Convexity

Convexity and Optimization – Part I

Lars-Åke Lindahl

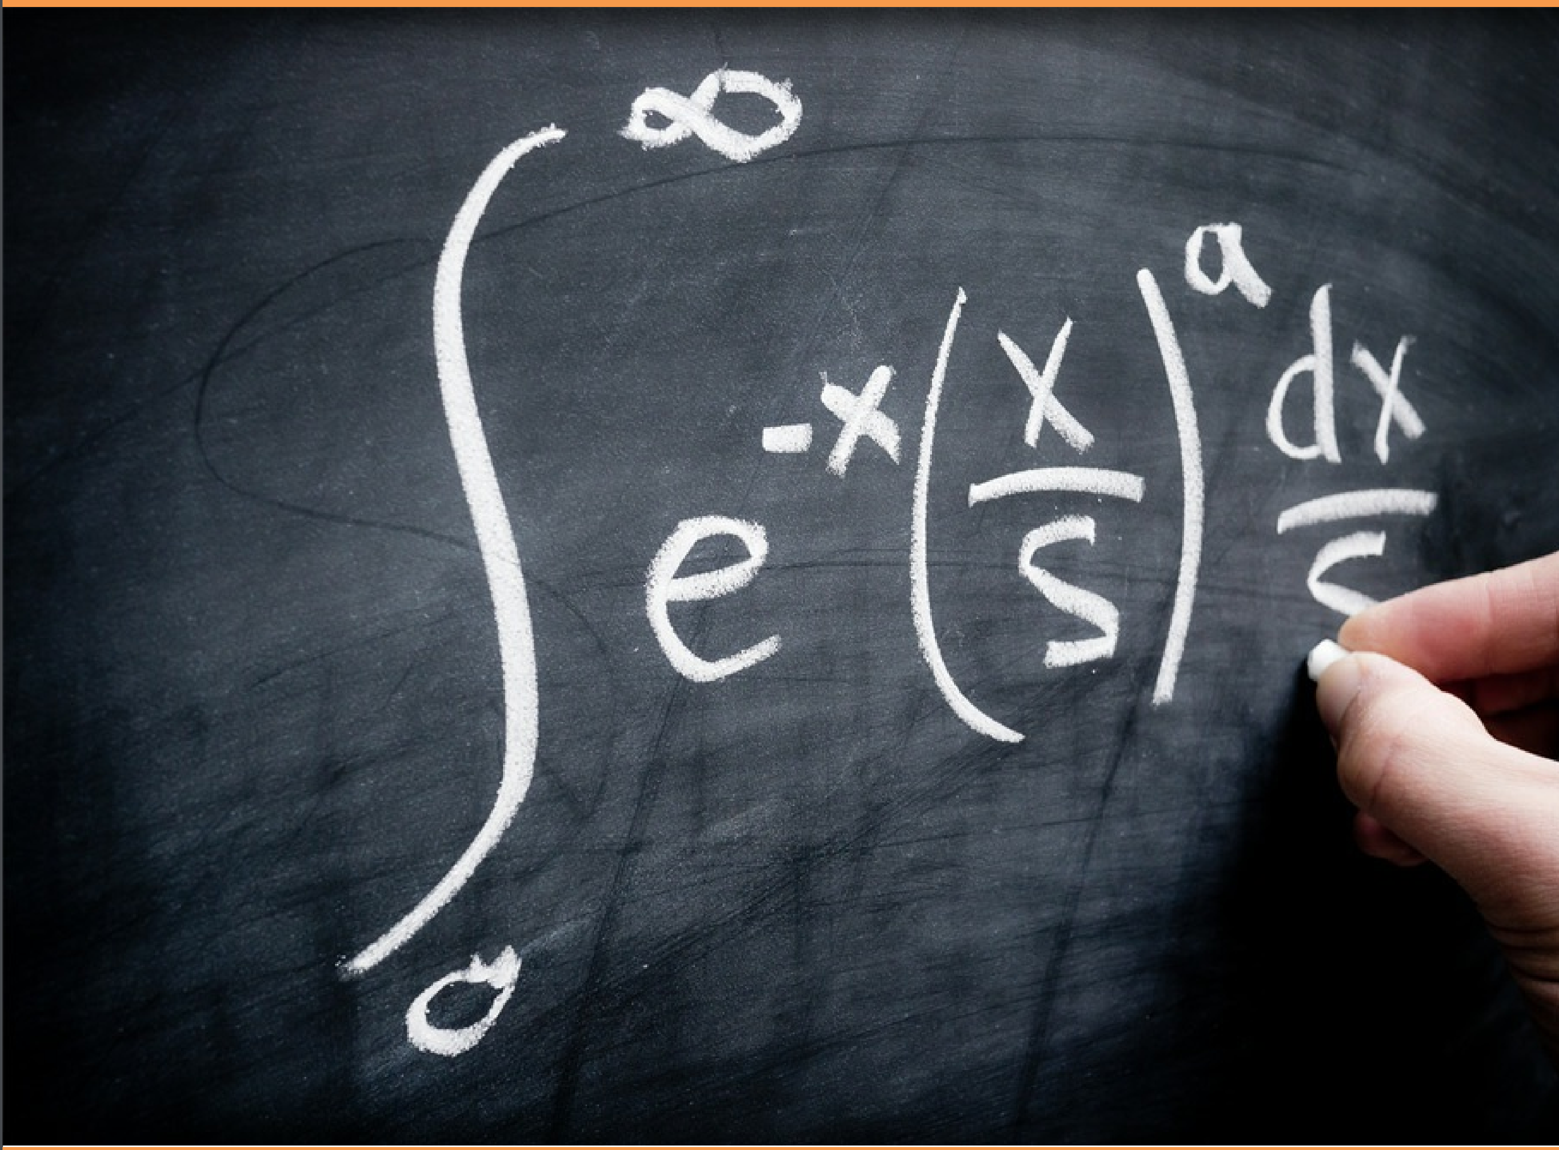

LARS-ÅKE LINDAHL

CONVEXITY

CONVEXITY AND OPTIMIZATION – PART I

Convexity: Convexity and Optimization – Part I
1st edition

ISBN 978-87-403-1382-6

CONTENTS

Preface

Mathematical optimization methods are today used routinely as a tool for economic and industrial planning, in production control and product design, in civil and military logistics, in medical image analysis, etc., and the development in the field of optimization has been tremendous since World War II. In 1945, George Stigler studied a diet problem with 77 foods and 9 constraints without being able to determine the optimal diet − today it is possible to solve optimization problems containing hundreds of thousands of variables and constraints. There are two factors that have made this possible − computers and efficient algorithms. It is the rapid development in the computer area that has been most visible to the common man, but the algorithm development has also been tremendous during the past 70 years, and computers would be of little use without efficient algorithms.

Maximization and minimization problems have of course been studied and solved since the beginning of the mathematical analysis, but optimization theory in the modern sense started around 1948 with George Dantzig, who introduced and popularized the concept of linear programming and proposed an efficient solution algorithm, the simplex algorithm, for such problems.

The type of optimization problems to be discussed by us are problems that can be formulated as the problem to maximize (or minimize) a given function over a somehow given subset of $\mathbf{R}^n$. In order to obtain general results of interest we need to make some assumptions about the function and the set, and it is here that convexity enters into the picture. The first part in this series of three on convexity and optimization therefore deals with finite dimensional convexity theory. Since convexity plays an important role in many areas of mathematics, significantly more about convexity is included than is used in the subsequent two parts on optimization, where Part II provides the basic classical theory for linear and convex optimization, and Part III describes Newton's algorithm, self-concordant functions and an interior point method with self-concordant barriers.

Parts II and III present a number of algorithms, but the emphasis is always on the mathematical theory, so we do not describe how the algorithms should be implemented numerically. Anyone who is interested in these important aspects should consult specialized literature in the field.

The embryo of this book is a compendium written by Christer Borell and myself 1978–79, but various additions, deletions and revisions over the years, have led to a completely different text, the most significant addition being Part III.

The presentation in this book is complete in the sense that all theorems are proved. Some of the proofs are quite technical, but none of them requires more previous knowledge than a good knowledge of linear algebra and calculus of several variables.

Uppsala, April 2016
Lars-Åke Lindahl

List of symbols

$\operatorname{aff} X$	affine hull of X, p. 22
$\operatorname{bdry} X$	boundary of X, p. 11
$\operatorname{cl} f$	closure of the function f, p. 172
$\operatorname{cl} X$	closure of X, p. 12
$\operatorname{con} X$	conic hull of X, p. 43
$\operatorname{cvx} X$	convex hull of X, p. 34
$\dim X$	dimension of X, p. 24
$\operatorname{dom} f$	the effective domain of f: $\{x \mid -\infty < f(x) < \infty\}$, p. 3
$\operatorname{epi} f$	epigraph of f, p. 104
$\operatorname{exr} X$	set of extreme rays of X, p. 79
$\operatorname{ext} X$	set of extreme points of X, p. 77
$\operatorname{int} X$	interior of X, p. 11
$\operatorname{lin} X$	recessive subspace of X, p. 51
$\operatorname{rbdry} X$	relative boundary of X, p. 37
$\operatorname{recc} X$	recession cone of X, p. 47
$\operatorname{rint} X$	relative interior of X, p. 37
$\operatorname{sublev}_\alpha f$	α-sublevel set of f, p. 104
$\mathbf{e}_i$	ith standard basis vector $(0,\dots,1,\dots,0)$, p. 4
f'	derivate or gradient of f, p. 17
$f'(x;v)$	direction derivate of f at x in direction v, p. 180
f''	second derivative or hessian of f, p. 19
f^*	conjugate function of f, p. 173
$B(a;r)$	open ball centered at a with radius r, p. 10
$\overline{B}(a;r)$	closed ball centered at a with radius r, p. 10
$Df(a)[v]$	differential of f at a, p. 17
$D^2f(a)[u,v]$	$\sum_{i,j=1}^n \frac{\partial^2 f}{\partial x_i \partial x_j}(a)u_iv_j$, p. 19
$D^3f(a)[u,v,w]$	$\sum_{i,j,k=1}^n \frac{\partial^3 f}{\partial x_i \partial x_j \partial x_k}(a)u_iv_jw_k$, p. 20
$\mathbf{R}_+$, $\mathbf{R}_{++}$	$\{x \in \mathbf{R} \mid x \geq 0\}$, $\{x \in \mathbf{R} \mid x > 0\}$, p. 1
$\mathbf{R}_-$	$\{x \in \mathbf{R} \mid x \leq 0\}$, p. 1
$\overline{\mathbf{R}}$, $\underline{\mathbf{R}}$, $\overline{\underline{\mathbf{R}}}$	$\mathbf{R} \cup \{\infty\}$, $\mathbf{R} \cup \{-\infty\}$, $\mathbf{R} \cup \{\infty, -\infty\}$, p. 1
S_X	support function of X, p. 137

$S_{\mu,L}(X)$	class of μ-strongly convex functions on X with L-Lipschitz continuous derivative, p. 157
X^+	dual cone of X, p. 65
$\mathbf{1}$	the vector $(1, 1, \dots, 1)$, p. 4
$\partial f(a)$	subdifferential of f at a, p. 163
ϕ_X	Minkowski functional of X, p. 140
∇f	gradient of f, p. 17
$\overrightarrow{x}$	ray from 0 through x, p. 40
$[x, y]$	line segment between x and y, p. 7
$]x, y[$	open line segment between x and y, p. 7
$\|\cdot\|_1$, $\|\cdot\|_2$, $\|\cdot\|_\infty$	ℓ^1-norm, Euclidean norm, maximum norm, p. 10

Chapter 1

Preliminaries

The purpose of this chapter is twofold – to explain certain notations and terminologies used throughout the book and to recall some fundamental concepts and results from calculus and linear algebra.

Real numbers

We use the standard notation $\mathbf{R}$ for the set of real numbers, and we let

$$\begin{aligned}\mathbf{R}_+ &= \{x \in \mathbf{R} \mid x \ge 0\},\\ \mathbf{R}_- &= \{x \in \mathbf{R} \mid x \le 0\},\\ \mathbf{R}_{++} &= \{x \in \mathbf{R} \mid x > 0\}.\end{aligned}$$

In other words, $\mathbf{R}_+$ consists of all nonnegative real numbers, and $\mathbf{R}_{++}$ denotes the set of all positive real numbers.

The extended real line

Each nonempty set A of real numbers that is bounded above has a least upper bound, denoted by $\sup A$, and each nonempty set A that is bounded below has a greatest lower bound, denoted by $\inf A$. In order to have these two objects defined for arbitrary subsets of $\mathbf{R}$ (and also for other reasons) we extend the set of real numbers with the two symbols $-\infty$ and ∞ and introduce the notation

$$\overline{\mathbf{R}} = \mathbf{R} \cup \{\infty\}, \quad \underline{\mathbf{R}} = \mathbf{R} \cup \{-\infty\} \quad \text{and} \quad \overline{\underline{\mathbf{R}}} = \mathbf{R} \cup \{-\infty, \infty\}.$$

We furthermore extend the order relation $<$ on $\mathbf{R}$ to the extended real line $\overline{\underline{\mathbf{R}}}$ by defining, for each real number x,

$$-\infty < x < \infty.$$

The arithmetic operations on $\mathbf{R}$ are partially extended by the following "natural" definitions, where x denotes an arbitrary real number:

$$
\begin{aligned}
&x+\infty=\infty+x=\infty+\infty=\infty\\
&x+(-\infty)=-\infty+x=-\infty+(-\infty)=-\infty\\
&x\cdot\infty=\infty\cdot x=\begin{cases}\infty & \text{if } x>0\\ 0 & \text{if } x=0\\ -\infty & \text{if } x<0\end{cases}\\
&x\cdot(-\infty)=-\infty\cdot x=\begin{cases}-\infty & \text{if } x>0\\ 0 & \text{if } x=0\\ \infty & \text{if } x<0\end{cases}\\
&\infty\cdot\infty=(-\infty)\cdot(-\infty)=\infty\\
&\infty\cdot(-\infty)=(-\infty)\cdot\infty=-\infty.
\end{aligned}
$$

It is now possible to define in a consistent way the least upper bound and the greatest lower bound of an arbitrary subset of the extended real line. For nonempty sets A which are not bounded above by any real number, we define $\sup A=\infty$, and for nonempty sets A which are not bounded below by any real number we define $\inf A=-\infty$. Finally, for the empty set $\emptyset$ we define $\inf\emptyset=\infty$ and $\sup\emptyset=-\infty$.

Sets and functions

We use standard notation for sets and set operations that are certainly well known to all readers, but the intersection and the union of an arbitrary family of sets may be new concepts for some readers.

So let $\{X_i \mid i\in I\}$ be an arbitrary family of sets X_i, indexed by the set I; their *intersection*, denoted by

$$\bigcap\{X_i \mid i\in I\} \quad\text{or}\quad \bigcap_{i\in I} X_i,$$

is by definition the set of elements that belong to all the sets X_i. The *union*

$$\bigcup\{X_i \mid i\in I\} \quad\text{or}\quad \bigcup_{i\in I} X_i$$

consists of the elements that belong to X_i for at least one $i\in I$.

We write $f\colon X\to Y$ to indicate that the function f is defined on the set X and takes its values in the set Y. The set X is then called the *domain* of the function and Y is called the *codomain*. Most functions in this book have domain equal to $\mathbf{R}^n$ or to some subset of $\mathbf{R}^n$, and their codomain is usually $\mathbf{R}$ or more generally $\mathbf{R}^m$ for some integer $m \geq 1$, but sometimes we also consider functions whose codomain equals $\overline{\mathbf{R}}$, $\underline{\mathbf{R}}$ or $\underline{\overline{\mathbf{R}}}$.

Let A be a subset of the domain X of the function f. The set

$$f(A) = \{f(x) \mid x \in A\}$$

is called the *image of* A under the function f. If B is a subset of the codomain of f, then

$$f^{-1}(B) = \{x \in X \mid f(x) \in B\}$$

is called the *inverse image of* B under f. There is no implication in the notation $f^{-1}(B)$ that the inverse f^{-1} exists.

For functions $f\colon X \to \underline{\overline{\mathbf{R}}}$ we use the notation $\operatorname{dom} f$ for the inverse image of $\mathbf{R}$, i.e.

$$\operatorname{dom} f = \{x \in X \mid -\infty < f(x) < \infty\}.$$

The set $\operatorname{dom} f$ thus consists of all $x \in X$ with finite function values $f(x)$, and it is called the *effective domain* of f.

The vector space $\mathbf{R}^n$

The reader is assumed to have a solid knowledge of elementary linear algebra and thus, in particular, to be familiar with basic vector space concepts such as linear subspace, linear independence, basis and dimension.

As usual, $\mathbf{R}^n$ denotes the vector space of all n-tuples $(x_1, x_2, \dots, x_n)$ of real numbers. The elements of $\mathbf{R}^n$, interchangeably called points and vectors, are denoted by lowercase letters from the beginning or the end of the alphabet, and if the letters are not numerous enough, we provide them with sub- or superindices. Subindices are also used to specify the coordinates of a vector, but there is no risk of confusion, because it will always be clear from the context whether for instance x_1 is a vector of its own or the first coordinate of the vector x.

Vectors in $\mathbf{R}^n$ will interchangeably be identified with *column matrices*. Thus, to us

$$(x_1, x_2, \dots, x_n) \qquad \text{and} \qquad \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

denote the same object.

The vectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ in $\mathbf{R}^n$, defined as

$$\mathbf{e}_1 = (1, 0, \dots, 0), \quad \mathbf{e}_2 = (0, 1, 0, \dots, 0), \quad \dots, \quad \mathbf{e}_n = (0, 0, \dots, 0, 1),$$

are called the natural basis vectors in $\mathbf{R}^n$, and $\mathbf{1}$ denotes the vector whose coordinates are all equal to one, so that

$$\mathbf{1} = (1, 1, \dots, 1).$$

The *standard scalar product* $\langle \cdot\,, \cdot \rangle$ on $\mathbf{R}^n$ is defined by the formula

$$\langle x, y \rangle = x_1 y_1 + x_2 y_2 + \dots + x_n y_n,$$

and, using matrix multiplication, we can write this as

$$\langle x, y \rangle = x^T y = y^T x,$$

where ^{T} denotes transposition. In general, A^T denotes the transpose of the matrix A.

The solution set to a homogeneous system of linear equations in n unknowns is a linear subspace of $\mathbf{R}^n$. Conversely, every linear subspace of $\mathbf{R}^n$

can be presented as the solution set to some homogeneous system of linear equations:

$$\begin{cases} a_{11}x_1 + \ a_{12}x_2 + \cdots + \ a_{1n}x_n = 0 \\ a_{21}x_1 + \ a_{22}x_2 + \cdots + \ a_{2n}x_n = 0 \\ \qquad\qquad\qquad\qquad\qquad\quad \vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n = 0 \end{cases}$$

Using matrices we can of course write the system above in a more compact form as

$$Ax = 0,$$

where the matrix A is called the *coefficient matrix* of the system.

The dimension of the solution set of the above system is given by the number $n - r$, where r equals the rank of the matrix A. Thus in particular, for each linear subspace X of $\mathbf{R}^n$ of dimension $n - 1$ there exists a nonzero vector $c = (c_1, c_2, \dots, c_n)$ such that

$$X = \{x \in \mathbf{R}^n \mid c_1x_1 + c_2x_2 + \cdots + c_nx_n = 0\}.$$

Sum of sets

If X and Y are nonempty subsets of $\mathbf{R}^n$ and α is a real number, we let

$$\begin{aligned} X + Y &= \{x + y \mid x \in X,\ y \in Y\}, \\ X - Y &= \{x - y \mid x \in X,\ y \in Y\}, \\ \alpha X &= \{\alpha x \mid x \in X\}. \end{aligned}$$

The set $X + Y$ is called the *(vector) sum* of X and Y, $X - Y$ is the *(vector) difference* and αX is the product of the number α and the set X.

It is convenient to have sums, differences and products defined for the empty set $\emptyset$, too. Therefore, we extend the above definitions by defining

$$X \pm \emptyset = \emptyset \pm X = \emptyset$$

for all sets X, and

$$\alpha\emptyset = \emptyset.$$

For singleton sets $\{a\}$ we write $a + X$ instead of $\{a\} + X$, and the set $a + X$ is called a *translation* of X.

It is now easy to verify that the following rules hold for arbitrary sets X,

Y and Z and arbitrary real numbers α and β:

$$\begin{aligned} X + Y &= Y + X \\ (X + Y) + Z &= X + (Y + Z) \\ \alpha X + \alpha Y &= \alpha(X + Y) \\ (\alpha + \beta)X &\subseteq \alpha X + \beta X\,. \end{aligned}$$

In connection with the last inclusion one should note that the converse inclusion $\alpha X + \beta X \subseteq (\alpha + \beta)X$ does **not** hold for general sets X.

Inequalites in $\mathbf{R}^n$

For vectors $x = (x_1, x_2, \dots, x_n)$ and $y = (y_1, y_2, \dots, y_n)$ in $\mathbf{R}^n$ we write $x \geq y$ if $x_j \geq y_j$ for all indices j, and we write $x > y$ if $x_j > y_j$ for all j. In particular, $x \geq 0$ means that all coordinates of x are nonnegative.

The set

$$\mathbf{R}^n_+ = \mathbf{R}_+ \times \mathbf{R}_+ \times \cdots \times \mathbf{R}_+ = \{x \in \mathbf{R}^n \mid x \geq 0\}$$

is called the *nonnegative orthant* of $\mathbf{R}^n$.

The order relation $\geq$ is a partial order on $\mathbf{R}^n$. It is thus, in other words, reflexive ($x \geq x$ for all x), transitive ($x \geq y$ & $y \geq z \Rightarrow x \geq z$) and antisymmetric ($x \geq y$ & $y \geq x \Rightarrow x = y$). However, the order is not a complete order when $n > 1$, since two vectors x and y may be unrelated.

Two important properties, which will be used now and then, are given by the following two trivial implications:

$$x \geq 0 \;\&\; y \geq 0 \;\Rightarrow\; \langle x, y\rangle \geq 0$$
$$x \geq 0 \;\&\; y \geq 0 \;\&\; \langle x, y\rangle = 0 \;\Rightarrow\; x = y = 0.$$

Line segments

Let x and y be points in $\mathbf{R}^n$. We define

$$[x, y] = \{(1-\lambda)x + \lambda y \mid 0 \leq \lambda \leq 1\}$$

and

$$]x, y[= \{(1-\lambda)x + \lambda y \mid 0 < \lambda < 1\},$$

and we call the set $[x, y]$ the *line segment* and the set $]x, y[$ the *open line segment* between x and y, if the two points are distinct. If the two points coincide, i.e. if $y = x$, then obviously $[x, x] =]x, x[= \{x\}$.

