to be taken over an empty set of vectors, and that $v \perp \varnothing$ holds for all $v \in \mathcal{X}$.)

There are alternative ways to describe positive semidefinite operators that are useful in different situations. In particular, the following statements are equivalent for every operator $P \in \mathrm{L}(\mathcal{X})$:

1. P is positive semidefinite.
2. $P = A^*A$ for an operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, for some choice of a complex Euclidean space $\mathcal{Y}$.
3. P is Hermitian and every eigenvalue of P is nonnegative.
4. $\langle u, Pu\rangle$ is a nonnegative real number for all $u \in \mathcal{X}$.
5. $\langle Q, P\rangle$ is a nonnegative real number for all $Q \in \mathrm{Pos}(\mathcal{X})$.
6. There exists a collection of vectors $\{u_a : a \in \Sigma\} \subset \mathcal{X}$ for which it holds that $P(a, b) = \langle u_a, u_b\rangle$ for all $a, b \in \Sigma$.

7. There exists a collection of vectors $\{u_a : a \in \Sigma\} \subset \mathcal{Y}$, for some choice of a complex Euclidean space $\mathcal{Y}$, for which it holds that $P(a,b) = \langle u_a, u_b \rangle$ for all $a, b \in \Sigma$.

Along similar lines, one has that the following statements are equivalent for every operator $P \in \mathrm{L}(\mathcal{X})$:

1. P is positive definite.
2. P is Hermitian, and every eigenvalue of P is positive.
3. $\langle u, Pu \rangle$ is a positive real number for every nonzero $u \in \mathcal{X}$.
4. $\langle Q, P \rangle$ is a positive real number for every nonzero $Q \in \mathrm{Pos}(\mathcal{X})$.
5. There exists a positive real number $\varepsilon > 0$ such that $P - \varepsilon \mathbb{1} \in \mathrm{Pos}(\mathcal{X})$.

The notations $P \geq 0$ and $0 \leq P$ indicate that P is positive semidefinite, while $P > 0$ and $0 < P$ indicate that P is positive definite. More generally, for Hermitian operators X and Y, one writes either $X \geq Y$ or $Y \leq X$ to indicate that $X - Y$ is positive semidefinite, and either $X > Y$ or $Y < X$ to indicate that $X - Y$ is positive definite.

Linear maps on square operators

Linear maps of the form

$$\Phi : \mathrm{L}(\mathcal{X}) \to \mathrm{L}(\mathcal{Y}), \tag{1.109}$$

for complex Euclidean spaces $\mathcal{X}$ and $\mathcal{Y}$, play a fundamental role in the theory of quantum information. The set of all such maps is denoted $\mathrm{T}(\mathcal{X}, \mathcal{Y})$, and is itself a complex vector space when addition and scalar multiplication are defined in the straightforward way:

1. Addition: given two maps $\Phi, \Psi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$, the map $\Phi + \Psi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is defined as
$$(\Phi + \Psi)(X) = \Phi(X) + \Psi(X) \tag{1.110}$$
for all $X \in \mathrm{L}(\mathcal{X})$.
2. Scalar multiplication: given a map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ and a scalar $\alpha \in \mathbb{C}$, the map $\alpha\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is defined as
$$(\alpha\Phi)(X) = \alpha\Phi(X) \tag{1.111}$$
for all $X \in \mathrm{L}(\mathcal{X})$.

For a given map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$, the *adjoint* of Φ is defined to be the unique map $\Phi^* \in \mathrm{T}(\mathcal{Y}, \mathcal{X})$ that satisfies

$$\langle \Phi^*(Y), X \rangle = \langle Y, \Phi(X) \rangle \tag{1.112}$$

for all $X \in \mathrm{L}(\mathcal{X})$ and $Y \in \mathrm{L}(\mathcal{Y})$.

Tensor products of maps of the form (1.109) are defined in a similar way to tensor products of operators. More specifically, for any choice of complex Euclidean spaces $\mathcal{X}_1, \ldots, \mathcal{X}_n$ and $\mathcal{Y}_1, \ldots, \mathcal{Y}_n$ and linear maps

$$\Phi_1 \in \mathrm{T}(\mathcal{X}_1, \mathcal{Y}_1), \ \ldots, \ \Phi_n \in \mathrm{T}(\mathcal{X}_n, \mathcal{Y}_n), \tag{1.113}$$

one defines the tensor product of these maps

$$\Phi_1 \otimes \cdots \otimes \Phi_n \in \mathrm{T}(\mathcal{X}_1 \otimes \cdots \otimes \mathcal{X}_n, \mathcal{Y}_1 \otimes \cdots \otimes \mathcal{Y}_n) \tag{1.114}$$

to be the unique linear map that satisfies the equation

$$(\Phi_1 \otimes \cdots \otimes \Phi_n)(X_1 \otimes \cdots \otimes X_n) = \Phi_1(X_1) \otimes \cdots \otimes \Phi_n(X_n) \tag{1.115}$$

for all operators $X_1 \in \mathrm{L}(\mathcal{X}_1), \ldots, X_n \in \mathrm{L}(\mathcal{X}_n)$. As for vectors and operators, the notation $\Phi^{\otimes n}$ denotes the n-fold tensor product of a map Φ with itself.

The notation $\mathrm{T}(\mathcal{X})$ is understood to be a shorthand for $\mathrm{T}(\mathcal{X}, \mathcal{X})$. The identity map $\mathbb{1}_{\mathrm{L}(\mathcal{X})} \in \mathrm{T}(\mathcal{X})$ is defined as

$$\mathbb{1}_{\mathrm{L}(\mathcal{X})}(X) = X \tag{1.116}$$

for all $X \in \mathrm{L}(\mathcal{X})$.

The trace function defined for square operators acting on $\mathcal{X}$ is a linear mapping of the form

$$\mathrm{Tr} : \mathrm{L}(\mathcal{X}) \to \mathbb{C}. \tag{1.117}$$

By making the identification $\mathrm{L}(\mathbb{C}) = \mathbb{C}$, one sees that the trace function is a linear map of the form

$$\mathrm{Tr} \in \mathrm{T}(\mathcal{X}, \mathbb{C}). \tag{1.118}$$

For a second complex Euclidean space $\mathcal{Y}$, one may consider the map

$$\mathrm{Tr} \otimes \mathbb{1}_{\mathrm{L}(\mathcal{Y})} \in \mathrm{T}(\mathcal{X} \otimes \mathcal{Y}, \mathcal{Y}). \tag{1.119}$$

By the definition of the tensor product of maps stated above, this is the unique map that satisfies the equation

$$(\mathrm{Tr} \otimes \mathbb{1}_{\mathrm{L}(\mathcal{Y})})(X \otimes Y) = \mathrm{Tr}(X)Y \tag{1.120}$$

for all operators $X \in \mathrm{L}(\mathcal{X})$ and $Y \in \mathrm{L}(\mathcal{Y})$. This map is called the *partial trace*, and is more commonly denoted $\mathrm{Tr}_{\mathcal{X}}$. Along similar lines, the map $\mathrm{Tr}_{\mathcal{Y}} \in \mathrm{T}(\mathcal{X} \otimes \mathcal{Y}, \mathcal{X})$ is defined as

$$\mathrm{Tr}_{\mathcal{Y}} = \mathbb{1}_{\mathrm{L}(\mathcal{X})} \otimes \mathrm{Tr}. \tag{1.121}$$

Generalizations of these maps may also be defined for tensor products of three or more complex Euclidean spaces.

The following classes of maps of the form (1.109) are among those that are discussed in greater detail later in this book:

1. *Hermitian-preserving maps.* A map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is *Hermitian-preserving* if it holds that
$$\Phi(H) \in \mathrm{Herm}(\mathcal{Y}) \tag{1.122}$$
for every Hermitian operator $H \in \mathrm{Herm}(\mathcal{X})$.
2. *Positive maps.* A map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is *positive* if it holds that
$$\Phi(P) \in \mathrm{Pos}(\mathcal{Y}) \tag{1.123}$$
for every positive semidefinite operator $P \in \mathrm{Pos}(\mathcal{X})$.
3. *Completely positive maps.* A map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is *completely positive* if it holds that
$$\Phi \otimes \mathbb{1}_{\mathrm{L}(\mathcal{Z})} \tag{1.124}$$
is a positive map for every complex Euclidean space $\mathcal{Z}$. The set of all completely positive maps of this form is denoted $\mathrm{CP}(\mathcal{X}, \mathcal{Y})$.
4. *Trace-preserving maps.* A map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is *trace-preserving* if it holds that
$$\mathrm{Tr}(\Phi(X)) = \mathrm{Tr}(X) \tag{1.125}$$
for all $X \in \mathrm{L}(\mathcal{X})$.
5. *Unital maps.* A map $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ is *unital* if
$$\Phi(\mathbb{1}_{\mathcal{X}}) = \mathbb{1}_{\mathcal{Y}}. \tag{1.126}$$

Maps of these sorts are discussed in greater detail in Chapters 2 and 4.

The operator-vector correspondence

There is a correspondence between the spaces $\mathrm{L}(\mathcal{Y}, \mathcal{X})$ and $\mathcal{X} \otimes \mathcal{Y}$, for any choice of complex Euclidean spaces $\mathcal{X} = \mathbb{C}^{\Sigma}$ and $\mathcal{Y} = \mathbb{C}^{\Gamma}$, that will be used repeatedly throughout this book. This correspondence is given by the linear mapping
$$\mathrm{vec} : \mathrm{L}(\mathcal{Y}, \mathcal{X}) \to \mathcal{X} \otimes \mathcal{Y}, \tag{1.127}$$
defined by the action
$$\mathrm{vec}(E_{a,b}) = e_a \otimes e_b \tag{1.128}$$
for all $a \in \Sigma$ and $b \in \Gamma$. In other words, this mapping is the change-of-basis taking the standard basis of $\mathrm{L}(\mathcal{Y}, \mathcal{X})$ to the standard basis of $\mathcal{X} \otimes \mathcal{Y}$. By linearity, it holds that
$$\mathrm{vec}(uv^*) = u \otimes \overline{v} \tag{1.129}$$

for $u \in \mathcal{X}$ and $v \in \mathcal{Y}$. This includes the special cases

$$\operatorname{vec}(u) = u \quad \text{and} \quad \operatorname{vec}(v^*) = \overline{v}, \tag{1.130}$$

obtained by setting $v = 1$ and $u = 1$, respectively.

The vec mapping is a linear bijection, which implies that every vector $u \in \mathcal{X} \otimes \mathcal{Y}$ uniquely determines an operator $A \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$ that satisfies $\operatorname{vec}(A) = u$. It is also an isometry, in the sense that

$$\langle A, B \rangle = \langle \operatorname{vec}(A), \operatorname{vec}(B) \rangle \tag{1.131}$$

for all $A, B \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$.

A few specific identities concerning the vec mapping will be especially useful throughout this book. One such identity is

$$(A_0 \otimes A_1) \operatorname{vec}(B) = \operatorname{vec}\bigl(A_0 B A_1^{\mathsf{T}}\bigr), \tag{1.132}$$

holding for all operators $A_0 \in \mathrm{L}(\mathcal{X}_0, \mathcal{Y}_0)$, $A_1 \in \mathrm{L}(\mathcal{X}_1, \mathcal{Y}_1)$, and $B \in \mathrm{L}(\mathcal{X}_1, \mathcal{X}_0)$, over all choices of complex Euclidean spaces $\mathcal{X}_0$, $\mathcal{X}_1$, $\mathcal{Y}_0$, and $\mathcal{Y}_1$. Two more such identities are

$$\operatorname{Tr}_{\mathcal{Y}}\bigl(\operatorname{vec}(A) \operatorname{vec}(B)^*\bigr) = AB^*, \tag{1.133}$$

$$\operatorname{Tr}_{\mathcal{X}}\bigl(\operatorname{vec}(A) \operatorname{vec}(B)^*\bigr) = A^{\mathsf{T}}\overline{B}, \tag{1.134}$$

which hold for all operators $A, B \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$, over all choices of complex Euclidean spaces $\mathcal{X}$ and $\mathcal{Y}$.

1.1.3 Operator decompositions and norms

Two decompositions of operators—the *spectral decomposition* and *singular value decomposition*—along with various related notions, are discussed in the present section. Among these related notions is a class of operator norms called *Schatten norms*, which include the trace norm, the Frobenius norm, and the spectral norm. These three norms are used frequently throughout this book.

The spectral theorem

The *spectral theorem* establishes that every normal operator can be expressed as a linear combination of projections onto pairwise orthogonal subspaces. A formal statement of the spectral theorem follows.

