1

Introduction

Convex polytopes, or simply polytopes, are geometric objects in some space $\mathbb{R}^d$; in fact, they are bounded intersections of finitely many closed halfspaces in $\mathbb{R}^d$. The space $\mathbb{R}^d$ can be regarded as a linear space or an affine space, and its linear or affine subspaces can be described by linear or affine equations. We introduce the basic concepts and results from linear algebra that allow the description and analysis of these subspaces. In particular, we describe the embedding of an affine space into a larger linear space, which often results in a clearer perspective of the initial space. In Chapter 2, we will embed an affine space into a projective space and so projective spaces are introduced in preparation for that embedding. Spaces are fundamental concepts in linear algebra, as are the maps that preserve their inherent structure; Section 1.4 revisits these maps. The principle of duality will surface often in this book, starting with an overview of dual spaces in Section 1.5 and dual sets in Section 1.11.

A polytope can alternatively be described as the convex hull of a finite set of points in $\mathbb{R}^d$ and so it is a convex set. Convex sets are therefore introduced, as well as the basic theorems of Carathéodory and Radon. Section 1.7 revisits topological properties of convex sets, with an emphasis on relative notions as these are based on a more natural setting, the affine hull of the set. We then review the separation and support of convex sets by hyperplanes. A convex set is formed by fitting together other polytopes of smaller dimensions, its faces; Section 1.9 discusses them. Finally, the chapter studies convex cones and lineality spaces of convex sets in $\mathbb{R}^d$; these sets are closely connected to the structure of unbounded convex sets.

1.1 Subspaces

From the outset, we want to make clear that all the spaces considered in this book are of finite dimension and 'concrete'; their underlying fields of scalars

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

are 'concrete': they are either the set $\mathbb{Q}$ of rational numbers, the set $\mathbb{R}$ of real numbers, or in rare occasions the set $\mathbb{C}$ of complex numbers.

The set of all d-tuples $(x_1, \ldots, x_d)$ with entries in $\mathbb{R}$ defines the *d-dimensional real linear space* $\mathbb{R}^d$ when endowed with vector addition and scalar multiplication. The elements of $\mathbb{R}^d$, called *vectors*, are always represented as column vectors, written in bold as $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$. When the vector coordinates are required, we will then write $\boldsymbol{x} = (x_1, \ldots, x_d)^t$ for typographical reasons, where X^t denotes the transpose of a matrix X. The all-one vector and the all-zero vector in $\mathbb{R}^d$ are denoted by $\mathbf{1}_d$ and $\mathbf{0}_d$, respectively; when there is no opportunity for confusion we write simply $\mathbf{0}$ and $\mathbf{1}$.

A *linear subspace* L of $\mathbb{R}^d$ is a nonempty subset that contains the linear combination of any two of its vectors; that is, a set of the form

$$L := \{\alpha_1 \boldsymbol{l}_1 + \alpha_2 \boldsymbol{l}_2 \mid \alpha_1, \alpha_2 \in \mathbb{R} \text{ and } \boldsymbol{l}_1, \boldsymbol{l}_2 \in L\}; \tag{1.1.1}$$

here $x := y$ and $y =: x$ define x as an object equal to y. The expression $\alpha_1 \boldsymbol{l}_1 + \alpha_2 \boldsymbol{l}_2$ is the *linear combination* of the vectors $\boldsymbol{l}_1$ and $\boldsymbol{l}_2$.

A point $\boldsymbol{x}$ in $\mathbb{R}^d$ lies in a line ℓ through distinct points $\boldsymbol{a}_1, \boldsymbol{a}_2 \in \mathbb{R}^d$ if $\boldsymbol{x}$ can be expressed as $\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2$ for some scalars $\alpha_1, \alpha_2 \in \mathbb{R}$ satisfying $\alpha_1 + \alpha_2 = 1$. The line ℓ can be defined as

$$\ell := \{\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2 \mid \alpha_1, \alpha_2 \in \mathbb{R}, \alpha_1 + \alpha_2 = 1\}.$$

In the particular case that the line ℓ passes through the origin, ℓ can be defined as $\ell := \{\alpha \boldsymbol{a} \mid \alpha \in \mathbb{R}, \boldsymbol{a} \in \ell\}$.

Example 1.1.2 (Linear subspaces) The following are examples of linear spaces.

(i) A line in $\mathbb{R}^d$ through the origin is a linear subspace of $\mathbb{R}^d$.
(ii) The set $\mathbb{R}^d$ is a linear subspace of $\mathbb{R}^d$.
(iii) The smallest linear space is $\{\mathbf{0}\}$.
(iv) The set of solutions of a system of homogeneous linear equations in d unknowns is a linear subspace of $\mathbb{R}^d$:

$$L := \left\{ \boldsymbol{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_d \end{pmatrix} \middle| \begin{array}{rcl} \alpha_{1,1}x_1 + \cdots + \alpha_{1,d}x_d & = & 0 \\ & \vdots & \\ \alpha_{n,1}x_1 + \cdots + \alpha_{n,d}x_d & = & 0 \end{array} \right\}.$$

A subset $L \subseteq \mathbb{R}^d$ is a linear subspace of $\mathbb{R}^d$ if and only if it is the solution set of a system of homogeneous linear equations in d unknowns (Problem 1.12.2).
(v) The set $\mathbb{R}^{d \times n}$ of $d \times n$ matrices with entries in $\mathbb{R}$ is a linear space with the usual matrix addition and scalar multiplication.

If we want to emphasise the affine properties of $\mathbb{R}^d$, we refer to it as the *d-dimensional real affine space* $\mathbb{A}^d$ and call its elements *points*. As with vectors, our points are always column vectors, written in bold.

A subset A of $\mathbb{R}^d$ is an *affine subspace* if it contains the line between any two of its elements:

$$A = \{\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2 \mid \alpha_1, \alpha_2 \in \mathbb{R}, \alpha_1 + \alpha_2 = 1, \text{ and } \boldsymbol{a}_1, \boldsymbol{a}_2 \in A\}. \qquad (1.1.3)$$

The expression $\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2$, where $\alpha_1, \alpha_2 \in \mathbb{R}$, $\alpha_1 + \alpha_2 = 1$, and $\boldsymbol{a}_1, \boldsymbol{a}_2 \in A$ is called the *affine combination* of the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$.

Example 1.1.4 (Affine subspaces) The following are examples of affine spaces.

(i) A line in $\mathbb{R}^d$ is an affine subspace of $\mathbb{R}^d$.
(ii) The set $\mathbb{R}^d$ is an affine subspace of $\mathbb{R}^d$.
(iii) The smallest affine space is $\varnothing$.
(iv) A set of solutions of a system of nonhomogeneous linear equations in d unknowns is an affine subspace of $\mathbb{R}^d$:

$$A := \left\{ \boldsymbol{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_d \end{pmatrix} \middle| \begin{array}{rcl} \alpha_{1,1}x_1 + \cdots + \alpha_{1,d}x_d & = & b_1 \\ & \vdots & \\ \alpha_{n,1}x_1 + \cdots + \alpha_{n,d}x_d & = & b_d \end{array} \right\}.$$

A subset $A \subseteq \mathbb{R}^d$ is an affine subspace of $\mathbb{R}^d$ if and only if it is the solution set of a system of nonhomogeneous linear equations in d unknowns (Problem 1.12.3).

Expressions (1.1.1) and (1.1.3) suggest a close relation between linear subspaces and nonempty affine subpaces. A nonempty affine subspace A is a *translate* $\boldsymbol{a}_0 + L$ of a linear subspace L by a point $\boldsymbol{a}_0 \in A$. Fixing any point $\boldsymbol{a}_0$ of A and translating any point of A by $-\boldsymbol{a}_0$ gives the linear space $L := \{\boldsymbol{l} \mid \boldsymbol{a}_0 + \boldsymbol{l} \in A\}$. In this setting, the point $\boldsymbol{a}_0$ plays the role of the origin of the subspace. We call the subspace L the *direction* of A and denote it by $\overrightarrow{A}$. Two nonempty affine subspaces are *parallel* if they have the same direction. Often we present a nonempty affine subspace A as the pair $(A, \overrightarrow{A})$ to keep track of its direction.

Theorem 1.1.5[1] *The nonempty affine spaces in $\mathbb{R}^d$ are precisely the translates of linear subspaces of $\mathbb{R}^d$. Furthermore, the direction of each nonempty affine subspace is unique.*

[1] A proof is available in Webster (1994, thm. 1.2.1).

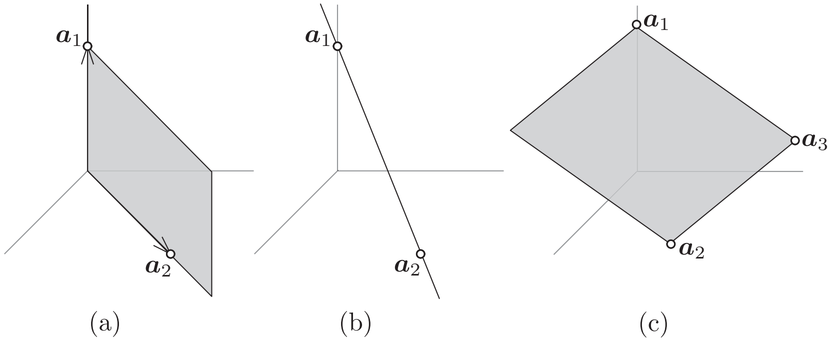

Figure 1.1.1 Affine and linear hulls in $\mathbb{R}^3$. (a) The linear hull of the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$, which is a plane in $\mathbb{R}^3$. (b) The affine hull of the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$, which is a line in $\mathbb{R}^3$. (c) The affine hull of the points $\boldsymbol{a}_1$, $\boldsymbol{a}_2$, and $\boldsymbol{a}_3$, which is a plane in $\mathbb{R}^3$.

It is obvious that any linear subspace admits an affine structure where the linear space itself acts as its direction. In view of this, it may not always be possible to make a clear distinction between points and vectors of a linear subspace.

Given a subset X of $\mathbb{R}^d$, it is often convenient to describe the smallest linear or affine subspace containing it. These new sets are the *linear hull* and *affine hull* of the set X, respectively. The linear hull is also called the *linear span* and, similarly, the affine hull is also called the *affine span*. The hulls of a set X can be described as the set of combinations of *finitely* many elements of X. Notationally,

$$\operatorname{lin} X := \{\alpha_1 \boldsymbol{a}_1 + \cdots + \alpha_n \boldsymbol{a}_n \mid \boldsymbol{a}_i \in X, \text{ and } \alpha_i \in \mathbb{R}\}.$$

$$\operatorname{aff} X := \left\{\alpha_1 \boldsymbol{a}_1 + \cdots + \alpha_n \boldsymbol{a}_n \,\middle|\, n \geqslant 1, \boldsymbol{a}_i \in X, \alpha_i \in \mathbb{R}, \text{ and } \sum_{i=1}^{n} \alpha_i = 1\right\}.$$

Note that $\operatorname{lin} \varnothing = \{\mathbf{0}\} = \operatorname{aff}(\varnothing \cup \{\mathbf{0}\})$, while $\operatorname{aff} \varnothing = \varnothing$. Figure 1.1.1 depicts examples of affine and linear hulls.

Dependence, Bases, and Dimension

Let $\boldsymbol{x}, \boldsymbol{y}$ be vectors or points of $\mathbb{R}^d$. We say that

$$(x_1, \ldots, x_d)^t = \boldsymbol{x} \leqslant \boldsymbol{y} = (y_1, \ldots, y_d)^t$$

if $x_i \leqslant y_i$ for every $1 \leqslant i \leqslant d$. Similarly, we define the relations $\boldsymbol{x} \geqslant \boldsymbol{y}$, $\boldsymbol{x} > \boldsymbol{y}$, $\boldsymbol{x} < \boldsymbol{y}$, and $\boldsymbol{x} = \boldsymbol{y}$. In contrast, the relation $\boldsymbol{x} \neq \boldsymbol{y}$ means that $x_i \neq y_i$, for some $1 \leqslant i \leqslant d$.

A set $L = \{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_n\}$ of vectors in $\mathbb{R}^d$ is said to be *linearly dependent* if and only if there exist distinct scalars $\alpha_1, \ldots, \alpha_n$, not all zero, such that

$$\alpha_1 \boldsymbol{l}_1 + \cdots + \alpha_n \boldsymbol{l}_n = \mathbf{0}_d.$$

If this equation is satisfied only when $\alpha_1 = \cdots = \alpha_n = 0$, the set is said to be *linearly independent*. A set $A = \{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n\}$ of points in $\mathbb{R}^d$ is said to be *affinely dependent* if and only if there exist distinct scalars $\alpha_1, \ldots, \alpha_n$, not all zero, such that

$$\alpha_1 \boldsymbol{a}_1 + \cdots + \alpha_n \boldsymbol{a}_n = \mathbf{0}_d \text{ and } \alpha_1 + \cdots + \alpha_n = 0.$$

If these equations are simultaneously satisfied only when $\alpha_1 = \cdots = \alpha_n = 0$, the set is said to be *affinely independent*.

A *linear basis* of a linear subspace L is a linearly independent set whose linear span is L, while an *affine basis* of an affine subspace A is an affinely independent set whose affine span is A. If an affine space A is written as $\boldsymbol{a}_0 + \overrightarrow{A}$, with $\boldsymbol{a}_0 \in A$, and the set $\{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_n\}$ is a linear basis of $\overrightarrow{A}$, then the set $\{\boldsymbol{a}_0, \boldsymbol{a}_0 + \boldsymbol{l}_1, \ldots, \boldsymbol{a}_0 + \boldsymbol{l}_n\}$ is an affine basis of A. Each point of a subspace can be expressed uniquely as an affine or linear combination of the elements of a basis of the subspace.

The *dimension* of a linear space L, denoted $\dim L$, is defined as the number of elements of any of its bases, while the *dimension* of a nonempty affine subspace is defined as the dimension of its direction. This definition ensures that when an affine space is a linear subspace, its affine dimension coincides with its linear dimension. The dimension of the empty affine subspace is -1.

The *dot product* $\cdot$ of two vectors $\boldsymbol{x} = (x_1, \ldots, x_d)^t$ and $\boldsymbol{y} = (y_1, \ldots, y_d)^t$ is defined as

$$\boldsymbol{x} \cdot \boldsymbol{y} := x_1 y_1 + \cdots + x_d y_d.$$

If the linear space $\mathbb{R}^d$ is equipped with a dot product, then it becomes a *Euclidean space*, the *d-dimensional Euclidean space*.

Example 1.1.6 recalls the definitions of five important subspaces.

Example 1.1.6 (Five important linear subspaces) Let M be an $n \times d$ matrix. The following subspaces can be defined from M.

(i) The (right) *nullspace* null M of M is the set of solutions of the homogeneous system $M\boldsymbol{x} = \mathbf{0}$. Notationally,

$$\operatorname{null} M := \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, M\boldsymbol{x} = \mathbf{0} \right\}. \tag{1.1.6.1}$$

(ii) The *left nullspace* of M is the nullspace of the transpose of M.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

(iii) The *column space* $\operatorname{col} M$ of M is the subspace of $\mathbb{R}^n$ spanned by the columns of M, seen as vectors in $\mathbb{R}^n$. Notationally,

$$\operatorname{col} M := \{\boldsymbol{b} \in \mathbb{R}^n \,|\, \text{The system } M\boldsymbol{x} = \boldsymbol{b} \text{ is solvable}\}. \tag{1.1.6.2}$$

(iv) The *row space* $\operatorname{row} M$ of M is the subspace of $\mathbb{R}^d$ spanned by the rows of M, seen as vectors in $\mathbb{R}^d$. Notationally,

$$\operatorname{row} M := \left\{\boldsymbol{b} \in \mathbb{R}^d \,\middle|\, \text{The system } M^t\boldsymbol{x} = \boldsymbol{b} \text{ is solvable}\right\}. \tag{1.1.6.3}$$

The dimension of $\operatorname{row} M$ equals the dimension of $\operatorname{col} M$, and either is called the *rank* of M.

(v) We say that two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ are *orthogonal* if $\boldsymbol{x} \cdot \boldsymbol{y} = 0$. Define the *orthogonal complement* of any linear subspace L as the set of vectors orthogonal to every vector in L. Notationally,

$$L^{\perp} := \left\{\boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \boldsymbol{x} \cdot \boldsymbol{l} = 0 \text{ for every } \boldsymbol{l} \in L\right\}. \tag{1.1.6.4}$$

The nullspace of M coincides with the orthogonal complement of its row space, whereas the left nullspace of M coincides with the orthogonal complement of its column space. These five subspaces are closely related, as Problems 1.12.4 and 1.12.5 attest.

Two linear subspaces L and L' of $\mathbb{R}^d$ are *orthogonal* if every vector of L is orthogonal to every vector of L'. A linear space and its orthogonal complement are clearly orthogonal. Two affine spaces are *orthogonal* if their corresponding directions are orthogonal.

We revisit the link between linear and affine spaces, and solutions to linear equations.

Example 1.1.7 (Subspaces and linear equations) In this example, given a linear subspace L and an affine subspace A of $\mathbb{R}^d$, we describe them as solutions of linear equations.