Linear maps and linear forms

Let us recall that a map $S\colon \mathbf{R}^n \to \mathbf{R}^m$ is called *linear* if

$$S(\alpha x + \beta y) = \alpha Sx + \beta Sy$$

for all vectors x, $y \in \mathbf{R}^n$ and all scalars (i.e. real numbers) α, β. A linear map $S\colon \mathbf{R}^n \to \mathbf{R}^n$ is also called a *linear operator* on $\mathbf{R}^n$.

Each linear map $S\colon \mathbf{R}^n \to \mathbf{R}^m$ gives rise to a unique $m \times n$-matrix $\tilde{S}$ such that

$$Sx = \tilde{S}x,$$

which means that the function value Sx of the map S at x is given by the matrixproduct $\tilde{S}x$. (Remember that vectors are identified with column matrices!) For this reason, the same letter will be used to denote a map and its matrix. We thus interchangeably consider Sx as the value of a map and as a matrix product.

By computing the scalar product $\langle x, Sy\rangle$ as a matrix product we obtain the following relation

$$\langle x, Sy\rangle = x^T Sy = (S^T x)^T y = \langle S^T x, y\rangle$$

between a linear map $S\colon \mathbf{R}^n \to \mathbf{R}^m$ (or $m \times n$-matrix S) and its *transposed map* $S^T\colon \mathbf{R}^m \to \mathbf{R}^n$ (or transposed matrix S^T).

An $n \times n$-matrix $A = [a_{ij}]$, and the corresponding linear map, is called *symmetric* if $A^T = A$, i.e. if $a_{ij} = a_{ji}$ for all indices i, j.

A linear map $f\colon \mathbf{R}^n \to \mathbf{R}$ with codomain $\mathbf{R}$ is called a *linear form.* A linear form on $\mathbf{R}^n$ is thus of the form

$$f(x) = c_1x_1 + c_2x_2 + \cdots + c_nx_n,$$

where $c = (c_1, c_2, \ldots, c_n)$ is a vector in $\mathbf{R}^n$. Using the standard scalar product we can write this more simply as

$$f(x) = \langle c, x\rangle,$$

and in matrix notation this becomes

$$f(x) = c^Tx.$$

Let $f(x) = \langle c, y\rangle$ be a linear form on $\mathbf{R}^m$ and let $S\colon \mathbf{R}^n \to \mathbf{R}^m$ be a linear map with codomain $\mathbf{R}^m$. The composition $f \circ S$ is then a linear form on $\mathbf{R}^n$, and we conclude that there exists a unique vector $d \in \mathbf{R}^n$ such that $(f \circ S)(x) = \langle d, x\rangle$ for all $x \in \mathbf{R}^n$. Since $f(Sx) = \langle c, Sx\rangle = \langle S^Tc, x\rangle$, it follows that $d = S^Tc$.

Quadratic forms

A function $q\colon \mathbf{R}^n \to \mathbf{R}$ is called a *quadratic form* if there exists a symmetric $n \times n$-matrix $Q = [q_{ij}]$ such that

$$q(x) = \sum_{i,j=1}^{n} q_{ij}x_ix_j,$$

or equivalently

$$q(x) = \langle x, Qx\rangle = x^TQx.$$

The quadratic form q determines the symmetric matrix Q uniquely, and this allows us to identify the form q with its matrix (or operator) Q.

An arbitrary quadratic polynomial $p(x)$ in n variables can now be written in the form

$$p(x) = \langle x, Ax\rangle + \langle b, x\rangle + c,$$

where $x \mapsto \langle x, Ax\rangle$ is a quadratic form determined by a symmetric operator (or matrix) A, $x \mapsto \langle b, x\rangle$ is a linear form determined by a vector b, and c is a real number.

Example. In order to write the quadratic polynomial

$$p(x_1, x_2, x_3) = x_1^2 + 4x_1x_2 - 2x_1x_3 + 5x_2^2 + 6x_2x_3 + 3x_1 + 2x_3 + 2$$

in this form we first replace the terms dx_ix_j for $i < j$ with $\frac{1}{2}dx_ix_j + \frac{1}{2}dx_jx_i$. This yields

$$\begin{aligned} p(x_1, x_2, x_3) &= (x_1^2 + 2x_1x_2 - x_1x_3 + 2x_2x_1 + 5x_2^2 + 3x_2x_3 - x_3x_1 + 3x_3x_2) \\ &\quad + (3x_1 + 2x_3) + 2 = \langle x, Ax\rangle + \langle b, x\rangle + c \end{aligned}$$

with $A = \begin{bmatrix} 1 & 2 & -1 \\ 2 & 5 & 3 \\ -1 & 3 & 0 \end{bmatrix}$, $b = \begin{bmatrix} 3 \\ 0 \\ 2 \end{bmatrix}$ and $c = 2$. □

A quadratic form q on $\mathbf{R}^n$ (and the corresponding symmetric operator and matrix) is called *positive semidefinite* if $q(x) \geq 0$ and *positive definite* if $q(x) > 0$ for all vectors $x \neq 0$ in $\mathbf{R}^n$.

Norms and balls

A *norm* $\|\cdot\|$ on $\mathbf{R}^n$ is a function $\mathbf{R}^n \to \mathbf{R}_+$ that satisfies the following three conditions:

(i)	$\|x+y\| \le \|x\| + \|y\|$	for all x, y
(ii)	$\|\lambda x\| = \|\lambda\|\,\|x\|$	for all $x \in \mathbf{R}^n$, $\lambda \in \mathbf{R}$
(iii)	$\|x\| = 0 \Leftrightarrow x = 0.$	

The most important norm to us is the *Euclidean norm*, defined via the standard scalar product as

$$\|x\| = \sqrt{\langle x, x\rangle} = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}.$$

This is the norm that we use unless the contrary is stated explicitely. We use the notation $\|\cdot\|_2$ for the Euclidean norm whenever we for some reason have to emphasize that the norm in question is the Euclidean one.

Other norms, that will occur now and then, are the *maximum norm*

$$\|x\|_\infty = \max_{1\le i\le n} |x_i|,$$

and the ℓ^1-*norm*

$$\|x\|_1 = \sum_{i=1}^{n} |x_i|.$$

It is easily verified that these really are norms, that is that conditions (i)–(iii) are satisfied.

All norms on $\mathbf{R}^n$ are equivalent in the following sense: If $\|\cdot\|$ and $\|\cdot\|'$ are two norms, then there exist two positive constants c and C such that

$$c\|x\|' \le \|x\| \le C\|x\|'$$

for all $x \in \mathbf{R}^n$.

For example, $\|x\|_\infty \le \|x\|_2 \le \sqrt{n}\,\|x\|_\infty$.

Given an arbitrary norm $\|\cdot\|$ we define the corresponding *distance* between two points x and a in $\mathbf{R}^n$ as $\|x - a\|$. The set

$$B(a; r) = \{x \in \mathbf{R}^n \mid \|x - a\| < r\},$$

consisting of all points x whose distance to a is less than r, is called the *open ball* centered at the point a and with radius r. Of course, we have to have $r > 0$ in order to get a nonempty ball. The set

$$\overline{B}(a; r) = \{x \in \mathbf{R}^n \mid \|x - a\| \le r\}$$

is the corresponding *closed ball*.

The geometric shape of the balls depends on the underlying norm. The ball $\overline{B}(0;1)$ in $\mathbf{R}^2$ is a square with corners at the points $(\pm1,\pm1)$ when the norm is the maximum norm, it is a square with corners at the points $(\pm1,0)$ and $(0,\pm1)$ when the norm is the ℓ^1-norm, and it is the unit disc when the norm is the Euclidean one.

If B denotes balls defined by one norm and B' denotes balls defined by a second norm, then there are positive constants c and C such that

$$B'(a;cr) \subseteq B(a;r) \subseteq B'(a;Cr) \tag{1.1}$$

for all $a \in \mathbf{R}^n$ and all $r > 0$. This follows easily from the equivalence of the two norms.

All balls that occur in the sequel are assumed to be Euclidean, i.e. defined with respect to the Euclidean norm, unless otherwise stated.

Topological concepts

We now use balls to define a number of topological concepts. Let X be an arbitrary subset of $\mathbf{R}^n$. A point $a \in \mathbf{R}^n$ is called

- an *interior point* of X if there exists an $r > 0$ such that $B(a;r) \subseteq X$;
- a *boundary point* of X if $X \cap B(a;r) \neq \emptyset$ and $\complement X \cap B(a;r) \neq \emptyset$ for all $r > 0$;
- an *exterior point* of X if there exists an $r > 0$ such that $X \cap B(a;r) = \emptyset$.

Observe that because of property (1.1), the above concepts do not depend on the kind of balls that we use.

A point is obviously either an interior point, a boundary point or an exterior point of X. Interior points belong to X, exterior points belong to the complement of X, while boundary points may belong to X but must not do so. Exterior points of X are interior points of the complement $\complement X$, and vice versa, and the two sets X and $\complement X$ have the same boundary points.

The set of all interior points of X is called the *interior* of X and is denoted by $\operatorname{int} X$. The set of all boundary points is called the *boundary* of X and is denoted by $\operatorname{bdry} X$.

A set X is called *open* if all points in X are interior points, i.e. if $\operatorname{int} X = X$.

It is easy to verify that the union of an arbitrary family of open sets is an open set and that the intersection of finitely many open sets is an open set. The empty set $\emptyset$ and $\mathbf{R}^n$ are open sets.

The interior $\operatorname{int} X$ is a (possibly empty) open set for each set X, and $\operatorname{int} X$ is the biggest open set that is included in X.

A set X is called *closed* if its complement $\complement X$ is an open set. It follows that X is closed if and only if X contains all its boundary points, i.e. if and only if $\operatorname{bdry} X \subseteq X$.

The intersection of an arbitrary family of closed sets is closed, the union of finitely many closed sets is closed, and $\mathbf{R}^n$ and $\emptyset$ are closed sets.

For arbitrary sets X we set

$$\operatorname{cl} X = X \cup \operatorname{bdry} X.$$

The set $\operatorname{cl} X$ is then a closed set that contains X, and it is called the *closure* (or *closed hull*) of X. The closure $\operatorname{cl} X$ is the smallest closed set that contains X as a subset.

For example, if $r > 0$ then

$$\operatorname{cl} B(a;r) = \{x \in \mathbf{R}^n \mid \|x - a\| \leq r\} = \overline{B}(a;r),$$

which makes it consistent to call the set $\overline{B}(a;r)$ a closed ball.

For nonempty subsets X of $\mathbf{R}^n$ and numbers $r > 0$ we define

$$X(r) = \{y \in \mathbf{R}^n \mid \exists x \in X \colon \|y - x\| < r\}.$$

The set $X(r)$ thus consists of all points whose distance to X is less than r.

A point x is an exterior point of X if and only if the distance from x to X is positive, i.e. if and only if there is an $r > 0$ such that $x \notin X(r)$. This means that a point x belongs to the closure $\operatorname{cl} X$, i.e. x is an interior point or a boundary point of X, if and only if x belongs to the sets $X(r)$ for all $r > 0$. In other words,

$$\operatorname{cl} X = \bigcap_{r>0} X(r).$$

A set X is said to be *bounded* if it is contained in some ball centered at 0, i.e. if there is a number $R > 0$ such that $X \subseteq B(0; R)$.

A set X that is both closed and bounded is called *compact*.

An important property of compact subsets X of $\mathbf{R}^n$ is given by the Bolzano–Weierstrass theorem: *Every infinite sequence $(x_n)_{n=1}^{\infty}$ of points x_n in a compact set X has a subsequence $(x_{n_k})_{k=1}^{\infty}$ that converges to a point in X.*

The cartesian product $X \times Y$ of a compact subset X of $\mathbf{R}^m$ and a compact subset Y of $\mathbf{R}^n$ is a compact subset of $\mathbf{R}^m \times \mathbf{R}^n \quad (= \mathbf{R}^{m+n})$.

Continuity

A function $f\colon X \to \mathbf{R}^m$, whose domain X is a subset of $\mathbf{R}^n$, is defined to be *continuous at the point* $a \in X$ if for each $\epsilon > 0$ there exists an $r > 0$ such that

$$f(X \cap B(a; r)) \subseteq B(f(a); \epsilon).$$

(Here, of course, the left B stands for balls in $\mathbf{R}^n$ and the right B stands for balls in $\mathbf{R}^m$.) The function is said to be *continuous on* X, or simply *continuous*, if it is continuous at all points $a \in X$.

The inverse image $f^{-1}(I)$ of an open interval under a continuous function $f\colon \mathbf{R}^n \to \mathbf{R}$ is an open set in $\mathbf{R}^n$. In particular, the sets $\{x \mid f(x) < a\}$ and $\{x \mid f(x) > a\}$, i.e. the sets $f^{-1}(]-\infty, a[)$ and $f^{-1}(]a, \infty[)$, are open for all $a \in \mathbf{R}$. Their complements, the sets $\{x \mid f(x) \geq a\}$ and $\{x \mid f(x) \leq a\}$, are thus closed.

Sums and (scalar) products of continuous functions are continuous, and quotients of real-valued continuous functions are continuous at all points where the quotients are well-defined. Compositions of continuous functions are continuous.

Compactness is preserved under continuous functions, that is the image $f(X)$ is compact if X is a compact subset of the domain of the continuous function f. For continuous functions f with codomain $\mathbf{R}$ this means that f is bounded on X and has a maximum and a minimum, i.e. there are two points $x_1, x_2 \in X$ such that $f(x_1) \leq f(x) \leq f(x_2)$ for all $x \in X$.

Lipschitz continuity

A function $f\colon X \to \mathbf{R}^m$ that is defined on a subset X of $\mathbf{R}^n$, is called *Lipschitz continuous* with Lipschitz constant L if

$$\|f(y) - f(x)\| \leq L\|y - x\| \qquad \text{for all } x, y \in X.$$

Note that the definition of Lipschitz continuity is norm independent, since all norms on $\mathbf{R}^n$ are equivalent, but the value of the Lipschitz constant L is obviously norm dependent.

Operator norms

Let $\|\cdot\|$ be a given norm on $\mathbf{R}^n$. Since the closed unit ball is compact and linear operators S on $\mathbf{R}^n$ are continuous, we get a finite number $\|S\|$, called the *operator norm*, by the definition

$$\|S\| = \sup_{\|x\|\leq 1} \|Sx\|.$$

That the operator norm really is a norm on the space of linear operators, i.e. that it satisfies conditions (i)–(iii) in the norm definition, follows immediately from the corresponding properties of the underlying norm on $\mathbf{R}^n$.

By definition, $S(x/\|x\|) \leq \|S\|$ for all $x \neq 0$, and consequently

$$\|Sx\| \leq \|S\|\|x\|$$

for all $x \in \mathbf{R}^n$.

From this inequality follows immediately that

$$\|STx\| \leq \|S\|\|Tx\| \leq \|S\|\|T\|\|x\|,$$

which gives us the important inequality

$$\|ST\| \leq \|S\|\|T\|$$

for the norm of a product of two operators.

The identity operator I on $\mathbf{R}^n$ clearly has norm equal to 1. Therefore, if the operator S is invertible, then, by choosing $T = S^{-1}$ in the above inequality, we obtain the inequality

$$\|S^{-1}\| \geq 1/\|S\|.$$

The operator norm obviously depends on the underlying norm on $\mathbf{R}^n$, but again, different norms on $\mathbf{R}^n$ give rise to equivalent norms on the space of operators. However, when speaking about the operator norm we shall in this book always assume that the underlying norm is the Euclidean norm even if this is not stated explicitely.

Symmetric operators, eigenvalues and norms

Every symmetric operator S on $\mathbf{R}^n$ is diagonizable according to the spectral theorem. This means that there is an ON-basis $e_1, e_2, \dots, e_n$ consisting of eigenvectors of S. Let $\lambda_1, \lambda_2, \dots, \lambda_n$ denote the corresponding eigenvalues.

The largest and the smallest eigenvalue $\lambda_{\max}$ and $\lambda_{\min}$ are obtained as maximum and minimum values, respectively, of the quadratic form $\langle x, Sx\rangle$ on the unit sphere $\|x\| = 1$:

$$\lambda_{\max} = \max_{\|x\|=1} \langle x, Sx\rangle \quad \text{and} \quad \lambda_{\min} = \min_{\|x\|=1} \langle x, Sx\rangle.$$

For, by using the expansion $x = \sum_{i=1}^n \xi_i e_i$ of x in the ON-basis of eigenvectors, we obtain the inequality

$$\langle x, Sx\rangle = \sum_{i=1}^n \lambda_i \xi_i^2 \le \lambda_{\max} \sum_{i=1}^n \xi_i^2 = \lambda_{\max} \|x\|^2,$$

and equality prevails when x is equal to the eigenvector e_i that corresponds to the eigenvalue $\lambda_{\max}$. An analogous inequality in the other direction holds for $\lambda_{\min}$, of course.

The operator norm (with respect to the Euclidean norm) moreover satisfies the equality

$$\|S\| = \max_{1\le i\le n} |\lambda_i| = \max\{|\lambda_{\max}|, |\lambda_{\min}|\}.$$

For, by using the above expansion of x, we have $Sx = \sum_{i=1}^n \lambda_i \xi_i e_i$, and consequently

$$\|Sx\|^2 = \sum_{i=1}^n \lambda_i^2 \xi_i^2 \le \max_{1\le i\le n} |\lambda_i|^2 \sum_{i=1}^n \xi_i^2 = (\max_{1\le i\le n} |\lambda_i|)^2 \, \|x\|^2,$$

with equality when x is the eigenvector that corresponds to $\max_i |\lambda_i|$.

If all eigenvalues of the symmetric operator S are nonzero, then S is invertible, and the inverse S^{-1} is symmetric with eigenvalues $\lambda_1^{-1}, \lambda_2^{-1}, \dots, \lambda_n^{-1}$. The norm of the inverse is given by

$$\|S^{-1}\| = 1/ \min_{1\le i\le n} |\lambda_i|.$$

A symmetric operator S is positive semidefinite if all its eigenvalues are nonnegative, and it is positive definite if all eigenvalues are positive. Hence, if S is positive definite, then

$$\|S\| = \lambda_{\max} \qquad \text{and} \qquad \|S^{-1}\| = 1/\lambda_{\min}.$$

It follows easily from the diagonizability of symmetric operators on $\mathbf{R}^n$ that every positive semidefinite symmetric operator S has a unique positive semidefinite symmetric square root $S^{1/2}$. Moreover, since

$$\langle x, Sx\rangle = \langle x, S^{1/2}(S^{1/2}x)\rangle = \langle S^{1/2}x, S^{1/2}x\rangle = \|S^{1/2}x\|$$

we conclude that the two operators S and $S^{1/2}$ have the same null space $\mathcal{N}(S)$ and that

$$\mathcal{N}(S) = \{x \in \mathbf{R}^n \mid Sx = 0\} = \{x \in \mathbf{R}^n \mid \langle x, Sx\rangle = 0\}.$$

Differentiability

A function $f\colon U \to \mathbf{R}$, which is defined on an open subset U of $\mathbf{R}^n$, is called *differentiable at the point* $a \in U$ if the partial derivatives $\frac{\partial f}{\partial x_i}$ exist at the point x and the equality

$$f(a+v) = f(a) + \sum_{i=1}^n \frac{\partial f}{\partial x_i}(a)\, v_i + r(v) \tag{1.2}$$

holds for all v in some neighborhood of the origin with a remainder term $r(v)$ that satisfies the condition

$$\lim_{v\to 0} \frac{r(v)}{\|v\|} = 0.$$

The linear form $Df(a)[v]$, defined by

$$Df(a)[v] = \sum_{i=1}^{n} \frac{\partial f}{\partial x_i}(a)\, v_i,$$

is called the *differential* of the function f at the point a. The coefficient vector

$$\Big(\frac{\partial f}{\partial x_1}(a), \frac{\partial f}{\partial x_2}(a), \dots, \frac{\partial f}{\partial x_n}(a)\Big)$$

of the differential is called the *derivative* or the *gradient* of f at the point a and is denoted by $f'(a)$ or $\nabla f(a)$. We shall mostly use the first mentioned notation.

The equation (1.2) can now be written in a compact form as

$$f(a+v) = f(a) + Df(a)[v] + r(v),$$

with

$$Df(a)[v] = \langle f'(a), v\rangle.$$

A function $f\colon U \to \mathbf{R}$ is called *differentiable (on U)* if it is differentiable at each point in U. In particular, this implies that U is an open set.

For functions of one variable, differentiability is clearly equivalent to the existence of the derivative, but for functions of several variables, the mere existence of the partial derivatives is no longer a guarantee for differentiability. However, if a function f has partial derivatives and these are continous on an open set U, then f is differentiable on U.

The Mean Value Theorem

Suppose $f\colon U \to \mathbf{R}$ is a differentiable function and that the line segment $[a, a+v]$ lies in U. Let $\phi(t) = f(a+tv)$. The function ϕ is then defined and differentiable on the interval $[0,1]$ with derivative

$$\phi'(t) = Df(a+tv)[v] = \langle f'(a+tv), v\rangle.$$

This is a special case of the chain rule but also follows easily from the definition of the derivative. By the usual mean value theorem for functions of one variable, there is a number $s \in]0,1[$ such that $\phi(1) - \phi(0) = \phi'(s)(1-0)$. Since $\phi(1) = f(a+v)$, $\phi(0) = f(a)$ and $a + sv$ is a point on the open line segment $]a, a+v[$, we have now deduced the following *mean value theorem* for functions of several variables.

Theorem 1.1.1. *Suppose the function $f\colon U \to \mathbf{R}$ is differentiable and that the line segment $[a, a+v]$ lies in U. Then there is a point $c \in]a, a+v[$ such that*

$$f(a+v) = f(a) + Df(c)[v].$$

Functions with Lipschitz continuous derivative

We shall sometimes need more precise information about the remainder term $r(v)$ in equation (1.2) than what follows from the definition of differentiability. We have the following result for functions with a Lipschitz continuous derivative.

Theorem 1.1.2. *Suppose the function $f\colon U \to \mathbf{R}$ is differentiable, that its derivative is Lipschitz continuous, i.e. that $\|f'(y) - f'(x)\| \le L\|y - x\|$ for all $x, y \in U$, and that the line segment $[a, a+v]$ lies in U. Then*

$$|f(a+v) - f(a) - Df(a)[v]| \le \frac{L}{2}\|v\|^2.$$

Proof. Define the function Φ on the interval $[0, 1]$ by

$$\Phi(t) = f(a+tv) - t\,Df(a)[v].$$

Then Φ is differentiable with derivative

$$\Phi'(t) = Df(a+tv)[v] - Df(a)[v] = \langle f'(a+tv) - f'(a), v\rangle,$$

and by using the Cauchy–Schwarz inequality and the Lipschitz continuity, we obtain the inequality

$$|\Phi'(t)| \le \|f'(a+tv) - f'(a)\| \cdot \|v\| \le Lt\,\|v\|^2.$$

Since $f(a+v) - f(a) - Df(a)[v] = \Phi(1) - \Phi(0) = \int_0^1 \Phi'(t)\,dt$, it now follows that

$$|f(a+v) - f(a) - Df(a)[v]| \le \int_0^1 |\Phi'(t)|\,dt \le L\|v\|^2 \int_0^1 t\,dt = \frac{L}{2}\,\|v\|^2.$$

$\square$

Two times differentiable functions

If the function f together with all its partial derivatives $\frac{\partial f}{\partial x_i}$ are differentiable on U, then f is said to be two times differentiable on U. The mixed partial second derivatives are then automatically equal, i.e.

$$\frac{\partial^2 f}{\partial x_i \partial x_j}(a) = \frac{\partial^2 f}{\partial x_j \partial x_i}(a)$$

for all i, j and all $a \in U$.