Theorem 1.3 (Spectral theorem) *Let $\mathcal{X}$ be a complex Euclidean space and let $X \in \mathrm{L}(\mathcal{X})$ be a normal operator. There exists a positive integer m, distinct complex numbers $\lambda_1, \ldots, \lambda_m \in \mathbb{C}$, and nonzero projection operators $\Pi_1, \ldots, \Pi_m \in \mathrm{Proj}(\mathcal{X})$ satisfying $\Pi_1 + \cdots + \Pi_m = \mathbb{1}_{\mathcal{X}}$, such that*

$$X = \sum_{k=1}^{m} \lambda_k \Pi_k. \tag{1.135}$$

The scalars $\lambda_1, \ldots, \lambda_m$ and projection operators $\Pi_1, \ldots, \Pi_m$ are unique, up to their ordering: each scalar λ_k is an eigenvalue of X with multiplicity equal to the rank of Π_k, and Π_k is the projection operator onto the space spanned by the eigenvectors of X corresponding to the eigenvalue λ_k.

The expression of a normal operator X in the form of the equation (1.135) is called a *spectral decomposition* of X.

A simple corollary of the spectral theorem follows. It expresses essentially the same fact as the spectral theorem, but in a slightly different form that will sometimes be convenient to refer to later in the book.

Corollary 1.4 *Let $\mathcal{X}$ be a complex Euclidean space having dimension n, let $X \in \mathrm{L}(\mathcal{X})$ be a normal operator, and assume that*

$$\mathrm{spec}(X) = \{\lambda_1, \ldots, \lambda_n\}. \tag{1.136}$$

There exists an orthonormal basis $\{x_1, \ldots, x_n\}$ of $\mathcal{X}$ such that

$$X = \sum_{k=1}^{n} \lambda_k x_k x_k^*. \tag{1.137}$$

It is evident from the expression (1.137), along with the requirement that the set $\{x_1, \ldots, x_n\}$ is an orthonormal basis, that each x_k is an eigenvector of X whose corresponding eigenvalue is λ_k. It is also evident that any operator X that is expressible in such a form as (1.137) is normal, implying that the condition of normality is equivalent to the existence of an orthonormal basis of eigenvectors.

On a few occasions later in the book, it will be convenient to index the eigenvectors and eigenvalues of a given normal operator $X \in \mathrm{L}(\mathbb{C}^{\Sigma})$ by symbols in the alphabet Σ rather than by integers in the set $\{1, \ldots, n\}$ for $n = |\Sigma|$. It follows immediately from Corollary 1.4 that a normal operator $X \in \mathrm{L}(\mathbb{C}^{\Sigma})$ may be expressed as

$$X = \sum_{a \in \Sigma} \lambda_a x_a x_a^* \tag{1.138}$$

for some choice of an orthonormal basis $\{x_a : a \in \Sigma\}$ of $\mathbb{C}^{\Sigma}$ and a collection

of complex numbers $\{\lambda_a : a \in \Sigma\}$. Indeed, such an expression may be derived from (1.137) by associating symbols in the alphabet Σ with integers in the set $\{1, \ldots, n\}$ with respect to an arbitrarily chosen bijection.

It is convenient to refer to expressions of operators having either of the forms (1.137) or (1.138) as *spectral decompositions*, despite the fact that they may differ slightly from the form (1.135). Unlike the form (1.135), the forms (1.137) and (1.138) are generally not unique. Along similar lines, the term *spectral theorem* is sometimes used to refer to the statement of Corollary 1.4, as opposed to the statement of Theorem 1.3. These conventions are followed throughout this book when there is no danger of any confusion resulting from their use.

The following important theorem states that the same orthonormal basis of eigenvectors $\{x_1, \ldots, x_n\}$ may be chosen for any two normal operators under the assumption that they commute.

Theorem 1.5 *Let $\mathcal{X}$ be a complex Euclidean space having dimension n and let $X, Y \in \mathrm{L}(\mathcal{X})$ be normal operators for which $[X, Y] = 0$. There exists an orthonormal basis $\{x_1, \ldots, x_n\}$ of $\mathcal{X}$ such that*

$$X = \sum_{k=1}^{n} \alpha_k x_k x_k^* \quad \text{and} \quad Y = \sum_{k=1}^{n} \beta_k x_k x_k^*, \tag{1.139}$$

for some choice of complex numbers $\alpha_1, \ldots, \alpha_n, \beta_1, \ldots, \beta_n$ satisfying

$$\operatorname{spec}(X) = \{\alpha_1, \ldots, \alpha_n\} \quad \text{and} \quad \operatorname{spec}(Y) = \{\beta_1, \ldots, \beta_n\}. \tag{1.140}$$

Jordan–Hahn decompositions

Every Hermitian operator is normal and has real eigenvalues. It therefore follows from the spectral theorem (Theorem 1.3) that, for every Hermitian operator $H \in \mathrm{Herm}(\mathcal{X})$, there exists a positive integer m, nonzero projection operators $\Pi_1, \ldots, \Pi_m$ satisfying

$$\Pi_1 + \cdots + \Pi_m = \mathbb{1}_{\mathcal{X}}, \tag{1.141}$$

and real numbers $\lambda_1, \ldots, \lambda_m$ such that

$$H = \sum_{k=1}^{m} \lambda_k \Pi_k. \tag{1.142}$$

By defining operators

$$P = \sum_{k=1}^{m} \max\{\lambda_k, 0\}\, \Pi_k \quad \text{and} \quad Q = \sum_{k=1}^{m} \max\{-\lambda_k, 0\}\, \Pi_k\,, \tag{1.143}$$

one finds that

$$H = P - Q \tag{1.144}$$

for $P, Q \in \mathrm{Pos}(\mathcal{X})$ satisfying $PQ = 0$. The expression (1.144) of a given Hermitian operator H in this form, for positive semidefinite operators P and Q satisfying $PQ = 0$, is called a *Jordan–Hahn decomposition*. There is only one such expression for a given operator $H \in \mathrm{Herm}(\mathcal{X})$; the operators P and Q are uniquely defined by the requirements that $P, Q \in \mathrm{Pos}(\mathcal{X})$, $PQ = 0$, and $H = P - Q$.

Functions of normal operators

Every function of the form $f : \mathbb{C} \to \mathbb{C}$ may be extended to the set of normal operators in $\mathrm{L}(\mathcal{X})$, for a given complex Euclidean space $\mathcal{X}$, by means of the spectral theorem (Theorem 1.3). In particular, if $X \in \mathrm{L}(\mathcal{X})$ is normal and has the spectral decomposition (1.135), then one defines

$$f(X) = \sum_{k=1}^{m} f(\lambda_k)\Pi_k. \tag{1.145}$$

Naturally, functions defined only on subsets of $\mathbb{C}$ may be extended to normal operators whose eigenvalues are restricted accordingly.

The following examples of scalar functions extended to operators will be important later in this book:

1. For $r > 0$, the function $\lambda \mapsto \lambda^r$ is defined for all $\lambda \in [0, \infty)$. For a positive semidefinite operator $P \in \mathrm{Pos}(\mathcal{X})$ having spectral decomposition

   $$P = \sum_{k=1}^{m} \lambda_k \Pi_k, \tag{1.146}$$

   for which it necessarily holds that $\lambda_k \geq 0$ for all $k \in \{1, \ldots, m\}$, one defines

   $$P^r = \sum_{k=1}^{m} \lambda_k^r \, \Pi_k. \tag{1.147}$$

   For positive integer values of r, it is evident that P^r coincides with the usual meaning of this expression given by operator multiplication.

   The case that $r = 1/2$ is particularly common, and in this case one may write $\sqrt{P}$ to denote $P^{1/2}$. The operator $\sqrt{P}$ is the unique positive semidefinite operator that satisfies the equation

   $$\sqrt{P}\sqrt{P} = P. \tag{1.148}$$

2. Along similar lines to the previous example, for any real number $r \in \mathbb{R}$, the function $\lambda \mapsto \lambda^r$ is defined for all $\lambda \in (0, \infty)$. For a given positive definite operator $P \in \mathrm{Pd}(\mathcal{X})$ having a spectral decomposition of the form (1.146), for which it holds that $\lambda_k > 0$ for all $k \in \{1, \ldots, m\}$, one defines P^r in a similar way to (1.147) above.
3. The (base-2) logarithm function $\lambda \mapsto \log(\lambda)$ is defined for all $\lambda \in (0, \infty)$. For a given positive definite operator $P \in \mathrm{Pd}(\mathcal{X})$, having a spectral decomposition (1.146) as above, one defines

$$\log(P) = \sum_{k=1}^{m} \log(\lambda_k) \Pi_k. \tag{1.149}$$

The singular value theorem

The *singular value theorem* has a close relationship to the spectral theorem. Unlike the spectral theorem, however, the singular value theorem holds for arbitrary (nonzero) operators, as opposed to just normal operators.

Theorem 1.6 (Singular value theorem) *Let $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ be a nonzero operator having rank equal to r, for complex Euclidean spaces $\mathcal{X}$ and $\mathcal{Y}$. There exist orthonormal sets $\{x_1, \ldots, x_r\} \subset \mathcal{X}$ and $\{y_1, \ldots, y_r\} \subset \mathcal{Y}$, along with positive real numbers $s_1, \ldots, s_r$, such that*

$$A = \sum_{k=1}^{r} s_k y_k x_k^*. \tag{1.150}$$

An expression of a given operator A in the form of (1.150) is said to be a *singular value decomposition* of A. The numbers $s_1, \ldots, s_r$ are called *singular values* and the vectors $x_1, \ldots, x_r$ and $y_1, \ldots, y_r$ are called *right* and *left singular vectors*, respectively.

The singular values $s_1, \ldots, s_r$ of an operator A are uniquely determined, up to their ordering. It will be assumed hereafter that singular values are always ordered from largest to smallest: $s_1 \geq \cdots \geq s_r$. When it is necessary to indicate the dependence of these singular values on the operator A, they are denoted $s_1(A), \ldots, s_r(A)$. Although 0 is not formally considered to be a singular value of any operator, it is convenient to also define $s_k(A) = 0$ for $k > \mathrm{rank}(A)$, and to take $s_k(A) = 0$ for all $k \geq 1$ when $A = 0$. The notation $s(A)$ is used to refer to the vector of singular values

$$s(A) = (s_1(A), \ldots, s_r(A)), \tag{1.151}$$

or to an extension of this vector

$$s(A) = (s_1(A), \ldots, s_m(A)) \tag{1.152}$$

when it is convenient to view it as an element of $\mathbb{R}^m$ for $m > \operatorname{rank}(A)$.

As suggested above, there is a close relationship between the singular value theorem and the spectral theorem. In particular, the singular value decomposition of an operator A and the spectral decompositions of the operators A^*A and AA^* are related in the following way: it holds that

$$s_k(A) = \sqrt{\lambda_k(AA^*)} = \sqrt{\lambda_k(A^*A)} \tag{1.153}$$

for $1 \leq k \leq \operatorname{rank}(A)$, and moreover the right singular vectors of A are eigenvectors of A^*A and the left singular vectors of A are eigenvectors of AA^*. One is free, in fact, to choose the left singular vectors of A to be any orthonormal collection of eigenvectors of AA^* for which the corresponding eigenvalues are nonzero—and once this is done the right singular vectors will be uniquely determined. Alternately, the right singular vectors of A may be chosen to be any orthonormal collection of eigenvectors of A^*A for which the corresponding eigenvalues are nonzero, which uniquely determines the left singular vectors.

In the special case that $X \in \mathrm{L}(\mathcal{X})$ is a normal operator, one may obtain a singular value decomposition of X directly from a spectral decomposition of the form

$$X = \sum_{k=1}^{n} \lambda_k x_k x_k^*. \tag{1.154}$$

In particular, one may define $S = \{k \in \{1, \ldots, n\} : \lambda_k \neq 0\}$, and set

$$s_k = |\lambda_k| \quad \text{and} \quad y_k = \frac{\lambda_k}{|\lambda_k|} x_k \tag{1.155}$$

for each $k \in S$. The expression

$$X = \sum_{k \in S} s_k y_k x_k^* \tag{1.156}$$

then represents a singular value decomposition of X, up to a relabeling of the terms in the sum.

The following corollary represents a reformulation of the singular value theorem that is useful in some situations.

Corollary 1.7 *Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces, let $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ be a nonzero operator, and let $r = \operatorname{rank}(A)$. There exists a diagonal and positive definite operator $D \in \mathrm{Pd}(\mathbb{C}^r)$ and isometries $U \in \mathrm{U}(\mathbb{C}^r, \mathcal{X})$ and $V \in \mathrm{U}(\mathbb{C}^r, \mathcal{Y})$ such that $A = VDU^*$.*

Polar decompositions

For every square operator $X \in \mathrm{L}(\mathcal{X})$, it is possible to choose a positive semidefinite operator $P \in \mathrm{Pos}(\mathcal{X})$ and a unitary operator $W \in \mathrm{U}(\mathcal{X})$ such that the equation

$$X = WP \tag{1.157}$$

holds; this follows from Corollary 1.7 by taking $W = VU^*$ and $P = UDU^*$. Alternatively, by similar reasoning it is possible to write

$$X = PW \tag{1.158}$$

for a (generally different) choice of operators $P \in \mathrm{Pos}(\mathcal{X})$ and $W \in \mathrm{U}(\mathcal{X})$. The expressions (1.157) and (1.158) are known as *polar decompositions* of X.