(i) (Linear case) Let $\{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_n\}$ be a set of vectors in $\mathbb{R}^d$ that form a basis of $L^{\perp}$. Letting M be the $n \times d$ matrix with the vectors $\boldsymbol{l}_1, \ldots, \boldsymbol{l}_n$ as rows, we find that L is the set of solutions of the homogeneous system $M\boldsymbol{x} = \boldsymbol{0}$. Notationally,

$$M := \begin{pmatrix} \boldsymbol{l}_1^t \\ \vdots \\ \boldsymbol{l}_n^t \end{pmatrix} \text{ and } L := \left\{\boldsymbol{x} \in \mathbb{R}^d \,\middle|\, M\boldsymbol{x} = \boldsymbol{0}\right\}.$$

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

(ii) (Affine case) Pick any point $\boldsymbol{a}_0 \in A$ and write $A = \boldsymbol{a}_0 + \overrightarrow{A}$ as the translation of its direction by $\boldsymbol{a}_0$. Represent $\overrightarrow{A}$ as the set of solutions of a homogeneous system $M\boldsymbol{x} = \boldsymbol{0}$ for some $n \times d$ matrix M, with $n \leqslant d$. Then

$$A = \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, M\boldsymbol{x} = \boldsymbol{b} \text{ with } \boldsymbol{b} = M\boldsymbol{a}_0 \right\}.$$

Examples 1.1.4, 1.1.6, and 1.1.7 contain all the ingredients to derive expressions for the dimensions of affine and linear spaces when given as solutions of systems of linear equations.

Proposition 1.1.8 *Let L and A be a linear subspace and an affine subspace of $\mathbb{R}^d$, respectively, given as solutions of systems of linear equations*

$$L = \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, M\boldsymbol{x} = \boldsymbol{0} \right\} \text{ and } A = \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, M\boldsymbol{x} = \boldsymbol{b} \right\},$$

for some $M \in \mathbb{R}^{n \times d}$ and $\boldsymbol{b} \in \mathbb{R}^n$. Then L is the direction of A and $\dim L = \dim A = d - \operatorname{rank} M$.

Interval Notation

The interval

$$[x, y] = \{z \in \mathbb{R} \mid x \leqslant z \leqslant y\}$$

is the set of real numbers between x and y. Similarly, we define the open interval (x, y) and the halfopen intervals $(x, y]$ and $[x, y)$ of real numbers from x to y. We will often consider intervals on the set $\mathbb{Z}$ of integers. The interval

$$[x...y] = \{z \in \mathbb{Z} \mid x \leqslant z \leqslant y\}$$

is the set of integers between x and y. Similarly, we define the open interval $(x...y)$ and the halfopen intervals $(x...y]$ and $[x \ldots y)$ of integers from x to y.

Sets in General Position

Another demand that will be imposed regularly on a subset of $\mathbb{R}^d$ is that of being in 'general position'. The precise meaning will depend on the nature of the underlying space, but the guiding principle is that the number of elements from the set in any hyperplane does not exceed the number of elements in any basis of the hyperplane. We say that a set of at least $d + 1$ points in $\mathbb{R}^d$ is in *general position* if and only if no $d + 1$ points of the set lie in an (affine) hyperplane; in other words, if and only if every subset of at most $d+1$ points is affinely independent. Similarly, we say that a set of at least d vectors in $\mathbb{R}^d$ is in *general position* if and only if no d vectors of the set lie in a linear hyperplane: if and only if every subset of at most d vectors is linearly independent.

An affine subspace in $\mathbb{R}^d$ of dimension $d-1$ is an *affine hyperplane* or just a *hyperplane*; this is geometrically defined as

$$H_d(\boldsymbol{a},\alpha) := \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \boldsymbol{a}\cdot\boldsymbol{x} = \alpha,\ \alpha \in \mathbb{R},\ \boldsymbol{a} \in \mathbb{R}^d \right\},$$

where $\boldsymbol{a}$ is a nonzero vector called the *normal* of the hyperplane. If $\alpha = 0$, the hyperplane is a linear subspace of dimension $d-1$ that we call a *linear hyperplane*.

1.2 Embedding Affine Spaces into Linear Spaces

In Section 1.1 we defined a nonempty r-dimensional affine subspace A of $\mathbb{A}^d$ as a set of the form $\boldsymbol{a}_0 + \overrightarrow{A}$ for some arbitrary element of $\boldsymbol{a}_0 \in A$ and a unique r-dimensional linear subspace $\overrightarrow{A}$. Here, we present A in another way, as a subspace of $\mathbb{R}^{d+1}$.

Every nonlinear hyperplane of $\mathbb{R}^{d+1}$ can be thought of as $\mathbb{A}^d$. Consider the concrete nonlinear hyperplane

$$H_A := \left\{ (x_1,\dots,x_{d+1})^t \in \mathbb{R}^{d+1} \,\middle|\, x_{d+1} = 1 \right\}$$

and an arbitrary point $\boldsymbol{a}_0 = (x_1,\dots,x_d,1)^t$ in H_A. The selection of $\boldsymbol{a}_0$ identifies the hyperplane H_A with its direction

$$\overrightarrow{H}_A := \left\{ (x_1,\dots,x_{d+1})^t \in \mathbb{R}^{d+1} \,\middle|\, x_{d+1} = 0 \right\}.$$

To every point $\boldsymbol{a} \in H_A$ there corresponds the vector $\boldsymbol{a}_0 + \boldsymbol{l}$ in $\mathbb{R}^{d+1}$ for some vector $\boldsymbol{l} \in \overrightarrow{H}_A$; that is, for an arbitrary $\boldsymbol{a}_0 \in H_A$ there is a bijection σ from H_A to the set of vectors in $\mathbb{R}^{d+1}$ with $x_{d+1} = 1$, which is defined as

$$\sigma(\boldsymbol{a}) = \boldsymbol{a}_0 + \boldsymbol{l}.$$

This embedding of $\mathbb{A}^d$ and its direction into $\mathbb{R}^{d+1}$ is called a *homogenisation* of $\mathbb{A}^d$. We may also say that $\mathbb{R}^{d+1}$ is a homogenisation of $\mathbb{A}^d$. See Fig. 1.2.1.

An *isomorphism* between linear spaces is a bijection that preserves vector addition and scalar multiplication; isomorphic spaces can be regarded as the *same* space, with the difference residing only in the nature or labelling of the vectors. Every d-dimensional linear or affine space over $\mathbb{R}$ is isomorphic to $\mathbb{R}^d$.

Starting from an affine subspace A and its direction $\overrightarrow{A}$, we can construct a linear subspace $\widehat{A}$ that contains both A and $\overrightarrow{A}$ and is unique up to isomorphism. As a consequence, we talk about the homogenisation $\widehat{A}$ of A. For one construction of the linear subspace $\widehat{A}$, check, for instance, Gallier (2011, ch. 4).

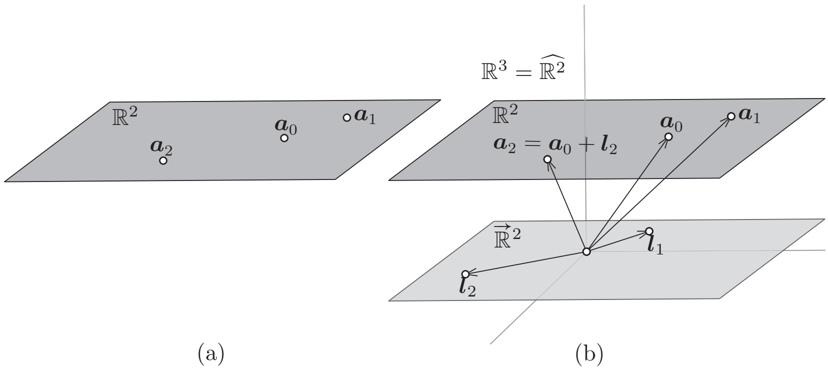

Figure 1.2.1 Homogenisation of the affine space $\mathbb{R}^2$. (a) The affine space $\mathbb{R}^2$. (b) The affine space $\mathbb{R}^2$ and its direction $\vec{\mathbb{R}}^2$ as subspaces of $\mathbb{R}^3$.

In the homogenisation of $\mathbb{A}^d$, a k-dimensional affine subspace A of $\mathbb{A}^d$ corresponds to a $(k+1)$-dimensional linear subspace $\widehat{A}$ of $\mathbb{R}^{d+1}$ spanned by an arbitrary point $\boldsymbol{a}_0$ in A and k linearly independent vectors $\boldsymbol{l}_1, \ldots, \boldsymbol{l}_k$ spanning the direction $\vec{A}$ of A in $\vec{H}_A$; that is, $A = H_A \cap \widehat{A}$ and $\vec{A} = \vec{H}_A \cap \widehat{A}$. Equivalently, we can see that $\widehat{A}$ is spanned by the $k+1$ linearly independent vectors $\boldsymbol{a}_0, \boldsymbol{a}_0 + \boldsymbol{l}_1, \ldots, \boldsymbol{a}_0 + \boldsymbol{l}_k$ and A is spanned by the $k+1$ affinely independent points $\sigma^{-1}(\boldsymbol{a}_0), \sigma^{-1}(\boldsymbol{a}_0 + \boldsymbol{l}_1), \ldots, \sigma^{-1}(\boldsymbol{a}_0 + \boldsymbol{l}_k)$. Affine properties of H_A then reduce to linear properties of $\mathbb{R}^{d+1}$. For instance, a set $\{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{k+1}\}$ of points in H_A is an affine basis of A if and only if the corresponding set $\{\sigma(\boldsymbol{a}_1), \ldots, \sigma(\boldsymbol{a}_{k+1})\}$ of vectors in $\mathbb{R}^{d+1}$ is a linear basis of $\widehat{A}$. Example 1.2.1 uses this embedding to give a concrete test for affine independence.

Example 1.2.1 (Criterion for checking affine independence) Let a set of affine points $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n$ in $\mathbb{A}^d$ be given, with $n \leqslant d+1$. Form the n vectors

$$\begin{pmatrix} \boldsymbol{a}_1 \\ 1 \end{pmatrix}, \ldots, \begin{pmatrix} \boldsymbol{a}_n \\ 1 \end{pmatrix}$$

and test their linear independence.

Concretely, verify that the points $\boldsymbol{a}_1 = (1,0,0)^t$, $\boldsymbol{a}_2 = (2,1,0)^t$, and $\boldsymbol{a}_3 = (1,1,0)^t$ are linearly independent. Form the matrix

$$M := \begin{pmatrix} 1 & 2 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix}$$

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

and compute the values of all possible determinants of order three, if necessary.

$$\begin{vmatrix}1 & 2 & 1\\0 & 1 & 1\\0 & 0 & 0\end{vmatrix} = 0, \begin{vmatrix}1 & 2 & 1\\0 & 1 & 1\\1 & 1 & 1\end{vmatrix} = 1, \begin{vmatrix}1 & 2 & 1\\0 & 0 & 0\\1 & 1 & 1\end{vmatrix} = 0, \begin{vmatrix}0 & 1 & 1\\0 & 0 & 0\\1 & 1 & 1\end{vmatrix} = 0.$$

Since there is a nonzero determinant, we conclude that the vectors $\begin{pmatrix}\boldsymbol{a}_1\\1\end{pmatrix}, \begin{pmatrix}\boldsymbol{a}_2\\1\end{pmatrix}, \begin{pmatrix}\boldsymbol{a}_3\\1\end{pmatrix}$ are linearly independent and, thus, that the points $\boldsymbol{a}_1$, $\boldsymbol{a}_2$, and $\boldsymbol{a}_3$ are affinely independent.

1.3 Projective Spaces

Our discussion of projective spaces is mostly utilitarian and not an aim in itself; in this book the objects live in affine spaces and the proofs mostly begin and end there. Our treatment will be analytic, based on linear algebra. For a synthetic treatment consult Hodge and Pedoe (1994, ch. VI).

The definition of a *d-dimensional projective space* $\mathbb{P}(\mathbb{R}^{d+1})$ over the linear space $\mathbb{R}^{d+1}$ is simple: it is the set of lines in $\mathbb{R}^{d+1}$ that pass through the origin. Each such line is a *projective point*. Projective points are our zero-dimensional projective subspaces. We prefer the notation $\mathbb{P}^d$ to $\mathbb{P}(\mathbb{R}^{d+1})$ when there is no risk of ambiguity. See Fig. 1.3.1(a) for a depiction of $\mathbb{P}(\mathbb{R}^3)$.

Since a line through the origin has the form $\alpha\boldsymbol{x}$ for some nonzero vector $\boldsymbol{x} \in \mathbb{R}^{d+1}$ and every scalar $\alpha \in \mathbb{R}$, we can define an equivalence relation $\sim$ on the nonzero vectors in $\mathbb{R}^{d+1}$ by relating two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ if they lie in the

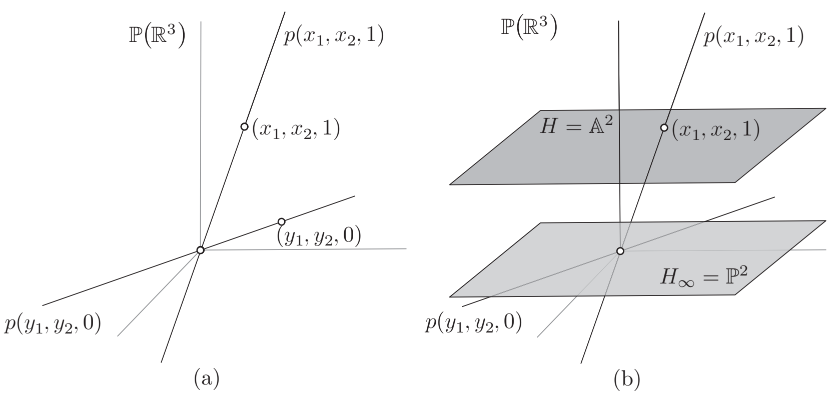

Figure 1.3.1 The projective space $\mathbb{P}(\mathbb{R}^3)$. (a) The affine space $\mathbb{R}^2$. (b) The projective space $\mathbb{P}(\mathbb{R}^3)$ as $H \cup H_\infty = \mathbb{A}^2 \cup \mathbb{P}(\mathbb{R}^2)$.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

same line through the origin; that is, if there exists a scalar α such that $\boldsymbol{x} = \alpha \boldsymbol{y}$. In this way, we get a map p that assigns to each nonzero vector $\boldsymbol{x} \in \mathbb{R}^{d+1}$ its equivalence class, or equivalently, the line $\ell_{\boldsymbol{x}} = \alpha \boldsymbol{x}$ that it spans. Notationally,

$$p(\boldsymbol{x}) = \ell_{\boldsymbol{x}}. \tag{1.3.1}$$

This equivalence relation on the vectors in $\mathbb{R}^{d+1}$ also implies that *uniqueness* in projective spaces must be understood as *uniqueness up to a scalar multiplication.*

Constructing projective subspaces from subspaces of $\mathbb{R}^{d+1}$ requires little effort. To every k-dimensional linear subspace L of $\mathbb{R}^{d+1}$ there corresponds a $(k-1)$-dimensional projective space $\mathbb{P}(L)$ defined as the set of lines through the origin that are spanned by the nonzero vectors in L. In other words,

$$p(L \backslash \{\mathbf{0}\}) = \mathbb{P}(L).$$

If $L = \{\mathbf{0}\}$, then $\mathbb{P}(L) = \varnothing$ and $\dim \mathbb{P}(L) = -1$. A one-dimensional projective subspace is a *projective line*, a two-dimensional projective subspace is a *projective plane*, and a $(d-1)$-dimensional projective subspace is a *projective hyperplane*.

An initial advantage of having the underlying linear space $\mathbb{R}^{d+1}$ in $\mathbb{P}(\mathbb{R}^{d+1})$ is that all its projective properties can be verified by linear properties of $\mathbb{R}^{d+1}$. We mention a couple of examples.

(P1) Every two projective lines intersect at a projective point. Every two linear planes (think of the ones defining the projective lines) intersect at a unique line through the origin.

(P2) Every two projective points determine a unique projective line. Every two lines through the origin (think of the ones defining the projective points) determine a unique linear plane.

Homogeneous Coordinates

Another advantage of the underlying linear space $\mathbb{R}^{d+1}$ in $\mathbb{P}(\mathbb{R}^{d+1})$ is the access to its linear bases. Consider the *standard basis* of $\mathbb{R}^{d+1}$, namely

$$\boldsymbol{e}_1 = (1, 0, \ldots, 0)^t, \ldots, \boldsymbol{e}_{d+1} = (0, \ldots, 0, 1)^t.$$

Via the standard basis, we have a bijection that maps the projective point $p(\boldsymbol{x})$ onto the set of coordinates of nonzero vectors of the form $\alpha \boldsymbol{x}$ where α is a scalar. This defines an equivalence class on all the nonzero vectors in $\mathbb{R}^{d+1}$. The equivalence class containing the coordinates $(\alpha_1, \ldots, \alpha_{d+1})$ of the vector $\boldsymbol{x}$ is denoted by $(\alpha_1 : \cdots : \alpha_{d+1})$ and defines the *homogeneous coordinates* of $p(\boldsymbol{x})$. That is,

$$(\beta_1, \ldots, \beta_{d+1}) \in (\alpha_1 : \cdots : \alpha_{d+1})$$

if and only if

$$(\alpha_1, \ldots, \alpha_{d+1})^t = \lambda(\beta_1, \ldots, \beta_{d+1})^t, \text{ for some nonzero } \lambda \in \mathbb{R}.$$

1.4 Maps

Now we revisit the maps that preserve the inherent character of affine, linear, or projective spaces.