A sufficient condition for the function f to be two times differentiable on U is that all partial derivatives of order up to two exist and are continuous on U.

If $f\colon U \to \mathbf{R}$ is a two times differentiable function and a is a point in U, we define a symmetric bilinear form $D^2 f(a)[u,v]$ on $\mathbf{R}^n$ by

$$D^2 f(a)[u,v] = \sum_{i,j=1}^{n} \frac{\partial^2 f}{\partial x_i \partial x_j}(a) u_i v_j, \quad u, v \in \mathbf{R}^n.$$

The corresponding symmetric linear operator is called the *second derivative* of f at the point a and it is denoted by $f''(a)$. The matrix of the second derivative, i.e. the matrix

$$\Big[\frac{\partial^2 f}{\partial x_i \partial x_j}(a)\Big]_{i,j=1}^{n},$$

is called the *hessian* of f (at the point a). Since we do not distinguish between matrices and operators, we also denote the hessian by $f''(a)$.

The above symmetric bilinear form can now be expressed in the form

$$D^2f(a)[u,v] = \langle u, f''(a)v\rangle = u^T f''(a)v,$$

depending on whether we interpret the second derivative as an operator or as a matrix.

Let us recall *Taylor's formula*, which reads as follows for two times differentiable functions.

Theorem 1.1.3. *Suppose the function f is two times differentiable in a neighborhood of the point a. Then*

$$f(a+v) = f(a) + Df(a)[v] + \tfrac{1}{2}D^2f(a)[v,v] + r(v)$$

with a remainder term that satisfies $\lim_{v\to 0} r(v)/\|v\|^2 = 0$.

Three times differentiable functions

To define self-concordance we also need to consider functions that are three times differentiable on some open subset U of $\mathbf{R}^n$. For such functions f and points $a \in U$ we define a trilinear form $D^3f(a)[u,v,w]$ in the vectors $u, v, w \in \mathbf{R}^n$ by

$$D^3f(a)[u,v,w] = \sum_{i,j,k=1}^{n} \frac{\partial^3 f}{\partial x_i \partial x_j \partial x_k}(a)u_i v_j w_k.$$

We leave to the reader to formulate Taylor's formula for functions that are three times differentiable. We have the following differentiation rules, which follow from the chain rule and will be used several times in the final chapters:

$$\begin{aligned}
\frac{d}{dt}f(x+tv) &= Df(x+tv)[v] \\
\frac{d}{dt}\Big(Df(x+tv)[u]\Big) &= D^2f(x+tv)[u,v], \\
\frac{d}{dt}\Big(D^2f(x+tw)[u,v]\Big) &= D^3f(x+tw)[u,v,w].
\end{aligned}$$

As a consequence we get the following expressions for the derivatives of the restriction ϕ of the function f to the line through the point x with the direction given by v:

$$\begin{aligned}
\phi(t) &= f(x+tv), \\
\phi'(t) &= Df(x+tv)[v], \\
\phi''(t) &= D^2f(x+tv)[v,v], \\
\phi'''(t) &= D^3f(x+tv)[v,v,v].
\end{aligned}$$

Chapter 2

Convex sets

2.1 Affine sets and affine maps

Affine sets

Definition. A subset of $\mathbf{R}^n$ is called *affine* if for each pair of distinct points in the set it contains the entire line through the points.

Thus, a set X is affine if and only if

$$x, y \in X,\ \lambda \in \mathbf{R} \Rightarrow \lambda x + (1 - \lambda)y \in X.$$

The empty set $\emptyset$, the entire space $\mathbf{R}^n$, linear subspaces of $\mathbf{R}^n$, singleton sets $\{x\}$ and lines are examples of affine sets.

Definition. A linear combination $y = \sum_{j=1}^m \alpha_j x_j$ of vectors $x_1, x_2, \dots, x_m$ is called an *affine combination* if $\sum_{j=1}^m \alpha_j = 1$.

Theorem 2.1.1. *An affine set contains all affine combination of its elements.*

Proof. We prove the theorem by induction on the number of elements in the affine combination. So let X be an affine set. An affine combination of one element is the element itself. Hence, X contains all affine combinations that can be formed by one element in the set.

Now assume inductively that X contains all affine combinations that can be formed out of $m - 1$ elements from X, where $m \geq 2$, and consider an arbitrary affine combination $x = \sum_{j=1}^m \alpha_j x_j$ of m elements $x_1, x_2, \dots, x_m$ in X. Since $\sum_{j=1}^m \alpha_j = 1$, at least one coefficient α_j must be different from 1; assume without loss of generality that $\alpha_m \neq 1$, and let $s = 1 - \alpha_m = \sum_{j=1}^{m-1} \alpha_j$.

Then $s \neq 0$ and $\sum_{j=1}^{m-1} \alpha_j/s = 1$, which means that the element

$$y = \sum_{j=1}^{m-1} \frac{\alpha_j}{s} x_j$$

is an affine combination of $m-1$ elements in X. Therefore, y belongs to X, by the induction assumption. But $x = sy + (1-s)x_m$, and it now follows from the definition of affine sets that x lies in X. This completes the induction step, and the theorem is proved. □

Definition. Let A be an arbitrary nonempty subset of $\mathbf{R}^n$. The set of all affine combinations $\lambda_1 a_1 + \lambda_2 a_2 + \cdots + \lambda_m a_m$ that can be formed of an arbitrary number of elements $a_1, a_2, \ldots, a_m$ from A, is called the *affine hull* of A and is denoted by $\operatorname{aff} A$.

In order to have the affine hull defined also for the empty set, we put $\operatorname{aff} \emptyset = \emptyset$.

Theorem 2.1.2. *The affine hull* $\operatorname{aff} A$ *is an affine set containing* A *as a subset, and it is the smallest affine subset with this property, i.e. if the set* X *is affine and* $A \subseteq X$*, then* $\operatorname{aff} A \subseteq X$.

Proof. The set $\operatorname{aff} A$ is an affine set, because any affine combination of two elements in $\operatorname{aff} A$ is obviously an affine combination of elements from A, and the set A is a subset of its affine hull, since any element is an affine combination of itself.

If X is an affine set, then $\operatorname{aff} X \subseteq X$, by Theorem 2.1.1, and if $A \subseteq X$, then obviously $\operatorname{aff} A \subseteq \operatorname{aff} X$. Thus, $\operatorname{aff} A \subseteq X$ whenever X is an affine set and A is a subset of X. □

Characterisation of affine sets

Nonempty affine sets are translations of linear subspaces. More precisely, we have the following theorem.

Theorem 2.1.3. *If* X *is an affine subset of* $\mathbf{R}^n$ *and* $a \in X$*, then* $-a + X$ *is a linear subspace of* $\mathbf{R}^n$*. Moreover, for each* $b \in X$ *we have* $-b + X = -a + X$.

Thus, to each nonempty affine set X *there corresponds a uniquely defined linear subspace* U *such that* $X = a + U$.

Proof. Let $U = -a + X$. If $u_1 = -a + x_1$ and $u_2 = -a + x_2$ are two elements in U and α_1, α_2 are arbitrary real numbers, then the linear combination

$$\alpha_1 u_1 + \alpha_2 u_2 = -a + (1 - \alpha_1 - \alpha_2)a + \alpha_1 x_1 + \alpha_2 x_2$$

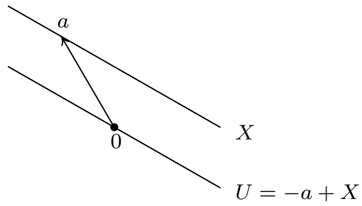

Figure 2.1. Illustration for Theorem 2.1.3: An affine set X and the corresponding linear subspace U.

is an element in U, because $(1-\alpha_1-\alpha_2)a+\alpha_1x_1+\alpha_2x_2$ is an affine combination of elements in X and hence belongs to X, according to Theorem 2.1.1. This proves that U is a linear subspace.

Now assume that $b \in X$, and let $v = -b + x$ be an arbitrary element in $-b + X$. By writing v as $v = -a + (a - b + x)$ we see that v belongs to $-a + X$, too, because $a - b + x$ is an affine combination of elements in X. This proves the inclusion $-b + X \subseteq -a + X$. The converse inclusion follows by symmetry. Thus, $-a + X = -b + X$. □

Dimension

The following definition is justified by Theorem 2.1.3.

Definition. The *dimension* $\dim X$ of a nonempty affine set X is defined as the dimension of the linear subspace $-a+X$, where a is an arbitrary element in X.

Since every nonempty affine set has a well-defined dimension, we can extend the dimension concept to arbitrary nonempty sets as follows.

Definition. The *(affine) dimension* $\dim A$ of a nonempty subset A of $\mathbf{R}^n$ is defined to be the dimension of its affine hull $\operatorname{aff} A$.

The dimension of an open ball $B(a;r)$ in $\mathbf{R}^n$ is n, and the dimension of a line segment $[x,y]$ is 1.

The dimension is *invariant under translation* i.e. if A is a nonempty subset of $\mathbf{R}^n$ and $a \in \mathbf{R}^n$ then

$$\dim(a+A) = \dim A,$$

and it is *increasing* in the following sense:

$$A \subseteq B \Rightarrow \dim A \leq \dim B.$$

Affine sets as solutions to systems of linear equations

Our next theorem gives a complete description of the affine subsets of $\mathbf{R}^n$.

Theorem 2.1.4. *Every affine subset of $\mathbf{R}^n$ is the solution set of a system of linear equations*

$$\begin{cases} c_{11}x_1 + c_{12}x_2 + \cdots + c_{1n}x_n = b_1 \\ c_{21}x_1 + c_{22}x_2 + \cdots + c_{2n}x_n = b_2 \\ \qquad\vdots \\ c_{m1}x_1 + c_{m2}x_2 + \cdots + c_{mn}x_n = b_m \end{cases}$$

and conversely. The dimension of a nonempty solution set equals $n-r$, where r is the rank of the coefficient matrix C.

Proof. The empty affine set is obtained as the solution set of an inconsistent system. Therefore, we only have to consider nonempty affine sets X, and these are of the form $X = x_0 + U$, where x_0 belongs to X and U is a linear subspace of $\mathbf{R}^n$. But each linear subspace is the solution set of a homogeneous system of linear equations. Hence there exists a matrix C such that

$$U = \{x \mid Cx = 0\},$$

and $\dim U = n - \operatorname{rank} C$. With $b = Cx_0$ it follows that $x \in X$ if and only if $Cx - Cx_0 = C(x - x_0) = 0$, i.e. if and only if x is a solution to the linear system $Cx = b$.

Conversely, if x_0 is a solution to the above linear system so that $Cx_0 = b$, then x is a solution to the same system if and only if the vector $z = x - x_0$ belongs to the solution set U of the homogeneous equation system $Cz = 0$. It follows that the solution set of the equation system $Cx = b$ is of the form $x_0 + U$, i.e. it is an affine set. □

Hyperplanes

Definition. Affine subsets of $\mathbf{R}^n$ of dimension $n - 1$ are called *hyperplanes*.

Theorem 2.1.4 has the following corollary:

Corollary 2.1.5. *A subset X of $\mathbf{R}^n$ is a hyperplane if and only if there exist a nonzero vector $c = (c_1, c_2, \dots, c_n)$ and a real number b so that*

$$X = \{x \in \mathbf{R}^n \mid \langle c, x\rangle = b\}.$$

It follows from Theorem 2.1.4 that every affine proper subset of $\mathbf{R}^n$ can be expressed as an intersection of hyperplanes.

Affine maps

Definition. Let X be an affine subset of $\mathbf{R}^n$. A map $T\colon X \to \mathbf{R}^m$ is called *affine* if

$$T(\lambda x + (1 - \lambda)y) = \lambda Tx + (1 - \lambda)Ty$$

for all $x, y \in X$ and all $\lambda \in \mathbf{R}$.

Using induction, it is easy to prove that if $T\colon X \to \mathbf{R}^m$ is an affine map and $x = \alpha_1 x_1 + \alpha_2 x_2 + \cdots + \alpha_m x_m$ is an affine combination of elements in X, then

$$Tx = \alpha_1 Tx_1 + \alpha_2 Tx_2 + \cdots + \alpha_m Tx_m.$$

Moreover, the image $T(Y)$ of an affine subset Y of X is an affine subset of $\mathbf{R}^m$, and the inverse image $T^{-1}(Z)$ of an affine subset Z of $\mathbf{R}^m$ is an affine subset of X.

The composition of two affine maps is affine. In particular, a linear map followed by a translation is an affine map, and our next theorem shows that each affine map can be written as such a composition.

Theorem 2.1.6. *Let X be an affine subset of $\mathbf{R}^n$, and suppose the map $T\colon X \to \mathbf{R}^m$ is affine. Then there exist a linear map $C\colon \mathbf{R}^n \to \mathbf{R}^m$ and a vector v in $\mathbf{R}^m$ so that $Tx = Cx + v$ for all $x \in X$.*

Proof. Write the domain of T in the form $X = x_0 + U$ with $x_0 \in X$ and U as a linear subspace of $\mathbf{R}^n$, and define the map C on the subspace U by

$$Cu = T(x_0 + u) - Tx_0.$$

Then, for each $u_1, u_2 \in U$ and $\alpha_1, \alpha_2 \in \mathbf{R}$ we have

$$\begin{aligned}
C(\alpha_1 u_1 + \alpha_2 u_2) &= T(x_0 + \alpha_1 u_1 + \alpha_2 u_2) - Tx_0 \\
&= T\big(\alpha_1(x_0 + u_1) + \alpha_2(x_0 + u_2) + (1 - \alpha_1 - \alpha_2)x_0\big) - Tx_0 \\
&= \alpha_1 T(x_0 + u_1) + \alpha_2 T(x_0 + u_2) + (1 - \alpha_1 - \alpha_2)Tx_0 - Tx_0 \\
&= \alpha_1\big(T(x_0 + u_1) - Tx_0\big) + \alpha_2\big(T(x_0 + u_2) - Tx_0\big) \\
&= \alpha_1 Cu_1 + \alpha_2 Cu_2.
\end{aligned}$$

So the map C is linear on U and it can, of course, be extended to a linear map on all of $\mathbf{R}^n$.

For $x \in X$ we now obtain, since $x - x_0$ belongs to U,

$$Tx = T(x_0 + (x - x_0)) = C(x - x_0) + Tx_0 = Cx - Cx_0 + Tx_0,$$

which proves the theorem with v equal to $Tx_0 - Cx_0$. □

2.2 Convex sets

Basic definitions and properties

Definition. A subset X of $\mathbf{R}^n$ is called *convex* if $[x, y] \subseteq X$ for all x, $y \in X$.

In other words, a set X is convex if and only if it contains the line segment between each pair of its points.

Figure 2.2. A convex set and a non-convex set

Example 2.2.1. Affine sets are obviously convex. In particular, the empty set $\emptyset$, the entire space $\mathbf{R}^n$ and linear subspaces are convex sets. Open line segments and closed line segments are clearly convex. □

Example 2.2.2. Open balls $B(a; r)$ (with respect to arbitrary norms $\|\cdot\|$) are convex sets. This follows from the triangle inequality and homogenouity, for if $x, y \in B(a; r)$ and $0 \le \lambda \le 1$, then

$$\begin{aligned}\|\lambda x + (1-\lambda)y - a\| &= \|\lambda(x-a) + (1-\lambda)(y-a)\| \\ &\le \lambda\|x-a\| + (1-\lambda)\|y-a\| < \lambda r + (1-\lambda)r = r,\end{aligned}$$

which means that each point $\lambda x+(1-\lambda)y$ on the segment $[x, y]$ lies in $B(a; r)$.

The corresponding closed balls $\overline{B}(a; r) = \{x \in \mathbf{R}^n \mid \|x - a\| \le r\}$ are of course convex, too. □

Definition. A linear combination $y = \sum_{j=1}^{m} \alpha_j x_j$ of vectors $x_1, x_2, \dots, x_m$ is called a *convex combination* if $\sum_{j=1}^{m} \alpha_j = 1$ and $\alpha_j \ge 0$ for all j.

Theorem 2.2.1. *A convex set contains all convex combinations of its elements.*

Proof. Let X be an arbitrary convex set. A convex combination of one element is the element itself, and hence X contains all convex combinations formed by just one element of the set. Now assume inductively that X contains all convex combinations that can be formed by $m-1$ elements of X, and consider an arbitrary convex combination $x = \sum_{j=1}^{m} \alpha_j x_j$ of $m \geq 2$ elements $x_1, x_2, \ldots, x_m$ in X. Since $\sum_{j=1}^{m} \alpha_j = 1$, some coefficient α_j must be strictly less than 1, and assume without loss of generality that $\alpha_m < 1$, and let $s = 1 - \alpha_m = \sum_{j=1}^{m-1} \alpha_j$. Then $s > 0$ and $\sum_{j=1}^{m-1} \alpha_j/s = 1$, which means that

$$y = \sum_{j=1}^{m-1} \frac{\alpha_j}{s} x_j$$

is a convex combination of $m-1$ elements in X. By the induction hypothesis, y belongs to X. But $x = sy + (1-s)x_m$, and it now follows from the convexity definition that x belongs to X. This completes the induction step and the proof of the theorem. □

2.3 Convexity preserving operations

We now describe a number of ways to construct new convex sets from given ones.

Image and inverse image under affine maps

Theorem 2.3.1. *Let $T : V \to \mathbf{R}^m$ be an affine map.*

(i) The image $T(X)$ of a convex subset X of V is convex.

(ii) The inverse image $T^{-1}(Y)$ of a convex subset Y of $\mathbf{R}^m$ is convex.

Proof. (i) Suppose y_1, $y_2 \in T(X)$ and $0 \leq \lambda \leq 1$. Let x_1, x_2 be points in X such that $y_i = T(x_i)$. Since

$$\lambda y_1 + (1-\lambda)y_2 = \lambda T x_1 + (1-\lambda)T x_2 = T(\lambda x_1 + (1-\lambda)x_2)$$

and $\lambda x_1 + (1-\lambda)x_2$ lies X, it follows that $\lambda y_1 + (1-\lambda)y_2$ lies in $T(X)$. This proves that the image set $T(X)$ is convex.

(ii) To prove the convexity of the inverse image $T^{-1}(Y)$ we instead assume that $x_1, x_2 \in T^{-1}(Y)$, i.e. that $Tx_1, Tx_2 \in Y$, and that $0 \leq \lambda \leq 1$. Since Y is a convex set,

$$T(\lambda x_1 + (1-\lambda)x_2) = \lambda T x_1 + (1-\lambda)T x_2$$

is an element of Y, and this means that $\lambda x_1 + (1-\lambda)x_2$ lies in $T^{-1}(Y)$. $\square$

As a special case of the preceding theorem it follows that translations $a + X$ of a convex set X are convex.

Example 2.3.1. The sets

$$\{x \in \mathbf{R}^n \mid \langle c, x\rangle \geq b\} \quad \text{and} \quad \{x \in \mathbf{R}^n \mid \langle c, x\rangle \leq b\},$$

where b is an arbitrary real number and $c = (c_1, c_2, \dots, c_n)$ is an arbirary nonzero vector, are called opposite *closed halfspaces*. Their complements, i.e.

$$\{x \in \mathbf{R}^n \mid \langle c, x\rangle < b\} \quad \text{and} \quad \{x \in \mathbf{R}^n \mid \langle c, x\rangle > b\},$$

are called *open halfspaces*.

The halfspaces $\{x \in \mathbf{R}^n \mid \langle c, x\rangle \geq b\}$ and $\{x \in \mathbf{R}^n \mid \langle c, x\rangle > b\}$ are inverse images of the real intervals $[b, \infty[$ and $]b, \infty[$, respectively, under the linear map $x \mapsto \langle c, x\rangle$. It therefore follows from Theorem 2.3.1 that *halfspaces are convex sets.* $\square$

Intersection and union

Theorem 2.3.2. *Let $\{X_i \mid i \in I\}$ be a family of convex subsets of $\mathbf{R}^n$. The intersection $\bigcap\{X_i \mid i \in I\}$ is a convex set.*

Proof. Suppose x, y are points in the intersection Y. The definition of an intersection implies that x and y lie in X_i for all indices $i \in I$, and convexity implies that $[x, y] \subseteq X_i$ for all $i \in I$. Therefore, $[x, y] \subseteq Y$, again by the definition of set intersection. This proves that the intersection is a convex set. □

A union of convex sets is, of course, in general not convex. However, there is a trivial case when convexity is preserved, namely when the sets can be ordered in such a way as to form an "increasing chain".

Theorem 2.3.3. *Suppose $\{X_i \mid i \in I\}$ is a family of convex sets X_i and that for each pair $i, j \in I$ either $X_i \subseteq X_j$ or $X_j \subseteq X_i$. The union $\bigcup\{X_i \mid i \in I\}$ is then a convex set.*

Proof. The assumptions imply that, for each pair of points x, y in the union there is an index $i \in I$ such that both points belong to X_i. By convexity, the entire segment $[x, y]$ lies in X_i, and thereby also in the union. □

Example 2.3.2. The convexity of closed balls follows from the convexity of open balls, because $\overline{B}(a; r_0) = \bigcap\{B(a; r) \mid r > r_0\}$.

Conversely, the convexity of open balls follows from the convexity of closed balls, since $B(a; r_0) = \bigcup\{\overline{B}(a; r) \mid r < r_0\}$ and the sets $\overline{B}(a; r)$ form an increasing chain. □

Definition. A subset X of $\mathbf{R}^n$ is called a *polyhedron* if X can be written as an intersection of finitely many closed halfspaces or if $X = \mathbf{R}^n$.†

Polyhedra are convex sets because of Theorem 2.3.2, and they can be represented as solution sets to systems of linear inequalities. By multiplying some of the inequalities by -1, if necessary, we may without loss of generality assume that all inequalities are of the form $c_1x_1 + c_2x_2 + \cdots + c_nx_n \geq d$. This means that every polyedron is the solution set to a system of the following

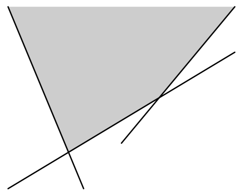

Figure 2.3. A polyhedron in $\mathbf{R}^2$

form

$$\begin{cases} c_{11}x_1 + \ c_{12}x_2 + \cdots + \ c_{1n}x_n \geq \ b_1 \\ c_{21}x_1 + \ c_{22}x_2 + \cdots + \ c_{2n}x_n \geq \ b_2 \\ \qquad\qquad\qquad\qquad\qquad\qquad \vdots \\ c_{m1}x_1 + c_{m2}x_2 + \cdots + c_{mn}x_n \geq b_m, \end{cases}$$

or in matrix notation

$$Cx \geq b.$$

The intersection of finitely many polyhedra is clearly a polyhedron. Since each hyperplane is the intersection of two opposite closed halfspaces, and each affine set (except the entire space) is the intersection of finitely many hyperplanes, it follows especially that affine sets are polyhedra. In particular, the empty set is a polyhedron.