The Moore–Penrose pseudo-inverse

For a given operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, one defines an operator $A^+ \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$, known as the *Moore–Penrose pseudo-inverse* of A, as the unique operator that possesses the following properties:

1. $AA^+A = A$,
2. $A^+AA^+ = A^+$, and
3. AA^+ and A^+A are both Hermitian.

It is evident that there is at least one such choice of A^+, for if

$$A = \sum_{k=1}^{r} s_k y_k x_k^* \tag{1.159}$$

is a singular value decomposition of a nonzero operator A, then

$$A^+ = \sum_{k=1}^{r} \frac{1}{s_k} x_k y_k^* \tag{1.160}$$

possesses the three properties listed above. One may observe that AA^+ and A^+A are projection operators, projecting onto the spaces spanned by the left singular vectors and right singular vectors of A, respectively.

The fact that A^+ is uniquely determined by the above equations may be verified as follows. Suppose that $B, C \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$ both possess the above properties:

1. $ABA = A = ACA$,
2. $BAB = B$ and $CAC = C$, and
3. AB, BA, AC, and CA are all Hermitian.

It follows that

$$\begin{aligned} B &= BAB = (BA)^*B = A^*B^*B = (ACA)^*B^*B \\ &= A^*C^*A^*B^*B = (CA)^*(BA)^*B = CABAB \\ &= CAB = CACAB = C(AC)^*(AB)^* = CC^*A^*B^*A^* \\ &= CC^*(ABA)^* = CC^*A^* = C(AC)^* = CAC = C, \end{aligned} \tag{1.161}$$

which shows that $B = C$.

Schmidt decompositions

Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces, and suppose that $u \in \mathcal{X} \otimes \mathcal{Y}$ is a nonzero vector. Given that the vec mapping is a bijection, there exists a unique operator $A \in \mathrm{L}(\mathcal{Y}, \mathcal{X})$ such that $u = \mathrm{vec}(A)$. For any singular value decomposition

$$A = \sum_{k=1}^{r} s_k x_k y_k^*, \tag{1.162}$$

it holds that

$$u = \mathrm{vec}(A) = \mathrm{vec}\left(\sum_{k=1}^{r} s_k x_k y_k^*\right) = \sum_{k=1}^{r} s_k x_k \otimes \overline{y_k}. \tag{1.163}$$

The orthonormality of $\{y_1, \ldots, y_r\}$ implies that $\{\overline{y_1}, \ldots, \overline{y_r}\}$ is orthonormal as well. It follows that every nonzero vector $u \in \mathcal{X} \otimes \mathcal{Y}$ can be expressed in the form

$$u = \sum_{k=1}^{r} s_k x_k \otimes z_k \tag{1.164}$$

for positive real numbers $s_1, \ldots, s_r$ and orthonormal sets $\{x_1, \ldots, x_r\} \subset \mathcal{X}$ and $\{z_1, \ldots, z_r\} \subset \mathcal{Y}$. An expression of u having this form is called a *Schmidt decomposition* of u.

Norms of operators

A *norm* on the space of operators $\mathrm{L}(\mathcal{X}, \mathcal{Y})$, for complex Euclidean spaces $\mathcal{X}$ and $\mathcal{Y}$, is a function $\|\cdot\|$ satisfying the following properties:

1. Positive definiteness: $\|A\| \geq 0$ for all $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, with $\|A\| = 0$ if and only if $A = 0$.
2. Positive scalability: $\|\alpha A\| = |\alpha| \|A\|$ for all $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ and $\alpha \in \mathbb{C}$.
3. The triangle inequality: $\|A + B\| \leq \|A\| + \|B\|$ for all $A, B \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$.

Many interesting and useful norms can be defined on spaces of operators, but this book will mostly be concerned with a single family of norms called *Schatten p-norms*. This family includes the three most commonly used norms in quantum information theory: the *spectral norm*, the *Frobenius norm*, and the *trace norm*.

For any operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ and any real number $p \geq 1$, one defines the Schatten p-norm of A as

$$\|A\|_p = \left(\mathrm{Tr}\left((A^*A)^{\frac{p}{2}}\right)\right)^{\frac{1}{p}}. \tag{1.165}$$

The Schatten ∞-norm is defined as

$$\|A\|_\infty = \max\left\{\|Au\| \,:\, u \in \mathcal{X},\ \|u\| \leq 1\right\}, \tag{1.166}$$

which coincides with $\lim_{p\to\infty}\|A\|_p$, explaining why the subscript ∞ is used. The Schatten p-norm of an operator A coincides with the ordinary vector p-norm of the vector of singular values of A:

$$\|A\|_p = \|s(A)\|_p. \tag{1.167}$$

The Schatten p-norms possess a variety of properties, including the ones summarized in the following list:

1. The Schatten p-norms are non-increasing in p: for every operator A and for $1 \leq p \leq q \leq \infty$, it holds that

$$\|A\|_p \geq \|A\|_q. \tag{1.168}$$

2. For every nonzero operator A and for $1 \leq p \leq q \leq \infty$, it holds that

$$\|A\|_p \leq \mathrm{rank}(A)^{\frac{1}{p}-\frac{1}{q}}\|A\|_q. \tag{1.169}$$

   In particular, one has

$$\|A\|_1 \leq \sqrt{\mathrm{rank}(A)}\|A\|_2 \quad \text{and} \quad \|A\|_2 \leq \sqrt{\mathrm{rank}(A)}\|A\|_\infty. \tag{1.170}$$

3. For every $p \in [1, \infty]$, the Schatten p-norm is isometrically invariant (and therefore unitarily invariant): for every $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, $U \in \mathrm{U}(\mathcal{Y}, \mathcal{Z})$, and $V \in \mathrm{U}(\mathcal{X}, \mathcal{W})$ it holds that

$$\|A\|_p = \|UAV^*\|_p. \tag{1.171}$$

4. For each $p \in [1, \infty]$, one defines $p^* \in [1, \infty]$ by the equation

$$\frac{1}{p} + \frac{1}{p^*} = 1. \tag{1.172}$$

For every operator $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, it holds that the Schatten p-norm and p^*-norm are dual, in the sense that

$$\|A\|_p = \max\bigl\{|\langle B, A\rangle| \,:\, B \in \mathrm{L}(\mathcal{X}, \mathcal{Y}),\ \|B\|_{p^*} \leq 1\bigr\}. \tag{1.173}$$

One consequence of (1.173) is the inequality

$$|\langle B, A\rangle| \leq \|A\|_p \|B\|_{p^*}, \tag{1.174}$$

which is known as the *Hölder inequality* for Schatten norms.

5. For operators $A \in \mathrm{L}(\mathcal{Z}, \mathcal{W})$, $B \in \mathrm{L}(\mathcal{Y}, \mathcal{Z})$, and $C \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, and any choice of $p \in [1, \infty]$, it holds that

$$\|ABC\|_p \leq \|A\|_\infty \|B\|_p \|C\|_\infty. \tag{1.175}$$

It follows that the Schatten p-norm is *submultiplicative*:

$$\|AB\|_p \leq \|A\|_p \|B\|_p. \tag{1.176}$$

6. For every $p \in [1, \infty]$ and every $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$, it holds that

$$\|A\|_p = \bigl\|A^*\bigr\|_p = \bigl\|A^{\mathsf{T}}\bigr\|_p = \bigl\|\overline{A}\bigr\|_p. \tag{1.177}$$

The Schatten 1-norm is commonly called the *trace norm*, the Schatten 2-norm is also known as the *Frobenius norm*, and the Schatten ∞-norm is called the *spectral norm* or *operator norm*. Some additional properties of these three norms are as follows:

1. *The spectral norm.* The spectral norm $\|\cdot\|_\infty$ is special in several respects. It is the operator norm *induced* by the Euclidean norm, which is its defining property (1.166). It also has the property that

$$\|A^*A\|_\infty = \|AA^*\|_\infty = \|A\|_\infty^2 \tag{1.178}$$

for every $A \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$. Hereafter in this book, the spectral norm of an operator A will be written $\|A\|$ rather than $\|A\|_\infty$, which reflects the fundamental importance of this norm.

2. *The Frobenius norm.* Substituting $p = 2$ into the definition of $\|\cdot\|_p$, one sees that the Frobenius norm $\|\cdot\|_2$ is given by

$$\|A\|_2 = \bigl(\mathrm{Tr}(A^*A)\bigr)^{\frac{1}{2}} = \sqrt{\langle A, A\rangle}, \tag{1.179}$$

and is therefore analogous to the Euclidean norm for vectors, but defined by the inner product on $\mathrm{L}(\mathcal{X}, \mathcal{Y})$.

In essence, the Frobenius norm corresponds to the Euclidean norm of an operator viewed as a vector:

$$\|A\|_2 = \|\mathrm{vec}(A)\| = \sqrt{\sum_{a,b} |A(a,b)|^2}, \tag{1.180}$$

where a and b range over the indices of the matrix representation of A.

3. *The trace norm.* Substituting $p = 1$ into the definition of $\|\cdot\|_p$, one has that the trace norm $\|\cdot\|_1$ is given by

$$\|A\|_1 = \mathrm{Tr}\Big(\sqrt{A^*A}\Big), \tag{1.181}$$

which is equal to the sum of the singular values of A. For two density operators $\rho, \sigma \in \mathrm{D}(\mathcal{X})$, the value $\|\rho - \sigma\|_1$ is typically referred to as the *trace distance* between ρ and σ.

A useful expression of $\|X\|_1$, for any square operator $X \in \mathrm{L}(\mathcal{X})$, is

$$\|X\|_1 = \max\{|\langle U, X\rangle| \,:\, U \in \mathrm{U}(\mathcal{X})\}, \tag{1.182}$$

which follows from (1.167) and the singular value theorem (Theorem 1.6). As a result, one has that the trace-norm is non-increasing under the action of partial tracing: for every operator $X \in \mathrm{L}(\mathcal{X} \otimes \mathcal{Y})$, it holds that

$$\begin{aligned} \|\mathrm{Tr}_{\mathcal{Y}}(X)\|_1 &= \max\{|\langle U \otimes \mathbb{1}_{\mathcal{Y}}, X\rangle| \,:\, U \in \mathrm{U}(\mathcal{X})\} \\ &\leq \max\{|\langle V, X\rangle| \,:\, V \in \mathrm{U}(\mathcal{X} \otimes \mathcal{Y})\} = \|X\|_1. \end{aligned} \tag{1.183}$$

The identity

$$\big\|\alpha uu^* - \beta vv^*\big\|_1 = \sqrt{(\alpha+\beta)^2 - 4\alpha\beta|\langle u, v\rangle|^2}, \tag{1.184}$$

which holds for all unit vectors u, v and nonnegative real numbers α, β, is used multiple times in this book. It may be proved by considering the spectrum of $\alpha uu^* - \beta vv^*$; this operator is Hermitian, and has at most two nonzero eigenvalues, represented by the expression

$$\frac{\alpha - \beta}{2} \pm \frac{1}{2}\sqrt{(\alpha+\beta)^2 - 4\alpha\beta|\langle u, v\rangle|^2}. \tag{1.185}$$

In particular, for unit vectors u and v, one has

$$\big\|uu^* - vv^*\big\|_1 = 2\sqrt{1 - |\langle u, v\rangle|^2}. \tag{1.186}$$

1.2 Analysis, convexity, and probability theory

Some of the proofs to be presented in this book will make use of concepts from analysis, convexity, and probability theory. The summary that follows provides an overview of these concepts, narrowly focused on the needs of this book.

1.2.1 Analysis and convexity

In the same spirit as the previous section on linear algebra, it is assumed that the reader is familiar with the most basic notions of mathematical analysis, including the supremum and infimum of sets of real numbers, sequences and limits, and standard univariate calculus over the real numbers.

The discussion below is limited to finite-dimensional real and complex vector spaces—and the reader is cautioned that some of the stated facts rely on the assumption that one is working with finite dimensional spaces. For the remainder of the subsection, $\mathcal{V}$ and $\mathcal{W}$ will denote finite dimensional real or complex vector spaces upon which some norm $\|\cdot\|$ is defined. Unless it is explicitly noted otherwise, the norm may be chosen arbitrarily—so the symbol $\|\cdot\|$ may not necessarily denote the Euclidean norm or spectral norm in this section.