A *linear map* is a function φ between two linear subspaces X and Y of $\mathbb{R}^d$ that satisfies

$$\varphi(\alpha_1 \boldsymbol{l}_1 + \alpha_2 \boldsymbol{l}_2) = \alpha_1 \varphi(\boldsymbol{l}_1) + \alpha_2 \varphi(\boldsymbol{l}_2),$$

for all $\alpha_1, \alpha_2 \in \mathbb{R}$ and $\boldsymbol{l}_1, \boldsymbol{l}_2 \in X$. Similarly, an *affine map* is a function ϱ between two affine spaces A and B of $\mathbb{R}^d$ that preserves affine combinations. In other words, it is a function $\varrho\colon A \to B$ that satisfies

$$\varrho(\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2) = \alpha_1 \varrho(\boldsymbol{a}_1) + \alpha_2 \varrho(\boldsymbol{a}_2),$$

for all $\alpha_1, \alpha_2 \in \mathbb{R}$ with $\alpha_1 + \alpha_2 = 1$ and for all $\boldsymbol{a}_1, \boldsymbol{a}_2 \in A$. As a consequence, every linear map is an affine map.

From the definitions it follows that an affine map ϱ is linear if and only if $\varrho(\boldsymbol{0}) = \boldsymbol{0}$. It also follows that linear maps send linear subspaces into linear subspaces and that affine maps do the same for affine subspaces. In the same vein, injective linear maps send linear subspaces of dimension r into linear subspaces of dimension r, and so do injective affine maps with affine subspaces. See Problem 1.12.6.

Bijective maps between two spaces that respect the structure of the spaces are called *isomorphisms*. Maps from a space to itself are called *transformations*, and bijective transformations are called *automorphisms*.

Projections

Orthogonal projections surface with some regularity in this book; we discuss them henceforth. Let L be a linear subspace of $\mathbb{R}^d$. Then every vector of $\mathbb{R}^d$ can be written uniquely as the sum of a vector in L and a vector in $L^\perp$ (Problem 1.12.4). Notationally, $\mathbb{R}^d = L + L^\perp$. The *orthogonal projection* π_L of any vector $\boldsymbol{x}$ of $\mathbb{R}^d$ onto L is the unique vector $\boldsymbol{l} \in L$ with the property that $\boldsymbol{x} - \boldsymbol{l} \in L^\perp$. See Fig. 1.4.1(a).

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

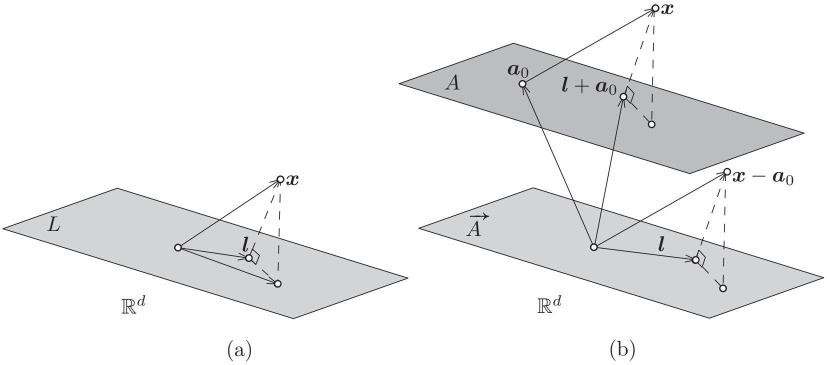

Figure 1.4.1 Orthogonal projections. (a) An orthogonal projection of a vector $\boldsymbol{x}$ onto a linear subspace L of $\mathbb{R}^d$. (b) An orthogonal projection of a vector $\boldsymbol{x}$ onto an affine subspace A of $\mathbb{R}^d$.

The projection $\pi_L(\boldsymbol{x})$ is the point in L closest to x. We provide three expressions for π_L, depending on the presentation of L.

Suppose that $\{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_r\}$ is a basis of L and write $M := (\boldsymbol{l}_1 \cdots \boldsymbol{l}_r)$. Then M is a $d \times r$ matrix whose columns are these basis vectors. The projection $\pi_L(\boldsymbol{x})$ is a point in L and so it can be written as $M\boldsymbol{y}$ for some vector $\boldsymbol{y} \in \mathbb{R}^r$. The vector $\boldsymbol{x} - \pi_L(\boldsymbol{x})$ is in $L^\perp$, which yields that $\boldsymbol{x} - \pi_L(\boldsymbol{x}) \in (\text{col}\, M)^\perp$ or, equivalently, that $\boldsymbol{x} - \pi_L(\boldsymbol{x}) \in \text{null}\, M^t$, as $(\text{col}\, M)^\perp = \text{null}\, M^t$ (Problem 1.12.5). Putting these elements together we get that

$$M^t(\boldsymbol{x} - \pi_L(\boldsymbol{x})) = \boldsymbol{0}$$
$$M^t(\boldsymbol{x} - M\boldsymbol{y}) = \boldsymbol{0}.$$

In addition, it can be shown that $M^t M$ is *nonsingular* – it has an inverse; see Problem 1.12.9. Hence, an expression for π_L is given by Equation (1.4.1) (Problem 1.12.9).

$$\pi_L(\boldsymbol{x}) = \left(M\left(M^t M\right)^{-1} M^t\right)\boldsymbol{x}. \tag{1.4.1}$$

We provide another expression for π_L, now in terms of orthogonal vectors. A set $X := \{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_r\}$ of vectors of $\mathbb{R}^d$ is *orthogonal* if $\boldsymbol{l}_i \cdot \boldsymbol{l}_j = 0$ whenever $i \neq j$. It follows that X is linearly independent. The set X is an *orthogonal basis* of L if it is both a linear basis of L and an orthogonal set. It is an *orthonormal basis* if it is an orthogonal basis consisting of unit vectors. A *unit vector* is a vector with norm one. Suppose $\{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_r\}$ is an orthogonal basis of L. Then

$$\pi_L(\boldsymbol{x}) = \sum_{i=1}^{r} \frac{\boldsymbol{x} \cdot \boldsymbol{l}_i}{\|\boldsymbol{l}_i\|^2} \boldsymbol{l}_i. \tag{1.4.2}$$

Define the *norm* $\|\cdot\|$ of a vector $\boldsymbol{x}$ in $\mathbb{R}^d$ as $\sqrt{\boldsymbol{x} \cdot \boldsymbol{x}}$.

Every linear basis $\{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_r\}$ of a linear subspace L can be transformed into an orthogonal basis or an orthonormal basis of L if we wish. To obtain an orthogonal basis $\{\boldsymbol{m}_1, \ldots, \boldsymbol{m}_r\}$, let

$$\begin{aligned} \boldsymbol{m}_1 &:= \boldsymbol{l}_1, \\ \boldsymbol{m}_i &:= \boldsymbol{l}_i - \frac{\boldsymbol{l}_i \cdot \boldsymbol{m}_1}{\|\boldsymbol{m}_1\|^2}\boldsymbol{m}_1 - \cdots - \frac{\boldsymbol{l}_i \cdot \boldsymbol{m}_{i-1}}{\|\boldsymbol{m}_{i-1}\|^2}\boldsymbol{m}_{i-1}, \text{ for each } i \in [2 \ldots r]. \end{aligned}$$

This is the *Gram–Schmidt orthogonalisation process*. If we are after an orthonormal basis, we divide each vector in $\{\boldsymbol{m}_1, \ldots, \boldsymbol{m}_r\}$ by its norm. Thus the set

$$\left\{ \frac{\boldsymbol{m}_1}{\|\boldsymbol{m}_1\|}, \ldots, \frac{\boldsymbol{m}_r}{\|\boldsymbol{m}_r\|} \right\}$$

is an orthonormal basis of L.

The final expression for π_L is obtained when a linear subspace L is given as a set of solutions of a homogeneous system:

$$L := \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, N\boldsymbol{x} = \boldsymbol{0} \right\}.$$

In this case, we get that

$$\pi_L(\boldsymbol{x}) = \boldsymbol{x} - \left(N^t \left(N N^t \right)^{-1} \right) (N\boldsymbol{x}). \tag{1.4.3}$$

In particular, if L is a linear hyperplane defined as $\left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \boldsymbol{a} \cdot \boldsymbol{x} = 0 \right\}$ then

$$\pi_L(\boldsymbol{x}) = \boldsymbol{x} - \frac{\boldsymbol{a} \cdot \boldsymbol{x}}{\|\boldsymbol{a}\|^2}\boldsymbol{a}. \tag{1.4.4}$$

Once we have defined projections onto linear spaces it is not difficult to extend the process to affine spaces via their directions. Let $(A, \overrightarrow{A})$ be an affine space and let $\boldsymbol{a}_0$ be a fixed point of A. Then the orthogonal projection π_A of any vector $\boldsymbol{x} \in \mathbb{R}^d$ onto A is defined as

$$\pi_A(\boldsymbol{x}) := \boldsymbol{a}_0 + \pi_{\overrightarrow{A}}(\boldsymbol{x} - \boldsymbol{a}_0). \tag{1.4.5}$$

See Fig. 1.4.1(b).

In analogy to the linear case, we obtain expressions for π_A depending on the presentation of A. If an affine subspace A is given as the set spanned by an affine basis $\{\boldsymbol{a}_0, \boldsymbol{a}_1, \ldots, \boldsymbol{a}_r\}$ then its direction is spanned by the linear basis

$\{\boldsymbol{a}_1 - \boldsymbol{a}_0, \ldots, \boldsymbol{a}_r - \boldsymbol{a}_0\}$. Writing $M := (\boldsymbol{a}_1 - \boldsymbol{a}_0 \cdots \boldsymbol{a}_r - \boldsymbol{a}_0)$, an application of (1.4.1) gives that

$$\pi_{\overrightarrow{A}}(\boldsymbol{x}) = \left(M \left(M^t M\right)^{-1} M^t\right) \boldsymbol{x},$$

wherefrom it follows that

$$\pi_A(\boldsymbol{x}) = \boldsymbol{a}_0 + \left(M \left(M^t M\right)^{-1} M^t\right) (\boldsymbol{x} - \boldsymbol{a}_0). \tag{1.4.6}$$

If, instead, an affine subspace A is given as a set of solutions of a nonhomogeneous system

$$A := \left\{\boldsymbol{x} \in \mathbb{R}^d \,\middle|\, N\boldsymbol{x} = N\boldsymbol{a}_0,\, \boldsymbol{a}_0 \in A\right\},$$

then

$$\pi_A(\boldsymbol{x}) = \boldsymbol{x} - \left(N^t \left(N N^t\right)^{-1}\right) (N\boldsymbol{x} - N\boldsymbol{a}_0). \tag{1.4.7}$$

In particular, if $A = \{\boldsymbol{x} \in \mathbb{R}^d \mid \boldsymbol{a} \cdot \boldsymbol{x} = \boldsymbol{a} \cdot \boldsymbol{a}_0,\, \boldsymbol{a}_0 \in A\}$ then

$$\pi_A(\boldsymbol{x}) = \boldsymbol{x} - \frac{\boldsymbol{a} \cdot \boldsymbol{x} - \boldsymbol{a} \cdot \boldsymbol{a}_0}{\|\boldsymbol{a}\|^2} \boldsymbol{a}. \tag{1.4.8}$$

Projective Maps

We define maps between projective spaces through linear maps between the underlying linear spaces. Recall that the map p assigns to each nonzero vector $\boldsymbol{x}$ in $\mathbb{R}^{d+1}$ the line $\alpha\boldsymbol{x}$ (see (1.3.1)). Let φ be a linear map between linear spaces X and Y. Since $\varphi(\alpha\boldsymbol{x}) = \alpha\varphi(\boldsymbol{x})$ for every nonzero vector $\boldsymbol{x}$ and scalar α, the map φ assigns lines through the origin to lines through the origin, provided $\boldsymbol{x} \notin \ker\varphi$. Here, $\ker\varphi$ denotes the *kernel* of φ, the subspace of X consisting of the vectors $\boldsymbol{x}$ for which $\varphi(\boldsymbol{x}) = \boldsymbol{0}$. It then follows that if φ is injective, namely $\ker\varphi = \{\boldsymbol{0}\}$, then it defines a projective map $p(\varphi)\colon \mathbb{P}(X) \to \mathbb{P}(Y)$ given by

$$p(\boldsymbol{x}) \mapsto p(\varphi(\boldsymbol{x})).$$

If φ is not injective, it defines a projective map from $\mathbb{P}(X)\backslash\mathbb{P}(\ker\varphi)$ to $\mathbb{P}(Y)$. Hence we hereafter restrict ourselves to injective linear maps. If a linear φ induces the map $p(\varphi)$, so does the map $\alpha\varphi$ for any nonzero α, and thus the linear map φ is determined up to multiplication by a nonzero scalar.

If there is a bijective projective map, a *projective isomorphism*, between two projective spaces, we say the spaces are (projectively) *isomorphic*. As before, projective maps from a space to itself are called *projective transformations* and bijective transformations are called *projective automorphisms*.

We next define a projective basis. We feel that a *projective basis* in $\mathbb{P}(\mathbb{R}^{d+1})$ is best defined as a set of points $\{p(\boldsymbol{x}_1), \ldots, p(\boldsymbol{x}_n)\}$ for which, given any

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

projective basis $\{p(\boldsymbol{y}_1), \ldots, p(\boldsymbol{y}_n)\}$ of $\mathbb{P}(\mathbb{R}^{d+1})$, there is a unique projective automorphism that takes $p(\boldsymbol{x}_i)$ to $p(\boldsymbol{y}_i)$ for all $i \in [1 \ldots n]$. This follows the general scheme that a *basis* in a linear or affine space Z ought to be a set B_1 of elements for which, given any basis B_2 in Z, there is a unique automorphism mapping B_1 onto B_2. This definition, while unusual, could have been equally used for linear and affine bases.

We say that a set of projective points $p(\boldsymbol{x}_1),\ldots,p(\boldsymbol{x}_n)$ is *projectively dependent* in $\mathbb{P}^d$ if the vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ are linearly dependent in $\mathbb{R}^{d+1}$; otherwise we say that the projective points are *projectively independent* in $\mathbb{P}^d$. In the same vein, we say that a set of projective points $p(\boldsymbol{x}_1), \ldots, p(\boldsymbol{x}_n)$ is in *general position* in $\mathbb{P}^d$ if the vectors $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ are in general position in $\mathbb{R}^{d+1}$; that is, if every subset of at most $d+1$ vectors is linearly independent.

Given the close connection between the notions of linear independence in $\mathbb{R}^{d+1}$ and projective independence in $\mathbb{P}(\mathbb{R}^{d+1})$, we may be tempted to say that a set $p(X) := \{p(\boldsymbol{x}_1), \ldots, p(\boldsymbol{x}_n)\}$ of points in $\mathbb{P}(\mathbb{R}^{d+1})$ is a basis if the set $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ is a basis in $\mathbb{R}^{d+1}$. But that would not satisfy our running definition, since no $d+1$ projective points determine a basis of $\mathbb{R}^{d+1}$, not even up to scalar multiplication; for this uniqueness, we need $d+2$ projective points.

Proposition 1.4.9[2] *Let* $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{d+1})$ *and* $(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{d+1})$ *be bases of* $\mathbb{R}^{d+1}$ *such that* $p(\boldsymbol{x}_i) = p(\boldsymbol{y}_i)$ *for* $i \in [1 \ldots d+1]$ *and* $p(\boldsymbol{x}_1 + \cdots + \boldsymbol{x}_{d+1}) = p(\boldsymbol{y}_1 + \cdots + \boldsymbol{y}_{d+1})$. *Then there is a nonzero scalar* α *such that* $\boldsymbol{x}_i = \alpha \boldsymbol{y}_i$, *for each* $i \in [1 \ldots d+1]$.

In view of Proposition 1.4.9, we say that a *projective basis* in $\mathbb{P}(\mathbb{R}^{d+1})$ is any set $\{p(\boldsymbol{x}_1), \ldots, p(\boldsymbol{x}_{d+1}), p(\boldsymbol{x}_{d+2})\}$ of $d+2$ projective points in general position. This definition is compatible with our initial definition as the following theorem attests.

Theorem 1.4.10 (Fundamental theorem of projective maps)[3] *Let* $\mathbb{P}(\mathbb{R}^r)$ *and* $\mathbb{P}(\mathbb{R}^s)$ *be projective spaces with corresponding projective bases* $(\boldsymbol{p}_1, \ldots, \boldsymbol{p}_{r+1})$ *and* $(\boldsymbol{q}_1, \ldots, \boldsymbol{q}_{s+1})$. *Then there exists a unique projective map that sends* $\boldsymbol{p}_i$ *to* $\boldsymbol{q}_i$, *for each* $i \in [1 \ldots r+1]$. *In the case* $r = s$, *the map is an isomorphism.*

Given an ordered linear basis $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{d+1})$ of $\mathbb{R}^{d+1}$, it is customary to take $\boldsymbol{x}_{d+2} = \boldsymbol{x}_1 + \cdots + \boldsymbol{x}_{d+1}$ to form the corresponding projective basis

[2] A proof is available in Berger (2009, lem. 4.4.2).
[3] A proof is available in Berger (2009, lem. 4.5.10).