Cartesian product

Theorem 2.3.4. *The Cartesian product $X \times Y$ of two convex sets X and Y is a convex set.*

Proof. Suppose X lies in $\mathbf{R}^n$ and Y lies in $\mathbf{R}^m$. The projections

$$P_1 \colon \mathbf{R}^n \times \mathbf{R}^m \to \mathbf{R}^n \quad \text{and} \quad P_2 \colon \mathbf{R}^n \times \mathbf{R}^m \to \mathbf{R}^m,$$

defined by $P_1(x, y) = x$ and $P_2(x, y) = y$, are linear maps, and

$$X \times Y = (X \times \mathbf{R}^m) \cap (\mathbf{R}^n \times Y) = P_1^{-1}(X) \cap P_2^{-1}(Y).$$

The assertion of the theorem is therefore a consequence of Theorem 2.3.1 and Theorem 2.3.2. □

Sum

Theorem 2.3.5. *The sum $X + Y$ of two convex subsets X and Y of $\mathbf{R}^n$ is convex, and the product αX of a number α and a convex set X is convex.*

Proof. The maps $S\colon \mathbf{R}^n \times \mathbf{R}^n \to \mathbf{R}^n$ and $T\colon \mathbf{R}^n \to \mathbf{R}^n$, defined by $S(x,y) = x + y$ and $Tx = \alpha x$, are linear. Since $X + Y = S(X \times Y)$ and $\alpha X = T(X)$, our assertions follow from Theorems 2.3.1 and 2.3.4. $\square$

Example 2.3.3. The set $X(r)$ of all points whose distance to a given set X is less than the positive number r, can be written as a sum, namely

$$X(r) = X + B(0; r).$$

Since open balls are convex, we conclude from Theorem 2.3.5 that the set $X(r)$ is convex if X is a convex set. $\square$

Image and inverse image under the perspective map

Definition. The *perspective map* $P\colon \mathbf{R}^n \times \mathbf{R}_{++} \to \mathbf{R}^n$ is defined by

$$P(x,t) = t^{-1}x$$

for $x \in \mathbf{R}^n$ and $t > 0$.

The perspective map thus first rescales points in $\mathbf{R}^n \times \mathbf{R}_{++}$ so that the last coordinate becomes 1 and then throws the last coordinate away. Figure 2.4 illustrates the process.

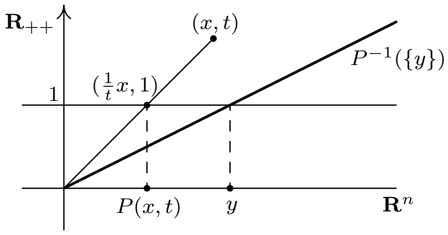

Figure 2.4. The perspective map P. The inverse image of a point $y \in \mathbf{R}^n$ is a halfline.

Theorem 2.3.6. *Let X be a convex subset of $\mathbf{R}^n \times \mathbf{R}_{++}$ and Y be a convex subset of $\mathbf{R}^n$. The image $P(X)$ of X and the inverse image $P^{-1}(Y)$ of Y under the perspective map $P\colon \mathbf{R}^n \times \mathbf{R}_{++} \to \mathbf{R}^n$ are convex sets.*

Proof. To prove that the image $P(X)$ is convex we assume that $y, y' \in P(X)$ and have to prove that the point $\lambda y + (1-\lambda)y'$ lies in $P(X)$ if $0 < \lambda < 1$. To achieve this we first note that there exist numbers $t, t' > 0$ such that the points (ty, t) and $(t'y', t')$ belong to X, and then define

$$\alpha = \frac{\lambda t'}{\lambda t' + (1-\lambda)t}.$$

Clearly $0 < \alpha < 1$, and it now follows from the convexity of X that the point

$$z = \alpha(ty,t) + (1-\alpha)(t'y',t') = \Big(\frac{tt'(\lambda y + (1-\lambda)y')}{\lambda t' + (1-\lambda)t}, \frac{tt'}{\lambda t' + (1-\lambda)t}\Big)$$

lies X. Thus, $P(z) \in P(X)$, and since $P(z) = \lambda y + (1-\lambda)y'$, we are done.

To prove that the inverse image $P^{-1}(Y)$ is convex, we instead assume that (x,t) and (x',t') are points in $P^{-1}(Y)$ and that $0 < \lambda < 1$. We will prove that the point $\lambda(x,t) + (1-\lambda)(x',t')$ lies in $P^{-1}(Y)$.

To this end we note that the the points $\frac{1}{t}x$ and $\frac{1}{t'}x'$ belong to Y and that

$$\alpha = \frac{\lambda t}{\lambda t + (1-\lambda)t'}$$

is a number between 0 and 1. Thus,

$$z = \alpha\frac{1}{t}x + (1-\alpha)\frac{1}{t'}x' = \frac{\lambda x + (1-\lambda)x'}{\lambda t + (1-\lambda)t'}$$

is a point in Y by convexity, and consequently $\big((\lambda t + (1-\lambda)t')z, \lambda t + (1-\lambda)t'\big)$ is a point in $P^{-1}(Y)$. But

$$\big((\lambda t + (1-\lambda)t')z, \lambda t + (1-\lambda)t'\big) = \lambda(x,t) + (1-\lambda)(x',t')$$

and this completes the proof. □

EXAMPLE 2.3.4. The set $\{(x, x_{n+1}) \in \mathbf{R}^n \times \mathbf{R} \mid \|x\| < x_{n+1}\}$ is the inverse image of the unit ball $B(0;1)$ under the perspective map, and it is therefore a convex set in $\mathbf{R}^{n+1}$ for each particular choice of norm $\|\cdot\|$. The following convex sets are obtained by choosing the ℓ^1-norm, the Euclidean norm and the maximum norm, respectively, as norm:

$$\{x \in \mathbf{R}^{n+1} \mid x_{n+1} > |x_1| + |x_2| + \cdots + |x_n|\,\},$$
$$\{x \in \mathbf{R}^{n+1} \mid x_{n+1} > (x_1^2 + x_2^2 + \cdots + x_n^2)^{1/2}\,\} \quad \text{and}$$
$$\{x \in \mathbf{R}^{n+1} \mid x_{n+1} > \max_{1 \le i \le n} |x_i|\,\}.$$
□

2.4 Convex hull

Definition. Let A be a nonempty set in $\mathbf{R}^n$. The set of all convex combinations $\lambda_1 a_1 + \lambda_2 a_2 + \cdots + \lambda_m a_m$ of an arbitrary number of elements $a_1, a_2, \dots, a_m$ in A is called the *convex hull* of A and is denoted by $\operatorname{cvx} A$.

Moreover, to have the convex hull defined for the empty set, we define $\operatorname{cvx} \emptyset = \emptyset$.

Theorem 2.4.1. *The convex hull* $\operatorname{cvx} A$ *is a convex set containing A, and it is the smallest set with this property, i.e. if X is a convex set and $A \subseteq X$, then* $\operatorname{cvx} A \subseteq X$.

Proof. $\operatorname{cvx} A$ is a convex set, because convex combinations of two elements of the type $\sum_{j=1}^m \lambda_j a_j$, where $m \ge 1$, $\lambda_1, \lambda_2, \dots, \lambda_m \ge 0$, $\sum_{j=1}^m \lambda_j = 1$ and $a_1, a_2, \dots, a_m \in A$, is obviously an element of the same type. Moreover, $A \subseteq \operatorname{cvx} A$, because each element in A is a convex combination of itself ($a = 1a$).

A convex set X contains all convex combinations of its elements, according to Theorem 2.2.1. If $A \subseteq X$, then in particular X contains all convex combinations of elements in A, which means that $\operatorname{cvx} A \subseteq X$. □

The convex hull of a set in $\mathbf{R}^n$ consists of all convex combinations of an arbitrary number of elements in the set, but each element of the hull is actually a convex combination of at most $n+1$ elements.

Theorem 2.4.2. *Let $A \subseteq \mathbf{R}^n$ and suppose that $x \in \operatorname{cvx} A$. Then A contains a subset B with at most $n+1$ elements such that $x \in \operatorname{cvx} B$.*

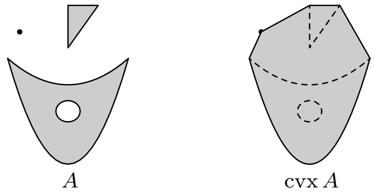

Figure 2.5. A set and its convex hull

Proof. According to the definition of convex hull there exists a finite subset B of A such that $x \in \operatorname{cvx} B$. Choose such a subset $B = \{b_1, b_2, \dots, b_m\}$ with as few elements as possible. By the minimality assumption, $x = \sum_{j=1}^{m} \lambda_j b_j$ with $\sum_{j=1}^{m} \lambda_j = 1$ and $\lambda_j > 0$ for all j.

Let $c_j = b_j - b_m$ for $j = 1, 2, \dots, m-1$. We will show that the set $C = \{c_1, c_2, \dots, c_{m-1}\}$ is a linearly independent subset of $\mathbf{R}^n$, and this obviously implies that $m \le n + 1$.

Suppose on the contrary that the set C is linearly dependent. Then there exist real numbers μ_j, not all of them equal to 0, such that $\sum_{j=1}^{m-1} \mu_j c_j = 0$. Now let $\mu_m = -\sum_{j=1}^{m-1} \mu_j$; then $\sum_{j=1}^{m} \mu_j = 0$ and $\sum_{j=1}^{m} \mu_j b_j = 0$. Moreover, at least one of the m numbers $\mu_1, \mu_2, \dots, \mu_m$ is positive.

Consider the numbers $\nu_j = \lambda_j - t\mu_j$ for $t > 0$. We note that

$$\sum_{j=1}^{m} \nu_j = \sum_{j=1}^{m} \lambda_j - t\sum_{j=1}^{m} \mu_j = 1 \quad \text{and} \quad \sum_{j=1}^{m} \nu_j b_j = \sum_{j=1}^{m} \lambda_j b_j - t\sum_{j=1}^{m} \mu_j b_j = x.$$

Moreover, $\nu_j \geq \lambda_j > 0$ if $\mu_j \leq 0$, and $\nu_j \geq 0$ if $\mu_j > 0$ and $t \leq \lambda_j/\mu_j$. Therefore, by choosing t as the smallest number of the numbers λ_j/μ_j with positive denominator μ_j, we obtain numbers ν_j such that $\nu_j \geq 0$ for all j and $\nu_{j_0} = 0$ for at least one index j_0. This means that x is a convex combination of elements in the set $B \setminus \{b_{j_0}\}$, which consists of $m - 1$ elements. This contradicts the minimality assumption concerning the set B, and our proof by contradiction is finished. □

2.5 Topological properties

Closure

Theorem 2.5.1. *The closure* $\operatorname{cl} X$ *of a convex set* X *is convex.*

Proof. We recall that $\operatorname{cl} X = \bigcap_{r>0} X(r)$, where $X(r)$ denotes the set of all points whose distance from X is less than r. The sets $X(r)$ are convex when the set X is convex (see Example 2.3.3), and an intersection of convex sets is convex. □

Interior and relative interior

The interior of a convex set may be empty. For example, line segments in $\mathbf{R}^n$ have no interior points when $n \geq 2$. A necessary and sufficient condition for nonempty interior is given by the following theorem.

Theorem 2.5.2. *A convex subset* X *of* $\mathbf{R}^n$ *has interior points if and only if* $\dim X = n$.

Proof. If X has an interior point a, then there exists an open ball $B = B(a; r)$ around a such that $B \subseteq X$, which implies that $\dim X \geq \dim B = n$, i.e. $\dim X = n$.

To prove the converse, let us assume that $\dim X = n$; we will prove that $\operatorname{int} X \neq \emptyset$. Since the dimension of a set is invariant under translations and $\operatorname{int}(a + X) = a + \operatorname{int} X$, we may assume that $0 \in X$.

Let $\{a_1, a_2, \dots, a_m\}$ be a maximal set of linearly independent vectors in X; then X is a subset of the linear subspace of dimension m which is spanned by these vectors, and it follows from the dimensionality assumption that $m = n$.

The set X contains the convex hull of the vectors $0, a_1, a_2, \dots, a_n$ as a subset, and, in particular, it thus contains the nonempty open set

$$\{\lambda_1 a_1 + \lambda_2 a_2 + \cdots + \lambda_n a_n \mid 0 < \lambda_1 + \cdots + \lambda_n < 1,\ \lambda_1 > 0, \dots, \lambda_n > 0\}.$$

All points in this latter set are interior points of X, so $\operatorname{int} X \neq \emptyset$. $\square$

A closed line segment $[a, b]$ in the two-dimensional space $\mathbf{R}^2$ has no interior points, but if we consider the line segment as a subset of a line and identify the line with the space $\mathbf{R}$, then it has interior points and its interior is equal to the corresponding open line segment $]a, b[$. A similar remark holds for the convex hull $T = \operatorname{cvx}\{a, b, c\}$ of three noncolinear points in three-space; the triangle T has interior points when viewed as a subset of $\mathbf{R}^2$, but it lacks interior points as a subset of $\mathbf{R}^3$. This conflict is unsatisfactory if we want a concept that is independent of the dimension of the surrounding space, but the dilemma disappears if we use the relative topology that the affine hull of the set inherits from the surrounding space $\mathbf{R}^n$.

Definition. Let X be an m-dimensional subset of $\mathbf{R}^n$, and let V denote the affine hull of X, i.e. V is the m-dimensional affine set that contains X.

A point $x \in X$ is called a *relative interior point* of X if there exists an $r > 0$ such that $B(x; r) \cap V \subseteq X$, and the set of all relative interior points of X is called the *relative interior* of X and is denoted by $\operatorname{rint} X$.

A point $x \in \mathbf{R}^n$ is called a *relative boundary point* of X if, for every $r > 0$, the intersection $B(x; r) \cap V$ contains at least one point from X and at least one point from $V \setminus X$. The set of all relative boundary points is called the *relative boundary* of X and is denoted by $\operatorname{rbdry} X$.

The relative interior of X is obviously a subset of X, and the relative boundary of X is a subset of the boundary of X. It follows that

$$\operatorname{rint} X \cup \operatorname{rbdry} X \subseteq X \cup \operatorname{bdry} X = \operatorname{cl} X.$$

Conversely, if x is a point in the closure $\operatorname{cl} X$, then for each $r > 0$

$$B(x, r) \cap V \cap X = B(x, r) \cap X \neq \emptyset.$$

Thus, x is either a relative boundary point or a relative interior point of X. This proves the converse inclusion, and we conclude that

$$\operatorname{rint} X \cup \operatorname{rbdry} X = \operatorname{cl} X,$$

or equivalently, that

$$\operatorname{rbdry} X = \operatorname{cl} X \setminus \operatorname{rint} X.$$

It follows from Theorem 2.5.2, with $\mathbf{R}^n$ replaced by $\operatorname{aff} X$, that every nonempty convex set has a nonempty relative interior.

Note that the relative interior of a line segment $[a, b]$ is the corresponding open line segment $]a, b[$, which is consistent with calling $]a, b[$ an open segment. The relative interior of a set $\{a\}$ consisting of just one point is the set itself.

Theorem 2.5.3. *The relative interior* $\operatorname{rint} X$ *of a convex set* X *is convex.*

Proof. The theorem follows as a corollary of the following theorem, since $\operatorname{rint} X \subseteq \operatorname{cl} X$. □

Theorem 2.5.4. *Suppose that* X *is a convex set, that* $a \in \operatorname{rint} X$ *and that* $b \in \operatorname{cl} X$*. The entire open line segment* $]a, b[$ *is then a subset of* $\operatorname{rint} X$.

Proof. Let $V = \operatorname{aff} X$ denote the affine set of least dimension that includes X, and let $c = \lambda a + (1 - \lambda)b$, where $0 < \lambda < 1$, be an arbitrary point on the open segment $]a, b[$. We will prove that c is a relative interior point of X by constructing an open ball B which contains c and whose intersection with V is contained in X.

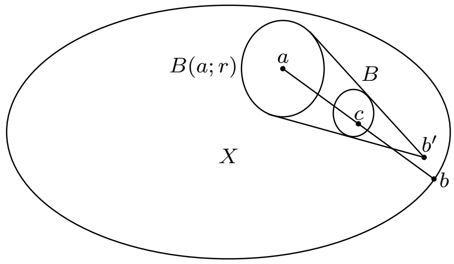

Figure 2.6. Illustration of the proof of Theorem 2.5.4. The convex hull of the ball $B(a;r)$ and the point b' forms a "cone" with nonempty interior that contains the point c.

To this end, we choose $r > 0$ such that $B(a;r) \cap V \subseteq X$ and a point $b' \in X$ such that $\|b' - b\| < \lambda r/(1-\lambda)$; this is possible since a is a relative interior point of X and b is a point in the closure of X. Let

$$B = \lambda B(a;r) + (1-\lambda)b',$$

and observe that $B = B(\lambda a + (1-\lambda)b'; \lambda r)$. The open ball B contains the point c because

$$\|c - (\lambda a + (1-\lambda)b')\| = \|(1-\lambda)(b-b')\| = (1-\lambda)\|b-b'\| < \lambda r.$$

Moreover, $B \cap V = \lambda(B(a;r)\cap V) + (1-\lambda)b' \subseteq \lambda X + (1-\lambda)X \subseteq X$, due to convexity. This completes the proof. $\square$

Theorem 2.5.5. *Let X be a convex set. Then*

(i) $\operatorname{cl}(\operatorname{rint} X) = \operatorname{cl} X$;

(ii) $\operatorname{rint}(\operatorname{cl} X) = \operatorname{rint} X$;

(iii) $\operatorname{rbdry}(\operatorname{cl} X) = \operatorname{rbdry}(\operatorname{rint} X) = \operatorname{rbdry} X$.

Proof. The equalities in (iii) for the relative boundaries follow from the other two and the definition of the relative boundary.

The inclusions $\operatorname{cl}(\operatorname{rint} X) \subseteq \operatorname{cl} X$ and $\operatorname{rint} X \subseteq \operatorname{rint}(\operatorname{cl} X)$ are both trivial, because it follows, for arbitrary sets A and B, that $A \subseteq B$ implies $\operatorname{cl} A \subseteq \operatorname{cl} B$ and $\operatorname{rint} A \subseteq \operatorname{rint} B$.

It thus only remains to prove the two inclusions

$$\operatorname{cl} X \subseteq \operatorname{cl}(\operatorname{rint} X) \quad \text{and} \quad \operatorname{rint}(\operatorname{cl} X) \subseteq \operatorname{rint} X.$$

So fix a point $x_0 \in \operatorname{rint} X$.

If $x \in \operatorname{cl} X$, then every point on the open segment $]x_0, x[$ lies in $\operatorname{rint} X$, by Theorem 2.5.4, and it follows from this that the point x is either an interior point or a boundary point of $\operatorname{rint} X$, i.e. a point in the closure $\operatorname{cl}(\operatorname{rint} X)$. This proves the inclusion $\operatorname{cl} X \subseteq \operatorname{cl}(\operatorname{rint} X)$.

To prove the remaining inclusion $\operatorname{rint}(\operatorname{cl} X) \subseteq \operatorname{rint} X$ we instead assume that $x \in \operatorname{rint}(\operatorname{cl} X)$ and define $y_t = (1-t)x_0 + tx$ for $t > 1$. Since $y_t \to x$ as $t \to 1$, the points y_t belong to $\operatorname{cl} X$ for all t sufficiently close to 1. Choose a number $t_0 > 1$ such that y_{t_0} belongs to $\operatorname{cl} X$. According to Theorem 2.5.4, all points on the open segment $]x_0, y_{t_0}[$ are relative interior points in X, and x is such a point since $x = \frac{1}{t_0} y_{t_0} + (1 - \frac{1}{t_0})x_0$. This proves the implication $x \in \operatorname{rint}(\operatorname{cl} X) \Rightarrow x \in \operatorname{rint} X$, and the proof is now complete. □

Compactness

Theorem 2.5.6. *The convex hull* $\operatorname{cvx} A$ *of a compact subset* A *in* $\mathbf{R}^n$ *is compact.*

Proof. Let $S = \{\lambda \in \mathbf{R}^{n+1} \mid \lambda_1, \lambda_2, \dots, \lambda_{n+1} \geq 0,\ \sum_{j=1}^{n+1} \lambda_j = 1\}$, and let $f\colon S \times \mathbf{R}^n \times \mathbf{R}^n \times \cdots \times \mathbf{R}^n \to \mathbf{R}^n$ be the function

$$f(\lambda, x_1, x_2, \dots, x_{n+1}) = \sum_{j=1}^{n+1} \lambda_j x_j.$$

The function f is of course continuous, and the set S is compact, since it is closed and bounded. According to Theorem 2.4.2, every element $x \in \operatorname{cvx} A$ can be written as a convex combination $x = \sum_{j=1}^{n+1} \lambda_j a_j$ of at most $n+1$ elements $a_1, a_2, \dots, a_{n+1}$ from the set A. This means that the convex hull $\operatorname{cvx} A$ coincides with the image $f(S \times A \times A \times \cdots \times A)$ under f of the compact set $S \times A \times A \times \cdots \times A$. Since compactness is preserved by continuous functions, we conclude that the convex hull $\operatorname{cvx} A$ is compact. □

2.6 Cones

Definition. Let x be a point in $\mathbf{R}^n$ different from 0. The set

$$\overrightarrow{x} = \{\lambda x \mid \lambda \geq 0\}$$

is called the *ray* through x, or the *halfline* from the origin through x.

A *cone* X in $\mathbf{R}^n$ is a non-empty set which contains the ray through each of its points.

A cone X is in other words a non-empty set which is closed under multiplication by nonnegative numbers, i.e. which satisfies the implication

$$x \in X, \lambda \geq 0 \Rightarrow \lambda x \in X.$$

In particular, all cones contain the point 0.

We shall study convex cones. Rays and linear subspaces of $\mathbf{R}^n$ are convex cones, of course. In particular, the entire space $\mathbf{R}^n$ and the trivial subspace $\{0\}$ are convex cones. Other simple examples of convex cones are provided by the following examples.

Example 2.6.1. A closed halfspace $\{x \in \mathbf{R}^n \mid \langle c, x\rangle \geq 0\}$, which is bounded by a hyperplane through the origin, is a convex cone and is called a *conic halfspace*.