Open and closed sets

A set $\mathcal{A} \subseteq \mathcal{V}$ is *open* if, for every $u \in \mathcal{A}$, there exists $\varepsilon > 0$ such that

$$\{v \in \mathcal{V} \,:\, \|u - v\| < \varepsilon\} \subseteq \mathcal{A}. \tag{1.187}$$

A set $\mathcal{A} \subseteq \mathcal{V}$ is *closed* if the complement of $\mathcal{A}$, defined as

$$\mathcal{V} \backslash \mathcal{A} = \{v \in \mathcal{V} \,:\, v \notin \mathcal{A}\}, \tag{1.188}$$

is open. Given subsets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}$, one defines that $\mathcal{A}$ is open or closed *relative to* $\mathcal{B}$ if $\mathcal{A}$ is the intersection of $\mathcal{B}$ with some set in $\mathcal{V}$ that is open or closed, respectively. Equivalently, $\mathcal{A}$ is open relative to $\mathcal{B}$ if, for every $u \in \mathcal{A}$, there exists a choice of $\varepsilon > 0$ such that

$$\{v \in \mathcal{B} \,:\, \|u - v\| < \varepsilon\} \subseteq \mathcal{A}; \tag{1.189}$$

and $\mathcal{A}$ is closed relative to $\mathcal{B}$ if $\mathcal{B} \backslash \mathcal{A}$ is open relative to $\mathcal{B}$.

For subsets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}$, one defines the *closure* of $\mathcal{A}$ relative to $\mathcal{B}$ as the intersection of all subsets $\mathcal{C}$ such that $\mathcal{A} \subseteq \mathcal{C} \subseteq \mathcal{B}$ and $\mathcal{C}$ is closed relative to $\mathcal{B}$. In other words, this is the smallest set that contains $\mathcal{A}$ and is closed relative to $\mathcal{B}$. The set $\mathcal{A}$ is *dense* in $\mathcal{B}$ if the closure of $\mathcal{A}$ relative to $\mathcal{B}$ is $\mathcal{B}$ itself.

Continuous functions

Let $f : \mathcal{A} \to \mathcal{W}$ be a function defined on some subset $\mathcal{A} \subseteq \mathcal{V}$. For any vector $u \in \mathcal{A}$, the function f is said to be *continuous* at u if the following holds: for every $\varepsilon > 0$ there exists $\delta > 0$ such that

$$\|f(v) - f(u)\| < \varepsilon \tag{1.190}$$

for all $v \in \mathcal{A}$ satisfying $\|u - v\| < \delta$. If f is continuous at every vector in $\mathcal{A}$, then one simply says that f is *continuous on* $\mathcal{A}$.

For a function $f : \mathcal{A} \to \mathcal{W}$ defined on some subset $\mathcal{A} \subseteq \mathcal{V}$, the *preimage* of a set $\mathcal{B} \subseteq \mathcal{W}$ is defined as

$$f^{-1}(\mathcal{B}) = \{u \in \mathcal{A} : f(u) \in \mathcal{B}\}. \tag{1.191}$$

Such a function f is continuous on $\mathcal{A}$ if and only if the preimage of every open set in $\mathcal{W}$ is open relative to $\mathcal{A}$. Equivalently, f is continuous on $\mathcal{A}$ if and only if the preimage of every closed set in $\mathcal{W}$ is closed relative to $\mathcal{A}$.

For a positive real number κ, a function $f : \mathcal{A} \to \mathcal{W}$ defined on a subset $A \subseteq \mathcal{V}$ is said to be a *κ-Lipschitz function* if

$$\|f(u) - f(v)\| \leq \kappa \|u - v\| \tag{1.192}$$

for all $u, v \in \mathcal{A}$. Every κ-Lipschitz function is necessarily continuous.

Compact sets

A set $\mathcal{A} \subseteq \mathcal{V}$ is *compact* if every sequence in $\mathcal{A}$ has a subsequence that converges to a vector $u \in \mathcal{A}$. As a consequence of the fact $\mathcal{V}$ is assumed to be finite dimensional, one has that a set $\mathcal{A} \subseteq \mathcal{V}$ is compact if and only if it is both closed and bounded—a fact known as the *Heine–Borel theorem*.

Two properties regarding continuous functions and compact sets that are particularly noteworthy for the purposes of this book are as follows:

1. If $\mathcal{A}$ is compact and $f : \mathcal{A} \to \mathbb{R}$ is continuous on $\mathcal{A}$, then f achieves both a maximum and minimum value on $\mathcal{A}$.
2. If $\mathcal{A} \subset \mathcal{V}$ is compact and $f : \mathcal{V} \to \mathcal{W}$ is continuous on $\mathcal{A}$, then

$$f(\mathcal{A}) = \{f(u) : u \in \mathcal{A}\} \tag{1.193}$$

is also compact. In words, continuous functions always map compact sets to compact sets.

Differentiation of multivariate real functions

Basic multivariate calculus will be employed in a few occasions later in this book, and in these cases it will be sufficient to consider only real-valued functions.

Suppose n is a positive integer, $f : \mathbb{R}^n \to \mathbb{R}$ is a function, and $u \in \mathbb{R}^n$ is a vector. Under the assumption that the partial derivative

$$\partial_k f(u) = \lim_{\alpha \to 0} \frac{f(u + \alpha e_k) - f(u)}{\alpha} \tag{1.194}$$

exists and is finite for each $k \in \{1, \ldots, n\}$, one defines the *gradient vector* of f at u as

$$\nabla f(u) = (\partial_1 f(u), \ldots, \partial_n f(u)). \tag{1.195}$$

A function $f : \mathbb{R}^n \to \mathbb{R}$ is *differentiable* at a vector $u \in \mathbb{R}^n$ if there exists a vector $v \in \mathbb{R}^n$ with the following property: for every sequence $(w_1, w_2, \ldots)$ of vectors in $\mathbb{R}^n$ that converges to 0, one has that

$$\lim_{k \to \infty} \frac{|f(u + w_k) - f(u) - \langle v, w_k \rangle|}{\|w_k\|} = 0 \tag{1.196}$$

(where here $\|\cdot\|$ denotes the Euclidean norm). In this case the vector v is necessarily unique, and one writes $v = (Df)(u)$. If f is differentiable at u, then it holds that

$$(Df)(u) = \nabla f(u). \tag{1.197}$$

It may be the case that the gradient vector $\nabla f(u)$ is defined for a vector u at which f is not differentiable, but if the function $u \mapsto \nabla f(u)$ is continuous at u, then f is necessarily differentiable at u.

If a function $f : \mathbb{R}^n \to \mathbb{R}$ is both differentiable and κ-Lipschitz, then for all $u \in \mathbb{R}^n$ and for $\|\cdot\|$ denoting the Euclidean norm, it must hold that

$$\|\nabla f(u)\| \leq \kappa. \tag{1.198}$$

Finally, suppose $g_1, \ldots, g_n : \mathbb{R} \to \mathbb{R}$ are functions that are differentiable at a real number $\alpha \in \mathbb{R}$ and $f : \mathbb{R}^n \to \mathbb{R}$ is a function that is differentiable at the vector $(g_1(\alpha), \ldots, g_n(\alpha))$. The *chain rule* for differentiation implies that the function $h : \mathbb{R} \to \mathbb{R}$ defined as

$$h(\beta) = f(g_1(\beta), \ldots, g_n(\beta)) \tag{1.199}$$

is differentiable at α, with its derivative being given by

$$h'(\alpha) = \big\langle \nabla f(g_1(\alpha), \ldots, g_n(\alpha)), (g_1'(\alpha), \ldots, g_n'(\alpha)) \big\rangle. \tag{1.200}$$

Nets

Let $\mathcal{V}$ be a real or complex vector space, let $\mathcal{A} \subseteq \mathcal{V}$ be a subset of $\mathcal{V}$, let $\|\cdot\|$ be a norm on $\mathcal{V}$, and let $\varepsilon > 0$ be a positive real number. A set of vectors $\mathcal{N} \subseteq \mathcal{V}$ is an *ε-net* for $\mathcal{A}$ if, for every vector $u \in \mathcal{A}$, there exists a vector $v \in \mathcal{N}$ such that $\|u - v\| \leq \varepsilon$. An ε-net $\mathcal{N}$ for $\mathcal{A}$ is *minimal* if $\mathcal{N}$ is finite and every ε-net of $\mathcal{A}$ contains at least $|\mathcal{N}|$ vectors.

The following theorem gives an upper bound for the number of elements in a minimal ε-net for the unit ball

$$\mathcal{B}(\mathcal{X}) = \{u \in \mathcal{X} : \|u\| \leq 1\} \tag{1.201}$$

in a complex Euclidean space, with respect to the Euclidean norm.

Theorem 1.8 (Pisier) *Let $\mathcal{X}$ be a complex Euclidean space of dimension n and let $\varepsilon > 0$ be a positive real number. With respect to the Euclidean norm on $\mathcal{X}$, there exists an ε-net $\mathcal{N} \subset \mathcal{B}(\mathcal{X})$ for the unit ball $\mathcal{B}(\mathcal{X})$ such that*

$$|\mathcal{N}| \leq \left(1 + \frac{2}{\varepsilon}\right)^{2n}. \tag{1.202}$$

The proof of this theorem does not require a complicated construction; one may take $\mathcal{N}$ to be any maximal set of vectors chosen from the unit ball for which it holds that $\|u - v\| \geq \varepsilon$ for all $u, v \in \mathcal{N}$ with $u \neq v$. Such a set is necessarily an ε-net for $\mathcal{B}(\mathcal{X})$, and the bound on $|\mathcal{N}|$ is obtained by comparing the volume of $\mathcal{B}(\mathcal{X})$ with the volume of the union of $\varepsilon/2$ balls around vectors in $\mathcal{N}$.

Borel sets and functions

Throughout this subsection, $\mathcal{A} \subseteq \mathcal{V}$ and $\mathcal{B} \subseteq \mathcal{W}$ will denote fixed subsets of finite-dimensional real or complex vector spaces $\mathcal{V}$ and $\mathcal{W}$.

A set $\mathcal{C} \subseteq \mathcal{A}$ is said to be a *Borel subset* of $\mathcal{A}$ if one or more of the following inductively defined properties holds:

1. $\mathcal{C}$ is an open set relative to $\mathcal{A}$.
2. $\mathcal{C}$ is the complement of a Borel subset of $\mathcal{A}$.
3. For $\{\mathcal{C}_1, \mathcal{C}_2, \ldots\}$ being a countable collection of Borel subsets of $\mathcal{A}$, it holds that $\mathcal{C}$ is equal to the union

$$\mathcal{C} = \bigcup_{k=1}^{\infty} \mathcal{C}_k. \tag{1.203}$$

The collection of all Borel subsets of $\mathcal{A}$ is denoted $\text{Borel}(\mathcal{A})$.

A function $f : \mathcal{A} \to \mathcal{B}$ is a *Borel function* if $f^{-1}(\mathcal{C}) \in \text{Borel}(\mathcal{A})$ for all $\mathcal{C} \in \text{Borel}(\mathcal{B})$. That is, Borel functions are functions for which the preimage

of every Borel subset is also a Borel subset. If f is a continuous function, then f is necessarily a Borel function. Another important type of Borel function is any function of the form

$$f(u) = \chi_{\mathcal{C}}(u)\, v \tag{1.204}$$

for any choice of $v \in \mathcal{B}$ and

$$\chi_{\mathcal{C}}(u) = \begin{cases} 1 & \text{if } u \in \mathcal{C} \\ 0 & \text{if } u \notin \mathcal{C} \end{cases} \tag{1.205}$$

being the characteristic function of a Borel subset $\mathcal{C} \in \text{Borel}(\mathcal{A})$.

The collection of all Borel functions $f : \mathcal{A} \to \mathcal{B}$ possesses a variety of closure properties, including the following properties:

1. If $\mathcal{B}$ is a vector space, $f, g : \mathcal{A} \to \mathcal{B}$ are Borel functions, and α is a scalar (either real or complex, depending on whether $\mathcal{B}$ is a real or complex vector space), then the functions αf and $f + g$ are also Borel functions.
2. If $\mathcal{B}$ is a subalgebra of $\text{L}(\mathcal{Z})$, for $\mathcal{Z}$ being a real or complex Euclidean space, and $f, g : \mathcal{A} \to \mathcal{B}$ are Borel functions, then the function $h : \mathcal{A} \to \mathcal{B}$ defined by

   $$h(u) = f(u)g(u) \tag{1.206}$$

   for all $u \in \mathcal{A}$ is also a Borel function. (This includes the special cases $f, g : \mathcal{A} \to \mathbb{R}$ and $f, g : \mathcal{A} \to \mathbb{C}$.)