$\{p(\boldsymbol{x}_1), \ldots, p(\boldsymbol{x}_{d+1}), p(\boldsymbol{x}_{d+2})\}$ of $\mathbb{P}(\mathbb{R}^{d+1})$. The projective basis arising from the standard basis is called the *projective standard basis*:

$$(1 : 0 : \cdots : 0), \ldots, (0 : 0 : \cdots : 1), (1 : 1 : \cdots : 1).$$

1.5 Dual Spaces

This section offers the first instance of the powerful principle of duality in this book; throughout the book, we will come in contact with many other incarnations of this principle. *Duality* roughly involves a pair of objects X and X^*, an involution $X \mapsto X^*$, and a correspondence, often order-reversing, between subsets or properties of X and X^* that translates results on X into results on X^*.

The set of all linear maps between two linear spaces is itself a linear space, where vector addition is defined as

$$(\varphi_1 + \varphi_2)(\boldsymbol{l}) = \varphi_1(\boldsymbol{l}) + \varphi_2(\boldsymbol{l})$$

and scalar multiplication as

$$(\alpha\varphi)(\boldsymbol{l}) = \alpha\varphi(\boldsymbol{l})$$

(Problem 1.12.6). In particular, the set of linear maps from $\mathbb{R}^d$ to $\mathbb{R}$ forms a linear space, called the *dual space* of $\mathbb{R}^d$ and denoted by $(\mathbb{R}^d)^*$.

Elements of a dual space are called *linear functionals*. Common examples of linear functionals include the *zero functional* on $\mathbb{R}^d$, which assigns zero to every vector in $\mathbb{R}^d$, and maps that define affine hyperplanes in $\mathbb{R}^d$, namely $\boldsymbol{x} \mapsto \boldsymbol{a} \cdot \boldsymbol{x}$ for every $\boldsymbol{x} \in \mathbb{R}^d$ and some $\boldsymbol{a} \in \mathbb{R}^d$.

The dual space of $\mathbb{R}^d$ is a linear space of dimension d. If $B = (\boldsymbol{l}_1, \ldots, \boldsymbol{l}_d)$ is an ordered basis of $\mathbb{R}^d$, there exists a uniquely determined basis $\varphi_1, \ldots, \varphi_d$ of $(\mathbb{R}^d)^*$ that is called the *dual basis* of B and satisfies

$$\varphi_i(\boldsymbol{l}_j) = \begin{cases} 1, & \text{if } i = j; \\ 0, & \text{otherwise.} \end{cases}$$

The dual basis of the standard basis of $\mathbb{R}^d$ can then be defined as the basis $\varphi_1, \ldots, \varphi_d$ given by $\varphi_i(x_1, \ldots, x_d) = x_i$.

There is an isomorphism between the space and its dual space. To each linear functional φ of $\mathbb{R}^d$ there corresponds a unique vector $\boldsymbol{a}$ of $\mathbb{R}^d$ such that $\varphi(\boldsymbol{x}) = \boldsymbol{a} \cdot \boldsymbol{x}$ for every $\boldsymbol{x} \in \mathbb{R}^d$. If φ is the zero functional then let $\boldsymbol{a} = \boldsymbol{0}$; otherwise, let $\boldsymbol{a} = \varphi(\boldsymbol{x}_0)\boldsymbol{x}_0$ for any nonzero unit vector $\boldsymbol{x}_0$ in $(\ker\varphi)^\perp$, which is not empty.

How does the dual space $(\mathbb{R}^d)^{**}$ of $(\mathbb{R}^d)^*$, the *double dual space* of $\mathbb{R}^d$ look? Is it what you would expect? An element of $(\mathbb{R}^d)^{**}$ is a linear functional

ϱ that sends a linear functional $\varphi\colon \mathbb{R}^d \to \mathbb{R}$ to an element in $\mathbb{R}$. There is also an isomorphism σ between $\mathbb{R}^d$ and $(\mathbb{R}^d)^{**}$. A common way to define σ is to first define a linear functional $\varrho_{\boldsymbol{x}_0}\colon (\mathbb{R}^d)^* \to \mathbb{R}$ as $\varrho_{\boldsymbol{x}_0}(\varphi) = \varphi(\boldsymbol{x}_0)$ for a fixed $\boldsymbol{x}_0 \in \mathbb{R}^d$ and every $\varphi \in (\mathbb{R}^d)^*$. Then define $\sigma(\boldsymbol{x}) = \varrho_{\boldsymbol{x}}$ for every $\boldsymbol{x} \in \mathbb{R}^d$. For finite-dimensional spaces it is customary to identify $\mathbb{R}^d$ with $(\mathbb{R}^d)^{**}$ via σ, and we do so in this book.

1.6 Convex Sets

If on the equation of a line through two points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$ of $\mathbb{R}^d$, namely $\alpha_1\boldsymbol{a}_1 + \alpha_2\boldsymbol{a}_2$ with $\alpha_1 + \alpha_2 = 1$, we impose the additional condition of $\alpha_1, \alpha_2 \geqslant 0$, then we obtain the segment $[\boldsymbol{a}_1, \boldsymbol{a}_2]$ joining the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$. Segments give rise to *convex sets*, sets that contain the segment between any of two of its points. That is, X is a *convex set* if

$$X = \{\alpha_1\boldsymbol{a}_1 + \alpha_2\boldsymbol{a}_2 |\ \alpha_1, \alpha_2 \geqslant 0, \alpha_1 + \alpha_2 = 1, \text{and } \boldsymbol{a}_1, \boldsymbol{a}_2 \in X\}.$$

The expression $\alpha_1\boldsymbol{a}_1 + \alpha_2\boldsymbol{a}_2$ where $\alpha_1, \alpha_2 \geqslant 0$, $\alpha_1 + \alpha_2 = 1$, and $\boldsymbol{a}_1, \boldsymbol{a}_2 \in X$ is the *convex combination* of the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$. The sets $\mathbb{R}^d$ and $\varnothing$ are convex sets, and so is any singleton in $\mathbb{R}^d$. Figure 1.6.1 shows examples of convex and nonconvex sets in $\mathbb{R}^3$.

In analogy to the affine and linear cases, given a set X in $\mathbb{R}^d$, the smallest convex set containing X is the *convex hull* of X and is denoted by $\operatorname{conv} X$. The convex hull of a set can be described as the set of all convex combinations of *finitely* many elements of X, namely

$$\operatorname{conv} X = \left\{ \alpha_1\boldsymbol{a}_1 + \cdots + \alpha_n\boldsymbol{a}_n \,\middle|\, \boldsymbol{a}_i \in X, \alpha_i \geqslant 0, \text{ and } \sum_{i=1}^{n} \alpha_i = 1 \right\}.$$

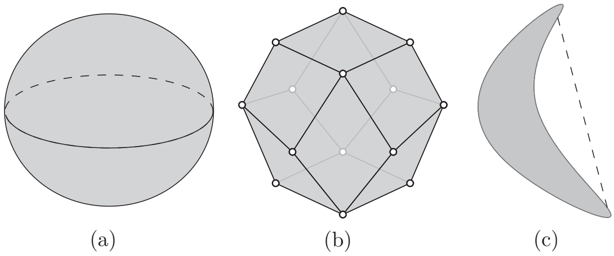

Figure 1.6.1 Convex and nonconvex sets. (a) A 3-dimensional ball, a convex set. (b) The rhombic dodecahedron, a convex 3-polytope. (c) A nonconvex set.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

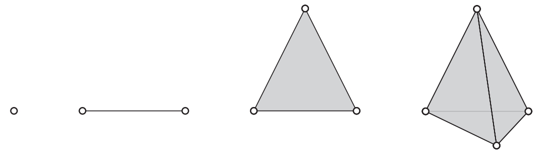

Figure 1.6.2 Simplices in $\mathbb{R}^3$.

Hyperplanes and the halfspaces in $\mathbb{R}^d$ bounded by a hyperplane are examples of convex sets. Each hyperplane $H_d(\boldsymbol{a},\alpha)$ in $\mathbb{R}^d$ with normal vector $\boldsymbol{a}$ and constant α determines or bounds four halfspaces in $\mathbb{R}^d$: the open halfspaces $H_d^+(\boldsymbol{a},\alpha)$ and $H_d^-(\boldsymbol{a},\alpha)$, and the closed halfspaces $H_d^+[\boldsymbol{a},\alpha]$ and $H_d^-[\boldsymbol{a},\alpha]$. In formulas we have that

$$H_d^+(\boldsymbol{a},\alpha) = \left\{\boldsymbol{x} \in \mathbb{R}^d : \boldsymbol{a}\cdot\boldsymbol{x} > \alpha\right\},\ H_d^-(\boldsymbol{a},\alpha) = \left\{\boldsymbol{x} \in \mathbb{R}^d : \boldsymbol{a}\cdot\boldsymbol{x} < \alpha\right\},$$
$$H_d^+[\boldsymbol{a},\alpha] = \left\{\boldsymbol{x} \in \mathbb{R}^d : \boldsymbol{a}\cdot\boldsymbol{x} \geqslant \alpha\right\},\ H_d^-[\boldsymbol{a},\alpha] = \left\{\boldsymbol{x} \in \mathbb{R}^d : \boldsymbol{a}\cdot\boldsymbol{x} \leqslant \alpha\right\}. \tag{1.6.1}$$

Define the Euclidean *distance* between any two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ in $\mathbb{R}^d$ as $\|\boldsymbol{x}-\boldsymbol{y}\|$. The distance function turns $\mathbb{R}^d$ into a metric space and enables more examples of convex sets such as the d-dimensional *open ball*, or simply the *open d-ball* $B_d(\boldsymbol{a},r)$, and the d-dimensional *closed ball*, or simply the *closed d-ball* $B_d[\boldsymbol{a},r]$, both with centre $\boldsymbol{a}$ and radius r:

$$B_d(\boldsymbol{a},r) = \left\{\boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \|\boldsymbol{x}-\boldsymbol{a}\| < r\right\},$$
$$B_d[\boldsymbol{a},r] = \left\{\boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \|\boldsymbol{x}-\boldsymbol{a}\| \leqslant r\right\}. \tag{1.6.2}$$

In all the formulas above, if there is no place for confusion we drop the subindex d. The ball $B_d[\boldsymbol{0},1]$ is referred to as the *closed unit d-ball*, while the ball $B_d(\boldsymbol{0},1)$ is referred to as the *open unit d-ball*.

Convex polytopes are other examples of convex sets. A *convex polytope* is the convex hull of a finite set of points in $\mathbb{R}^d$ (see Fig. 1.6.1(b)). In this book we speak only of convex polytopes. Hence we drop the adjective 'convex' hereafter. *Simplices* are an important class of polytopes; they are convex hulls of affinely independent points in $\mathbb{R}^d$. Figure 1.6.2 depicts all the simplices in $\mathbb{R}^3$.

It turns out that to get the convex hull of a set X, we do not require the finite convex combinations of all the points of X but rather the finite combinations of all collections of affinely independent points of X. Carathéordory's theorem elucidates these remarks.

Theorem 1.6.3 (Carathéodory, 1907)[4] *The convex hull of a set X in $\mathbb{R}^d$ is formed by* all *the convex combinations of at most $d+1$ affinely independent points of X. Furthermore, if* conv X *is not a simplex then no fixed collection of affinely independent points from X suffices to span* conv X.

Carathéordory's theorem shows that a simplex possesses a proper *convex basis*: a fixed subset of affinely independent points that uniquely generate each element of the simplex. Since the set $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r\}$ of a simplex $T := \operatorname{conv}\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r\}$ is affinely independent, every point $\boldsymbol{x}$ of T has a unique representation as an affine combination of $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r\}$ and, in particular, as a convex combination of $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_r\}$.

Remark 1.6.4 If, for a set $X \subseteq \mathbb{R}^d$, conv X is not a simplex, then the notion of 'convex basis' is not available for (at least) two reasons:

(i) no fixed, finite collection of affinely independent points from X would suffice to generate conv X (by Carathéordory's theorem), and
(ii) some points in conv X have no unique representation as the convex combination of a fixed subset of affinely independent points from X.

Let X be a d-dimensional *sphere*, or simply *d-sphere* $S_d(\boldsymbol{a}, r)$, namely a set of the form $\{\boldsymbol{x} \in \mathbb{R}^{d+1} \mid \|\boldsymbol{x} - \boldsymbol{a}\| = r\}$. The convex hull of X is the closed ball $B_d[\boldsymbol{a}, r]$. The ball $B_d[\boldsymbol{a}, r]$ exemplifies Remark 1.6.4(ii) in the case of X being an infinite set. Every point in $B_d[\boldsymbol{a}, r]$ is a convex combination of at most two points from X, every point in the open ball admits more than one representation, and no finite subset of $S_d(\boldsymbol{a}, r)$ spans $B_d[\boldsymbol{a}, r]$. The justification for (ii) in the case of X being a finite set is given by Radon's theorem.

We will be mostly dealing with the sphere $S_d(\boldsymbol{0}, 1)$, which is referred to as the *unit d-sphere* and is denoted by $\mathbb{S}^d$.

Theorem 1.6.5 (Radon, 1921)[5] *Let X be a finite set of affinely dependent points in $\mathbb{R}^d$. Then the set X can be partitioned into subsets X_+ and X_- whose convex hulls intersect. Furthermore,* conv X_+ *and* conv X_- *can be assumed to be simplices.*

A *Radon partition* $\{X_+, X_-\}$ of a set $X \subseteq \mathbb{R}^d$ is a partition of X such that $\operatorname{conv} X_+ \cap \operatorname{conv} X_- \neq \emptyset$. A *Radon point* is a point in $\operatorname{conv} X_+ \cap \operatorname{conv} X_-$, a point that admits more than one representation as a convex hull of points of X. Figure 1.6.3 shows Radon partitions of four points in $\mathbb{R}^2$.

[4] A proof is available in Webster (1994, thm. 2.2.4).
[5] A proof is available in Webster (1994, thm. 2.2.5).

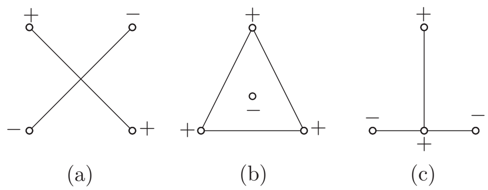

Figure 1.6.3 Radon partitions of four points in $\mathbb{R}^2$. Elements of the sets X_+ and X_- in Radon's theorem (1.6.5) are labelled with a $+$ and $-$ sign, respectively. Then $\operatorname{conv} X_+ \cap \operatorname{conv} X_- \neq \emptyset$. (a) Two line segments intersecting at a point; the convex hull of the four points is a quadrilateral. (b) A triangle and a point contained in the triangle; the convex hull of the four points is a triangle. (c) Two line segments intersecting at a point; the convex hull of the four points is a triangle.

An implication of the lack of a 'basis' for a convex set is that its *dimension* is defined as the dimension of its affine hull. It follows that $\dim \emptyset = -1$, since $\operatorname{aff} \emptyset = \emptyset$. As it should be, the dimension of a convex set is an invariant that does not dependent on the space in which the set is embedded.

Convexity is preserved by a number of operations; the final theorem of this section gathers some of these. Before stating the theorem, we require some basic definitions.

Definition 1.6.6 The *Minkowski sum* or *sum* $X + Y$ of two sets X and Y is the set

$$X + Y := \{x + y \mid x \in X, y \in Y\}.$$

The sets X and Y are the *summands* of $X + Y$.

The *scalar multiple* αX of a set X is the set

$$\alpha X := \{\alpha \boldsymbol{x} \mid \boldsymbol{x} \in X, \alpha \in \mathbb{R}\}.$$

Minkowski sums and scalar multiples of sets are illustrated in Fig. 1.6.4.

Let $\varphi\colon X \to Y$ be a function and let $B \subseteq Y$. The *preimage* $\varphi^{-1}(B)$ of B is the subset $\{x \in X \mid \varphi(x) \in B\}$ of X. The preimage of a function is well defined even if the function is not a bijection. In the case that φ is a bijection, $\varphi^{-1}(B)$ coincides with the image of B under the inverse function φ^{-1} of φ.

Theorem 1.6.7[6] *The following operations in $\mathbb{R}^d$ all return convex sets.*

(i) *The intersection of an arbitrary family of convex sets.*
(ii) *The Minkowski sum of convex sets.*

[6] A proof is available in Webster (1994, sec. 2.1).