The union $\{x \in \mathbf{R}^n \mid \langle c, x\rangle > 0\} \cup \{0\}$ of the corresponding open halfspace and the origin is also a convex cone. □

Example 2.6.2. The nonnegative orthant

$$\mathbf{R}^n_+ = \{x = (x_1, \dots, x_n) \in \mathbf{R}^n \mid x_1 \geq 0, \dots, x_n \geq 0\}$$

in $\mathbf{R}^n$ is a convex cone. □

Definition. A cone that does not contain any line through 0, is called a *proper* cone.‡

That a cone X does not contain any line through 0 is equivalent to the condition

$$x, -x \in X \Rightarrow x = 0.$$

In other words, a cone X is a proper cone if and only if $X \cap (-X) = \{0\}$.

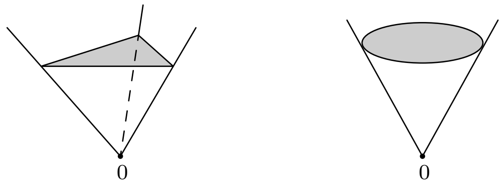

Figure 2.7. A plane cut through two proper convex cones in $\mathbf{R}^3$

Closed conic halfspaces in $\mathbf{R}^n$ are non-proper cones if $n \geq 2$. The non-negative orthant $\mathbf{R}^n_+$ is a proper cone. The cones $\{x \in \mathbf{R}^n \mid \langle c, x\rangle > 0\} \cup \{0\}$ are also proper cones.

We now give two alternative ways to express that a set is a convex cone.

Theorem 2.6.1. *The following three conditions are equivalent for a nonempty subset X of $\mathbf{R}^n$:*

(i) X is a convex cone.

(ii) X is a cone and $x + y \in X$ for all $x, y \in X$.

(iii) $\lambda x + \mu y \in X$ for alla $x, y \in X$ and all $\lambda, \mu \in \mathbf{R}_+$.

Proof. (i) $\Rightarrow$ (ii): If X is a convex cone and $x, y \in X$, then $z = \frac{1}{2}x + \frac{1}{2}y$ belongs to X because of convexity, and $x + y$ $(= 2z)$ belongs to X since X is cone.

(ii) $\Rightarrow$ (iii): If (ii) holds, $x, y \in X$ and $\lambda, \mu \in \mathbf{R}_+$, then λx and μy belong to X by the cone condition, and $\lambda x + \mu y \in X$ by additivity.

(iii) $\Rightarrow$ (i): If (iii) holds, then we conclude that X is a cone by choosing $y = x$ and $\mu = 0$, and that the cone is convex by choosing $\lambda + \mu = 1$. $\square$

Definition. A linear combination $\sum_{j=1}^m \lambda_j x_j$ of vectors $x_1, x_2, \dots, x_m$ in $\mathbf{R}^n$ is called a *conic combination* if all coefficients $\lambda_1, \lambda_2, \dots, \lambda_m$ are nonnegative.

Theorem 2.6.2. *A convex cone contains all conic combinations of its elements.*

Proof. Follows immediately by induction from the characterization of convex cones in Theorem 2.6.1 (iii). □

Cone preserving operations

The proofs of the four theorems below are analogous to the proofs of the corresponding theorems on convex sets, and they are therefore left as exercises.

Theorem 2.6.3. *Let $T\colon \mathbf{R}^n \to \mathbf{R}^m$ be a linear map.*
(i) The image $T(X)$ of a convex cone X in $\mathbf{R}^n$ is a convex cone.
(ii) The inverse image $T^{-1}(Y)$ of a convex cone in $\mathbf{R}^m$ is a convex cone.

Theorem 2.6.4. *The intersection $\bigcap_{i\in I} X_i$ of an arbitrary family of convex cones X_i in $\mathbf{R}^n$ is a convex cone.*

Theorem 2.6.5. *The Cartesian product $X \times Y$ of two convex cones X and Y is a convex cone.*

Theorem 2.6.6. *The sum $X + Y$ of two convex cones X and Y in $\mathbf{R}^n$ is a convex cone, and $-X$ is a convex cone if X is a convex cone.*

Example 2.6.3. An intersection

$$X = \bigcap_{i=1}^{m} \{x \in \mathbf{R}^n \mid \langle c_i, x\rangle \geq 0\}$$

of finitely many closed conic halfspaces is called a *polyhedral cone* or a *conic polyhedron*.

By defining C as the $m \times n$-matrix with rows c_i^T, $i = 1, 2, \dots, m$, we can write the above polyhedral cone X in a more compact way as

$$X = \{x \in \mathbf{R}^n \mid Cx \geq 0\}.$$

A polyhedral cone is in other words the solution set of a system of homogeneous linear inequalities. □

Conic hull

Definition. Let A be an arbitrary nonempty subset of $\mathbf{R}^n$. The set of all conic combinations of elements of A is called the *conic hull* of A, and it is denoted by $\operatorname{con} A$. The elements of A are called *generators* of $\operatorname{con} A$.

We extend the concept to the empty set by defining $\operatorname{con}\emptyset = \{0\}$.

Theorem 2.6.7. *The set* $\operatorname{con} A$ *is a convex cone that contains* A *as a subset, and it is the smallest convex cone with this property, i.e. if* X *is a convex cone and* $A \subseteq X$*, then* $\operatorname{con} A \subseteq X$.

Proof. A conic combination of two conic combinations of elements from A is clearly a new conic combination of elements from A, and hence $\operatorname{con} A$ is a convex cone. The inclusion $A \subseteq \operatorname{con} A$ is obvious. Since a convex cone contains all conic combinations of its elements, a convex cone X that contains A as a subset must in particular contain all conic combinations of elements from A, which means that $\operatorname{con} A$ is a subset of X. $\square$

Theorem 2.6.8. *Let* $X = \operatorname{con} A$ *be a cone in* $\mathbf{R}^n$, $Y = \operatorname{con} B$ *be a cone in* $\mathbf{R}^m$ *and* $T\colon \mathbf{R}^n \to \mathbf{R}^m$ *be a linear map. Then*

(i) $T(X) = \operatorname{con} T(A)$;

(ii) $X \times Y = \operatorname{con}\big((A \times \{0\}) \cup (\{0\} \times B)\big)$;

(iii) $X + Y = \operatorname{con}(A \cup B)$, *provided that* $m = n$ *so that the sum* $X + Y$ *is well-defined.*

Proof. (i) The cone X consists of all conic combinations $x = \sum_{j=1}^{p} \lambda_j a_j$ of elements a_j in A. For such a conic combination $Tx = \sum_{j=1}^{p} \lambda_j T a_j$. The image cone $T(X)$ thus consists of all conic combinations of the elements $Ta_j \in T(A)$, which means that $T(X) = \operatorname{con} T(A)$.

(ii) The cone $X \times Y$ consists of all pairs $(x, y) = \bigl(\sum_{j=1}^{p} \lambda_j a_j, \sum_{k=1}^{q} \mu_k b_k\bigr)$ of conic combinations of elements in A and B, respectively. But

$$(x, y) = \sum_{j=1}^{p} \lambda_j (a_j, 0) + \sum_{k=1}^{q} \mu_k (0, b_k),$$

and hence (x, y) is a conic combination of elements in $(A \times \{0\}) \cup (\{0\} \times B)$, that is (x, y) is an element of the cone $Z = \operatorname{con}\bigl((A \times \{0\}) \cup (\{0\} \times B)\bigr)$. This proves the inclusion $X \times Y \subseteq Z$.

The converse inclusion $Z \subseteq X \times Y$ follows at once from the trivial inclusion $(A \times \{0\}) \cup (\{0\} \times B) \subseteq X \times Y$, and the fact that $X \times Y$ is a cone.

(iii) A typical element of $X + Y$ has the form $\sum_{j=1}^{p} \lambda_j a_j + \sum_{k=1}^{q} \mu_k b_k$, which is a conic combination of elements in $A \cup B$. This proves the assertion. □

Finitely generated cones

Definition. A convex cone X is called *finitely generated* if $X = \operatorname{con} A$ for some finite set A.

EXAMPLE 2.6.4. The nonnegative orthant $\mathbf{R}^n_+$ is finitely generated by the standard basis $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ of $\mathbf{R}^n$. □

Theorem 2.6.8 has the following immediate corollary.

Corollary 2.6.9. *Cartesian products $X \times Y$, sums $X + Y$ and images $T(X)$ under linear maps T of finitely generated cones X and Y, are themselves finitely generated cones.*

Intersections $X \cap Y$ and inverse images $T^{-1}(Y)$ of finitely generated cones are finitely generated, too, but the proof of this fact has to wait until we have shown that finitely generated cones are polyhedral, and vice versa. See Chapter 5.

Theorem 2.6.10. *Suppose that $x \in \operatorname{con} A$, where A is a subset of $\mathbf{R}^n$. Then $x \in \operatorname{con} B$ for some linearly independent subset B of A. The number of elements in B is thus at most equal to n.*

Proof. Since x is a conic combination of elements of A, x is per definition a conic combination of finitely many elements chosen from A. Now choose a subset B of A with as few elements as possible and such that $x \in \operatorname{con} B$. We will prove that the set B is linearly independent.

If $B = \emptyset$ (i.e. if $x = 0$), then we are done, because the empty set is linearly independent. So assume that $B = \{b_1, b_2, \dots, b_m\}$, where $m \geq 1$. Then $x = \sum_{j=1}^{m} \lambda_j b_j$, where each $\lambda_j > 0$ due to the minimality assumption.

We will prove our assertion by contradiction. So, suppose that the set B is linearly dependent. Then there exist scalars $\mu_1, \mu_2, \dots, \mu_m$, at least one of them being positive, such that $\sum_{j=1}^{m} \mu_j b_j = 0$, and it follows that $x = \sum_{j=1}^{m} (\lambda_j - t\mu_j) b_j$ for all $t \in \mathbf{R}$.

Now let $t_0 = \min \lambda_j/\mu_j$, where the minimum is taken over all indices such that $\mu_j > 0$, and let j_0 be an index where the minimum is achieved. Then $\lambda_j - t_0\mu_j \geq 0$ for all indices j, and $\lambda_{j_0} - t_0\mu_{j_0} = 0$. This means that x belongs to the cone generated by the set $B \setminus \{b_{j_0}\}$, which contradicts the minimality assumption about B. Thus, B is linearly independent. □

Theorem 2.6.11. *Every finitely generated cone is closed.*

Proof. Let X be a finitely generated cone in $\mathbf{R}^n$ so that $X = \operatorname{con} A$ for some finite set A.

We first treat the case when $A = \{a_1, a_2, \dots, a_m\}$ is a linearly independent set. Then $m \leq n$, and it is possible to extend the set A, if necessary, with vectors $a_{m+1}, \dots, a_n$ to a basis for $\mathbf{R}^n$. Let $(c_1(x), c_2(x), \dots, c_n(x))$ denote the coordinates of the vector x with respect to the basis $a_1, a_2, \dots, a_n$, so that $x = \sum_{j=1}^{n} c_j(x)\, a_j$. The coordinate functions $c_j(x)$ are linear forms on $\mathbf{R}^n$.

A vector x belongs to X if and only if x is a conic combination of the first m basis vectors, and this means that

$$X = \{x \in \mathbf{R}^n \mid c_1(x) \geq 0, \dots, c_m(x) \geq 0,\ c_{m+1}(x) = \dots = c_n(x) = 0\}.$$

We conclude that X is equal to the intersection of the closed halfspaces $\{x \in \mathbf{R}^n \mid c_j(x) \geq 0\}$, $1 \leq j \leq m$, and the hyperplanes $\{x \in \mathbf{R}^n \mid c_j(x) = 0\}$, $m + 1 \leq j \leq n$. This proves that X is a closed cone in the present case and indeed a polyhedral cone.

Let us now turn to the general case. Let A be an arbitrary finite set. By the previous theorem, there corresponds to each $x \in \operatorname{con} A$ a linearly independent subset B of A such that $x \in \operatorname{con} B$, and this fact implies that $\operatorname{con} A = \bigcup \operatorname{con} B$, where the union is to be taken over all linearly independent subsets B of A. Of course, there are only finitely many such subsets, and hence $\operatorname{con} A$ is a union of finitely many cones $\operatorname{con} B$, each of them being closed, by the first part of the proof. A union of finitely many closed sets is closed. Hence, $\operatorname{con} A$ is a closed cone. □

2.7 The recession cone

The recession cone of a set consists of the directions in which the set is unbounded and in this way provides information about the behavior of the set at infinity. Here is the formal definition of the concept.

Definition. Let X be a subset of $\mathbf{R}^n$ and let v be a nonzero vector in $\mathbf{R}^n$. We say that the set X *recedes* in the direction of v and that v is a *recession vector* of X if X contains all halflines with direction vector v that start from an arbitrary point of X.

The set consisting of all recession vectors of X and the zero vector is called the *recession cone* and is denoted by $\operatorname{recc} X$. Hence, if X is a nonempty set then

$$\operatorname{recc} X = \{v \in \mathbf{R}^n \mid x + tv \in X \text{ for all } x \in X \text{ and all } t > 0\},$$

whereas $\operatorname{recc} \emptyset = \{0\}$.

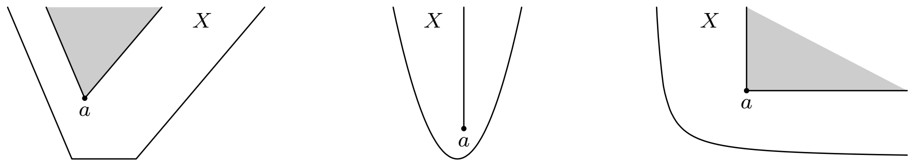

Figure 2.8. Three convex sets X and the corresponding translated recession cones $a + \operatorname{recc} X$.

Theorem 2.7.1. *The recession cone of an arbitrary set X is a convex cone and*

$$X = X + \operatorname{recc} X.$$

Proof. That $\operatorname{recc} X$ is a cone follows immediately from the very definition of recession cones, and the same holds for the inclusion $X + \operatorname{recc} X \subseteq X$. The converse inclusion $X \subseteq X + Y$ is trivially true for all sets Y containing 0 and thus in particular for the cone $\operatorname{recc} X$.

If v and w are two recession vectors of X, x is an arbitrary point in X and t is an arbitrary positive number, then first $x + tv$ belongs to X by definition, and then $x + t(v + w) = (x + tv) + tw$ belongs to X. This means that the sum $v + w$ is also a recession vector. So the recession cone has the additivity property $v, w \in \operatorname{recc} X \Rightarrow v + w \in \operatorname{recc} X$, which implies that the cone is convex according to Theorem 2.6.1. □

Example 2.7.1. Here are some simple examples of recession cones:

$$\begin{aligned}
&\operatorname{recc}(\mathbf{R}_+ \times [0,1]) = \operatorname{recc}(\mathbf{R}_+ \times \,]0,1[) = \mathbf{R}_+ \times \{0\},\\
&\operatorname{recc}(\mathbf{R}_+ \times \,]0,1[\cup \{(0,0)\}) = \{(0,0)\},\\
&\operatorname{recc}\{x \in \mathbf{R}^2 \mid x_1^2 + x_2^2 \leq 1\} = \{(0,0)\},\\
&\operatorname{recc}\{x \in \mathbf{R}^2 \mid x_2 \geq x_1^2\} = \{0\} \times \mathbf{R}_+,\\
&\operatorname{recc}\{x \in \mathbf{R}^2 \mid x_2 \geq 1/x_1,\ x_1 > 0\} = \mathbf{R}_+^2.
\end{aligned}$$ □

The computation of the recession cone of a convex set is simplified by the following theorem.

Theorem 2.7.2. *A vector v is a recession vector of a nonempty convex set X if and only if $x + v \in X$ for all $x \in X$.*

Proof. If v is a recession vector, then obviously $x + v \in X$ for all $x \in X$.

To prove the converse, assume that $x + v \in X$ for all $x \in X$, and let x be an arbitrary point in X. Then, $x + nv \in X$ for all natural numbers n, by induction, and since X is a convex set, we conclude that the closed line segment $[x, x+nv]$ lies in X for all n. Of course, this implies that $x+tv \in X$ for all positive numbers t, and hence v is a recession vector of X. $\square$

Corollary 2.7.3. *If X is a convex cone, then* $\operatorname{recc} X = X$.

Proof. The inclusion $\operatorname{recc} X \subseteq X$ holds for all sets X containing 0 and thus in particular for cones X. The converse inclusion $X \subseteq \operatorname{recc} X$ is according to Theorem 2.7.2 a consequence of the additivity property $x, v \in X \Rightarrow x+v \in X$ for convex cones. $\square$

EXAMPLE 2.7.2. $\operatorname{recc} \mathbf{R}^2_+ = \mathbf{R}^2_+$, $\operatorname{recc}(\mathbf{R}^2_{++} \cup \{(0,0)\}) = \mathbf{R}^2_{++} \cup \{(0,0)\}$. $\square$

The recession vectors of a closed convex set are characterized by the following theorem.

Theorem 2.7.4. *Let X be a nonempty closed convex set. The following three conditions are equivalent for a vector v.*

(i) v is a recession vector of X.
(ii) There exists a point $x \in X$ such that $x + nv \in X$ for all $n \in \mathbf{Z}_+$.
(iii) There exist a sequence $(x_n)_1^\infty$ of points x_n in X and a sequence $(\lambda_n)_1^\infty$ of positive numbers such that $\lambda_n \to 0$ and $\lambda_n x_n \to v$ as $n \to \infty$.

Proof. (i) $\Rightarrow$ (ii): Trivial, since $x + tv \in X$ for all $x \in X$ and all $t \in \mathbf{R}_+$, if v is a recession vector of X.

(ii) $\Rightarrow$ (iii): If (ii) holds, then condition (iii) is satisfied by the points $x_n = x + nv$ and the numbers $\lambda_n = 1/n$.

(iii) $\Rightarrow$ (i): Assume that $(x_n)_1^\infty$ and $(\lambda_n)_1^\infty$ are sequences of points in X and positive numbers such that $\lambda_n \to 0$ and $\lambda_n x_n \to v$ as $n \to \infty$, and let x be an arbitrary point in X. The points $z_n = (1-\lambda_n)x + \lambda_n x_n$ then lie in X for all sufficiently large n, and since $z_n \to x + v$ as $n \to \infty$ and X is a closed set, it follows that $x + v \in X$. Hence, v is a recession vector of X according to Theorem 2.7.2. $\square$

Theorem 2.7.5. *The recession cone* $\operatorname{recc} X$ *of a closed convex set X is a closed convex cone.*

Proof. The case $X = \emptyset$ is trivial, so assume that X is a nonempty closed convex set. To prove that the recession cone $\operatorname{recc} X$ is closed, we assume that v is a boundary point of the cone and choose a sequence $(v_n)_1^\infty$ of recession vectors that converges to v as $n \to \infty$. If x is an arbitrary point in X, then the points $x + v_n$ lie in X for each natural number n, and this implies that their limit point $x + v$ lies in X, since X is a closed set. Hence, v is a recession vector, i.e. v belongs to the recession cone $\operatorname{recc} X$. This proves that the recession cone contains all its boundary points. $\square$

Theorem 2.7.6. *Let $\{X_i \mid i \in I\}$ be a family of closed convex sets, and assume that their intersection is nonempty. Then* $\operatorname{recc}(\bigcap_{i\in I} X_i) = \bigcap_{i\in I} \operatorname{recc} X_i$.

Proof. Let x_0 be a point in $\bigcap_i X_i$. By Theorem 2.7.4, $v \in \operatorname{recc}(\bigcap_i X_i)$ if and only if $x_0 + nv$ lies in X_i for all positive integers n and all $i \in I$, and this holds if and only if $v \in \operatorname{recc} X_i$ for all $i \in I$. $\square$

The recession cone of a polyhedron is given by the following theorem.

Theorem 2.7.7. *If $X = \{x \in \mathbf{R}^n \mid Cx \geq b\}$ is a nonempty polyhedron, then* $\operatorname{recc} X = \{x \in \mathbf{R}^n \mid Cx \geq 0\}$.

Proof. The recession cone of a closed halfspace is obviously equal to the corresponding conical halfspace. The theorem is thus an immediate consequence of Theorem 2.7.6. $\square$

Note that the recession cone of a subset Y of X can be bigger than the recession cone of X. For example,

$$\operatorname{recc} \mathbf{R}^2_{++} = \mathbf{R}^2_{+} \supsetneq \mathbf{R}^2_{++} \cup \{(0,0)\} = \operatorname{recc}(\mathbf{R}^2_{++} \cup \{(0,0)\}).$$

However, this cannot occur if the superset X is closed.

Theorem 2.7.8. *(i) Suppose that X is a closed convex set and that $Y \subseteq X$. Then* $\operatorname{recc} Y \subseteq \operatorname{recc} X$.

(ii) If X is a convex set, then $\operatorname{recc}(\operatorname{rint} X) = \operatorname{recc}(\operatorname{cl} X)$.

Proof. (i) The case $Y = \emptyset$ being trivial, we assume that Y is a nonempty subset of X and that y is an arbitrary point in Y. If v is a recession vector of Y, then $y + nv$ are points in Y and thereby also in X for all natural numbers n. We conclude from Theorem 2.7.4 that v is a recession vector of X.

(ii) The inclusion $\operatorname{recc}(\operatorname{rint} X) \subseteq \operatorname{recc}(\operatorname{cl} X)$ follows from part (i), because $\operatorname{cl} X$ is a closed convex subset.

To prove the converse inclusion, assume that v is a recession vector of $\operatorname{cl} X$, and let x be an arbitrary point in $\operatorname{rint} X$. Then $x + 2v$ belongs to $\operatorname{cl} X$, so it follows from Theorem 2.5.4 that the open line segment $]x, x + 2v[$ is a subset of $\operatorname{rint} X$, and this implies that the point $x + v$ belongs to $\operatorname{rint} X$. Thus, $x \in \operatorname{rint} X \Rightarrow x + v \in \operatorname{rint} X$, and we conclude from Theorem 2.7.2 that v is a recession vector of $\operatorname{rint} X$. □

Theorem 2.7.9. *Let X be a closed convex set. Then X is bounded if and only if* $\operatorname{recc} X = \{0\}$.

Proof. Obviously, $\operatorname{recc} X = \{0\}$ if X is a bounded set. So assume that X is unbounded. Then there exists a sequence $(x_n)_1^\infty$ of points in X such that $\|x_n\| \to \infty$ as $n \to \infty$. The bounded sequence $(x_n/\|x_n\|)_1^\infty$ has a convergent subsequence, and by deleting elements, if necessary, we may as well assume that the original sequence is convergent. The limit v is a vector of norm 1, which guarantees that $v \neq 0$. With $\lambda_n = 1/\|x_n\|$ we now have a sequence of points x_n in X and a sequence of positive numbers λ_n such that $\lambda_n \to 0$ and $\lambda_n x_n \to v$ as $n \to \infty$, and this means that v is a recession vector of X according to Theorem 2.7.4. Hence, $\operatorname{recc} X \neq \{0\}$. □

Definition. Let X be an arbitary set. The intersection $\operatorname{recc} X \cap (-\operatorname{recc} X)$ is a linear subspace, which is called the *recessive subspace of* X and is denoted $\operatorname{lin} X$.