Measures on Borel sets

A *Borel measure* (or simply a *measure*) defined on $\text{Borel}(\mathcal{A})$ is a function

$$\mu : \text{Borel}(\mathcal{A}) \to [0, \infty] \tag{1.207}$$

that possesses two properties:

1. $\mu(\varnothing) = 0$.
2. For any countable collection $\{\mathcal{C}_1, \mathcal{C}_2, \ldots\} \subseteq \text{Borel}(\mathcal{A})$ of pairwise disjoint Borel subsets of $\mathcal{A}$, it holds that

   $$\mu\left(\bigcup_{k=1}^{\infty} \mathcal{C}_k\right) = \sum_{k=1}^{\infty} \mu(\mathcal{C}_k). \tag{1.208}$$

A measure μ defined on $\text{Borel}(\mathcal{A})$ is said to be *normalized* if it holds that $\mu(\mathcal{A}) = 1$. The term *probability measure* is also used to refer to a normalized measure.

There exists a measure ν defined on Borel($\mathbb{R}$), known as the *standard Borel measure*,[3] that has the property

$$\nu([\alpha, \beta]) = \beta - \alpha \tag{1.209}$$

for all choices of $\alpha, \beta \in \mathbb{R}$ with $\alpha \leq \beta$.

If $\mathcal{A}_1, \ldots, \mathcal{A}_n$ are subsets of (not necessarily equal) finite-dimensional real or complex vector spaces, and

$$\mu_k : \text{Borel}(\mathcal{A}_k) \to [0, \infty] \tag{1.210}$$

is a measure for each $k \in \{1, \ldots, n\}$, then there is a uniquely defined *product measure*

$$\mu_1 \times \cdots \times \mu_n : \text{Borel}(\mathcal{A}_1 \times \cdots \times \mathcal{A}_n) \to [0, \infty] \tag{1.211}$$

for which

$$(\mu_1 \times \cdots \times \mu_n)(\mathcal{B}_1 \times \cdots \times \mathcal{B}_n) = \mu_1(\mathcal{B}_1) \cdots \mu_n(\mathcal{B}_n) \tag{1.212}$$

for all $\mathcal{B}_1 \in \text{Borel}(\mathcal{A}_1), \ldots, \mathcal{B}_n \in \text{Borel}(\mathcal{A}_n)$.

Integration of Borel functions

For some (but not all) Borel functions $f : \mathcal{A} \to \mathcal{B}$, and for μ being a Borel measure of the form $\mu : \text{Borel}(\mathcal{A}) \to [0, \infty]$, one may define the integral

$$\int f(u) \, \mathrm{d}\mu(u), \tag{1.213}$$

which is an element of $\mathcal{B}$ when it is defined.

An understanding of the specifics of the definition through which such an integral is defined is not critical within the context of this book, but some readers may find that a high-level overview of the definition is helpful in associating an intuitive meaning to the integrals that do arise. In short, one defines what is meant by the integral of an increasingly large collection of functions, beginning with functions taking nonnegative real values, and then proceeding to vector (or operator) valued functions by taking linear combinations.

1. *Nonnegative simple functions.* A function $g : \mathcal{A} \to [0, \infty)$ is a *nonnegative simple function* if it may be written as

$$g(u) = \sum_{k=1}^{m} \alpha_k \, \chi_k(u) \tag{1.214}$$

[3] The standard Borel measure agrees with the well-known *Lebesgue measure* on every Borel subset of $\mathbb{R}$. The Lebesgue measure is also defined for some subsets of $\mathbb{R}$ that are not Borel subsets, which endows it with additional properties that happen not to be relevant within the context of this book.

for a nonnegative integer m, distinct positive real numbers $\alpha_1, \ldots, \alpha_m$, and characteristic functions $\chi_1, \ldots, \chi_m$ given by

$$\chi_k(u) = \begin{cases} 1 & \text{if } u \in \mathcal{C}_k \\ 0 & \text{if } u \notin \mathcal{C}_k \end{cases} \tag{1.215}$$

for disjoint Borel sets $\mathcal{C}_1, \ldots, \mathcal{C}_m \in \text{Borel}(\mathcal{A})$. (It is to be understood that the sum is empty when $m = 0$, which corresponds to g being identically zero.)

A nonnegative simple function g of the form (1.214) is *integrable* with respect to a measure $\mu : \text{Borel}(\mathcal{A}) \to [0, \infty]$ if $\mu(\mathcal{C}_k)$ is finite for every $k \in \{1, \ldots, m\}$, and in this case the integral of g with respect to μ is defined as

$$\int g(u) \, \mathrm{d}\mu(u) = \sum_{k=1}^{m} \alpha_k \, \mu(\mathcal{C}_k). \tag{1.216}$$

This is a well-defined quantity, by virtue of the fact that the expression (1.214) happens to be unique for a given simple function g.

2. *Nonnegative Borel functions.* The integral of a Borel function of the form $f : \mathcal{A} \to [0, \infty)$, with respect to a given measure $\mu : \text{Borel}(\mathcal{A}) \to [0, \infty]$, is defined as

$$\int f(u) \, \mathrm{d}\mu(u) = \sup \int g(u) \, \mathrm{d}\mu(u), \tag{1.217}$$

where the supremum is taken over all nonnegative simple functions of the form $g : \mathcal{A} \to [0, \infty)$ for which it holds that $g(u) \leq f(u)$ for all $u \in \mathcal{A}$. It is said that f is *integrable* if the supremum value in (1.217) is finite.

3. *Real and complex Borel functions.* A Borel function $g : \mathcal{A} \to \mathbb{R}$ is *integrable* with respect to a measure $\mu : \text{Borel}(\mathcal{A}) \to [0, \infty]$ if there exist integrable Borel functions $f_0, f_1 : \mathcal{A} \to [0, \infty)$ such that $g = f_0 - f_1$, and in this case the integral of g with respect to μ is defined as

$$\int g(u) \, \mathrm{d}\mu(u) = \int f_0(u) \, \mathrm{d}\mu(u) - \int f_1(u) \, \mathrm{d}\mu(u). \tag{1.218}$$

Similarly, a Borel function $h : \mathcal{A} \to \mathbb{C}$ is *integrable* with respect to a measure $\mu : \text{Borel}(\mathcal{A}) \to [0, \infty]$ if there exist integrable Borel functions $g_0, g_1 : \mathcal{A} \to \mathbb{R}$ such that $h = g_0 + i g_1$, and in this case the integral of h with respect to μ is defined as

$$\int h(u) \, \mathrm{d}\mu(u) = \int g_0(u) \, \mathrm{d}\mu(u) + i \int g_1(u) \, \mathrm{d}\mu(u). \tag{1.219}$$

4. *Arbitrary Borel functions.* An arbitrary Borel function $f : \mathcal{A} \to \mathcal{B}$ is *integrable* with respect to a given measure $\mu : \text{Borel}(\mathcal{A}) \to [0, \infty]$ if there exists a finite-dimensional vector space $\mathcal{W}$ such that $\mathcal{B} \subseteq \mathcal{W}$, a basis $\{w_1, \ldots, w_m\}$ of $\mathcal{W}$, and integrable functions $g_1, \ldots, g_m : \mathcal{A} \to \mathbb{R}$ or $g_1, \ldots, g_m : \mathcal{A} \to \mathbb{C}$ (depending on whether $\mathcal{W}$ is a real or complex vector space) such that

$$f(u) = \sum_{k=1}^{m} g_k(u) w_k. \tag{1.220}$$

In this case, the integral of f with respect to μ is defined as

$$\int f(u) \, \mathrm{d}\mu(u) = \sum_{k=1}^{m} \left(\int g_k(u) \, \mathrm{d}\mu(u) \right) w_k. \tag{1.221}$$

The fact that the third and fourth items in this list lead to uniquely defined integrals of integrable functions is not immediate and requires a proof.

A selection of properties and conventions regarding integrals defined in this way, targeted to the specific needs of this book, follows.

1. *Linearity.* For integrable functions f and g, and scalar values α and β, one has

$$\int (\alpha f(u) + \beta g(u)) \, \mathrm{d}\mu(u) = \alpha \int f(u) \, \mathrm{d}\mu(u) + \beta \int g(u) \, \mathrm{d}\mu(u). \tag{1.222}$$

2. *Standard Borel measure as the default.* Hereafter in this book, whenever $f : \mathbb{R} \to \mathbb{R}$ is an integrable function, and ν denotes the standard Borel measure on $\mathbb{R}$, the shorthand notation

$$\int f(\alpha) \, \mathrm{d}\alpha = \int f(\alpha) \, \mathrm{d}\nu(\alpha) \tag{1.223}$$

will be used. It is the case that, whenever f is an integrable function for which the commonly studied *Riemann integral* is defined, the Riemann integral will be in agreement with the integral defined as above for the standard Borel measure—so this shorthand notation is not likely to lead to confusion or ambiguity.

3. *Integration over subsets.* For an integrable function $f : \mathcal{A} \to \mathcal{B}$ and a Borel subset $\mathcal{C} \in \text{Borel}(\mathcal{A})$, one defines

$$\int_{\mathcal{C}} f(u) \, \mathrm{d}\mu(u) = \int f(u) \chi_{\mathcal{C}}(u) \, \mathrm{d}\mu(u), \tag{1.224}$$

for $\chi_{\mathcal{C}}$ being the characteristic function of $\mathcal{C}$. The notation

$$\int_{\beta}^{\gamma} f(\alpha)\,\mathrm{d}\alpha = \int_{[\beta,\gamma]} f(\alpha)\,\mathrm{d}\alpha \tag{1.225}$$

is also used in the case that f takes the form $f : \mathbb{R} \to \mathcal{B}$ and $\beta, \gamma \in \mathbb{R}$ satisfy $\beta \leq \gamma$.

4. *Order of integration.* Suppose that $\mathcal{A}_0 \subseteq \mathcal{V}_0$, $\mathcal{A}_1 \subseteq \mathcal{V}_1$, and $\mathcal{B} \subseteq \mathcal{W}$ are subsets of finite-dimensional real or complex vector spaces, where it is to be assumed that $\mathcal{V}_0$ and $\mathcal{V}_1$ are either both real or both complex for simplicity. If $\mu_0 : \mathrm{Borel}(\mathcal{A}_0) \to [0,\infty]$ and $\mu_1 : \mathrm{Borel}(\mathcal{A}_1) \to [0,\infty]$ are Borel measures, $f : \mathcal{A}_0 \times \mathcal{A}_1 \to \mathcal{B}$ is a Borel function, and f is integrable with respect to the product measure $\mu_0 \times \mu_1$, then it holds (by a theorem known as *Fubini's theorem*) that

$$\begin{aligned} \int \biggl(\int f(u,v)\,\mathrm{d}\mu_0(u) \biggr) \mathrm{d}\mu_1(v) &= \int f(u,v)\,\mathrm{d}(\mu_0 \times \mu_1)(u,v) \\ &= \int \biggl(\int f(u,v)\,\mathrm{d}\mu_1(v) \biggr) \mathrm{d}\mu_0(u). \end{aligned} \tag{1.226}$$

Convex sets, cones, and functions

Let $\mathcal{V}$ be a vector space over the real or complex numbers. A subset $\mathcal{C}$ of $\mathcal{V}$ is *convex* if, for all vectors $u, v \in \mathcal{C}$ and scalars $\lambda \in [0,1]$, it holds that

$$\lambda u + (1-\lambda)v \in \mathcal{C}. \tag{1.227}$$

Intuitively speaking, this means that for any two distinct elements u and v of $\mathcal{C}$, the line segment whose endpoints are u and v lies entirely within $\mathcal{C}$. The intersection of any collection of convex sets is also convex.

If $\mathcal{V}$ and $\mathcal{W}$ are vector spaces, either both over the real numbers or both over the complex numbers, and $\mathcal{A} \subseteq \mathcal{V}$ and $\mathcal{B} \subseteq \mathcal{W}$ are convex sets, then the set

$$\{u \oplus v \,:\, u \in \mathcal{A},\, v \in \mathcal{B}\} \subseteq \mathcal{V} \oplus \mathcal{W} \tag{1.228}$$

is also convex. Moreover, if $A \in \mathrm{L}(\mathcal{V}, \mathcal{W})$ is an operator, then the set

$$\{Au \,:\, u \in \mathcal{A}\} \subseteq \mathcal{W} \tag{1.229}$$

is convex as well.