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

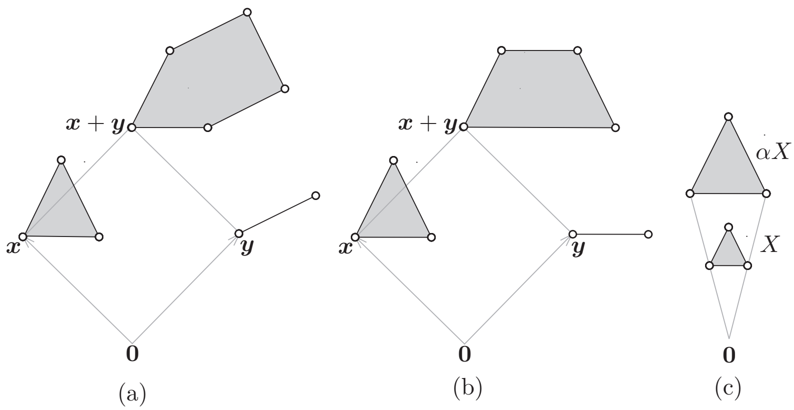

Figure 1.6.4 Minkowski sums and scalar multiple of sets. (a)–(b) Minkowski sums of a triangle and a segment. (c) Scalar multiple of a triangle.

(iii) *The scalar multiple of a convex set.*
(iv) *The image of a convex set under an affine map.*
(v) *The preimage of a convex set under an affine map.*

A function is *convexity-preserving* if it maps convex sets to convex sets. Every affine map is a convexity-preserving function (Theorem 1.6.7(iv)), but not every convexity-preserving function is an affine map (Problem 1.12.13). It is, however, true that every injective convexity-preserving function $\varphi\colon \mathbb{R}^r \to \mathbb{R}^s$ with $r \geqslant 2$ is an affine map (Meyer and Kay, 1973).

1.7 Interior, Boundary, and Closure

This section considers the topological notions of interior, boundary, and closure in the context of convex sets.

The *interior* of a set X in $\mathbb{R}^d$ is defined as the set of all points of X that are centres of some open d-dimensional ball that lies in X. Notationally,

$$\operatorname{int} X := \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \text{there exists } r > 0 \text{ such that } B_d(\boldsymbol{x}, r) \subseteq X \right\}. \qquad (1.7.1)$$

The points in the interior of X are its *interior points*.

An immediate consequence of the above definition is that a set X in $\mathbb{R}^d$ of dimension less than d has an empty interior. We benefit from studying the interior of X relative to the smallest affine space containing X; this is the *relative interior* of X. Notationally,

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

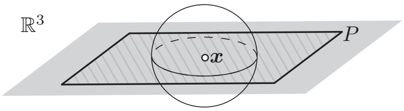

Figure 1.7.1 Difference between the relative interior and the interior of a 2-polytope, or *polygon*, P in $\mathbb{R}^3$. The affine hull of P is a plane; its relative interior is highlighted in a tiling pattern. However, the interior of P is empty: no ball in $\mathbb{R}^3$ centred at a point of P is fully contained in P.

$$\operatorname{rint} X := \left\{ \boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \text{there exists } r > 0 \text{ such that } (B_d(\boldsymbol{x}, r) \cap \operatorname{aff} X) \subseteq X \right\}. \tag{1.7.2}$$

The points in the relative interior of X are its *relative interior points*.

The relative interior of a set is a more natural concept than that of its interior and, as such, we discuss it in more detail. The geometric difference between the relative interior and the interior of a set is captured in Fig. 1.7.1.

For a set $X \subseteq \mathbb{R}^d$, it is obvious that $\operatorname{int} X \subseteq \operatorname{rint} X \subseteq X$ (Problem 1.12.14). In the case of X being a nonempty convex set, it is not straightforward but it is true that $\operatorname{int} X = \operatorname{rint} X$ if and only if $\operatorname{int} X \neq \emptyset$ (Problem 1.12.15(ii)); the necessity of the statement relies on the nontrivial assertion that $\operatorname{rint} X \neq \emptyset$.

Theorem 1.7.3[7] *If X is a nonempty convex set in $\mathbb{R}^d$, then* $\operatorname{rint} X \neq \emptyset$.

We unveil some geometric properties of the relative interior of a convex set.

Theorem 1.7.4[8] *Let X be a convex set in $\mathbb{R}^d$ and let $\boldsymbol{a} \in \mathbb{R}^d$. The following statements hold.*

(i) *If $\boldsymbol{a} \in \operatorname{rint} X$, then the halfopen segment $[\boldsymbol{a}, \boldsymbol{b})$ lies in $\operatorname{rint} X$ for every $\boldsymbol{b} \in X$.*
(ii) *If $\boldsymbol{a} \in \operatorname{rint} X$ then, for each $\boldsymbol{b} \in \operatorname{aff} X$, there exists $\alpha > 1$ such that $\alpha\boldsymbol{a} + (1-\alpha)\boldsymbol{b} \in X$.*
(iii) *If, for each $\boldsymbol{b} \in X$, there exists $\alpha > 1$ such that $\alpha\boldsymbol{a} + (1-\alpha)\boldsymbol{b} \in X$, then $\boldsymbol{a} \in \operatorname{rint} X$.*

As a corollary of Theorem 1.7.4 we get another geometric property of relative interior points.

Corollary 1.7.5 *Let X be a convex set in $\mathbb{R}^d$ and let $\boldsymbol{a} \in \operatorname{rint} X$. Then, for each $\boldsymbol{b} \in \operatorname{aff} X \setminus \{\boldsymbol{a}\}$, there is a point $\boldsymbol{c} \in X \setminus \{\boldsymbol{a}\}$ such that $\boldsymbol{a} \in (\boldsymbol{b}, \boldsymbol{c})$.*

[7] A proof is available in Webster (1994, thm. 2.3.1).
[8] A proof is available in Webster (1994, sec. 2.3).

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

Proof According to Theorem 1.7.4(ii), there exists an $\alpha > 1$ such that $\boldsymbol{c} := \alpha\boldsymbol{a} + (1-\alpha)\boldsymbol{b} \in X$. Wherefrom it follows that

$$\boldsymbol{a} = \frac{1}{\alpha}\boldsymbol{c} + \left(1 - \frac{1}{\alpha}\right)\boldsymbol{b},$$

and with this equality, the corollary. □

Perhaps the most useful geometric property of the relative interior is the following.

Theorem 1.7.6[9] *Let $X \subseteq \mathbb{R}^d$ be defined as* $\operatorname{conv}\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ *and let $\boldsymbol{a} \in X$. The point $\boldsymbol{a}$ is in* $\operatorname{rint} X$ *if and only if there exist scalars $\alpha_1, \ldots, \alpha_n > 0$ with $\sum_{i=1}^n \alpha_i = 1$ such that*

$$\boldsymbol{a} = \sum_{i=1}^n \alpha_i \boldsymbol{x}_i.$$

A *closure point* of a set $X \subseteq \mathbb{R}^d$ is a point $\boldsymbol{x}$ in which every open d-dimensional ball with centre at $\boldsymbol{x}$ meets X. The set of closure points is the *closure* of X and is denoted $\operatorname{cl} X$. Notationally,

$$\operatorname{cl} X := \left\{\boldsymbol{x} \in \mathbb{R}^d \,\middle|\, \text{for every } r > 0,\, B_d(\boldsymbol{x}, r) \cap X \neq \emptyset\right\}.$$

A sequence $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n, \ldots$ of points in $\mathbb{R}^d$, written as $(\boldsymbol{x}_n)$, is said to *converge* to a point $\boldsymbol{x}$ if $\|\boldsymbol{x}_n - \boldsymbol{x}\|$ tends to zero as n tends to infinity. Sequences provide a characterisation of closure points.

Theorem 1.7.7[10] *Let $X \subseteq \mathbb{R}^d$ be a set in $\mathbb{R}^d$. A point $\boldsymbol{x} \in \mathbb{R}^d$ is a closure point of X if and only if there exists a sequence of points of X converging to $\boldsymbol{x}$.*

A *boundary point* of a convex set $X \subseteq \mathbb{R}^d$ is a point $\boldsymbol{x}$ in which every open ball with centre at $\boldsymbol{x}$ meets both X and $\mathbb{R}^d \setminus X$. The set of boundary points is the *boundary* of X and is denoted $\operatorname{bd} X$. The definition of a boundary point gives the following at once.

Proposition 1.7.8 *Let $X \subseteq \mathbb{R}^d$ be a set and let $\boldsymbol{x} \in \operatorname{bd} X$. The following assertions hold.*

(i) $\operatorname{bd} X = \operatorname{cl} X \setminus \operatorname{int} X$.
(ii) *There exists a sequence of points of X converging to $\boldsymbol{x}$.*
(iii) *There exists a sequence of points of $\mathbb{R}^d \setminus X$ converging to $\boldsymbol{x}$.*

9 A proof is available in Webster (1994, thm. 2.3.7).
10 A proof is available in Webster (1994, thm. 1.8.2).

The boundary of a convex set X with respect to aff X defines its relative boundary. Formally, a *relative boundary point* of a convex set $X \subseteq \mathbb{R}^d$ is a point in $\operatorname{cl} X \setminus \operatorname{rint} X$. The set of boundary points is the *relative boundary* of X and is denoted rbd X. If $X \subseteq \mathbb{R}^d$ is full dimensional then $\operatorname{rbd} X = \operatorname{bd} X$, while if $\dim X < d$ then $\operatorname{rbd} X \subset \operatorname{bd} X = \operatorname{cl} X$.

Theorem 1.7.9[11] *Let $X \subseteq \mathbb{R}^d$ be a convex set. Then* rint X, int X, *and* cl X *are all convex sets.*

1.8 Separation and Support

Two disjoint nonempty compact convex sets can be 'separated' by a hyperplane (Theorem 1.8.5); this is an intuitive and fundamental result in convexity. Also intuitive and fundamental is the result that a closed convex set is the intersection of (possibly infinitely many) halfspaces (Theorem 1.8.3) that 'support' the set. This section explores this kind of result.

Every nonempty closed set X contains a point $\boldsymbol{x}_0$ closest to a given point $\boldsymbol{a}$ of $\mathbb{R}^d$. If we add convexity then the point $\boldsymbol{x}_0$ is unique and the angle between the vectors $\boldsymbol{a} - \boldsymbol{x}$ and $\boldsymbol{x} - \boldsymbol{x}_0$ is nonacute for every $\boldsymbol{x} \in X$. The details are captured in Theorem 1.8.1.

Theorem 1.8.1[12] *Let X be a nonempty closed convex set in $\mathbb{R}^d$ and let $\boldsymbol{a} \in \mathbb{R}^d$. Then there exists a unique point $\boldsymbol{x}_0 \in X$ that is closest to $\boldsymbol{a}$; notationally,*

$$\|\boldsymbol{a} - \boldsymbol{x}_0\| = \inf\{\|\boldsymbol{a} - \boldsymbol{x}\| \mid \boldsymbol{x} \in X\}.$$

Moreover, the angle between the vectors $\boldsymbol{a} - \boldsymbol{x}_0$ and $\boldsymbol{x} - \boldsymbol{x}_0$ is nonacute for every $\boldsymbol{x} \in X$; notationally,

$$(\boldsymbol{a} - \boldsymbol{x}_0) \cdot (\boldsymbol{x} - \boldsymbol{x}_0) \leqslant 0 \text{ for every } \boldsymbol{x} \in X.$$

Figure 1.8.1 depicts a geometric description of the theorem.

We proceed with a number of definitions. Let X be a set in $\mathbb{R}^d$. A *supporting halfspace* of X is a closed halfspace in $\mathbb{R}^d$ that contains X and whose bounding hyperplane meets cl X. A *supporting hyperplane* of X in $\mathbb{R}^d$ is a hyperplane H that bounds a supporting halfspace of X. In the case $X \subseteq H$, the supporting hyperplane is said to be *trivial*; otherwise it is said to be *nontrivial*. Notationally, a hyperplane $H(\boldsymbol{a}, \alpha)$ is a supporting hyperplane of X if and only if

[11] A proof is available in Webster (1994, thm. 2.3.5).
[12] A proof is available in Webster (1994, thm. 2.4.1).

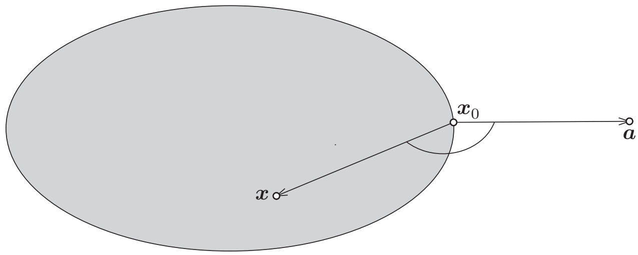

Figure 1.8.1 The point $\boldsymbol{x}_0$ in a closed convex set that is closest to a point $\boldsymbol{a} \in \mathbb{R}^d$; the nonacute angle between $\boldsymbol{a} - \boldsymbol{x}_0$ and $\boldsymbol{x} - \boldsymbol{x}_0$ is also shown.

$$\text{either } \alpha = \sup_{\boldsymbol{x} \in X} \boldsymbol{a} \cdot \boldsymbol{x} \quad \text{or} \quad \alpha = \inf_{\boldsymbol{x} \in X} \boldsymbol{a} \cdot \boldsymbol{x}.$$

And a supporting hyperplane $H(\boldsymbol{a}, \alpha)$ is nontrivial if and only if

$$\text{either } \inf_{\boldsymbol{x} \in X} \boldsymbol{a} \cdot \boldsymbol{x} < \sup_{\boldsymbol{x} \in X} \boldsymbol{a} \cdot \boldsymbol{x} = \alpha \quad \text{or} \quad \sup_{\boldsymbol{x} \in X} \boldsymbol{a} \cdot \boldsymbol{x} > \inf_{\boldsymbol{x} \in X} \boldsymbol{a} \cdot \boldsymbol{x} = \alpha.$$

A supporting hyperplane of X is said to *support* X at the points where it intersects $\operatorname{cl} X$.

Theorem 1.8.1 facilitates the proof of a number of results; the first such result comes next.

Corollary 1.8.2 *Let $X \subseteq \mathbb{R}^d$ be a nonempty closed convex set and let $\boldsymbol{a} \notin X$. Then there exists a hyperplane that does not contain $\boldsymbol{a}$ and that supports X at the point $\boldsymbol{x}_0 \in X$ closest to $\boldsymbol{a}$.*

Proof Thanks to Theorem 1.8.1, we find a point $\boldsymbol{x}_0 \in X$ closest to $\boldsymbol{a}$ such that

$$(\boldsymbol{a} - \boldsymbol{x}_0) \cdot (\boldsymbol{x} - \boldsymbol{x}_0) \leqslant 0 \text{ for every } \boldsymbol{x} \in X$$

or, equivalently, that

$$(\boldsymbol{a} - \boldsymbol{x}_0) \cdot \boldsymbol{x} \leqslant (\boldsymbol{a} - \boldsymbol{x}_0) \cdot \boldsymbol{x}_0 \text{ for every } \boldsymbol{x} \in X.$$

Define $\boldsymbol{b} := \boldsymbol{a} - \boldsymbol{x}_0$ and $\beta := (\boldsymbol{a} - \boldsymbol{x}_0) \cdot \boldsymbol{x}_0$. It immediately follows that $H^-[\boldsymbol{b}, \beta]$ is a supporting halfspace of X and that $H(\boldsymbol{b}, \beta)$ supports X at $\boldsymbol{x}_0$.

Suppose that $\boldsymbol{a} \in H^-[\boldsymbol{b}, \beta]$. Then

$$0 \geqslant (\boldsymbol{a} - \boldsymbol{x}_0) \cdot \boldsymbol{a} - (\boldsymbol{a} - \boldsymbol{x}_0) \cdot \boldsymbol{x}_0 = \|\boldsymbol{a} - \boldsymbol{x}_0\|^2,$$

which implies that $\boldsymbol{a} = \boldsymbol{x}_0$ (Problem 1.12.1(ii)). This is a contradiction because $\boldsymbol{x}_0 \in X$ and $\boldsymbol{a} \notin X$. □

Another consequence of Theorem 1.8.1 is that a closed convex set is the set of solutions of a system of (possibly infinitely many) linear inequalities.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

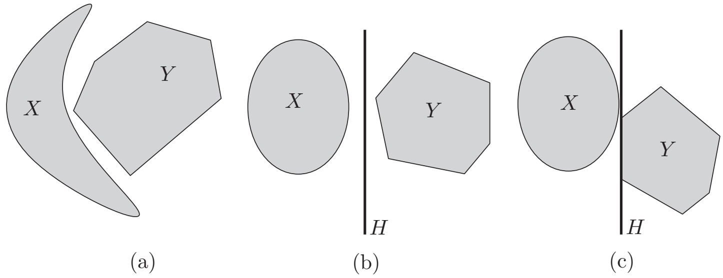

Figure 1.8.2 Examples of the separation of sets $X, Y \subseteq \mathbb{R}^2$. (a) Sets X and Y that cannot be separated. (b) Sets X and Y that are strictly separated by the hyperplane H. (c) Sets X and Y that are separated but not strictly by the hyperplane H.

Theorem 1.8.3[13] *A nonempty closed convex set in $\mathbb{R}^d$ is the intersection of its supporting halfspaces.*

It seems intuitive that through each relative boundary point of a closed convex set passes a nontrivial supporting hyperplane; this intuition is correct.