A closed convex set is called *line-free* if $\operatorname{lin} X = \{0\}$. The set X is in other words line-free if and only if $\operatorname{recc} X$ is a proper cone.

If X is a nonempty closed convex subset of $\mathbf{R}^n$ and $x \in X$ is arbitrary, then clearly

$$\operatorname{lin} X = \{x \in \mathbf{R}^n \mid a + tx \in X \text{ for all } t \in \mathbf{R}\}.$$

The image $T(X)$ of a closed convex set X under a linear map T is not necessarily closed. A counterexample is given by $X = \{x \in \mathbf{R}^2_+ \mid x_1x_2 \geq 1\}$ and the projection $T(x_1, x_2) = x_1$ of $\mathbf{R}^2$ onto the first factor, the image being $T(X) =]0\,,\infty[$. The reason why the image is not closed in this case is the fact that X has a recession vector $v = (0, 1)$ which is mapped on 0 by T.

However, we have the following general result, where $\mathcal{N}(T)$ denotes the null space of the map T, i.e. $\mathcal{N}(T) = \{x \mid Tx = 0\}$.

Theorem 2.7.10. *Let $T\colon \mathbf{R}^n \to \mathbf{R}^m$ be a linear map, let X be a closed convex subset of $\mathbf{R}^n$, and suppose that*

$$\mathcal{N}(T) \cap \operatorname{recc} X \subseteq \operatorname{lin} X.$$

The image $T(X)$ is then a closed set, and

$$\operatorname{recc} T(X) = T(\operatorname{recc} X).$$

In particular, the image $T(X)$ is closed if X is a closed convex set and $x = 0$ is the only vector in $\operatorname{recc} X$ such that $Tx = 0$.

Proof. The intersection

$$L = \mathcal{N}(T) \cap \operatorname{lin} X = \mathcal{N}(T) \cap \operatorname{recc} X$$

is a linear subspace of $\mathbf{R}^n$. Let $L^\perp$ denote its orthogonal complement. Then $X = X \cap L + X \cap L^\perp$, and

$$T(X) = T(X \cap L^\perp),$$

since $Tx = 0$ for all $x \in L$.

Let y be an arbitrary boundary point of the image $T(X)$. Due to the equality above, there exists a sequence $(x_n)_1^\infty$ of points $x_n \in X \cap L^\perp$ such that $\lim_{n\to\infty} Tx_n = y$.

We claim that the sequence $(x_n)_1^\infty$ is bounded. Assume the contrary. The sequence $(x_n)_1^\infty$ has then a subsequence $(x_{n_k})_1^\infty$ such that $\|x_{n_k}\| \to \infty$ as $k \to \infty$ and the bounded sequence $(x_{n_k}/\|x_{n_k}\|)_1^\infty$ converges. The limit v is, of course, a vector of norm 1 in the linear subspace $L^\perp$. Moreover, since $x_{n_k} \in X$ and $1/\|x_{n_k}\| \to 0$, it follows from Theorem 2.7.4 that $v \in \operatorname{recc} X$. Finally,

$$Tv = \lim_{k\to\infty} T(x_{n_k}/\|x_{n_k}\|) = \lim_{k\to\infty} \|x_{n_k}\|^{-1} Tx_{n_k} = 0 \cdot y = 0,$$

and hence v belongs to $\mathcal{N}(T)$, and thereby also to L. This means that $v \in L \cap L^\perp$, which contradicts tha fact that $v \neq 0$, since $L \cap L^\perp = \{0\}$.

The sequence $(x_n)_1^\infty$ is thus bounded. Let $(x_{n_k})_1^\infty$ be a convergent subsequence, and let $x = \lim_{k\to\infty} x_{n_k}$. The limit x lies in X since X is closed, and $y = \lim_{k\to\infty} Tx_{n_k} = Tx$, which implies that $y \in T(X)$. This proves that the image $T(X)$ contains its boundary points, so it is a closed set.

The inclusion $T(\operatorname{recc} X) \subseteq \operatorname{recc} T(X)$ holds for all sets X. To prove this, assume that v is a recession vector of X and let y be an arbitrary point in $T(X)$. Then $y = Tx$ for some point $x \in X$, and since $x + tv \in X$ for all $t > 0$ and $y + tTv = T(x + tv)$, we conclude that the points $y + tTv$ lie in $T(X)$ for all $t > 0$, and this means that Tv is a recession vector of $T(X)$.

To prove the converse inclusion $\operatorname{recc} T(X) \subseteq T(\operatorname{recc} X)$ for closed convex sets X and linear maps T fulfilling the assumptions of the theorem, we assume that $w \in \operatorname{recc} T(X)$ and shall prove that there is a vector $v \in \operatorname{recc} X$ such that $w = Tv$.

We first note that there exists a sequence $(y_n)_1^\infty$ of points $y_n \in T(X)$ and a sequence $(\lambda_n)_1^\infty$ of positive numbers such that $\lambda_n \to 0$ and $\lambda_n y_n \to w$ as $n \to \infty$. For each n, choose a point $x_n \in X \cap L^\perp$ such that $y_n = T(x_n)$.

The sequence $(\lambda_n x_n)_1^\infty$ is bounded. Because assume the contrary; then there is a subsequence such that $\|\lambda_{n_k} x_{n_k}\| \to \infty$ and $(x_{n_k}/\|x_{n_k}\|)_1^\infty$ converges to a vector z as $k \to \infty$. It follows from Theorem 2.7.4 that $z \in \operatorname{recc} X$, because the x_{n_k} are points in X and $\|x_{n_k}\| \to \infty$ as $k \to \infty$. The limit z

belongs to the subspace $L^\perp$, and since

$$\begin{aligned} Tz &= \lim_{k\to\infty} T(x_{n_k}/\|x_{n_k}\|) = \lim_{k\to\infty} T(\lambda_{n_k} x_{n_k}/\|\lambda_{n_k} x_{n_k}\|) \\ &= \lim_{k\to\infty} \lambda_{n_k} y_{n_k}/\|\lambda_{n_k} x_{n_k}\| = 0\cdot w = 0, \end{aligned}$$

we also have $z \in \mathcal{N}(T) \cap \operatorname{recc} X = L$. Hence, $z \in L \cap L^\perp = \{0\}$, which contradicts the fact that $\|z\| = 1$.

The sequence $(\lambda_n x_n)_1^\infty$, being bounded, has a subsequence that converges to a vector v, which belongs to $\operatorname{recc} X$ according to Theorem 2.7.4. Since $T(\lambda_n x_n) = \lambda_n y_n \to w$, we conclude that $Tv = w$. Hence, $w \in T(\operatorname{recc} X)$. □

Theorem 2.7.11. *Let X and Y be nonempty closed convex subsets of $\mathbf{R}^n$ and suppose that*

$$x \in \operatorname{recc} X \ \&\ y \in \operatorname{recc} Y \ \&\ x + y = 0 \quad\Rightarrow\quad x \in \operatorname{lin} X \ \&\ y \in \operatorname{lin} Y.$$

The sum $X + Y$ is then a closed convex set and

$$\operatorname{recc}(X + Y) = \operatorname{recc} X + \operatorname{recc} Y.$$

Remark. The assumption of the theorem is fulfilled if $\operatorname{recc} X$ and $-\operatorname{recc} Y$ have no common vector other than the zero vector.

Proof. Let $T\colon \mathbf{R}^n \times \mathbf{R}^n \to \mathbf{R}^n$ be the linear map $T(x, y) = x + y$. We leave as an easy exercise to show that $\operatorname{recc}(X \times Y) = \operatorname{recc} X \times \operatorname{recc} Y$ and that $\operatorname{lin}(X\times Y) = \operatorname{lin} X \times \operatorname{lin} Y$. Since $\mathcal{N}(T) = \{(x, y) \mid x+y = 0\}$, the assumption of the theorem yields the inclusion $\mathcal{N}(T) \cap \operatorname{recc}(X \times Y) \subseteq \operatorname{lin}(X \times Y)$, and it now follows from Theorem 2.7.10 that $T(X \times Y)$, i.e. the sum $X + Y$, is closed and that

$$\begin{aligned} \operatorname{recc}(X + Y) &= \operatorname{recc} T(X \times Y) = T(\operatorname{recc}(X \times Y)) = T(\operatorname{recc} X \times \operatorname{recc} Y) \\ &= \operatorname{recc} X + \operatorname{recc} Y. \end{aligned}$$
□

Corollary 2.7.12. *The sum $X + Y$ of a nonempty closed convex set X and a nonempty compact convex set Y is a closed convex set and*

$$\operatorname{recc}(X + Y) = \operatorname{recc} X.$$

Proof. The assumptions of Theorem 2.7.11 are trivially fulfilled, because $\operatorname{recc} Y = \{0\}$. □

Theorem 2.7.13. *Suppose that C is a closed convex cone and that Y is a nonempty compact convex set. Then* $\operatorname{recc}(C + Y) = C$.

Proof. The corollary is a special case of Corollary 2.7.12 since $\operatorname{recc} C = C$. □

Exercises

2.1 Prove that the set $\{x \in \mathbf{R}^2_+ \mid x_1x_2 \geq a\}$ is convex and, more generally, that the set $\{x \in \mathbf{R}^n_+ \mid x_1x_2 \cdots x_n \geq a\}$ is convex.
Hint: Use the inequality $x_i^{\lambda}y_i^{1-\lambda} \leq \lambda x_i + (1-\lambda)y_i$; see Theorem 6.4.1.

2.2 Determine the convex hull $\operatorname{cvx} A$ for the following subsets A of $\mathbf{R}^2$:
a) $A = \{(0,0),(1,0),(0,1)\}$ b) $A = \{x \in \mathbf{R}^2 \mid \|x\| = 1\}$
c) $A = \{x \in \mathbf{R}^2_+ \mid x_1x_2 = 1\} \cup \{(0,0)\}$.

2.3 Give an example of a closed set with a non-closed convex hull.

2.4 Find the inverse image $P^{-1}(X)$ of the convex set $X = \{x \in \mathbf{R}^2_+ \mid x_1x_2 \geq 1\}$ under the perspective map $P\colon \mathbf{R}^2 \times \mathbf{R}_{++} \to \mathbf{R}^2$.

2.5 Prove that the set $\{x \in \mathbf{R}^{n+1} \mid \big(\sum_{j=1}^n x_j^2\big)^{1/2} \leq x_{n+1}\}$ is a cone.

2.6 Let $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ denote the standard basis in $\mathbf{R}^n$ and let $\mathbf{e}_0 = -\sum_1^n \mathbf{e}_j$. Prove that $\mathbf{R}^n$ is generated as a cone by the $n+1$ vectors $\mathbf{e}_0, \mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$.

2.7 Prove that each conical halfspace in $\mathbf{R}^n$ is the conic hull of a set consisting of $n+1$ elements.

2.8 Prove that each closed cone in $\mathbf{R}^2$ is the conic hull of a set consisting of at most three elements.

2.9 Prove that the sum of two closed cones in $\mathbf{R}^2$ is a closed cone.

2.10 Find $\operatorname{recc} X$ and $\operatorname{lin} X$ for the following convex sets:

a) $X = \{x \in \mathbf{R}^2 \mid -x_1 + x_2 \leq 2,\ x_1 + 2x_2 \geq 2,\ x_2 \geq -1\}$
b) $X = \{x \in \mathbf{R}^2 \mid -x_1 + x_2 \leq 2,\ x_1 + 2x_2 \leq 2,\ x_2 \geq -1\}$
c) $X = \{x \in \mathbf{R}^3 \mid 2x_1 + x_2 + x_3 \leq 4,\ x_1 + 2x_2 + x_3 \leq 4\}$
d) $X = \{x \in \mathbf{R}^3 \mid x_1^2 - x_2^2 \geq 1,\ x_1 \geq 0\}$.

2.11 Let X and Y be arbitrary nonempty sets. Prove that

$$\operatorname{recc}(X \times Y) = \operatorname{recc} X \times \operatorname{recc} Y$$

and that

$$\operatorname{lin}(X \times Y) = \operatorname{lin} X \times \operatorname{lin} Y.$$

2.12 Let $P\colon \mathbf{R}^n \times \mathbf{R}_{++} \to \mathbf{R}^n$ be the perspective map. Suppose X is a convex subset of $\mathbf{R}^n$, and let $c(X) = P^{-1}(X) \cup \{(0,0)\}$.

a) Prove that $c(X)$ is a cone and, more precisely, that $c(X) = \operatorname{con}(X \times \{1\})$.
b) Find an explicit expression for the cones $c(X)$ and $\operatorname{cl}(c(X))$ if
(i) $n = 1$ and $X = [2,3]$;
(ii) $n = 1$ and $X = [2,\infty[$;
(iii) $n = 2$ and $X = \{x \in \mathbf{R}^2 \mid x_1 \geq x_2^2\}$.

c) Find $c(X)$ if $X = \{x \in \mathbf{R}^n \mid \|x\| \le 1\}$ and $\|\cdot\|$ is an arbitrary norm on $\mathbf{R}^n$.

d) Prove that $\operatorname{cl}(c(X)) = c(\operatorname{cl} X) \cup (\operatorname{recc}(\operatorname{cl} X) \times \{0\})$.

e) Prove that $\operatorname{cl}(c(X)) = c(\operatorname{cl} X)$ if and only if X is a bounded set.

f) Prove that the cone $c(X)$ is closed if and only if X is compact.

2.13 $Y = \{x \in \mathbf{R}^3 \mid x_1x_3 \ge x_2^2,\ x_3 > 0\} \cup \{x \in \mathbf{R}^3 \mid x_1 \ge 0,\ x_2 = x_3 = 0\}$ is a closed cone. (Cf. problem 2.12 b) (iii)). Put

$$Z = \{x \in \mathbf{R}^3 \mid x_1 \le 0,\ x_2 = x_3 = 0\}.$$

Show that

$$Y + Z = \{x \in \mathbf{R}^3 \mid x_3 > 0\} \cup \{x \in \mathbf{R}^3 \mid x_2 = x_3 = 0\},$$

with the conclusion that the sum of two closed cones in $\mathbf{R}^3$ is not necessarily a closed cone.

2.14 Prove that the sum $X + Y$ of an arbitrary closed set X and an arbitrary compact set Y is closed.

Chapter 3

Separation

3.1 Separating hyperplanes

Definition. Let X and Y be two sets in $\mathbf{R}^n$. We say that the hyperplane H *separarates* the two sets if the following two conditions† are satisfied:

(i) X is contained in one of the two opposite closed halvspaces defined by H and Y is contained in the other closed halfspace;
(ii) X and Y are not both subsets of the hyperplane H.

The separation is called *strict* if there exist two parallell hyperplanes to H, one on each side of H, that separates X and Y.

The hyperplane $H = \{x \mid \langle c, x\rangle = b\}$ thus separates the two sets X and Y, if $\langle c, x\rangle \le b \le \langle c, y\rangle$ for all $x \in X$ and all $y \in Y$ and $\langle c, x\rangle \ne b$ for some element $x \in X \cup Y$.

The separation is strict if there exist numbers $b_1 < b < b_2$ such that $\langle c, x\rangle \le b_1 < b_2 \le \langle c, y\rangle$ for all $x \in X$, $y \in Y$.

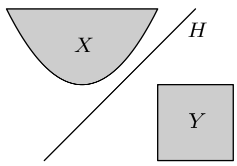

Figure 3.1. A strictly separating hyperplane H

The existence of separating hyperplanes is in a natural way connected to extreme values of linear functions.

Theorem 3.1.1. *Let X and Y be two nonempty subsets of $\mathbf{R}^n$.*

(i) There exists a hyperplane that separates X and Y if and only if there exists a vector c such that

$$\sup_{x\in X}\langle c,x\rangle \leq \inf_{y\in Y}\langle c,y\rangle \quad \textit{and} \qquad \inf_{x\in X}\langle c,x\rangle < \sup_{y\in Y}\langle c,y\rangle.$$

(ii) There exists a hyperplane that separates X and Y strictly if and only if there exists a vector c such that

$$\sup_{x\in X}\langle c,x\rangle < \inf_{y\in Y}\langle c,y\rangle.$$

Proof. A vector c that satisfies the conditions in (i) or (ii) is nonzero, of course.

Suppose that c satisfies the conditions in (i) and choose the number b so that $\sup_{x\in X}\langle c,x\rangle \leq b \leq \inf_{y\in Y}\langle c,y\rangle$. Then $\langle c,x\rangle \leq b$ for all $x \in X$ and $\langle c,y\rangle \geq b$ for all $y \in Y$. Moreover, $\langle c,x\rangle \neq b$ for some $x \in X \cup Y$ because of the second inequality in (i). The hyperplane $H = \{x \mid \langle c,x\rangle = b\}$ thus separates the two sets X and Y.

If c satisfies the condition in (ii), we choose instead b so that

$$\sup_{x\in X}\langle c,x\rangle < b < \inf_{y\in Y}\langle c,y\rangle,$$

and now conclude that the hyperplane H separates X and Y strictly.

Conversely, if the hyperplane H separates X and Y, then, by changing the signs of c and b if necessary, we may assume that $\langle c,x\rangle \leq b$ for all $x \in X$ and $\langle c,y\rangle \geq b$ for all $y \in Y$, and this implies that $\sup_{x\in X}\langle c,x\rangle \leq b \leq \inf_{y\in Y}\langle c,y\rangle$. Since H does not contain both X and Y, there are points $x_1 \in X$ and $y_1 \in Y$ with $\langle c,x_1\rangle < \langle c,y_1\rangle$, and this gives us the second inequality in (i).

If the separation is strict, then there exist two parallel hyperplanes $H_i = \{x \mid \langle c,x\rangle = b_i\}$ with $b_1 < b < b_2$ that separate X and Y. Assuming that X lies in the halfspace $\{x \mid \langle c,x\rangle \leq b_1\}$, we conclude that

$$\sup_{x\in X}\langle c,x\rangle \leq b_1 < b < b_2 \leq \inf_{y\in Y}\langle c,y\rangle,$$

i.e. the vector c satisfies the condition in (ii). □

The following simple lemma reduces the problem of separating two sets to the case when one of the sets consists of just one point.

Lemma 3.1.2. *Let X and Y be two nonempty sets.*

(i) If there exists a hyperplane that separates 0 from the set $X - Y$, then there exists a hyperplane that separates X and Y.

(ii) If there exists a hyperplane that strictly separates 0 *from the set* $X - Y$*, then there exists a hyperplane that strictly separates* X *and* Y*.*

Proof. (i) If there exists a hyperplane that separates 0 from $X - Y$, then by Theorem 3.1.1 there exists a vector c such that

$$\begin{cases} 0 = \langle c, 0\rangle \le \inf\limits_{x\in X,\, y\in Y} \langle c, x - y\rangle = \inf\limits_{x\in X} \langle c, x\rangle - \sup\limits_{y\in Y} \langle c, y\rangle \\ 0 = \langle c, 0\rangle < \sup\limits_{x\in X,\, y\in Y} \langle c, x - y\rangle = \sup\limits_{x\in X} \langle c, x\rangle - \inf\limits_{y\in Y} \langle c, y\rangle \end{cases}$$

i.e. $\sup_{y\in Y}\langle c, y\rangle \le \inf_{x\in X}\langle c, x\rangle$ and $\inf_{y\in Y}\langle c, y\rangle < \sup_{x\in X}\langle c, x\rangle$, and we conclude that there exists a hyperplane that separates X and Y.

(ii) If instead there exists a hyperplane that strictly separates 0 from $X - Y$, then there exists a vector c such that

$$0 = \langle c, 0\rangle < \inf_{x\in X,\, y\in Y} \langle c, x - y\rangle = \inf_{x\in X} \langle c, x\rangle - \sup_{y\in Y} \langle c, y\rangle$$

and it now follows that $\sup_{y\in Y}\langle c, y\rangle < \inf_{x\in X}\langle c, x\rangle$, which shows that Y and X can be strictly separated by a hyperplane. $\square$

Our next theorem is the basis for our results on separation of convex sets.

Theorem 3.1.3. *Suppose X is a convex set and that $a \notin \operatorname{cl} X$. Then there exists a hyperplane H that strictly separates a and X.*

Proof. The set $\operatorname{cl} X$ is closed and convex, and a hyperplane that strictly separates a and $\operatorname{cl} X$ will, of course, also strictly separate a and X, since X is a subset of $\operatorname{cl} X$. Hence, it suffices to prove that we can strictly separate a point a from each closed convex set that does not contain the point.

We may therefore assume, without loss of generality, that the convex set X is closed and nonempty. Define $d(x) = \|x - a\|^2$, i.e. $d(x)$ is the square of the Euclidean distance between x and a. We start by proving that the restriction of the continuous function $d(\cdot)$ to X has a minimum point.

To this end, choose a positive real number r so big that the closed ball $\overline{B}(a;r)$ intersects the set X. Then $d(x) > r^2$ for all $x \in X \setminus \overline{B}(a;r)$, and $d(x) \leq r^2$ for all $x \in X \cap \overline{B}(a;r)$. The restriction of d to the compact set $X \cap \overline{B}(a;r)$ has a minimum point $x_0 \in X \cap \overline{B}(a;r)$, and this point is clearly also a minimum point for d restricted to X, i.e. the inequality $d(x_0) \leq d(x)$ holds for all $x \in X$.

Now, let $c = a - x_0$. We claim that $\langle c, x - x_0\rangle \leq 0$ for all $x \in X$. Therefore, assume the contrary, i.e. that there is a point $x_1 \in X$ such that $\langle c, x_1 - x_0\rangle > 0$. We will prove that this assumption yields a contradiction.

Consider the points $x_t = tx_1 + (1-t)x_0$. They belong to X when $0 \leq t \leq 1$ because of convexity. Let $f(t) = d(x_t) = \|x_t - a\|^2$. Then

$$\begin{aligned} f(t) &= \|x_t - a\|^2 = \|t(x_1 - x_0) + (x_0 - a)\|^2 = \|t(x_1 - x_0) - c\|^2 \\ &= t^2\|x_1 - x_0\|^2 - 2t\,\langle c, x_1 - x_0\rangle + \|c\|^2. \end{aligned}$$

The function $f(t)$ is a quadratic polynomial in t, and its derivative at 0 satisfies $f'(0) = -2\,\langle c, y_1 - y_0\rangle < 0$. Hence, $f(t)$ is strictly decreasing in a neighbourhood of $t = 0$, which means that $d(x_t) < d(x_0)$ for all sufficiently small positive numbers t.

This is a contradiction to x_0 being the minimum point of the function and proves our assertion that $\langle c, x - x_0\rangle \leq 0$ for all $x \in X$. Consequently,

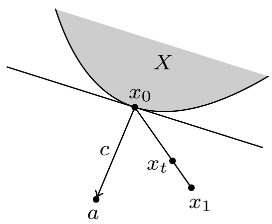

Figure 3.2. Illustration for the proof of Theorem 3.1.3.