A set $\mathcal{K} \subseteq \mathcal{V}$ is a *cone* if, for all choices of $u \in \mathcal{K}$ and $\lambda \geq 0$, one has that $\lambda u \in \mathcal{K}$. The cone *generated* by a set $\mathcal{A} \subseteq \mathcal{V}$ is defined as

$$\mathrm{cone}(\mathcal{A}) = \{\lambda u \,:\, u \in \mathcal{A},\, \lambda \geq 0\}. \tag{1.230}$$

If $\mathcal{A}$ is a compact set that does not include 0, then $\operatorname{cone}(\mathcal{A})$ is necessarily a closed set. A *convex cone* is simply a cone that is also convex. A cone $\mathcal{K}$ is convex if and only if it is closed under addition, meaning that $u + v \in \mathcal{K}$ for every choice of $u, v \in \mathcal{K}$.

A function $f : \mathcal{C} \to \mathbb{R}$ defined on a convex set $\mathcal{C} \subseteq \mathcal{V}$ is a *convex function* if the inequality

$$f(\lambda u + (1-\lambda)v) \leq \lambda f(u) + (1-\lambda) f(v) \tag{1.231}$$

holds for all $u, v \in \mathcal{C}$ and $\lambda \in [0, 1]$. A function $f : \mathcal{C} \to \mathbb{R}$ defined on a convex set $\mathcal{C} \subseteq \mathcal{V}$ is a *midpoint convex function* if the inequality

$$f\Big(\frac{u+v}{2}\Big) \leq \frac{f(u) + f(v)}{2} \tag{1.232}$$

holds for all $u, v \in \mathcal{C}$. Every continuous midpoint convex function is convex.

A function $f : \mathcal{C} \to \mathbb{R}$ defined on a convex set $\mathcal{C} \subseteq \mathcal{V}$ is a *concave function* if $-f$ is convex. Equivalently, f is concave if the reverse of the inequality (1.231) holds for all $u, v \in \mathcal{C}$ and $\lambda \in [0, 1]$. Similarly, a function $f : \mathcal{C} \to \mathbb{R}$ defined on a convex set $\mathcal{C} \subseteq \mathcal{V}$ is a *midpoint concave function* if $-f$ is a midpoint convex function, and therefore every continuous midpoint concave function is concave.

Convex hulls

For any alphabet Σ, a vector $p \in \mathbb{R}^{\Sigma}$ is said to be a *probability vector* if it holds that $p(a) \geq 0$ for all $a \in \Sigma$ and

$$\sum_{a \in \Sigma} p(a) = 1. \tag{1.233}$$

The set of all such vectors will be denoted $\mathcal{P}(\Sigma)$.

For any vector space $\mathcal{V}$ and any subset $\mathcal{A} \subseteq \mathcal{V}$, a *convex combination* of vectors in $\mathcal{A}$ is any expression of the form

$$\sum_{a \in \Sigma} p(a) u_a, \tag{1.234}$$

for some choice of an alphabet Σ, a probability vector $p \in \mathcal{P}(\Sigma)$, and a collection

$$\{u_a : a \in \Sigma\} \subseteq \mathcal{A} \tag{1.235}$$

of vectors in $\mathcal{A}$.

The *convex hull* of a set $\mathcal{A} \subseteq \mathcal{V}$, denoted $\operatorname{conv}(\mathcal{A})$, is the intersection of all convex sets containing $\mathcal{A}$. The set $\operatorname{conv}(\mathcal{A})$ is equal to the set of all vectors that may be written as a convex combination of elements of $\mathcal{A}$. (This is true

even in the case that $\mathcal{A}$ is infinite.) The convex hull $\operatorname{conv}(\mathcal{A})$ of a closed set $\mathcal{A}$ need not itself be closed. However, if $\mathcal{A}$ is compact, then so too is $\operatorname{conv}(\mathcal{A})$.

The theorem that follows provides an upper bound on the number of elements over which one must take convex combinations in order to generate every point in the convex hull of a given set. The theorem refers to the notion of an *affine subspace*: a set $\mathcal{U} \subseteq \mathcal{V}$ is an affine subspace of $\mathcal{V}$ having dimension n if there exists a subspace $\mathcal{W} \subseteq \mathcal{V}$ of dimension n and a vector $u \in \mathcal{V}$ such that

$$\mathcal{U} = \{u + v \,:\, v \in \mathcal{W}\}. \tag{1.236}$$

Theorem 1.9 (Carathéodory's theorem) *Let $\mathcal{V}$ be a real vector space and let $\mathcal{A}$ be a subset of $\mathcal{V}$. Assume, moreover, that $\mathcal{A}$ is contained in an affine subspace of $\mathcal{V}$ having dimension n. For every vector $v \in \operatorname{conv}(\mathcal{A})$ in the convex hull of $\mathcal{A}$, there exist $m \leq n + 1$ vectors $u_1, \ldots, u_m \in \mathcal{A}$ such that $v \in \operatorname{conv}(\{u_1, \ldots, u_m\})$.*

Extreme points

A point $w \in \mathcal{C}$ in a convex set $\mathcal{C}$ is said to be an *extreme point* of $\mathcal{C}$ if, for every expression

$$w = \lambda u + (1 - \lambda)v \tag{1.237}$$

for which $u, v \in \mathcal{C}$ and $\lambda \in (0, 1)$, it holds that $u = v = w$. In words, the extreme points are those elements of $\mathcal{C}$ that do not lie properly between two distinct points of $\mathcal{C}$.

The following theorem states that every convex and compact subset of a finite-dimensional vector space, over the real or complex numbers, is equal to the convex hull of its extreme points.

Theorem 1.10 (Minkowski) *Let $\mathcal{V}$ be a finite-dimensional vector space over the real or complex numbers, let $\mathcal{C} \subseteq \mathcal{V}$ be a compact and convex set, and let $\mathcal{A} \subseteq \mathcal{C}$ be the set of extreme points of $\mathcal{C}$. It holds that $\mathcal{C} = \operatorname{conv}(\mathcal{A})$.*

A few examples of convex and compact sets, along with an identification of their extreme points, follow.

1. *The spectral norm unit ball.* For any complex Euclidean space $\mathcal{X}$, the set

   $$\{X \in \mathrm{L}(\mathcal{X}) \,:\, \|X\| \leq 1\} \tag{1.238}$$

   is a convex and compact set. The extreme points of this set are the unitary operators $\mathrm{U}(\mathcal{X})$.

2. *The trace norm unit ball.* For any complex Euclidean space $\mathcal{X}$, the set

$$\{X \in \mathrm{L}(\mathcal{X}) : \|X\|_1 \leq 1\} \tag{1.239}$$

is a convex and compact set. The extreme points of this set are those operators of the form uv^* for $u, v \in \mathcal{S}(\mathcal{X})$ unit vectors.

3. *Density operators.* For any complex Euclidean space $\mathcal{X}$, the set $\mathrm{D}(\mathcal{X})$ of density operators acting on $\mathcal{X}$ is convex and compact. The extreme points of $\mathrm{D}(\mathcal{X})$ coincide with the rank-one projection operators. These are the operators of the form uu^* for $u \in \mathcal{S}(\mathcal{X})$ being a unit vector.

4. *Probability vectors.* For any alphabet Σ, the set of probability vectors $\mathcal{P}(\Sigma)$ is convex and compact. The extreme points of this set are the elements of the standard basis $\{e_a : a \in \Sigma\}$ of $\mathbb{R}^\Sigma$.

Hyperplane separation and min-max theorems

Convex sets in real Euclidean spaces possess a fundamentally important property: every vector lying outside of a given convex set in a real Euclidean space can be separated from that convex set by a *hyperplane*. That is, if the underlying real Euclidean space has dimension n, then there exists an affine subspace of that space having dimension $n - 1$ that divides the entire space into two half-spaces: one contains the convex set and the other contains the chosen point lying outside of the convex set. The following theorem represents one specific formulation of this fact.

Theorem 1.11 (Hyperplane separation theorem) *Let $\mathcal{V}$ be a real Euclidean space, let $\mathcal{C} \subset \mathcal{V}$ be a closed, convex subset of $\mathcal{V}$, and let $u \in \mathcal{V}$ be a vector with $u \notin \mathcal{C}$. There exists a vector $v \in \mathcal{V}$ and a scalar $\alpha \in \mathbb{R}$ such that*

$$\langle v, u \rangle < \alpha \leq \langle v, w \rangle \tag{1.240}$$

for all $w \in \mathcal{C}$. If $\mathcal{C}$ is a cone, then v may be chosen so that (1.240) *holds for $\alpha = 0$.*

Another theorem concerning convex sets that finds uses in the theory of quantum information is the following theorem.

Theorem 1.12 (Sion's min-max theorem) *Let $\mathcal{X}$ and $\mathcal{Y}$ be real or complex Euclidean spaces, let $\mathcal{A} \subseteq \mathcal{X}$ and $\mathcal{B} \subseteq \mathcal{Y}$ be convex sets with $\mathcal{B}$ compact, and let $f : \mathcal{A} \times \mathcal{B} \to \mathbb{R}$ be a continuous function such that*

1. *$u \mapsto f(u, v)$ is a convex function on $\mathcal{A}$ for all $v \in \mathcal{B}$, and*
2. *$v \mapsto f(u, v)$ is a concave function on $\mathcal{B}$ for all $u \in \mathcal{A}$.*

It holds that

$$\inf_{u\in\mathcal{A}} \max_{v\in\mathcal{B}} f(u,v) = \max_{v\in\mathcal{B}} \inf_{u\in\mathcal{A}} f(u,v). \tag{1.241}$$

1.2.2 Probability theory

Concepts from probability theory will play an important role throughout much of this book. Probability distributions over alphabets or other finite sets will be viewed as having fundamental importance; they arise naturally when information-theoretic tasks and settings are considered. The reader is assumed to have familiarity with basic probability theory for distributions over sets with finitely many elements. It will also be convenient to use the language of probability theory to discuss properties of Borel measures.

Random variables distributed with respect to probability measures

Suppose $\mathcal{A}$ is a subset of a finite-dimensional real or complex vector space $\mathcal{V}$ and $\mu : \text{Borel}(\mathcal{A}) \to [0,1]$ is a probability measure (by which it is meant that μ is a normalized Borel measure). A *random variable* X distributed with respect to μ is a real-valued, integrable Borel function of the form

$$X : \mathcal{A} \to \mathbb{R}, \tag{1.242}$$

which is typically viewed as representing an outcome of a random process of some sort.

For every Borel subset $\mathcal{B} \subseteq \mathbb{R}$ of the real numbers, the probability that X takes a value in $\mathcal{B}$ is defined as

$$\Pr(X \in \mathcal{B}) = \mu\big(\{u \in \mathcal{A} : X(u) \in \mathcal{B}\}\big). \tag{1.243}$$

As a matter of notational convenience, one often writes expressions such as

$$\Pr(X \geq \beta) \quad \text{and} \quad \Pr(|X - \beta| \geq \varepsilon), \tag{1.244}$$

which are to be understood as meaning $\Pr(X \in \mathcal{B})$ for

$$\mathcal{B} = \{\alpha \in \mathbb{R} : \alpha \geq \beta\} \quad \text{and} \quad \mathcal{B} = \{\alpha \in \mathbb{R} : |\alpha - \beta| \geq \varepsilon\}, \tag{1.245}$$

respectively. Other expressions of this form are interpreted in an analogous way.

The *union bound* states, for any random variable X and arbitrary Borel subsets $\mathcal{B}_1, \ldots, \mathcal{B}_n$ of $\mathbb{R}$, that

$$\Pr\big(X \in \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_n\big) \leq \Pr(X \in \mathcal{B}_1) + \cdots + \Pr(X \in \mathcal{B}_n). \tag{1.246}$$

The *expected value* (or *mean value*) of a random variable X, distributed with respect to a probability measure $\mu : \text{Borel}(\mathcal{A}) \to [0,1]$, is defined as

$$\mathrm{E}(X) = \int X(u)\,\mathrm{d}\mu(u). \tag{1.247}$$

If X is a random variable taking nonnegative real values, then it holds that

$$\mathrm{E}(X) = \int_0^\infty \Pr(X \geq \lambda)\,\mathrm{d}\lambda. \tag{1.248}$$

Random variables for discrete distributions

For a given alphabet Σ and a probability vector $p \in \mathcal{P}(\Sigma)$, one may also define a random variable X, distributed with respect to p, in an analogous way to a random variable distributed with respect to a Borel measure. In particular, such a random variable is a function of the form

$$X : \Sigma \to \mathbb{R}, \tag{1.249}$$

and for every subset $\Gamma \subseteq \Sigma$ one writes

$$\Pr(X \in \Gamma) = \sum_{a \in \Gamma} p(a). \tag{1.250}$$

In this case, the *expected value* (or *mean value*) of X is

$$\mathrm{E}(X) = \sum_{a \in \Sigma} p(a) X(a). \tag{1.251}$$

It is, in some sense, not necessary for random variables distributed with respect to probability vectors of the form $p \in \mathcal{P}(\Sigma)$ to be viewed as being fundamentally different from random variables distributed with respect to Borel probability measures. Indeed, one may consider the set

$$\{1, \ldots, n\} \subset \mathbb{R}, \tag{1.252}$$

for some choice of a positive integer n, and observe that every subset of $\{1, \ldots, n\}$ is a Borel subset of this set. The Borel probability measures

$$\mu : \text{Borel}(\{1, \ldots, n\}) \to [0,1] \tag{1.253}$$

coincide precisely with the set of all probability vectors $p \in \mathcal{P}(\{1, \ldots, n\})$ through the equations

$$\mu(\mathcal{B}) = \sum_{b \in \mathcal{B}} p(b) \quad \text{and} \quad p(a) = \mu(\{a\}), \tag{1.254}$$

for every $\mathcal{B} \subseteq \{1, \ldots, n\}$ and $a \in \{1, \ldots, n\}$.