Theorem 1.8.4[14] *Let X be a nonempty convex set in $\mathbb{R}^d$ and let $\boldsymbol{x}_0 \in \operatorname{bd} X$. Then there exists a hyperplane in $\mathbb{R}^d$ supporting X at $\boldsymbol{x}_0$. In the case $\boldsymbol{x}_0 \in \operatorname{rbd} X$, the hyperplane can be assumed to be nontrivial.*

A hyperplane in $\mathbb{R}^d$ separates the space into two closed halfspaces. This fact gives rise to the important concept of separation of convex sets. Let X and Y be two sets in $\mathbb{R}^d$ and let H be a hyperplane in $\mathbb{R}^d$. The sets X and Y are said to be *separated* by H if X and Y lie in opposite closed halfspaces defined by H. And the sets X and Y are said to be *strictly separated* by H if X and Y lie in opposite open halfspaces defined by H. Figure 1.8.2(b)–(c) exemplifies these new notions.

Not every two disjoint sets in $\mathbb{R}^d$ can be separated by a hyperplane, as Fig. 1.8.2(a) shows. But every two disjoint convex sets can. We, however, content ourselves with an instance of this assertion.

Theorem 1.8.5 (Separation theorem)[15] *Let X and Y be disjoint nonempty convex sets in $\mathbb{R}^d$. Suppose that X is closed and Y is compact. Then X and Y can be strictly separated by a hyperplane in $\mathbb{R}^d$.*

Recall that a set X in $\mathbb{R}^d$ is said to be *compact* if each sequence of its points contains a subsequence that converges to a point of X. Compact sets in $\mathbb{R}^d$

[13] A proof is available in Webster (1994, cor. 2.4.8).
[14] A proof is available in Webster (1994, thm. 2.4.12).
[15] A proof is available in Webster (1994, thm. 2.4.6).

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

have a neat characterisation: they are precisely the closed and bounded sets in $\mathbb{R}^d$ (Webster, 1994, thm. 1.8.4).

We state a useful, particular case of Theorem 1.8.5.

Corollary 1.8.6 *Let $X \subseteq \mathbb{R}^d$ be a nonempty closed convex set and let $\boldsymbol{a} \notin X$. Then X can be strictly separated from $\boldsymbol{a}$ by a hyperplane.*

1.9 Faces

Convex sets are structured around other convex sets, their faces; this feature makes their mathematics amenable to inductive arguments. A convex subset F of a convex set $X \subseteq \mathbb{R}^d$ is a *face* of X if each time the relative interior of a segment in X meets F then the segment is fully contained in F. By the convexity of F, a segment is in F if its endpoints are. Hence the definition of a face amounts to the following.

$$\text{If } \alpha \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2 \in F \text{ with } \boldsymbol{x}_1, \boldsymbol{x}_2 \in X \text{ and } \alpha \in (0,1), \text{ then } \boldsymbol{x}_1, \boldsymbol{x}_2 \in F. \tag{1.9.1}$$

It is clear that $\emptyset$ and X itself are faces of X; these are its *improper* faces. Every other face of X is a *proper* face.

A k-dimensional face of a convex set X is referred to as a *k-face*. The set of k-faces of X is denoted by $\mathcal{F}_k(X)$ and the set of all faces of X is denoted by $\mathcal{F}(X)$. The 0-faces are called *extreme points*, and the set of extreme points of X is denoted by $\operatorname{ext} X$.

Proposition 1.9.2 *Let X be a convex set. An extreme point $\boldsymbol{v}$ of X is not in any segment $(\boldsymbol{x}_1, \boldsymbol{x}_2)$ with $\boldsymbol{x}_1, \boldsymbol{x}_2 \in X$ or, equivalently, the set $X \backslash \{\boldsymbol{v}\}$ is again convex.*

Another immediate consequence of (1.9.1) is that 'is a face of' is a transitive relation on the faces of a convex set X.

Proposition 1.9.3 *Let X be a convex set in $\mathbb{R}^d$. If F is a face of X and I is a face of F, then I is a face of X.*

We now characterise faces.

Theorem 1.9.4[16] *Let X be a convex set in $\mathbb{R}^d$ and let F be a convex subset of X. Then F is a face of X if and only if $X \backslash F$ is convex and $F = X \cap \operatorname{aff} F$.*

[16] A proof is available in Webster (1994, thm. 2.6.2).

We list a corollary of Theorem 1.9.4.

Corollary 1.9.5 *If X is a convex set in $\mathbb{R}^d$ and F is a proper face of X, then* $\dim F < \dim X$.

Proof Suppose $\dim F = \dim X$. Then, $\operatorname{aff} F = \operatorname{aff} X$ and, by Theorem 1.9.4, $F = X \cap \operatorname{aff} X = X$. Hence, a proper face F of X must satisfy $\dim F < \dim X$. □

The next theorem reveals five ways in which faces appear.

Theorem 1.9.6 *Let $X \subseteq \mathbb{R}^d$ be a convex set. Then the following assertions hold.*

(i) *The intersection of any nonempty family of faces of X is a face of X.*
(ii) *The intersection of any nonempty family of faces of X can be expressed as the intersection of at most $d+1$ members of the family.*
(iii) *The intersection of X and any of its supporting hyperplanes is a face of X.*
(iv) *Let $\varphi\colon \mathbb{R}^d \to \mathbb{R}^d$ be an affine transformation. If F is a face of X, then the set $\varphi^{-1}(F)$ is a face of $\varphi^{-1}(X)$.*
(v) *Let $\varphi\colon \mathbb{R}^d \to \mathbb{R}^d$ be an injective affine transformation and let $F \subseteq X$. The set F is a face of X if and only if the set $\varphi(F)$ is a face of $\varphi(X)$.*

Proof Take $\boldsymbol{x}_1, \boldsymbol{x}_2 \in X$ and $\alpha \in (0,1)$.

(i) Let I be some index set and let $\{F_i \mid i \in I\}$ be a family of faces of X. Let $F := \bigcap_{i \in I} F_i$. Then F is a convex subset of X (Theorem 1.6.7). Suppose that $\alpha \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2 \in F$. It follows that $\alpha \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2 \in F_i$ for each $i \in I$. Since F_i is a face, we get that $\boldsymbol{x}_1, \boldsymbol{x}_2 \in F_i$ for each $i \in I$. This confirms that $\boldsymbol{x}_1, \boldsymbol{x}_2 \in F$ and that F is a face of X.

(ii) Let I be some index set; let $\{F_i \mid i \in I\}$ be a family of distinct faces of X and let $F := \bigcap_{i \in I} F_i$. We show that there are elements $i_1, \ldots, i_n$ of I, with $n \leqslant d+1$, such that

$$F = F_{i_1} \cap \cdots \cap F_{i_n}. \tag{1.9.6.1}$$

Pick any element $i_1 \in I$. If $F = F_{i_1}$ then we are done; otherwise $F \subset F_{i_1}$, $\dim F < \dim F_{i_1}$ by Corollary 1.9.5, and there is another index i_2 of I such that

$$F \subseteq F_{i_1} \cap F_{i_2} \subset F_{i_1}.$$

The selection of i_2 is possible because the intersection of all these faces F_i (with $i \in I$) is F and $\dim F < \dim F_{i_1}$. If the statement is false, then there are indices $i_1, \ldots, i_n$ of I such that $n \geqslant d+2$, and

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

$$F \subset F_{i_1} \cap \cdots \cap F_{i_n} \subset \cdots \subset F_{i_1} \cap F_{i_2} \subset F_{i_1}.$$

Additionally, as a consequence of Corollary 1.9.5, we have that $\dim F_{i_1} \leqslant d$ and

$$\dim F < \dim(F_{i_1} \cap F_{i_2} \cap \cdots \cap F_{i_n}) < \cdots < \dim(F_{i_1} \cap F_{i_2}) < \dim F_{i_1}.$$

But from this series of inequalities, it follows that

$$-1 \leqslant \dim F < \dim(F_{i_1} \cap F_{i_2} \cap \cdots \cap F_{i_n}) < \dim F_{i_1} - (n-1) \leqslant d - n + 1,$$

which yields that $-1 < d - n + 1$, a contradiction for $n \geqslant d + 2$. Hence (1.9.6.1) holds and $n \leqslant d + 1$, as desired.

(iii) Let $H := H(\boldsymbol{b}, \beta)$ be a supporting hyperplane of X with $X \subseteq H^-[\boldsymbol{b}, \beta]$. Then $\boldsymbol{b} \cdot \boldsymbol{x} \leqslant \beta$ for every $\boldsymbol{x} \in X$. Let $F := X \cap H$. It follows that F is a convex subset of X. Now suppose that $\alpha \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2 \in F$. Since $\alpha \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2 \in H$, we must have that

$$\beta = \boldsymbol{b} \cdot (\alpha \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{x}_2) = \alpha \boldsymbol{b} \cdot \boldsymbol{x}_1 + (1-\alpha)\boldsymbol{b} \cdot \boldsymbol{x}_2 \leqslant \alpha\beta + (1-\alpha)\beta = \beta.$$

Hence, $\boldsymbol{b} \cdot \boldsymbol{x}_1 = \boldsymbol{b} \cdot \boldsymbol{x}_2 = \beta$ since $\alpha \in (0,1)$. This proves that $\boldsymbol{x}_1, \boldsymbol{x}_2 \in X \cap H$ and that F is a face of X.

(iv) Take $\boldsymbol{y}_1, \boldsymbol{y}_2 \in \varphi^{-1}(X)$. Then $\varphi(\boldsymbol{y}_1), \varphi(\boldsymbol{y}_2) \in X$. From F being a convex set of X, it follows that $\varphi^{-1}(F)$ is a convex subset of $\varphi^{-1}(X)$ (Theorem 1.6.7(v)). Suppose that

$$\alpha \boldsymbol{y}_1 + (1-\alpha)\boldsymbol{y}_2 \in \varphi^{-1}(F).$$

Then $\varphi(\alpha \boldsymbol{y}_1 + (1-\alpha)\boldsymbol{y}_2) \in F$. Since φ is an affine transformation, we have that

$$\varphi(\alpha \boldsymbol{y}_1 + (1-\alpha)\boldsymbol{y}_2) = \alpha\varphi(\boldsymbol{y}_1) + (1-\alpha)\varphi(\boldsymbol{y}_2),$$

and since F is a face of X we have that $\varphi(\boldsymbol{y}_1), \varphi(\boldsymbol{y}_2) \in F$. Hence $\boldsymbol{y}_1, \boldsymbol{y}_2 \in \varphi^{-1}(F)$.

(v) The proof goes along the lines of that for (iv) and is left to the reader. □

The condition of φ being injective in Theorem 1.9.6(v) cannot be removed. We can find many examples of convex sets X, faces F of X, and noninjective affine transformations φ for which $\varphi(F)$ is not a face of $\varphi(X)$. Consider, for instance, the two-dimensional unit ball $X := B[\boldsymbol{0}, 1]$ and the orthogonal projection φ of $\mathbb{R}^2$ onto the x-axis. Then $\varphi(X) = [-\boldsymbol{1}, \boldsymbol{1}]$. It is now clear that, for any extreme point $\boldsymbol{x}$ of X other than $\pm\boldsymbol{1}$, $\varphi(\boldsymbol{x})$ is not a face of $[-\boldsymbol{1}, \boldsymbol{1}]$.

A face of a convex set X is *exposed* if it is either an improper face of X or is of the form $X \cap H$ where H is a supporting hyperplane of X (Theorem 1.9.6(iii)); Figure 1.9.1(b) gives examples of exposed faces. Suppose

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

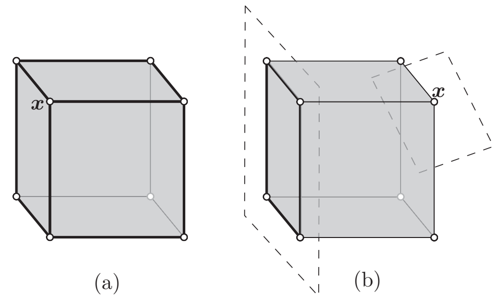

Figure 1.9.1 Faces of convex sets. (a) An extreme point $\boldsymbol{x}$ as the intersection of three 2-faces. (b) A 2-face and an extreme point $\boldsymbol{x}$ as exposed faces.

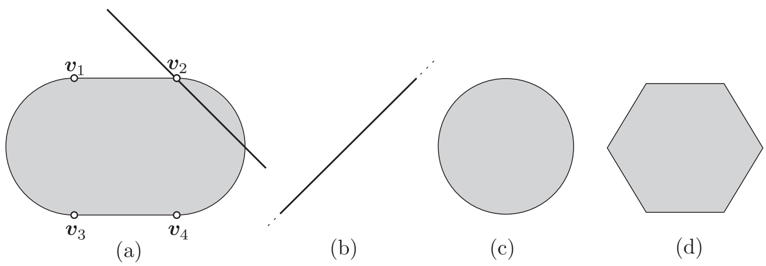

Figure 1.9.2 Faces in closed convex sets. (a) A two-dimensional closed convex set X with faces that are not exposed. The set X is the convex hull of two closed discs; its boundary consists of two closed segments $[\boldsymbol{v}_1, \boldsymbol{v}_2]$ and $[\boldsymbol{v}_3, \boldsymbol{v}_4]$ and two open halfcircles. The extreme points are the points $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4$ and the points on each of the two open halfcircles, and the 1-faces are the two aforementioned segments. The extreme points $\boldsymbol{v}_1, \boldsymbol{v}_2, \boldsymbol{v}_3, \boldsymbol{v}_4$ are not exposed: a supporting hyperplane of X passing through precisely one of them must also meet one of the two open halfcircles; every other face of X is exposed. (b) A convex set with no extreme points. (c) A two-dimensional closed convex set with extreme points but no 1-faces. (d) A two-dimensional closed convex set with extreme points and 1-faces.

that $H = \{\boldsymbol{x} \in \mathbb{R}^d \mid \boldsymbol{a} \cdot \boldsymbol{x} = \alpha\}$ and that X lies in the supporting halfspace $H^-[\boldsymbol{a}, \alpha]$. From the definition it follows that an exposed face $F := X \cap H$ maximises the linear functional $\boldsymbol{a} \cdot \boldsymbol{x}$ over X:

$$\alpha = \max\{\boldsymbol{a} \cdot \boldsymbol{x} \mid \boldsymbol{x} \in X\} \text{ and } F = \{\boldsymbol{x} \in X \mid \boldsymbol{a} \cdot \boldsymbol{x} = \alpha\}. \tag{1.9.7}$$

This face is often denoted by $F(X, \boldsymbol{a})$. An inequality $\boldsymbol{c} \cdot \boldsymbol{x} \leqslant \gamma$ is *valid* for a set in $\mathbb{R}^d$ if it is satisfied for every point $\boldsymbol{x}$ in the set. The inequality $\boldsymbol{a} \cdot \boldsymbol{x} \leqslant \alpha$ that defines $H^-[\boldsymbol{a}, \alpha]$ is valid for X.

Not every face of a convex set X is exposed, as Fig. 1.9.2(a) shows. However, every intersection of exposed faces of X is exposed.

Proposition 1.9.8[17] *Let X be a convex set in $\mathbb{R}^d$. Then the intersection of any nonempty family of exposed faces of X is also exposed.*

Theorem 1.8.4 established that, through each relative boundary point of a nonempty convex set X, it passes a nontrivial supporting hyperplane of X, while Theorem 1.9.6(iii) established that the intersection of X and a nontrivial supporting hyperplane of it is a proper face of X. Consequently, some proper faces of X meet its relative boundary. More is true: proper faces of X are contained in the relative boundary of X.

Theorem 1.9.9[18] *If X is a closed convex set in $\mathbb{R}^d$ and F is a proper face of X, then $F \subseteq \operatorname{rbd} X$.*

According to Theorem 1.9.6(i), for any subset Y of a convex set X there is a smallest face of X containing Y, the intersection of all faces containing Y. In the case of Y being a point, a simple characterisation of the smallest face containing it follows from Theorem 1.9.9.

Theorem 1.9.10[19] *Let X be a closed convex set in $\mathbb{R}^d$, let F be a proper face of X, and let $\boldsymbol{x} \in F$. Then F is the smallest face of X that contains $\boldsymbol{x}$ if and only if $\boldsymbol{x} \in \operatorname{rint} F$.*

Figure 1.9.2(b) depicts a closed convex set in $\mathbb{R}^2$ with no extreme points. But every bounded, closed, and convex set in $\mathbb{R}^d$ has extreme points and, moreover, it is spanned by them. This is perhaps the most important result on the facial structure of compact convex sets.

Theorem 1.9.11 (Minkowski–Krein–Milman's theorem)[20] *Let X be a compact convex set in $\mathbb{R}^d$. Then X is the convex hull of its extreme points. Notationally,*

$$X = \operatorname{conv}(\operatorname{ext} X).$$

A simple consequence of Minkowski–Krein–Milman's theorem (1.9.11) is that every compact convex set has extreme points, but this does not extend to higher dimensional faces. Figure 1.9.2(c) depicts a compact convex set with no 1-faces. In fact, a compact convex set in $\mathbb{R}^d$ can have a very diverse face-dimension pattern. The *face-dimension pattern* of a compact convex set X is an increasing sequence $(d_1, \ldots, d_n)$ of positive integers that encode all possible positive dimensions of faces in X. We have all the possible patterns in $\mathbb{R}^2$:

[17] A proof is available in Webster (1994, thm. 2.6.17).
[18] A proof is available in Brøndsted (1983, thm. 5.3).
[19] A proof is available in Brøndsted (1983, thm. 5.6).
[20] A proof is available in Brøndsted (1983, thm. 5.10).