$\langle c, x\rangle \le \langle c, x_0\rangle = \langle c, a-c\rangle = \langle c, a\rangle - \|c\|^2$ for all $x \in X$, and this implies that $\sup_{x\in X}\langle c, x\rangle < \langle c, a\rangle$. So there exists a hyperplane that strictly separates a from X according to Theorem 3.1.1. □

Definition. Let X be a subset of $\mathbf{R}^n$ and let x_0 be a point in X. A hyperplane H through x_0 is called a *supporting hyperplane* of X, if it separates x_0 and X. We then say that the hyperplane H *supports* X *at the point* x_0.

The existence of a supporting hyperplane of X at x_0 is clearly equivalent to the condition that there exists a vector c such that

$$\langle c, x_0\rangle = \inf_{x\in X}\langle c, x\rangle \qquad \text{and} \qquad \langle c, x_0\rangle < \sup_{x\in X}\langle c, x\rangle.$$

The hyperplane $\{x \mid \langle c, x\rangle = \langle c, x_0\rangle\}$ is then a supporting hyperplane.

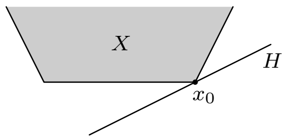

Figure 3.3. A supporting hyperplane of X at the point x_0

If a hyperplane supports the set X at the point x_0, then x_0 is necessarily a relative boundary point of X. For convex sets the following converse holds.

Theorem 3.1.4. *Suppose that X is a convex set and that $x_0 \in X$ is a relative boundary point of X. Then there exists a hyperplane H that supports X at the point x_0.*

Proof. First suppose that the dimension of X equals the dimension of the surrounding space $\mathbf{R}^n$. Since x_0 is then a boundary point of X, there exists a sequence $(x_n)_1^\infty$ of points $x_n \notin \operatorname{cl} X$ that converges to x_0 as $n \to \infty$, and by Theorem 3.1.3 there exists, for each $n \ge 1$, a hyperplane which strictly separates x_n and X. Theorem 3.1.1 thus gives us a sequence $(c_n)_1^\infty$ of vectors such that

$$\langle c_n, x_n\rangle < \langle c_n, x\rangle \qquad \text{for all } x \in X \tag{3.1}$$

and all $n \ge 1$, and we can obviously normalize the vectors c_n so that $\|c_n\| = 1$ for all n.

The unit sphere $\{x \in \mathbf{R}^n \mid \|x\| = 1\}$ is compact. Hence, by the Bolzano–Weierstrass theorem, the sequence $(c_n)_1^\infty$ has a subsequence $(c_{n_k})_{k=1}^\infty$ which converges to some vector c of length $\|c\| = 1$. Clearly $\lim_{k\to\infty} x_{n_k} = x_0$, so by going to the limit in the inequality (3.1) we conclude that $\langle c, x_0\rangle \le \langle c, x\rangle$ for all $x \in X$. The set X is therefore a subset of one of the two closed halfspaces determined by the hyperplane $H = \{x \in \mathbf{R}^n \mid \langle c, x\rangle = \langle c, x_0\rangle\}$, and X is not a subset of H, since $\dim X = n$. The hyperplane H is consequently a supporting hyperplane of X at the point x_0.

Next suppose that $\dim X < n$. Then there exists an affine subspace $a + U$ that contains X, where U is a linear subspace of $\mathbf{R}^n$ and $\dim U = \dim X$. Consider the set $Y = X + U^\perp$, where $U^\perp$ is the orthogonal complement of U. Compare with figure 3.4. Y is a "cylinder" with X as "base", and each $y \in Y$ has a unique decomposition of the form $y = x + v$ with $x \in X$ and $v \in U^\perp$.

The set Y is a convex set of dimension n with x_0 as a boundary point. By the already proven case of the theorem, there exists a hyperplane which supports Y at the point x_0, i.e. there exists a vector c such that

$$\langle c, x_0\rangle = \inf_{y\in Y}\langle c, y\rangle = \inf_{x\in X,\, v\in U^\perp}\langle c, x+v\rangle = \inf_{x\in X}\langle c, x\rangle + \inf_{v\in U^\perp}\langle c, v\rangle$$

and

$$\langle c, x_0\rangle < \sup_{y\in Y}\langle c, y\rangle = \sup_{x\in X,\, v\in U^\perp}\langle c, x+v\rangle = \sup_{x\in X}\langle c, x\rangle + \sup_{v\in U^\perp}\langle c, v\rangle.$$

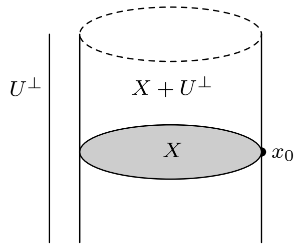

Figure 3.4. Illustration for the proof of Theorem 3.1.4.

It follows from the first equation that $\inf_{v\in U^\perp}\langle c, v\rangle$ is a finite number, and since $U^\perp$ is a vector space, this is possible if and only if $\langle c, v\rangle = 0$ for all $v \in U^\perp$. The conditions above are therefore reduced to the conditions

$$\langle c, x_0\rangle = \inf_{x\in X}\langle c, x\rangle \quad \text{and} \quad \langle c, x_0\rangle < \sup_{x\in X}\langle c, x\rangle,$$

which show that X has indeed a supporting hyperplane at x_0. □

We are now able to prove the following necessary and sufficient condition for separation of convex sets.

Theorem 3.1.5. *Two convex sets X and Y can be separated by a hyperplane if and only if their relative interiors are disjoint.*

Proof. A hyperplane that separates two sets A and B clearly also separates their closures $\operatorname{cl} A$ and $\operatorname{cl} B$ and thereby also all sets C and D that satisfy the inclusions $A \subseteq C \subseteq \operatorname{cl} A$ and $B \subseteq D \subseteq \operatorname{cl} B$.

To prove the existence of a hyperplane that separates the two convex sets X and Y provided $\operatorname{rint} X \cap \operatorname{rint} Y = \emptyset$, it hence suffices to prove that there exists a hyperplane that separates the two convex sets $A = \operatorname{rint} X$ and $B = \operatorname{rint} Y$, because $\operatorname{rint} X \subseteq X \subseteq \operatorname{cl}(\operatorname{rint} X) = \operatorname{cl} X$, and the corresponding inclusions are of course also true for Y.

Since the sets A and B are disjoint, 0 does not belong to the convex set $A - B$. Thus, the point 0 either lies in the complement of $\operatorname{cl}(A - B)$ or belongs to $\operatorname{cl}(A - B)$ and is a relative boundary point of $\operatorname{cl}(A - B)$, because

$$\begin{aligned}\operatorname{cl}(A - B) \setminus (A - B) &\subseteq \operatorname{cl}(A - B) \setminus \operatorname{rint}(A - B)\\ &= \operatorname{rbdry}(A - B) = \operatorname{rbdry}(\operatorname{cl}(A - B)).\end{aligned}$$

In the first case it follows from Theorem 3.1.3 that there is a hyperplane that strictly separates 0 and $A - B$, and in the latter case Theorem 3.1.4 gives us a hyperplane that separates 0 from the set $\operatorname{cl}(A - B)$, and thereby afortiori also 0 from $A - B$. The existence of a hyperplane that separates A and B then follows from Lemma 3.1.2.

Now, let us turn to the converse. Assume that the hyperplane H separates the two convex sets X and Y. We will prove that there is no point that is a relative interior point of both sets. To this end, let us assume that x_0 is a point in the intersection $X \cap Y$. Then, x_0 lies in the hyperplane H because X and Y are subsets of opposite closed halfspaces determined by H. According to the separability definition, at least one of the two convex sets, X say, has points that lie outside H, and this clearly implies that the affine hull $V = \operatorname{aff} X$ is not a subset of H. Hence, there are points in V from each side of H. Therefore, the intersection $V \cap B(x_0; r)$ between V and an arbitrary open ball $B(x_0; r)$ centered at x_0 also contains points from both sides of H, and consequently surely points that do not belong to X. This means that x_0 must be a relative boundary point of X.

Hence, every point in the intersection $X \cap Y$ is a relative boundary point of either of the two sets X and Y. The intersection $\operatorname{rint} X \cap \operatorname{rint} Y$ is thus empty. □

Let us now consider the possibility of strict separation. A hyperplane that strictly separates two sets obviously also strictly separates their closures, so it suffices to examine when two closed convex subsets X and Y can be strictly separated. Of course, the two sets have to be disjoint, i.e. $0 \notin X - Y$ is a necessary condition, and Lemma 3.1.2 now reduces the problem of separating X strictly from Y to the problem of separating 0 strictly from $X - Y$. So it follows at once from Theorem 3.1.3 that there exists a separating hyperplane if the set $X - Y$ is closed. This gives us the following theorem, where the sufficient conditions follow from Theorem 2.7.11 and Corollary 2.7.12.

Theorem 3.1.6. *Two disjoint closed convex sets X and Y can be strictly separated by a hyperplane if the set $X - Y$ is closed, and a sufficient condition for this to be the case is $\operatorname{recc} X \cap \operatorname{recc} Y = \{0\}$. In particular, two disjoint closed convex set can be separated strictly by a hyperplane if one of the sets is bounded.*

We conclude this section with a result that shows that proper convex cones are proper subsets of conic halfspaces. More precisely, we have:

Theorem 3.1.7. *Let $X \neq \{0\}$ be a proper convex cone in $\mathbf{R}^n$, where $n \geq 2$. Then X is a proper subset of some conic halfspace $\{x \in \mathbf{R}^n \mid \langle c, x\rangle \geq 0\}$, whose boundary $\{x \in \mathbf{R}^n \mid \langle c, x\rangle = 0\}$ does not contain X as a subset.*

Proof. The point 0 is a relative boundary point of X, because no point on the line segment $]0, -a[$ belongs to X when a is a point $\neq 0$ in X. Hence, by Theorem 3.1.4, there exists a hyperplane $H = \{x \in \mathbf{R}^n \mid \langle c, x\rangle = 0\}$ through 0 such that X lies in the closed halfspace $K = \{x \in \mathbf{R}^n \mid \langle c, x\rangle \geq 0\}$ without X being a subset of H. K is a conic halfspace, and the proper cone X must be different from K, since no conic halfspaces in $\mathbf{R}^n$ are proper cones when $n \geq 2$. □

3.2 The dual cone

To each subset A of $\mathbf{R}^n$ we associate a new subset A^+ of $\mathbf{R}^n$ by letting

$$A^+ = \{x \in \mathbf{R}^n \mid \langle a, x\rangle \geq 0 \text{ for all } a \in A\}.$$

In particular, for sets $\{a\}$ consisting of just one point we have

$$\{a\}^+ = \{x \in \mathbf{R}^n \mid \langle a, x\rangle \geq 0\},$$

which is a conic closed halfspace. For general sets A, $A^+ = \bigcap_{a\in A}\{a\}^+$, and this is an intersection of conic closed halfspaces. The set A^+ is thus in general a closed convex cone, and it is a polyhedral cone if A is a finite set.

Definition. The closed convex cone A^+ is called the *dual cone* of the set A.

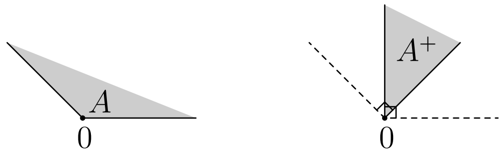

Figure 3.5. A cone A and its dual cone A^+.

The dual cone A^+ of a set A in $\mathbf{R}^n$ has an obvious geometric interpretation when $n \le 3$; it consists of all vectors that form an acute angle or are perpendicular to all vectors in A.

Theorem 3.2.1. *The following properties hold for subsets A and B of $\mathbf{R}^n$.*

(i) $A \subseteq B \;\Rightarrow\; B^+ \subseteq A^+$;

(ii) $A^+ = (\operatorname{con} A)^+$;

(iii) $A^+ = (\operatorname{cl} A)^+$.

Proof. Property (i) is an immediate consequence of the definition of the dual cone.

To prove (ii) and (iii), we first observe that

$$(\operatorname{con} A)^+ \subseteq A^+ \text{ and } (\operatorname{cl} A)^+ \subseteq A^+,$$

because of property (i) and the obvious inclusions $A \subseteq \operatorname{con} A$ and $A \subseteq \operatorname{cl} A$.

It thus only remains to prove the converse inclusions. So let us assume that $x \in A^+$. Then

$$\langle \lambda_1 a_1 + \cdots + \lambda_k a_k, x\rangle = \lambda_1\langle a_1, x\rangle + \cdots + \lambda_k\langle a_k, x\rangle \ge 0$$

for all conic combinations of elements a_i in A. This proves the implication $x \in A^+ \Rightarrow x \in (\operatorname{con} A)^+$, i.e. the inclusion $A^+ \subseteq (\operatorname{con} A)^+$.

For each $a \in \operatorname{cl} A$ there exists a sequence $(a_k)_1^\infty$ of elements in A such that $a_k \to a$ as $k \to \infty$. If $x \in A^+$, then $\langle a_k, x\rangle \ge 0$ for all k, and it follows, by passing to the limit, that $\langle a, x\rangle \ge 0$. Since $a \in \operatorname{cl} A$ is arbitrary, this proves the implication $x \in A^+ \Rightarrow x \in (\operatorname{cl} A)^+$ and the inclusion $A^+ \subseteq (\operatorname{cl} A)^+$. □

Example 3.2.1. Clearly, $(\mathbf{R}^n)^+ = \{0\}$ and $\{0\}^+ = \mathbf{R}^n$. □

Example 3.2.2. Let, as usual, $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_n$ denote the standard basis of $\mathbf{R}^n$. Then

$$\{\mathbf{e}_j\}^+ = \{x \in \mathbf{R}^n \mid \langle \mathbf{e}_j, x\rangle \ge 0\} = \{x \in \mathbf{R}^n \mid x_j \ge 0\}.$$

Since $\mathbf{R}^n_+ = \operatorname{con}\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$, it follows that

$$(\mathbf{R}^n_+)^+ = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}^+ = \bigcap_{j=1}^n \{\mathbf{e}_j\}^+ = \{x \in \mathbf{R}^n \mid x_1 \geq 0, \dots, x_n \geq 0\} = \mathbf{R}^n_+.$$

□

The bidual cone

Definition. The dual cone A^+ of a set A in $\mathbf{R}^n$ is a new set in $\mathbf{R}^n$, and we may therefore form the dual cone $(A^+)^+$ of A^+. The cone $(A^+)^+$ is called the *bidual cone* of A, and we write A^{++} instead of $(A^+)^+$.

Theorem 3.2.2. *Let A be an arbitrary set in $\mathbf{R}^n$. Then*

$$A \subseteq \operatorname{con} A \subseteq A^{++}.$$

Proof. The definitions of dual and bidual cones give us the implications

$$\begin{aligned} a \in A \quad &\Rightarrow \quad \langle x, a \rangle = \langle a, x \rangle \geq 0 \text{ for all } x \in A^+ \\ &\Rightarrow \quad a \in A^{++}, \end{aligned}$$

which show that $A \subseteq A^{++}$. Since A^{++} is a cone and $\operatorname{con} A$ is the smallest cone containing A, we conclude that $\operatorname{con} A \subseteq A^{++}$. □

Because of the previous theorem, it is natural to ask when $\operatorname{con} A = A^{++}$. Since A^{++} is a closed cone, a necessary condition for this to be the case is that the cone $\operatorname{con} A$ be closed. Our next theorem shows that this condition is also sufficient.

Theorem 3.2.3. *Let X be a convex cone. Then $X^{++} = \operatorname{cl} X$, and consequently, $X^{++} = X$ if and only if the cone X is a closed.*

Proof. It follows from the inclusion $X \subseteq X^{++}$ and the closedness of the bidual cone X^{++} that $\operatorname{cl} X \subseteq X^{++}$.

To prove the converse inclusion $X^{++} \subseteq \operatorname{cl} X$, we assume that $x_0 \notin \operatorname{cl} X$ and will prove that $x_0 \notin X^{++}$.

By Theorem 3.1.3, there exists a hyperplane that strictly separates x_0 from $\operatorname{cl} X$. Hence, there exist a vector $c \in \mathbf{R}^n$ and a real number b such that the inequality $\langle c, x \rangle \geq b > \langle c, x_0 \rangle$ holds for all $x \in X$. In particular, $t\langle c, x \rangle = \langle c, tx \rangle \geq b$ for all $x \in X$ and all numbers $t \geq 0$, since X is a cone, and this clearly implies that $b \leq 0$ and that $\langle c, x \rangle \geq 0$ for all $x \in X$. Hence, $c \in X^+$, and since $\langle c, x_0 \rangle < b \leq 0$, we conclude that $x_0 \notin X^{++}$. □

By Theorem 2.6.11, finitely generated cones are closed, so we have the following immediate corollary.

Corollary 3.2.4. *If the cone X is finitely generated, then $X^{++} = X$.*

Example 3.2.3. The dual cone of the polyhedral cone

$$X = \bigcap_{i=1}^{m} \{x \in \mathbf{R}^n \mid \langle a_i, x\rangle \geq 0\}$$

is the cone

$$X^+ = \operatorname{con}\{a_1, a_2, \dots, a_m\}.$$

This follows from the above corollary and Theorem 3.2.1, for

$$X = \{a_1, a_2, \dots, a_m\}^+ = (\operatorname{con}\{a_1, a_2, \dots, a_m\})^+.$$

If we use matrices to write the above cone X as $\{x \in \mathbf{R}^n \mid Ax \geq 0\}$, then the vector a_i corresponds to the ith column of the *transposed* matrix A^T (cf. Example 2.6.3), and the dual cone X^+ is consequently generated by the *columns* of A^T. Thus,

$$\{x \in \mathbf{R}^n \mid Ax \geq 0\}^+ = \{A^T y \mid y \in \mathbf{R}^m_+\}. \qquad \square$$

3.3 Solvability of systems of linear inequalities

Corollary 3.2.4 can be reformulated as a criterion for the solvability of systems of linear inequalities. The proof of this criterion uses the following lemma about dual cones.

Lemma 3.3.1. *Let X and Y be closed convex cones in $\mathbf{R}^n$. Then*

(i) $X \cap Y = (X^+ + Y^+)^+$*;*

(ii) $X + Y = (X^+ \cap Y^+)^+$*, provided that the cone $X + Y$ is closed.*

Proof. We have $X^+ \subseteq (X \cap Y)^+$ and $Y^+ \subseteq (X \cap Y)^+$, by Theorem 3.2.1 (i). Hence, $X^+ + Y^+ \subseteq (X \cap Y)^+ + (X \cap Y)^+ = (X \cap Y)^+$.

Another application of Theorem 3.2.1 in combination with Theorem 3.2.3 now yields $X \cap Y = (X \cap Y)^{++} \subseteq (X^+ + Y^+)^+$.

To obtain the converse inclusion we first deduce from $X^+ \subseteq X^+ + Y^+$ that $(X^+ + Y^+)^+ \subseteq X^{++} = X$, and the inclusion $(X^+ + Y^+)^+ \subseteq Y$ is of course obtained in the same way. Consequently, $(X^+ + Y^+)^+ \subseteq X \cap Y$. This completes the proof of property (i).

By replacing X and Y in (i) by the closed cones X^+ and Y^+, we obtain the equality $X^+ \cap Y^+ = (X^{++} + Y^{++})^+ = (X + Y)^+$, and since the cone $X + Y$ is assumed to be closed, we conclude that

$$X + Y = (X + Y)^{++} = (X^+ \cap Y^+)^+. \qquad \square$$

We are now ready for the promised result on the solvability of certain systems of linear inequalities, a result that will be used in our proof of the duality theorem in linear programming.

Theorem 3.3.2. *Let U be a finitely generated cone in $\mathbf{R}^n$, V be a finitely generated cone in $\mathbf{R}^m$, A be an $m \times n$-matrix and c be an $n \times 1$-matrix. Then the system*

$$\text{(S)} \qquad \begin{cases} Ax \in V^+ \\ x \in U^+ \\ c^T x < 0 \end{cases}$$

has a solution x if and only if the system

$$\text{(S}^*\text{)} \qquad \begin{cases} c - A^T y \in U \\ y \in V \end{cases}$$

has no solution y.

Proof. The system (S^*) clearly has a solution if and only if $c \in (A^T(V) + U)$, and consequently, there is no solution if and only if $c \notin (A^T(V) + U)$. Therefore, it is worthwhile to take a closer look at the cone $A^T(V) + U$.

The cones $A^T(V)$, U and $A^T(V) + U$ are closed, since they are finitely generated. We may therefore apply Lemma 3.3.1 with

$$A^T(V) + U = \big(A^T(V)^+ \cap U^+\big)^+$$

as conclusion. The condition $c \notin (A^T(V) + U)$ is now seen to be equivalent to the existence of a vector $x \in A^T(V)^+ \cap U^+$ satisfying the inequality $c^T x < 0$, i.e. to the existence of an x such that

$$(\dagger) \qquad \begin{cases} \quad x \in A^T(V)^+ \\ \quad x \in U^+ \\ c^T x < 0 \,. \end{cases}$$

It now only remains to translate the condition $x \in A^T(V)^+$; it is equivalent to the condition

$$\langle y, Ax \rangle = \langle A^T y, x \rangle \geq 0 \qquad \text{for all } y \in V,$$

i.e. to $Ax \in V^+$. The two systems (†) and (S) are therefore equivalent, and this observation completes the proof. □

By choosing $U = \{0\}$ and $V = \mathbf{R}^m_+$ with dual cones $U^+ = \mathbf{R}^n$ and $V^+ = \mathbf{R}^m_+$, we get the following special case of Theorem 3.3.2.

Corollary 3.3.3 (Farkas's lemma). *Let A be an $m \times n$-matrix and c be an $n \times 1$-matrix, and consider the two systems:*

$$\text{(S)} \quad \begin{cases} Ax \geq 0 \\ c^T x < 0 \end{cases} \qquad \textit{and} \qquad (\mathrm{S}^*) \quad \begin{cases} A^T y = c \\ \quad y \geq 0 \end{cases}$$

The system (S) *has a solution if and only if the system* (S^*) *has no solution.*

EXAMPLE 3.3.1. The system

$$\begin{cases} x_1 - x_2 + 2x_3 \geq 0 \\ -x_1 + x_2 - x_3 \geq 0 \\ 2x_1 - x_2 + 3x_3 \geq 0 \\ 4x_1 - x_2 + 10x_3 < 0 \end{cases}$$

has no solution, because the dual system

$$\begin{cases} y_1 - y_2 + 2y_3 = 4 \\ -y_1 + y_2 - y_3 = -1 \\ 2y_1 - y_2 + 3y_3 = 10 \end{cases}$$

has a nonnegative solution $y = (3, 5, 3)$. □

Example 3.3.2. The system

$$\begin{cases} 2x_1 + x_2 - x_3 \geq 0 \\ x_1 + 2x_2 - 2x_3 \geq 0 \\ x_1 - x_2 + x_3 \geq 0 \\ x_1 - 4x_2 + 4x_3 < 0 \end{cases}$$

is solvable, because the solutions of the dual system

$$\begin{cases} 2y_1 + y_2 + y_3 = 1 \\ y_1 + 2y_2 - y_3 = -4 \\ -y_1 - 2y_2 + y_3 = 4 \end{cases}$$

are of the form $y = (2-t, -3+t, t)$, $t \in \mathbf{R}$, and none of those is nonnegative since $y_1 < 0$ for $t > 2$ and $y_2 < 0$ for $t < 3$. □

The following generalization of Example 3.2.3 will be needed in Chapter 10 in Part II.