Thus, by associating an arbitrary alphabet Σ with the set $\{1, \ldots, n\}$, one finds that a random variable distributed with respect to a probability vector $p \in \mathcal{P}(\Sigma)$ is represented by a random variable distributed with respect to a Borel probability measure.

Vector and operator valued random variables

It is sometimes convenient to define random variables that take vector or operator values, rather than real number values. Random variables of this sort will always be specified explicitly in terms of ordinary random variables (i.e., ones that take real values) in this book. For example, given random variables $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$, for some choice of a positive integer n, one may refer to the vector-valued random variables

$$(X_1, \ldots, X_n) \in \mathbb{R}^n \quad \text{and} \quad (X_1 + iY_1, \ldots, X_n + iY_n) \in \mathbb{C}^n. \tag{1.255}$$

The default meaning of the term *random variable* should be understood as referring to real-valued random variables, and the term *vector-valued random variable* or *operator-valued random variable* will be used when referring to random variables obtained in the manner just described.

Independent and identically distributed random variables

Two random variables X and Y are said to be *independent* if

$$\Pr((X, Y) \in \mathcal{A} \times \mathcal{B}) = \Pr(X \in \mathcal{A}) \Pr(Y \in \mathcal{B}) \tag{1.256}$$

for every choice of Borel subsets $\mathcal{A}, \mathcal{B} \subseteq \mathbb{R}$, and are said to be *identically distributed* if

$$\Pr(X \in \mathcal{A}) = \Pr(Y \in \mathcal{A}) \tag{1.257}$$

for every Borel subset $\mathcal{A} \subseteq \mathbb{R}$. In general, these conditions do not require that X and Y are defined with respect to the same Borel measure. In both cases, these notions may be extended to more than two random variables, as well as to vector-valued random variables, in a straightforward way.

Suppose that $\mathcal{A}$ is a subset of a finite-dimensional real or complex vector space, $\mu : \text{Borel}(\mathcal{A}) \to [0, 1]$ is a probability measure, and $Y : \mathcal{A} \to \mathbb{R}$ is a random variable distributed with respect to μ. For any choice of a positive integer n, one may consider *independent and identically distributed* random variables $X_1, \ldots, X_n$, each being distributed in the same way as Y. For the purposes of this book, one may assume without a loss of generality that this means that $X_1, \ldots, X_n$ are Borel functions, taking the form

$$X_k : \mathcal{A}^n \to \mathbb{R} \tag{1.258}$$

and being defined as

$$X_k(u_1, \ldots, u_n) = Y(u_k) \tag{1.259}$$

for each k and each $(u_1, \ldots, u_n) \in \mathcal{A}^n$. Moreover, each X_k is understood to be distributed with respect to the n-fold product measure $\mu \times \cdots \times \mu$ on $\mathcal{A}^n$. In essence, this formal specification represents the simple and intuitive notion that $X_1, \ldots, X_n$ are uncorrelated copies of the random variable Y.

A few fundamental theorems

A few fundamental theorems concerning random variables will be used later in this book. While these theorem do hold for more general notions of random variables, the theorem statements that follow should be understood to apply to random variables distributed with respect to Borel probability measures (including random variables distributed with respect to probability vectors of the form $p \in \mathcal{P}(\Sigma)$ as a special case, as described above).

The first theorem to be stated in this subsection is *Markov's inequality*, which provides a sometimes coarse upper bound on the probability that a nonnegative random variable exceeds a given threshold value.

Theorem 1.13 (Markov's inequality) *Let X be a random variable taking nonnegative real values, and let $\varepsilon > 0$ be a positive real number. It holds that*

$$\Pr(X \geq \varepsilon) \leq \frac{\mathrm{E}(X)}{\varepsilon}. \tag{1.260}$$

The next theorem, known as *Jensen's inequality*, concerns the expected value of a convex function applied to a random variable.

Theorem 1.14 (Jensen's inequality) *Suppose that X is a random variable and $f : \mathbb{R} \to \mathbb{R}$ is a convex function. It holds that*

$$f(\mathrm{E}(X)) \leq \mathrm{E}(f(X)). \tag{1.261}$$

Two additional theorems—known as the *weak law of large numbers* and *Hoeffding's inequality*—provide bounds on the deviation of the average value of a collection of independent and identically distributed random variables from their mean value.

Theorem 1.15 (Weak law of large numbers) *Let X be a random variable and let $\alpha = \mathrm{E}(X)$. Assume, moreover, for every positive integer n, that $X_1, \ldots, X_n$ are independent random variables identically distributed to X. For every positive real number $\varepsilon > 0$, it holds that*

$$\lim_{n \to \infty} \Pr\left(\left| \frac{X_1 + \cdots + X_n}{n} - \alpha \right| \geq \varepsilon \right) = 0. \tag{1.262}$$

Theorem 1.16 (Hoeffding's inequality) *Let $X_1, \ldots, X_n$ be independent and identically distributed random variables taking values in the interval $[0, 1]$ and having mean value α. For every positive real number $\varepsilon > 0$ it holds that*

$$\Pr\left(\left|\frac{X_1 + \cdots + X_n}{n} - \alpha\right| \geq \varepsilon\right) \leq 2\exp(-2n\varepsilon^2). \tag{1.263}$$

Gaussian measure and normally distributed random variables

The *standard Gaussian measure* on $\mathbb{R}$ is the Borel probability measure

$$\gamma : \text{Borel}(\mathbb{R}) \to [0, 1] \tag{1.264}$$

defined as

$$\gamma(\mathcal{A}) = \frac{1}{\sqrt{2\pi}} \int_{\mathcal{A}} \exp\left(-\frac{\alpha^2}{2}\right) \mathrm{d}\alpha \tag{1.265}$$

for every $\mathcal{A} \in \text{Borel}(\mathbb{R})$, where the integral is to be taken with respect to the standard Borel measure on $\mathbb{R}$. The fact that this is a well-defined measure follows from the observation that the function

$$\alpha \mapsto \begin{cases} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\alpha^2}{2}\right) & \text{if } \alpha \in \mathcal{A} \\ 0 & \text{otherwise} \end{cases} \tag{1.266}$$

is an integrable Borel function for every Borel subset $\mathcal{A} \subseteq \mathbb{R}$, and the fact that it is a probability measure follows from the Gaussian integral

$$\int \exp\left(-\frac{\alpha^2}{2}\right) \mathrm{d}\alpha = \sqrt{2\pi}. \tag{1.267}$$

A random variable X is a *standard normal random variable* if it holds that $\Pr(X \in \mathcal{A}) = \gamma(\mathcal{A})$ for every $A \in \text{Borel}(\mathbb{R})$. This is equivalent to saying that X is identically distributed to a random variable $Y(\alpha) = \alpha$ distributed with respect to the standard Gaussian measure γ on $\mathbb{R}$.

The following integrals are among many integrals of a similar sort that are useful when reasoning about standard normal random variables:

1. For every positive real number $\lambda > 0$ and every real number $\beta \in \mathbb{R}$ it holds that

$$\int \exp(-\lambda\alpha^2 + \beta\alpha) \, \mathrm{d}\alpha = \sqrt{\frac{\pi}{\lambda}} \exp\left(\frac{\beta^2}{4\lambda}\right). \tag{1.268}$$

2. For every positive integer n, it holds that

$$\int_0^\infty \alpha^n \, \mathrm{d}\gamma(\alpha) = \frac{2^{\frac{n}{2}} \Gamma\left(\frac{n+1}{2}\right)}{2\sqrt{\pi}}, \tag{1.269}$$

where the Γ-function may be defined at positive half-integer points as follows:

$$\Gamma\Big(\frac{m+1}{2}\Big) = \begin{cases} \sqrt{\pi} & \text{if } m = 0 \\ 1 & \text{if } m = 1 \\ \frac{m-1}{2} \, \Gamma\left(\frac{m-1}{2}\right) & \text{if } m \geq 2. \end{cases} \tag{1.270}$$

3. For every positive real number $\lambda > 0$ and every pair of real numbers $\beta_0, \beta_1 \in \mathbb{R}$ with $\beta_0 \leq \beta_1$ it holds that

$$\int_{\beta_0}^{\beta_1} \alpha \exp(-\lambda \alpha^2) \, \mathrm{d}\alpha = \frac{1}{2\lambda} \exp(-\lambda \beta_0^2) - \frac{1}{2\lambda} \exp(-\lambda \beta_1^2). \tag{1.271}$$

This formula also holds for infinite values of β_0 and β_1, with the natural interpretation $\exp(-\infty) = 0$.

For every positive integer n, the *standard Gaussian measure* on $\mathbb{R}^n$ is the Borel probability measure

$$\gamma_n : \mathrm{Borel}(\mathbb{R}^n) \to [0, 1] \tag{1.272}$$

obtained by taking the n-fold product measure of γ with itself. Equivalently,

$$\gamma_n(\mathcal{A}) = (2\pi)^{-\frac{n}{2}} \int_{\mathcal{A}} \exp\left(-\frac{\|u\|^2}{2}\right) \mathrm{d}\nu_n(u), \tag{1.273}$$

where ν_n denotes the n-fold product measure of the standard Borel measure ν with itself and the norm is the Euclidean norm.

The standard Gaussian measure on $\mathbb{R}^n$ is invariant under orthogonal transformations (which include rotations):

$$\gamma_n(U\mathcal{A}) = \gamma_n(\mathcal{A}) \tag{1.274}$$

for every Borel set $\mathcal{A} \subseteq \mathbb{R}^n$ and every orthogonal operator $U \in \mathrm{L}(\mathbb{R}^n)$, meaning one that satisfies $UU^{\mathsf{T}} = \mathbb{1}$. Therefore, for independent standard normal random variables $X_1, \ldots, X_n$, one has that the vector valued random variable $(X_1, \ldots, X_n)$ is identically distributed to the vector-valued random variable $(Y_1, \ldots, Y_n)$ obtained by defining

$$Y_k = \sum_{j=1}^{n} U(k, j) X_j \tag{1.275}$$

for each $k \in \{1, \ldots, n\}$, for $U \in \mathrm{L}(\mathbb{R}^n)$ being any orthogonal operator. As a consequence of this fact, one has that if the standard Gaussian measure is projected onto a subspace, it is equivalent to the standard Gaussian measure on that subspace.

Proposition 1.17 *Let m and n be positive integers satisfying $m < n$ and let $V \in \mathrm{L}(\mathbb{R}^m, \mathbb{R}^n)$ satisfy $V^{\mathsf{T}}V = \mathbb{1}$. For every Borel set $\mathcal{A} \subseteq \mathbb{R}^m$, one has*

$$\gamma_m(\mathcal{A}) = \gamma_n\big(\{u \in \mathbb{R}^n \,:\, V^{\mathsf{T}}u \in \mathcal{A}\}\big). \tag{1.276}$$

It follows from this proposition that the standard Gaussian measure $\gamma_n(\mathcal{V})$ of any proper subspace $\mathcal{V}$ of $\mathbb{R}^n$ is zero.

Finally, for independent standard normal random variables $X_1, \ldots, X_n$, one may define a random variable

$$Y = \sqrt{X_1^2 + \cdots + X_n^2}. \tag{1.277}$$

The distribution of Y is known as the *χ-distribution*. The mean value of Y has the following closed-form expression:

$$\mathrm{E}(Y) = \frac{\sqrt{2}\,\Gamma\big(\frac{n+1}{2}\big)}{\Gamma\big(\frac{n}{2}\big)}. \tag{1.278}$$

From this expression, it may be proved that

$$\mathrm{E}(Y) = \upsilon_n \sqrt{n}, \tag{1.279}$$

where $(\upsilon_1, \upsilon_2, \ldots)$ is a strictly increasing sequence that begins

$$\upsilon_1 = \sqrt{\frac{2}{\pi}}, \quad \upsilon_2 = \frac{\sqrt{\pi}}{2}, \quad \upsilon_3 = \sqrt{\frac{8}{3\pi}}, \quad \ldots \tag{1.280}$$

and converges to 1 in the limit as n goes to infinity.