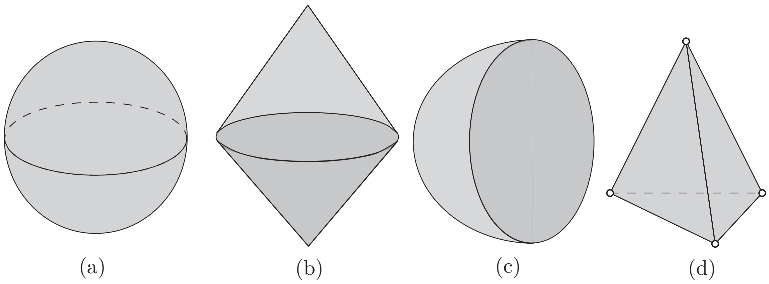

Figure 1.9.3 All the possible face-dimension patterns for three-dimensional compact convex sets in $\mathbb{R}^3$. (a) The three-dimensional unit ball, which has pattern (3). (b) A set obtained as the convex hull of a circle on a plane and two points at opposite sides of the plane; the set has pattern $(1,3)$. (c) A set obtained by intersecting a closed halfspace with the three-dimensional, unit ball; the set has pattern $(2,3)$. (d) A 3-simplex, which has pattern $(1,2,3)$.

the pattern (), exemplified by a singleton; the pattern (1), exemplified by a line segment; the pattern (2), exemplified by Fig. 1.9.2(c); and the pattern $(1,2)$, exemplified by Fig. 1.9.2(d). This gives a total of 2^2 patterns in $\mathbb{R}^2$. Figure 1.9.3 shows examples with all the possible face-dimension patterns in three-dimensional compact convex sets: (3), $(1,3)$, $(2,3)$, and $(1,2,3)$; these examples together with the lower dimensional examples give a total of 2^3 patterns in $\mathbb{R}^3$.

A result of Roshchina et al. (2018) states that for any finite, increasing sequence of positive integers, there exists a compact convex set in $\mathbb{R}^d$ that has extreme points and faces with dimensions only from this prescribed sequence; in other words, for any of the 2^d possible face-dimension patterns, there is a compact convex set in $\mathbb{R}^d$ exhibiting that pattern.

1.10 Cones and Lineality Spaces

If on the line through the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_1 + \boldsymbol{a}_2$, namely the set

$$\left\{\boldsymbol{a}_1 + \alpha \boldsymbol{a}_2 \,\middle|\, \boldsymbol{a}_1 \in \mathbb{R}^d, \boldsymbol{0} \neq \boldsymbol{a}_2 \in \mathbb{R}^d, \text{and every } \alpha \in \mathbb{R}\right\},$$

we impose the condition of $\alpha \geqslant 0$, we arrive at the definition of a *ray*.

A subset X of $\mathbb{R}^d$ is a convex *cone* if it is convex and it contains the ray passing through any of its points and the origin; that is, X is a set of the form

$$\{\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2 |\; \alpha_1, \alpha_2 \geqslant 0 \text{ and } \boldsymbol{a}_1, \boldsymbol{a}_2 \in X\}. \tag{1.10.1}$$

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

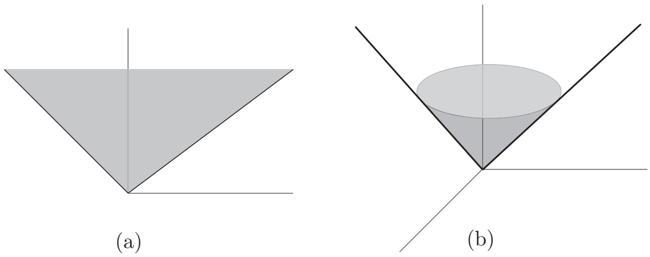

Figure 1.10.1 Parts of cones. (a) Part of a cone in $\mathbb{R}^2$. (b) Part of a cone in $\mathbb{R}^3$.

The expression $\alpha_1 \boldsymbol{a}_1 + \alpha_2 \boldsymbol{a}_2$ where $\alpha_1, \alpha_2 \geqslant 0$ is the *positive combination* or *conical combination* of the points $\boldsymbol{a}_1$ and $\boldsymbol{a}_2$. We consider only convex cones, so we will drop the adjective convex hereafter, unless we want to reinforce the convexity of the cone. Figure 1.10.1 depicts examples of cones in $\mathbb{R}^3$.

The *positive hull* or *conical hull* of a set $X \subseteq \mathbb{R}^d$, denoted cone X, is the smallest cone containing X. It is also defined as the set of all positive combinations of finitely many points of X. That is,

$$\operatorname{cone} X = \{\alpha_1 \boldsymbol{a}_1 + \cdots + \alpha_n \boldsymbol{a}_n \mid \boldsymbol{a}_i \in X \text{ and } \alpha_i \geqslant 0\}.$$

In this case, we say that cone X is *generated* by X and call the elements of X *generators*. And if X is finite, then we call the cone a *V-cone*.

A set $X := \{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n\}$ is *positively dependent* if some $\boldsymbol{a}_i$ is a positive combination of the others; otherwise it is *positively independent*. The set X *positively spans* a linear subspace L if any vector of the subpace can be expressed as a positive combination of elements of X. And the set X is a *positive basis* of L if it is positively independent and positively spans the subspace. The next characterisation of positively spanning sets is well known.

Theorem 1.10.2 (Davis, 1954, thm. 3.6) *Let $X := \{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n\}$ with each $\boldsymbol{a}_i \neq \mathbf{0}$ such that X linearly spans $\mathbb{R}^d$. Then the following statements are equivalent.*

(i) *The set X positively spans $\mathbb{R}^d$.*
(ii) *For every $i \in [1 \ldots n]$, the point $-\boldsymbol{a}_i \in \operatorname{cone}(X \setminus \{\boldsymbol{a}_i\})$.*
(iii) $\mathbf{0} = \sum_{i=1}^{n} \alpha_i \boldsymbol{a}_i$, *with $\alpha_i > 0$ for $i \in [1 \ldots n]$.*

It will often be useful to call translations of convex cones *affine convex cones*, following the same analogy between linear spaces and affine spaces. An *affine convex cone* A in $\mathbb{R}^d$ is a set of the form $\boldsymbol{a}_0 + C$ where $\boldsymbol{a}_0 \in \mathbb{R}^d$ and C is a convex cone in $\mathbb{R}^d$, the translation of C by $\boldsymbol{a}_0$. We often call the point

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

$\boldsymbol{a}_0$ the *apex* of the cone, or we say that the affine convex cone is *based at* $\boldsymbol{a}_0$. Every (standard) convex cone is an affine convex cone with $\mathbf{0}$ as the apex. In this way, we have that hyperplanes and halfspaces are affine convex cones.

Lineality Spaces and Recession Cones

The dimension of the dual set of a set (Section 1.11) and the structure of an unbounded convex set can be described by the lines and rays that are contained in the set, and specifically by the lineality space and recession cone of the set.

Let X be a nonempty convex set in $\mathbb{R}^d$. The *lineality space* of X is defined as

$$\operatorname{lineal} X := \left\{ \boldsymbol{y} \in \mathbb{R}^d \,\middle|\, \boldsymbol{x} + \alpha \boldsymbol{y} \in X \text{ for all } \boldsymbol{x} \in X \text{ and all } \alpha \in \mathbb{R} \right\}, \qquad (1.10.3)$$

and the *recession cone* of X is defined as

$$\operatorname{rec} X := \left\{ \boldsymbol{y} \in \mathbb{R}^d \,\middle|\, \boldsymbol{x} + \alpha \boldsymbol{y} \in X \text{ for all } \boldsymbol{x} \in X \text{ and all } \alpha \geqslant 0 \right\}. \qquad (1.10.4)$$

From the definitions, it easily follows that

$$\operatorname{lineal} X = \operatorname{rec} X \cap (-\operatorname{rec} X) \qquad (1.10.5)$$

and that the set lineal X is a linear subspace of $\mathbb{R}^d$.

A convex set X is *pointed* if it contains no lines, which together with the definition of lineal X implies the following.

Proposition 1.10.6 *A nonempty convex set X in $\mathbb{R}^d$ is pointed if and only if* lineal $X = \{\mathbf{0}\}$.

Another basic property of lineality spaces is stated below.

Proposition 1.10.7[21] *If X is a nonempty closed convex set in $\mathbb{R}^d$, then*

$$X = \operatorname{lineal} X + \left(X \cap (\operatorname{lineal} X)^{\perp} \right),$$

where $X \cap (\operatorname{lineal} X)^{\perp}$ is a pointed closed convex set.

According to Proposition 1.10.7, every point $\boldsymbol{x}$ of a closed convex set X can be uniquely written as a sum of a point in lineal X and one in a pointed closed convex set. This decomposition often makes it possible to focus on pointed closed convex sets.

Definitions (1.10.3) and (1.10.4) also ensure, albeit not at once, that the recession cone is a convex cone and that if X is a cone, then $\operatorname{rec} X = X$ (Problem 1.12.17).

[21] A proof is available in Webster (1994, thm. 2.5.8).

1.11 Dual Sets

This section explores another manifestation of duality.

Definition 1.11.1 (Dual set) With each set $X \subseteq \mathbb{R}^d$, we associate the set $X^* \subseteq \mathbb{R}^d$ defined as

$$X^* := \left\{ \boldsymbol{y} \in \mathbb{R}^d \,\middle|\, \boldsymbol{y} \cdot \boldsymbol{x} \leqslant 1 \text{ for every } \boldsymbol{x} \in X \right\}.$$

The set X^* is said to be the *dual* of X.

From this definition, it is plain that the dual of $\varnothing$ and $\{\mathbf{0}\}$ is $\mathbb{R}^d$ and that the dual of $\mathbb{R}^d$ is $\{\mathbf{0}\}$. Moreover, it follows that the set X^* can also be expressed as

$$X^* = \bigcap_{\boldsymbol{x} \in X} \{\boldsymbol{x}\}^* = \bigcap_{\boldsymbol{x} \in X} H^-[\boldsymbol{x}, 1]. \tag{1.11.2}$$

The geometric relation between $\boldsymbol{x} \in X$ and $\{\boldsymbol{x}\}^* = H^-[\boldsymbol{x}, 1]$ is depicted in Fig. 1.11.1. The next proposition is also immediate from the definition of dual sets.

Proposition 1.11.3 *Let X be a subset in $\mathbb{R}^d$. Then X^* is a closed convex set in $\mathbb{R}^d$ that contains the origin.*

Example 1.11.4 We find the dual of a closed ball. Let $r > 0$. Each nonzero point $\boldsymbol{y} := (y_1, \ldots, y_d)^t$ of $B_d^*[\mathbf{0}_d, r]$ satisfies $\boldsymbol{y} \cdot \boldsymbol{x} \leqslant 1$ for every point $\boldsymbol{x} \in B_d[\mathbf{0}_d, r]$ and, in particular, for the point $\boldsymbol{x}_{\boldsymbol{y}} := r\boldsymbol{y}/\|\boldsymbol{y}\|$ of $B_d[\mathbf{0}_d, r]$. Then

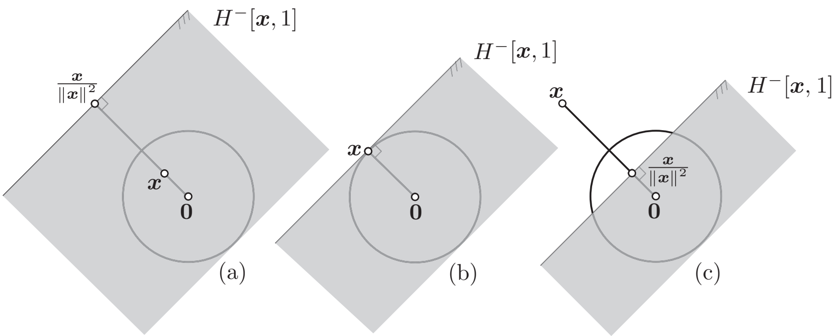

Figure 1.11.1 Geometric relation between $\boldsymbol{x}$ and $H^-[\boldsymbol{x}, 1]$ in $\mathbb{R}^2$. The halfspace $H^-[\boldsymbol{x}, 1]$ is highlighted in grey. The hyperplane bounding $H^-[\boldsymbol{x}, 1]$ passes through $\boldsymbol{x}/\|\boldsymbol{x}\|^2$. Also, the unit ball is depicted to better see the relative Euclidean distance between $\boldsymbol{x}$ and $\mathbf{0}$. (a) $\|\boldsymbol{x}\| < 1$. (b) $\|\boldsymbol{x}\| = 1$. (c) $\|\boldsymbol{x}\| > 1$.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

$$\boldsymbol{y} \cdot \boldsymbol{x}_y = \boldsymbol{y} \cdot \frac{r\boldsymbol{y}}{\|\boldsymbol{y}\|} = r\|\boldsymbol{y}\| \leqslant 1.$$

Hence

$$B_d^*[\mathbf{0}_d, r] \subseteq \left\{ \boldsymbol{z} \in \mathbb{R}^d \,\middle|\, \|\boldsymbol{z}\| \leqslant 1/r \right\}.$$

Take $\boldsymbol{z} \in \mathbb{R}^d$ such that $\|\boldsymbol{z}\| \leqslant 1/r$. Then, for every point $\boldsymbol{x}$ in $B_d[\mathbf{0}_d, r]$, the Cauchy–Schwarz inequality (Problem 1.12.1(iii)) ensures that

$$\boldsymbol{z} \cdot \boldsymbol{x} \leqslant |\boldsymbol{z} \cdot \boldsymbol{x}| \leqslant \|\boldsymbol{z}\|\|\boldsymbol{x}\| \leqslant \frac{1}{r} r = 1;$$

here $|x|$ denotes the absolute value of the real number x. Hence, $\boldsymbol{z} \in B_d^*[\mathbf{0}_d, r]$ and

$$B_d^*[\mathbf{0}_d, r] = \left\{ \boldsymbol{z} \in \mathbb{R}^d \,\middle|\, \|\boldsymbol{z}\| \leqslant 1/r \right\} = B_d[\mathbf{0}_d, 1/r]. \tag{1.11.4.1}$$

This concludes the example.

Example 1.11.4 illustrates another consequence of Definition 1.11.1, the order-reversing inclusion between subsets of X and subsets of X^*:

$$Y \subseteq X \implies X^* \subseteq Y^*. \tag{1.11.5}$$

For any set $X \subseteq \mathbb{R}^d$, the dual set X^{**} of X^* is well defined. Additionally, if $\boldsymbol{x} \in X$ then, for every $\boldsymbol{y} \in X^*$, we have that $\boldsymbol{x} \cdot \boldsymbol{y} \leqslant 1$ by Definition 1.11.1. Hence $\boldsymbol{x} \in X^{**}$ and

$$X \subseteq X^{**}. \tag{1.11.6}$$

Combining (1.11.2), (1.11.5), and (1.11.6), we get that the set $X^{**} \subseteq \mathbb{R}^d$ is a closed convex set containing X and the origin. Thus, X^{**} contains the smallest closed convex set in $\mathbb{R}^d$ that contains X and the origin, namely $\operatorname{cl}(\operatorname{conv}(X \cup \{\mathbf{0}\}))$. But more is true: $X^{**} = \operatorname{cl}(\operatorname{conv}(X \cup \{\mathbf{0}\}))$ (Webster, 1994, thm. 2.8.3). The next theorem follows from this discussion.

Theorem 1.11.7[22] *If $X \subseteq \mathbb{R}^d$ is a closed convex set that contains the origin, then $X^{**} = X$.*

Because of Theorem 1.11.7, we are mostly interested in dual sets of closed convex sets that contains the origin.

While taking duals is an involution in the class of closed convex sets in $\mathbb{R}^d$ that contain the origin (Theorem 1.11.7), compactness is not necessarily preserved by this involution. Compactness is, however, preserved in a special subclass.

[22] A proof is available in Webster (1994, thm. 2.8.3).

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

Theorem 1.11.8[23] *Let $X \subseteq \mathbb{R}^d$ be a closed convex set that contains the origin. The set X^* is bounded if and only if the set X contains the origin in its interior, and vice versa.*

We state a consequence of Theorems 1.11.7 and 1.11.8.

Corollary 1.11.9 *If $X \subseteq \mathbb{R}^d$ is a compact convex set that contains the origin in its interior, then so is X^*. In addition, we have that $X^{**} = X$.*

Linear Subspaces and Cones

In the case of linear subspaces and convex cones, Definition 1.11.1 for the dual set can be sharpened.

Let L be a linear subspace of $\mathbb{R}^d$. We show that if $\boldsymbol{y} \in L^*$, then $\boldsymbol{y} \cdot \boldsymbol{x} = 0$ for every $\boldsymbol{x} \in L$. Suppose otherwise: $\boldsymbol{y} \cdot \boldsymbol{x} \neq 0$ for some $\boldsymbol{x} \in L$. Because L is a linear space, $\alpha \boldsymbol{x} \in L$ for each $\alpha \in \mathbb{R}$. In the case $\boldsymbol{y} \cdot \boldsymbol{x} > 0$, we choose $\alpha > 0$ sufficiently large, and in the case $\boldsymbol{y} \cdot \boldsymbol{x} < 0$, we choose $\alpha < 0$ with $|\alpha|$ sufficiently large. In either case, our chosen α would cause $\boldsymbol{y} \cdot (\alpha \boldsymbol{x}) > 1$. This contradiction validates our claim. Hence,

$$L^* = \left\{ \boldsymbol{y} \in \mathbb{R}^d \,\middle|\, \boldsymbol{y} \cdot \boldsymbol{x} = 0 \text{ for every } \boldsymbol{x} \in L \right\} = L^{\perp}. \tag{1.11.10}$$

In other words, the dual of a linear subspace coincides with its orthogonal complement (Example 1.1.6). Therefore, we will use (1.11.10) as the definition of the *dual linear subspace* L^* of L. It is instructive to compare this discussion on dual linear subspaces with our discussion on dual spaces in Section 1.5.

Let C be a convex cone. The analysis in the previous paragraph also proves that if $\boldsymbol{y} \in C^*$, then $\boldsymbol{y} \cdot \boldsymbol{x} \leqslant 0$ for every $\boldsymbol{x} \in C$. Suppose otherwise: $\boldsymbol{y} \cdot \boldsymbol{x} > 0$ for some $\boldsymbol{x} \in C$. Then, since $\alpha \boldsymbol{x} \in C$ for any $\alpha > 0$, choosing α sufficiently large would cause $\boldsymbol{y} \cdot (\alpha \boldsymbol{x}) > 1$. This contradiction validates our claim. Hence,

$$C^* = \left\{ \boldsymbol{y} \in \mathbb{R}^d \,\middle|\, \boldsymbol{y} \cdot \boldsymbol{x} \leqslant 0 \text{ for every } \boldsymbol{x} \in C \right\} = \bigcap_{\boldsymbol{x} \in C} H^-[\boldsymbol{x}, 0]. \tag{1.11.11}$$

It is clear that $\alpha_1 \boldsymbol{y}_1 + \alpha_2 \boldsymbol{y}_2 \in C^*$ for every $\boldsymbol{y}_1, \boldsymbol{y}_2 \in C^*$ and every $\alpha_1, \alpha_2 \geqslant 0$. Hence, C^* is a closed convex cone by (1.10.1) and (1.11.11). As a consequence, we will use (1.11.11) as the definition of the *dual cone* C^* of C. Note that our dual cone is sometimes called the polar cone of C; see, for instance, Lauritzen (2013, sec. 3.4)

[23] A proof is available in Webster (1994, thm. 2.8.4).

1.12 Problems

1.12.1 Let $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$ be three vectors in $\mathbb{R}^d$ and let $\alpha \in \mathbb{R}$. Prove the following properties of the dot product.

(i) $\boldsymbol{x} \cdot \boldsymbol{y} = \boldsymbol{y} \cdot \boldsymbol{x}$.
(ii) $\boldsymbol{x} \cdot \boldsymbol{x} = \|\boldsymbol{x}\|^2 \geqslant 0$, and $\boldsymbol{x} \cdot \boldsymbol{x} = 0$ if and only if $\boldsymbol{x} = \boldsymbol{0}$.
(iii) $|\boldsymbol{x} \cdot \boldsymbol{y}| \leqslant \|\boldsymbol{x}\| \|\boldsymbol{y}\|$, and equality holds if and only if one vector is a scalar multiple of the other (Cauchy–Schwarz inequality).
(iv) $\|\boldsymbol{x} + \boldsymbol{y}\| \leqslant \|\boldsymbol{x}\| + \|\boldsymbol{y}\|$ (triangle inequality).
(v) $|\,\|\boldsymbol{x}\| - \|\boldsymbol{y}\|\,| \leqslant \|\boldsymbol{x} - \boldsymbol{y}\|$.
(vi) $(\alpha \boldsymbol{x}) \cdot \boldsymbol{y} = \alpha(\boldsymbol{x} \cdot \boldsymbol{y})$.
(vii) $\boldsymbol{x} \cdot (\boldsymbol{y} + \boldsymbol{z}) = \boldsymbol{x} \cdot \boldsymbol{y} + \boldsymbol{x} \cdot \boldsymbol{z}$.

1.12.2 Prove that a set of $\mathbb{R}^d$ is a linear subspace if and only if it is the solution of a system of homogenous linear equations.

1.12.3 Prove that a set of $\mathbb{A}^d$ is an affine subspace if and only if it is the solution of a system of linear equations.

1.12.4 Let L be a linear subspace of $\mathbb{R}^d$. Prove the following assertions.

(i) $(L^\perp)^\perp = L$.
(ii) $L^\perp \cap L = \{\boldsymbol{0}\}$.
(iii) $\dim L^\perp = d - \dim L$.
(iv) Every vector of $\mathbb{R}^d$ can be written uniquely as the sum of a vector in L and a vector in $L^\perp$.

1.12.5 Let $M \in \mathbb{R}^{n \times d}$. Prove the following assertions.

(i) $\operatorname{null} M = (\operatorname{row} M)^\perp$.
(ii) $\operatorname{null} M^t = (\operatorname{col} M)^\perp$.
(iii) $\dim(\operatorname{null} M) + \operatorname{rank} M = d$ (nullity–rank theorem).

1.12.6 Prove the following assertions related to linear and affine maps.

(i) Not every affine map is a linear map.
(ii) An affine map ϱ is linear if and only if $\varrho(\boldsymbol{0}) = \boldsymbol{0}$.
(iii) Any affine map can be obtained as a translation of some unique linear map.
(iv) Linear maps send linear subspaces into linear subspaces, and affine maps do the same for affine subspaces.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

(v) Injective linear maps send linear subspaces of dimension r into linear subspaces of dimension r, and so do injective affine maps with affine subspaces.
(vi) The set of linear automorphisms form a group, called the *general linear group*, under composition of functions.
(vii) The set of affine automorphisms form a group, called the *affine group*, under composition of functions.
(viii) The set of all linear maps between two linear spaces is itself a linear space.

1.12.7 Define the *image* of a linear map $\varphi\colon X \to Y$ as the linear subspace of Y consisting of the images of X under φ, and define the *rank* of φ as the dimension of its image. Prove that the rank of φ coincides with the rank of any matrix representing it.

1.12.8 (Continuity of linear and affine maps) This exercise explores continuity of linear and affine maps.

A map $\varphi\colon X \to \mathbb{R}^s$ defined on a nonempty set $X \subseteq \mathbb{R}^r$ is said to satisfy a *Lipschitz condition* on X if there exists a real number α such that, for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in X$, we have that

$$\|\varphi(\boldsymbol{x}_1) - \varphi(\boldsymbol{x}_2)\| \leqslant \alpha \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|.$$

Prove the following.

(i) If a map $\varphi\colon X \to \mathbb{R}^s$ satisfies the Lipschitz condition on a subset $X \subseteq \mathbb{R}^r$, then φ is continuous on X.
(ii) Every linear functional $\varphi(\boldsymbol{x}) := \boldsymbol{a} \cdot \boldsymbol{x}$ for some $\boldsymbol{a} \in \mathbb{R}^d$ is continuous on $\mathbb{R}^d$.
(iii) Every linear map $\varphi\colon \mathbb{R}^r \to \mathbb{R}^s$ is continuous on $\mathbb{R}^r$.
(iv) Every affine map $\varphi\colon \mathbb{R}^r \to \mathbb{R}^s$ is continuous on $\mathbb{R}^r$.

1.12.9 Suppose that $B := \{\boldsymbol{l}_1, \ldots, \boldsymbol{l}_r\}$ is a basis of a linear subspace L of $\mathbb{R}^d$ and let $\boldsymbol{x}$ be any vector of $\mathbb{R}^d$. Write $M := (\boldsymbol{l}_1 \cdots \boldsymbol{l}_r)$. Then M is a $d \times r$ matrix whose columns are these basis vectors. Prove the following.

(i) (1.4.1): $\pi_L(\boldsymbol{x}) := (M(M^t M)^{-1} M^t)\boldsymbol{x}$.
(ii) (1.4.2): if B is an orthogonal basis of L, then $\pi_L(\boldsymbol{x}) := \sum_{i=1}^{r} \frac{\boldsymbol{x} \cdot \boldsymbol{l}_i}{\|\boldsymbol{l}_i\|^2} \boldsymbol{l}_i$.
(iii) The orthogonal projection is a linear transformation.
(iv) $M^t M$ is nonsingular.

https://doi.org/10.1017/9781009257794.002 Published online by Cambridge University Press

1.12.10 Let $L := \left\{ \boldsymbol{x} \in \mathbb{R}^d \middle| \, N\boldsymbol{x} = \boldsymbol{0} \right\}$ be a linear subspace of $\mathbb{R}^d$ and let $\boldsymbol{x}$ be any vector of $\mathbb{R}^d$. Prove (1.4.3): $\pi_L(\boldsymbol{x}) = \boldsymbol{x} - (N^t(NN^t)^{-1})(N\boldsymbol{x})$.

1.12.11 Let A be an affine subspace of $\mathbb{R}^d$ and let $\boldsymbol{a}_0 \in A$. Suppose that $B := \{\boldsymbol{a}_0, \boldsymbol{a}_1, \ldots, \boldsymbol{a}_r\}$ is a basis of A of $\mathbb{R}^d$ and let $\boldsymbol{x}$ be any vector of $\mathbb{R}^d$. Prove (1.4.6): write $M := (\boldsymbol{a}_1 - \boldsymbol{a}_0 \cdots \boldsymbol{a}_r - \boldsymbol{a}_0)$; then

$$\pi_A(\boldsymbol{x}) = \boldsymbol{a}_0 + \left(M \left(M^t M \right)^{-1} M^t \right) (\boldsymbol{x} - \boldsymbol{a}_0).$$

1.12.12 Let $A := \left\{ \boldsymbol{x} \in \mathbb{R}^d \middle| \, N\boldsymbol{x} = N\boldsymbol{a}_0, \, \boldsymbol{a}_0 \in A \right\}$ be an affine subspace of $\mathbb{R}^d$ and let $\boldsymbol{x}$ be any vector of $\mathbb{R}^d$. Prove (1.4.7), namely

$$\pi_A(\boldsymbol{x}) = \boldsymbol{x} - (N^t(NN^t)^{-1})(N\boldsymbol{x} - N\boldsymbol{a}_0).$$

1.12.13 Find convexity-preserving functions that are not affine.

1.12.14 Let $X, Y \subseteq \mathbb{R}^d$ be sets. Prove the following.

(i) $\operatorname{int} X \subseteq \operatorname{rint} X \subseteq X$.
(ii) If $X \subseteq Y$ and $\operatorname{aff} X = \operatorname{aff} Y$, then $\operatorname{rint} X \subseteq \operatorname{rint} Y$.
(iii) $\operatorname{rint}(\operatorname{rint} X) = \operatorname{rint} X$.
(iv) If $\operatorname{rint} X \subseteq Y \subseteq X$ then $\operatorname{rint} X = \operatorname{rint} Y$.
(v) $\operatorname{rint}(X + \boldsymbol{x}) = \operatorname{rint} X + \boldsymbol{x}$ for every $\boldsymbol{x} \in \mathbb{R}^d$.

1.12.15 Let $X \subseteq \mathbb{R}^d$ be a convex set. Prove that $\operatorname{int} X = \operatorname{rint} X$ if and only if $\operatorname{int} X \neq \varnothing$.

1.12.16 (Supporting function) Let $X \subseteq \mathbb{R}^d$ be a nonempty convex set. For every $\boldsymbol{y} \in \mathbb{R}^d$, define the *supporting function* h of X as

$$h(X, \boldsymbol{y}) := \sup \{\boldsymbol{x} \cdot \boldsymbol{y} | \, \boldsymbol{x} \in X\}. \qquad (1.12.0.1)$$

It follows that, if $h(X, \boldsymbol{a}) < \infty$ and $\boldsymbol{a} \neq \boldsymbol{0}_d$, then the hyperplane

$$\left\{ \boldsymbol{z} \in \mathbb{R}^d \middle| \, \boldsymbol{a} \cdot \boldsymbol{z} = h(X, \boldsymbol{a}) \right\}$$

is a supporting hyperplane of X with normal vector $\boldsymbol{a}$.

Prove that, for any nonempty convex sets $X \subseteq \mathbb{R}^d$ and $Y \subseteq \mathbb{R}^d$, the function h satisfies the following.

(i) $X = \{\boldsymbol{x} \in \mathbb{R}^d \mid \boldsymbol{x} \cdot \boldsymbol{a} \leqslant h(X, \boldsymbol{a})$ for all $\boldsymbol{a} \in \mathbb{R}^d\}$.
(ii) $h(X, \alpha\boldsymbol{x}) = \alpha h(X, \boldsymbol{x})$ for all $\alpha \geqslant 0$ and all $\boldsymbol{x} \in \mathbb{R}^d$.
(iii) $h(\alpha X, \boldsymbol{x}) = \alpha h(X, \boldsymbol{x})$ for all $\alpha \geqslant 0$ and all $\boldsymbol{x} \in \mathbb{R}^d$.
(iv) $h(X, \boldsymbol{x} + \boldsymbol{y}) \leqslant h(X, \boldsymbol{x}) + h(X, \boldsymbol{y})$ for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^d$.
(v) $h(X, \boldsymbol{x}) \leqslant h(Y, \boldsymbol{x})$ if and only if $X \subseteq Y$ for all $\boldsymbol{x} \in \mathbb{R}^d$.

(vi) $h(X+Y,\boldsymbol{x}) = h(X,\boldsymbol{x}) + h(Y,\boldsymbol{x})$.
(vii) If, in addition, X and Y are closed sets that satisfy $h(X,\boldsymbol{x}) = h(Y,\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathbb{R}^d$, then $X = Y$.

1.12.17 Let X be a nonempty closed convex set set in $\mathbb{R}^d$. Prove the following.

(i) lineal X is a linear subspace of $\mathbb{R}^d$.
(ii) rec X is a convex cone.
(iii) If X is a convex cone, then rec $X = X$.
(iv) lineal $\left(X \cap (\text{lineal } X)^{\perp}\right) = \{\mathbf{0}\}$.

1.13 Postscript

The information related to linear subspaces and linear maps (Sections 1.1 and 1.4) can be found in most linear algebra books; for instance, in Shifrin and Adams (2011). For the material on affine subspaces, while fairly standard, one may need look outside linear algebra books; for instance in Webster (1994, chap. 1) or Lauritzen (2013, chap. 2). The presentation in Section 1.2 on the embedding of affine spaces into linear spaces follows that of Berger (2009, ch. 3) and Gallier (2011, ch. 4). The material on projective spaces and projective maps can be found in Berger (2009, ch. 4), Gallier (2011, ch. 5), and Reid and Szendroi (2005, ch. 5); if there is a need for a concise review of results in projective geometry that we did not cover, albeit with no proofs, we recommend Fortuna et al. (2016, ch. 1). The section on dual spaces is based on Halmos (1974, secs. 13–20,67–69).

The presentation of basic convexity in Sections 1.6 to 1.11 is standard and can be found elsewhere, for instance, in Webster (1994); Soltan (2015); Lauritzen (2013); Brøndsted (1983).

Carathéodory's theorem (1.6.3), Radon's theorem (1.6.5), and a theorem of Helly (1923) all appeared in the first half of the 20th century (Carathéodory, 1907; Radon, 1921). Since then they have sparked a great deal of interest in intersection and covering patterns of convex sets. Equally influential was the second wave of such theorems that appeared in the second half of the 20th century, including the colourful Carathéodory theorem (Bárány, 1982) and Tveberg's theorem (Tverberg, 1966). All these results are covered with care in Bárány (2021).

Radon's theorem (1.6.5) can be restated in terms of affine maps: for every $n \geqslant d+1$ and every affine map φ from the $(n-1)$-simplex to $\mathbb{R}^d$, there exists a pair of disjoint faces of the simplex whose φ-images intersect. The equivalence between the statement in Theorem 1.6.5 and this affine formulation stems from noticing that every set $X \subseteq \mathbb{R}^d$ with cardinality n determines a unique affine map φ that takes the extreme points of the $(n-1)$-simplex to the elements of X and, in this way, each face of the simplex is mapped to the convex hull of the images of its extreme points. A topological version of Radon's theorem, due to Bajmóczy and Bárány (1979), replaces the adjective 'affine' with the adjective 'continuous', and thus relaxes the condition on φ. These topics and their topological versions are also presented in Bárány (2021).

The main concepts related to separation and support originated in Minkowski (1896), including the separation theorem (1.8.5). The converse of Theorem 1.8.1 is also true: if $X \subseteq \mathbb{R}^d$ is a nonempty set in which, to each point $\boldsymbol{a}$ in $\mathbb{R}^d$, there is a unique point in X closest to $\boldsymbol{a}$, then X must be closed and convex; this was shown independently by Bunt (1934) and Motzkin (1935).

A more general version of Minkowski–Krein–Milman's theorem (1.9.11) appeared in Krein and Milman (1940). Theorem 1.9.11 is often called Minkowski's theorem because Minkowski (1911, pp. 131–229) proved the finite-dimensional version that we presented.