1.2.3 Semidefinite programming

The paradigm of *semidefinite programming* finds numerous applications in the theory of quantum information, both analytical and computational. This section describes a formulation of semidefinite programming that is well-suited to its (primarily analytical) applications found in this book.

Definitions associated with semidefinite programs

Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces, let $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ be a Hermitian-preserving map, and let $A \in \mathrm{Herm}(\mathcal{X})$ and $B \in \mathrm{Herm}(\mathcal{Y})$ be Hermitian operators. A *semidefinite program* is a triple (Φ, A, B), with which the following pair of optimization problems is associated:

Primal problem		Dual problem	
maximize:	$\langle A, X\rangle$	minimize:	$\langle B, Y\rangle$
subject to:	$\Phi(X) = B,$	subject to:	$\Phi^*(Y) \geq A,$
	$X \in \mathrm{Pos}(\mathcal{X}).$		$Y \in \mathrm{Herm}(\mathcal{Y}).$

With these problems in mind, one defines the *primal feasible* set $\mathcal{A}$ and the *dual feasible* set $\mathcal{B}$ of (Φ, A, B) as follows:

$$\begin{aligned} \mathcal{A} &= \{X \in \mathrm{Pos}(\mathcal{X}) \,:\, \Phi(X) = B\}, \\ \mathcal{B} &= \{Y \in \mathrm{Herm}(\mathcal{Y}) \,:\, \Phi^*(Y) \geq A\}. \end{aligned} \tag{1.281}$$

Operators $X \in \mathcal{A}$ and $Y \in \mathcal{B}$ are also said to be *primal feasible* and *dual feasible*, respectively.

The function $X \mapsto \langle A, X\rangle$, from $\mathrm{Herm}(\mathcal{X})$ to $\mathbb{R}$, is the *primal objective function*, while the function $Y \mapsto \langle B, Y\rangle$, from $\mathrm{Herm}(\mathcal{Y})$ to $\mathbb{R}$, is the *dual objective function* of (Φ, A, B). The *optimum values* associated with the primal and dual problems are defined as

$$\alpha = \sup\{\langle A, X\rangle \,:\, X \in \mathcal{A}\} \quad \text{and} \quad \beta = \inf\{\langle B, Y\rangle \,:\, Y \in \mathcal{B}\}, \tag{1.282}$$

respectively. (If it is the case that $\mathcal{A} = \varnothing$ or $\mathcal{B} = \varnothing$, then one defines $\alpha = -\infty$ and $\beta = \infty$, respectively.)

Semidefinite programming duality

Semidefinite programs have associated with them a notion of *duality*, which refers to the special relationship between the primal and dual problems.

The property of *weak duality*, which holds for all semidefinite programs, is that the primal optimum can never exceed the dual optimum. In more succinct terms, it necessarily holds that $\alpha \leq \beta$. This implies that every dual feasible operator $Y \in \mathcal{B}$ provides an upper bound of $\langle B, Y\rangle$ on the value $\langle A, X\rangle$ that is achievable over all choices of a primal feasible $X \in \mathcal{A}$, and likewise every primal feasible operator $X \in \mathcal{A}$ provides a lower bound of $\langle A, X\rangle$ on the value $\langle B, Y\rangle$ that is achievable over all dual feasible operators $Y \in \mathcal{B}$.

It is not always the case that the primal optimum and dual optimum of a semidefinite program (Φ, A, B) agree, but for many semidefinite programs that arise naturally in applications, the primal optimum and dual optimum will be equal. This situation is called *strong duality*. The following theorem provides one set of conditions under which strong duality is guaranteed.

Theorem 1.18 (Slater's theorem for semidefinite programs) *Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces, let $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ be a Hermitian-preserving map, and let $A \in \mathrm{Herm}(\mathcal{X})$ and $B \in \mathrm{Herm}(\mathcal{Y})$ be Hermitian operators. Letting $\mathcal{A}$, $\mathcal{B}$, α, and β be as defined above for the semidefinite program (Φ, A, B), one has the following two implications:*

1. *If α is finite and there exists a Hermitian operator $Y \in \mathrm{Herm}(\mathcal{Y})$ such that $\Phi^*(Y) > A$, then $\alpha = \beta$, and moreover there exists a primal-feasible operator $X \in \mathcal{A}$ such that $\langle A, X \rangle = \alpha$.*
2. *If β is finite and there exists a positive definite operator $X \in \mathrm{Pd}(\mathcal{X})$ such that $\Phi(X) = B$, then $\alpha = \beta$, and moreover there exists a dual-feasible operator $Y \in \mathcal{B}$ such that $\langle B, Y \rangle = \beta$.*

In the situation that the optimum primal and dual values are equal, and are both achieved for some choice of feasible operators, a simple relationship between these operators, known as *complementary slackness*, must hold.

Proposition 1.19 (Complementary slackness for semidefinite programs) *Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces, let $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ be a Hermitian-preserving map, and let $A \in \mathrm{Herm}(\mathcal{X})$ and $B \in \mathrm{Herm}(\mathcal{Y})$ be Hermitian operators. Let $\mathcal{A}$ and $\mathcal{B}$ be the primal-feasible and dual-feasible sets associated with the semidefinite program (Φ, A, B), and suppose that $X \in \mathcal{A}$ and $Y \in \mathcal{B}$ are operators satisfying $\langle A, X \rangle = \langle B, Y \rangle$. It holds that*

$$\Phi^*(Y)X = AX. \tag{1.283}$$

Simplified forms and alternative expressions of semidefinite programs

Semidefinite programs are typically presented in a way that is somewhat less formal than a precise specification of a triple (Φ, A, B), for $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ being a Hermitian-preserving map and $A \in \mathrm{Herm}(\mathcal{X})$ and $B \in \mathrm{Herm}(\mathcal{Y})$ being Hermitian operators. Rather, the primal and dual problems are stated directly, often in a simplified form, and it is sometimes left to the reader to formulate a triple (Φ, A, B) that corresponds to the simplified problem statements.

Two examples of semidefinite programs follow, in both cases including their formal specifications and simplified forms.

Example 1.20 (Semidefinite program for the trace norm) Let $\mathcal{X}$ and $\mathcal{Y}$ be complex Euclidean spaces and let $K \in \mathrm{L}(\mathcal{X}, \mathcal{Y})$ be any operator. Define a Hermitian-preserving map $\Phi \in \mathrm{T}(\mathcal{X} \oplus \mathcal{Y})$ as

$$\Phi \begin{pmatrix} X & \cdot \\ \cdot & Y \end{pmatrix} = \begin{pmatrix} X & 0 \\ 0 & Y \end{pmatrix} \tag{1.284}$$

for all $X \in \mathrm{L}(\mathcal{X})$ and $Y \in \mathrm{L}(\mathcal{Y})$, where the dots represent elements of $\mathrm{L}(\mathcal{X}, \mathcal{Y})$ and $\mathrm{L}(\mathcal{Y}, \mathcal{X})$ that are effectively zeroed out by Φ. The map Φ is self-adjoint: $\Phi^* = \Phi$. Also define $A, B \in \mathrm{Herm}(\mathcal{X} \oplus \mathcal{Y})$ as

$$A = \frac{1}{2}\begin{pmatrix} 0 & K^* \\ K & 0 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} \mathbb{1}_{\mathcal{X}} & 0 \\ 0 & \mathbb{1}_{\mathcal{Y}} \end{pmatrix}. \tag{1.285}$$

The primal and dual problems associated with the semidefinite program (Φ, A, B) may, after some simplifications, be expressed as follows:

Primal problem		Dual problem	
maximize:	$\frac{1}{2}\langle K, Z\rangle + \frac{1}{2}\langle K^*, Z^*\rangle$	minimize:	$\frac{1}{2}\mathrm{Tr}(X) + \frac{1}{2}\mathrm{Tr}(Y)$
subject to:	$\begin{pmatrix} \mathbb{1}_{\mathcal{X}} & Z^* \\ Z & \mathbb{1}_{\mathcal{Y}} \end{pmatrix} \geq 0,$	subject to:	$\begin{pmatrix} X & -K^* \\ -K & Y \end{pmatrix} \geq 0,$
	$Z \in \mathrm{L}(\mathcal{X}, \mathcal{Y}).$		$X \in \mathrm{Pos}(\mathcal{X}),$
			$Y \in \mathrm{Pos}(\mathcal{Y}).$

The primal and dual optima are equal for all choices of K, and given by $\|K\|_1$. (Given a singular value decomposition of K, one can construct both a primal feasible and dual feasible solution achieving this value.)

A standard way of expressing this semidefinite program would be to list only the simplified primal and dual problems given above, letting the triple (Φ, A, B) be specified implicitly.

Example 1.21 (Semidefinite programs with inequality constraints) Let $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{Z}$ be complex Euclidean spaces, let $\Phi \in \mathrm{T}(\mathcal{X}, \mathcal{Y})$ and $\Psi \in \mathrm{T}(\mathcal{X}, \mathcal{Z})$ be Hermitian-preserving maps, and let $A \in \mathrm{Herm}(\mathcal{X})$, $B \in \mathrm{Herm}(\mathcal{Y})$, and $C \in \mathrm{Herm}(\mathcal{Z})$ be Hermitian operators. Define a map

$$\Xi \in \mathrm{T}(\mathcal{X} \oplus \mathcal{Z}, \mathcal{Y} \oplus \mathcal{Z}) \tag{1.286}$$

as

$$\Xi \begin{pmatrix} X & \cdot \\ \cdot & Z \end{pmatrix} = \begin{pmatrix} \Phi(X) & 0 \\ 0 & \Psi(X) + Z \end{pmatrix} \tag{1.287}$$

for all $X \in \mathrm{L}(\mathcal{X})$ and $Z \in \mathrm{L}(\mathcal{Z})$. (Similar to the previous example, the dots in the argument to Ξ represent arbitrary elements of $\mathrm{L}(\mathcal{X}, \mathcal{Z})$ and $\mathrm{L}(\mathcal{Z}, \mathcal{X})$ upon which Ξ does not depend.) The adjoint map

$$\Xi^* \in \mathrm{T}(\mathcal{Y} \oplus \mathcal{Z}, \mathcal{X} \oplus \mathcal{Z}) \tag{1.288}$$

to Ξ is given by

$$\Xi^*\begin{pmatrix} Y & \cdot \\ \cdot & Z \end{pmatrix} = \begin{pmatrix} \Phi^*(Y) + \Psi^*(Z) & 0 \\ 0 & Z \end{pmatrix}. \tag{1.289}$$

The primal and dual problems of the semidefinite program specified by the map Ξ, together with the Hermitian operators

$$\begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix} \in \mathrm{Herm}(\mathcal{X} \oplus \mathcal{Z}) \quad \text{and} \quad \begin{pmatrix} B & 0 \\ 0 & C \end{pmatrix} \in \mathrm{Herm}(\mathcal{Y} \oplus \mathcal{Z}), \tag{1.290}$$

may be expressed in the following simplified form:

Primal problem		Dual problem	
maximize:	$\langle A, X \rangle$	minimize:	$\langle B, Y \rangle + \langle C, Z \rangle$
subject to:	$\Phi(X) = B,$	subject to:	$\Phi^*(Y) + \Psi^*(Z) \geq A,$
	$\Psi(X) \leq C,$		$Y \in \mathrm{Herm}(\mathcal{Y}),$
	$X \in \mathrm{Pos}(\mathcal{X}).$		$Z \in \mathrm{Pos}(\mathcal{Z}).$

It is sometimes convenient to consider semidefinite programming problems of this form, that include both equality and inequality constraints in the primal problem, as opposed to just equality constraints.

1.3 Suggested references

Several textbooks cover the material on linear algebra summarized in this chapter; the classic books of Halmos (1978) and Hoffman and Kunze (1971) are two examples. Readers interested in a more modern development of linear algebra for finite dimensional spaces are referred to the book of Axler (1997). The books of Horn and Johnson (1985) and Bhatia (1997) also cover much of the material on linear algebra that has been summarized in this chapter (and a great deal more, including relevant theorems to be proved in subsequent chapters of this book), with a focus on the matrix-theoretic aspects of the subject.

There are also many textbooks on mathematical analysis, including the classic texts of Rudin (1964) and Apostol (1974), as well as the books of Bartle (1966) and Halmos (1974) that focus on measure theory. The book of Rockafellar (1970) is a standard reference on convex analysis, and the two volume collection of Feller (1968, 1971) is a standard reference on probability theory. Semidefinite programming is discussed by Wolkowicz, Saigal, and Vandenberge (2000).