Theorem 3.3.4. *Let $a_1, a_2, \dots, a_m$ be vectors in $\mathbf{R}^n$, and let I, J be a partition of the index set $\{1, 2, \dots, m\}$, i.e. $I \cap J = \emptyset$ and $I \cup J = \{1, 2, \dots, m\}$. Let*

$$X = \bigcap_{i \in I} \{x \in \mathbf{R}^n \mid \langle a_i, x\rangle \geq 0\} \cap \bigcap_{i \in J} \{x \in \mathbf{R}^n \mid \langle a_i, x\rangle > 0\},$$

and suppose that $X \neq \emptyset$. Then

$$X^+ = \operatorname{con}\{a_1, a_2, \dots, a_m\}.$$

Proof. Let

$$Y = \bigcap_{i=1}^{m} \{x \in \mathbf{R}^n \mid \langle a_i, x\rangle \geq 0\}.$$

The set Y is closed and contains X, and we will prove that $Y = \operatorname{cl} X$ by showing that every neighborhood of an arbitrary point $y \in Y$ contains points from X.

So, fix a point $x_0 \in X$, and consider the points $y + tx_0$ for $t > 0$. These points lie in X, for

$$\langle a_i, y + tx_0\rangle = \langle a_i, y\rangle + t\langle a_i, x_0\rangle = \begin{cases} \geq 0 & \text{if } i \in I \\ > 0 & \text{if } i \in J, \end{cases}$$

and since $y + tx_0 \to y$ as $t \to 0$, there are indeed points in X arbitrarily close to y.

Hence $X^+ = (\operatorname{cl} X)^+ = Y^+$, by Theorem 3.2.1, and the conclusion of the theorem now follows from the result in Example 3.2.3. □

How do we decide whether the set X in Theorem 3.3.4 is nonempty? If just one of the m linear inequalities that define X is strict (i.e. if the index set J consists of one element), then Farkas's lemma gives a necessary and sufficient condition for X to be nonempty. A generalization to the general case reads as follows.

Theorem 3.3.5. *The set X in Theorem 3.3.4 is nonempty if and only if*

$$\begin{cases} \sum_{i=1}^{m} \lambda_i a_i = 0 \\ \lambda_i \geq 0 \ \textit{for all } i \end{cases} \quad \Rightarrow \quad \lambda_i = 0 \ \textit{for all } i \in J.$$

Proof. Let the vectors $\hat{a}_i$ in $\mathbf{R}^{n+1}$ $(= \mathbf{R}^n \times \mathbf{R})$ be defined by

$$\hat{a}_i = \begin{cases} (a_i, 0) & \text{if } i \in I \\ (a_i, 1) & \text{if } i \in J. \end{cases}$$

Write $\tilde{x} = (x, x_{n+1})$, and let $\tilde{X}$ be polyhedral cone

$$\tilde{X} = \bigcap_{i=1}^{m} \{\tilde{x} \in \mathbf{R}^{n+1} \mid \langle \hat{a}_i, \tilde{x} \rangle \geq 0\} = (\operatorname{con}\{\hat{a}_1, \dots, \hat{a}_m\})^+.$$

Since

$$\langle \hat{a}_i, (x, x_{n+1}) \rangle = \begin{cases} \langle a_i, x \rangle & \text{if } i \in I \\ \langle a_i, x \rangle + x_{n+1} & \text{if } i \in J, \end{cases}$$

and $\langle a_i, x \rangle > 0$ for all $i \in J$ if and only if there exists a negative real number x_{n+1} such that $\langle a_i, x \rangle + x_{n+1} \geq 0$ for all $i \in J$, we conclude that the point x lies in X if and only if there exists a negative number x_{n+1} such that $\langle \hat{a}_i, (x, x_{n+1}) \rangle \geq 0$ for all i, i.e. if and only if there exists a negative number x_{n+1} such that $(x, x_{n+1}) \in \tilde{X}$. This is equivalent to saying that the set X is empty if and only if the implication $\tilde{x} \in \tilde{X} \Rightarrow x_{n+1} \geq 0$ is true, i.e. if and only if $\tilde{X} \subseteq \mathbf{R}^n \times \mathbf{R}_+$. Using the results on dual cones in Theorems 3.2.1 and 3.2.3 we thus obtain the following chain of equivalences:

$$\begin{aligned} X = \emptyset \;\; &\Leftrightarrow \;\; \tilde{X} \subseteq \mathbf{R}^n \times \mathbf{R}_+ \\ &\Leftrightarrow \;\; \{0\} \times \mathbf{R}_+ = (\mathbf{R}^n \times \mathbf{R}_+)^+ \subseteq \tilde{X}^+ = \operatorname{con}\{\hat{a}_1, \hat{a}_2, \dots, \hat{a}_m\} \\ &\Leftrightarrow \;\; (0, 1) \in \operatorname{con}\{\hat{a}_1, \hat{a}_2, \dots, \hat{a}_m\} \\ &\Leftrightarrow \;\; \text{there are numbers } \lambda_i \geq 0 \text{ such that } \sum_{i=1}^{m} \lambda_i a_i = 0 \text{ and } \sum_{i \in J} \lambda_i = 1 \\ &\Leftrightarrow \;\; \text{there are numbers } \lambda_i \geq 0 \text{ such that } \sum_{i=1}^{m} \lambda_i a_i = 0 \text{ and } \lambda_i > 0 \text{ for} \\ &\qquad \text{some } i \in J. \end{aligned}$$

(The last equivalence holds because of the homogenouity of the condition $\sum_{i=1}^{m} \lambda_i a_i = 0$. If the condition is fulfilled for a set of nonnegative numbers λ_i with $\lambda_i > 0$ for at least one $i \in J$, then we can certainly arrange so that $\sum_{i \in J} \lambda_i = 1$ by multiplying with a suitable constant.)

Since the first and the last assertion in the above chain of equivalences are equivalent, so are their negations, and this is the statement of the theorem. □

The following corollary is an immediate consequence of Theorem 3.3.5.

Corollary 3.3.6. *The set X in Theorem 3.3.4 is nonempty if the vectors $a_1, a_2, \dots, a_m$ are linearly independent.*

The following equivalent matrix version of Theorem 3.3.5 is obtained by considering the vectors a_i, $i \in I$ and a_i, $i \in J$ in Theorems 3.3.4 and 3.3.5 as rows in two matrices A and C, respectively.

Theorem 3.3.7. *Let A be a $p \times n$-matrix and C be $q \times n$-matris. Then exactly one of the two dual systems*

$$\begin{cases} Ax \geq 0 \\ Cx > 0 \end{cases} \quad \text{and} \quad \begin{cases} A^T y + C^T z = 0 \\ y, z \geq 0,\ z \neq 0 \end{cases}$$

has a solution.

Theorem 3.3.7 will be generalized in Chapter 6.5, where we prove a theorem on the solvability of systems of convex and affine inequalities.

Exercises

3.1 Find two disjoint closed convex sets in $\mathbf{R}^2$ that are not strictly separable by a hyperplane (i.e. by a line in $\mathbf{R}^2$).

3.2 Let X be a convex proper subset of $\mathbf{R}^n$. Show that X is an intersection of closed halfspaces if X is closed, and an intersection of open halfspaces if X is open.

3.3 Prove the following converse of Lemma 3.1.2: If two sets X and Y are (strictly) separable, then $X - Y$ and 0 are (strictly) separable.

3.4 Find the dual cones of the following cones in $\mathbf{R}^2$:

a) $\mathbf{R}_+ \times \{0\}$ b) $\mathbf{R} \times \{0\}$ c) $\mathbf{R} \times \mathbf{R}_+$ d) $(\mathbf{R}_{++} \times \mathbf{R}_{++}) \cup \{(0,0)\}$

e) $\{x \in \mathbf{R}^2 \mid x_1 + x_2 \geq 0,\ x_2 \geq 0\}$

3.5 Prove for arbitrary sets X and Y that $(X \times Y)^+ = X^+ \times Y^+$.

3.6 Determine the cones X, X^+ and X^{++}, if $X = \operatorname{con} A$ and

a) $A = \{(1,0), (1,1), (-1,1)\}$ b) $A = \{(1,0), (-1,1), (-1,-1)\}$

c) $A = \{x \in \mathbf{R}^2 \mid x_1 x_2 = 1,\ x_1 > 0\}$.

3.7 Let $A = \{a_1, a_2, \ldots, a_m\}$ be a subset of $\mathbf{R}^n$, and suppose $0 \notin A$. Prove that the following three conditions are equivalent:

(i) $\operatorname{con} A$ is a proper cone.

(ii) $\sum_{j=1}^m \lambda_j a_j = 0$, $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m) \geq 0 \;\Rightarrow\; \lambda = 0$.

(iii) There is a vector c such that $\langle c, a \rangle > 0$ for all $a \in A$.

3.8 Is the following system consistent?

$$\begin{cases} x_1 - 2x_2 - 7x_3 \geq 0 \\ 5x_1 + x_2 - 2x_3 \geq 0 \\ x_1 + 2x_2 + 5x_3 \geq 0 \\ 18x_1 + 5x_2 - 3x_3 < 0 \end{cases}$$

3.9 Show that

$$\begin{cases} x_1 + x_2 - x_3 \geq 2 \\ x_1 - x_2 \geq 1 \\ x_1 + x_3 \geq 3 \end{cases} \quad\Rightarrow\quad 6x_1 - 2x_2 + x_3 \geq 11.$$

3.10 For which values of the parameter $\alpha \in \mathbf{R}$ is the system

$$\begin{cases} x_1 + x_2 + \alpha x_3 \geq 0 \\ x_1 + \alpha x_2 + x_3 \geq 0 \\ \alpha x_1 + x_2 + x_3 \geq 0 \\ x_1 + \alpha x_2 + \alpha^2 x_3 < 0 \end{cases}$$

solvable?

3.11 Let A be an $m \times n$-matrix. Prove that exactly one of the two systems (S) and (S*) has a solution if

a) (S) $\begin{cases} Ax = 0 \\ x \geq 0 \\ x \neq 0 \end{cases}$ and (S*) $A^T y > 0$

b) (S) $\begin{cases} Ax = 0 \\ x > 0 \end{cases}$ and (S*) $\begin{cases} A^T y \geq 0 \\ A^T y \neq 0\,. \end{cases}$

3.12 Prove that the following system of linear inequalities is solvable:

$$\begin{cases} Ax = 0 \\ x \geq 0 \\ A^T y \geq 0 \\ A^T y + x > 0\,. \end{cases}$$

Chapter 4

More on convex sets

4.1 Extreme points and faces

Extreme point

Polyhedra, like the one in figure 4.1, have vertices. A vertex is characterized by the fact that it is not an interior point of any line segment that lies entirely in the polyhedron. This property is meaningful for arbitrary convex sets.

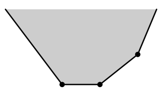

Figure 4.1. A polyhedron with vertices.

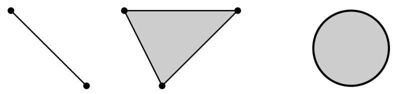

Figure 4.2. Extreme points of a line segment, a triangle and a circular disk.

Definition. A point x in a convex set X is called an *extreme point* of the set if it does not lie in any open line segment joining two points of X, i.e. if

$$a_1, a_2 \in X \ \& \ a_1 \neq a_2 \quad \Rightarrow \quad x \notin]a_1, a_2[.$$

The set of all extreme points of X will be denoted by $\operatorname{ext} X$.

A point in the relative interior of a convex set is clearly never an extreme point, except when the convex set consists just one point.† With an exception for this trivial case, $\operatorname{ext} X$ is consequently a subset of the relative boundary of X. In particular, open convex sets have no extreme points.

EXAMPLE 4.1.1. The two endpoints are the extreme points of a closed line segment. The three vertices are the extreme points of a triangle. All points on the boundary $\{x \mid \|x\|_2 = 1\}$ are extreme points of the Euclidean closed unit ball $\overline{B}(0;1)$ in $\mathbf{R}^n$. □

Extreme ray

The extreme point concept is of no interest for convex cones, because non-proper cones have no extreme points, and proper cones have 0 as their only extreme point. Instead, for cones the correct extreme concept is about rays, and in order to define it properly we first have to define what it means for a ray to lie between to rays.

Definition. We say that the ray $R = \overrightarrow{a}$ *lies between* the two rays $R_1 = \overrightarrow{a_1}$ and $R_2 = \overrightarrow{a_2}$ if the two vectors a_1 and a_2 are linearly independent and there exist two positive numbers λ_1 and λ_2 so that $a = \lambda_1 a_1 + \lambda_2 a_2$.

It is easy to convince oneself that the concept "lie between" only depends on the rays R, R_1 and R_2, and not on the vectors a, a_1 and a_2 chosen to represent them. Furthermore, a_1 and a_2 are linearly independent if and only if the rays R_1 and R_2 are different and not opposite to each other, i.e. if and only if $R_1 \neq \pm R_2$.

Definition. A ray R in a convex cone X is called an *extreme ray* of the cone if the following two conditions are satisfied:

(i) the ray R does not lie between any rays in the cone X;

(ii) the opposite ray $-R$ does not lie in X.

The set of all extreme rays of X is denoted by $\operatorname{exr} X$.

The second condition (ii) is automatically satisfied for all proper cones, and it implies, as we shall see later (Theorem 4.2.4), that non-proper cones have no extreme rays.

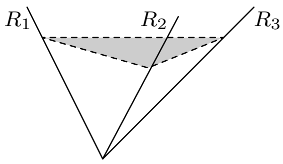

Figure 4.3. A polyhedral cone in $\mathbf{R}^3$ with three extreme rays.

It follows from the definition that no extreme ray of a convex cone with dimension greater than 1 can pass through a relative interior point of the cone. The extreme rays of a cone of dimension greater than 1 are in other words subsets of the relative boundary of the cone.

EXAMPLE 4.1.2. The extreme rays of the four subcones of $\mathbf{R}$ are as follows: $\operatorname{exr}\{0\} = \operatorname{exr} \mathbf{R} = \emptyset$, $\operatorname{exr} \mathbf{R}_+ = \mathbf{R}_+$ and $\operatorname{exr} \mathbf{R}_- = \mathbf{R}_-$.

The non-proper cone $\mathbf{R} \times \mathbf{R}_+$ in $\mathbf{R}^2$ (the ”upper halfplane”) has no extreme rays, since the two boundary rays $\mathbf{R}_+ \times \{0\}$ and $\mathbf{R}_- \times \{0\}$ are disqualified by condition (ii) of the extreme ray definition. □

Face

Definition. A subset F of a convex set X is called a *proper face* of X if $F = X \cap H$ for some supporting hyperplane H of X. In addition, the set X itself and the empty set $\emptyset$ are called *non-proper faces* of X.‡

The reason for including the set itself and and the empty set among the faces is that it simplifies the wording of some theorems and proofs.

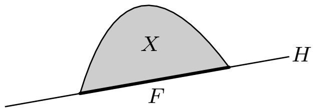

Figure 4.4. A convex set X with F as one its faces.

The faces of a convex set are obviously convex sets. And the proper faces of a convex cone are cones, since the supporting hyperplanes of a cone must pass through the origin and thus be linear subspaces.

EXAMPLE 4.1.3. Every point on the boundary $\{x \mid \|x\|_2 = 1\}$ is a face of the closed unit ball $\overline{B}(0;1)$, because the tangent plane at a boundary point is a supporting hyperplane and does not intersect the unit ball in any other point. □

EXAMPLE 4.1.4. A cube in $\mathbf{R}^3$ has 26 proper faces: 8 vertices, 12 edges and 6 sides. □

Theorem 4.1.1. *The relative boundary of a closed convex set X is equal to the union of all proper faces of X.*

Proof. We have to prove that $\operatorname{rbdry} X = \bigcup F$, where the union is taken over all proper faces F of X. So suppose that $x_0 \in F$, where $F = X \cap H$ is a proper face of X, and H is a supporting hyperplane. Since H supports X at x_0, and since, by definition, X is not contained in H, it follows that x_0 is a relative boundary point of X. This proves the inclusion $\bigcup F \subseteq \operatorname{rbdry} X$.

Conversely, if x_0 is a relative boundary point of X, then there exists a hyperplane H that supports X at x_0, and this means that x_0 lies in the proper face $X \cap H$. □

Theorem 4.1.2. *The intersection of two faces of a convex set is a face of the set.*

Proof. Let F_1 and F_2 be two faces of the convex set X, and let $F = F_1 \cap F_2$. That F is a face is trivial if the two faces F_1 and F_2 are identical, or if they are disjoint, or if one of them is non-proper.

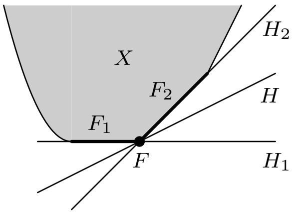

Figure 4.5. Illustration for the proof of Theorem 4.1.2.

So suppose that the two faces F_1 and F_2 are distinct and proper, i.e. that they are of the form $F_i = X \cap H_i$, where H_1 och H_2 are distinct supporting hyperplanes of the set X, and $F \neq \emptyset$. Let

$$H_i = \{x \in \mathbf{R}^n \mid \langle c_i, x\rangle = b_i\},$$

where the normal vectors c_i of the hyperplanes are chosen so that X lies in the two halfspaces $\{x \in \mathbf{R}^n \mid \langle c_i, x\rangle \leq b_i\}$, and let $x_1 \in X$ be a point satisfying the condition $\langle c_1, x_1\rangle < b_1$.

The hyperplanes H_1 and H_2 must be non-parallel, since $X \cap H_1 \cap H_2 = F \neq \emptyset$. Hence $c_2 \neq -c_1$, and we obtain a new hyperplane

$$H = \{x \in \mathbf{R}^n \mid \langle c, x \rangle = b\}$$

by defining $c = c_1 + c_2$ and $b = b_1 + b_2$. We will show that H is a supporting hyperplane of X and that $F = X \cap H$, which proves our claim that the intersection F is a face of X.

For all $x \in X$, we have the inequality

$$\langle c, x \rangle = \langle c_1, x \rangle + \langle c_2, x \rangle \leq b_1 + b_2 = b,$$

and the inequality is strict for the particular point $x_1 \in X$, since

$$\langle c, x_1 \rangle = \langle c_1, x_1 \rangle + \langle c_2, x_1 \rangle < b_1 + b_2 = b.$$

So X lies in one of the two closed halfspaces determined by H without being a subset of H. Moreover, for all $x \in F = X \cap H_1 \cap H_2$,

$$\langle c, x \rangle = \langle c_1, x \rangle + \langle c_2, x \rangle = b_1 + b_2 = b.$$

which implies that H is a supporting hyperplane of X and that $F \subseteq X \cap H$.

Conversely, if $x \in X \cap H$, then $\langle c_1, x \rangle \leq b_1$, $\langle c_2, x \rangle \leq b_2$ and

$$\langle c_1, x \rangle + \langle c_2, x \rangle = b_1 + b_2,$$

and this implies that $\langle c_1, x \rangle = b_1$ and $\langle c_2, x \rangle = b_2$. Hence, $x \in X \cap H_1 \cap H_2 = F$. This proves the inclusion $X \cap H \subseteq F$. $\square$

Theorem 4.1.3. *(i) Suppose F is a face of the convex set X. A point x in F is an extreme point of F if and only if x is an extreme point of X.*

(ii) Suppose F is a face of the convex cone X. A ray R in F is an extreme ray of F if and only if R is an extreme ray of X.

Proof. Since the assertions are trivial for non-proper faces, we may assume that $F = X \cap H$ for some supporting hyperplane H of X.

No point in a hyperplane lies in the interior of a line segment whose endpoints both lie in the same halfspace, unless both endpoints lie in the hyperplane, i.e. unless the line segment lies entirely in the hyperplane.

Analogously, no ray in a hyperplane H (through the origin) lies between two rays in the same closed halfspace determined by H, unless both these rays lie in the hyperplane H. And the opposite ray $-R$ of a ray R in a hyperplane clearly lies in the same hyperplane.

(i) If $x \in F$ is an interior point of some line segment with both endpoints belonging to X, then x is in fact an interior point of a line segment whose endpoints both belong to F. This proves the implication

$$x \notin \operatorname{ext} X \Rightarrow x \notin \operatorname{ext} F.$$

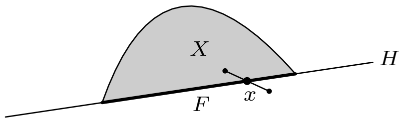

Figure 4.6. The two endpoints of an open line segment that intersects the hyperplane H must either both belong to H or else lie in opposite open halfspaces.

The converse implication is trivial, since every line segment in F is a line segment i X. Hence, $x \notin \operatorname{ext} X \Leftrightarrow x \notin \operatorname{ext} F$, and this is of course equivalent to assertion (i).

(ii) Suppose R is a ray in F and that R is not an extreme ray of the cone X. Then there are two possibilities: R lies between two rays R_1 and R_2 in X, or the opposite ray $-R$ lies in X. In the first case, R_1 and R_2 will necessarily lie in F, too. In the second case, the ray $-R$ will lie in F. Thus, both cases lead to the conclusion that R is not an extreme ray of the cone F, and this proves the implication $R \notin \operatorname{exr} X \Rightarrow R \notin \operatorname{exr} F$.

The converse implication is again trivial, and this observation concludes the proof of assertion (ii). □

4.2 Structure theorems for convex sets

Theorem 4.2.1. *Let X be a line-free closed convex set with* $\dim X \geq 2$. *Then*

$$X = \operatorname{cvx}(\operatorname{rbdry} X).$$

Proof. Let $n = \dim X$. By identifying the affine hull of X with $\mathbf{R}^n$, we may without loss of generality assume that X is a subset of $\mathbf{R}^n$ of full dimension, so that $\operatorname{rbdry} X = \operatorname{bdry} X$. To prove the theorem it is now enough to prove that every point in X lies in the convex hull of the boundary $\operatorname{bdry} X$, because the inclusionen $\operatorname{cvx}(\operatorname{bdry} X) \subseteq X$ is trivially true.

The recession cone $C = \operatorname{recc} X$ is a proper cone, since X is supposed to be line-free. Hence, there exists a closed halfspace

$$K = \{x \in \mathbf{R}^n \mid \langle c, x\rangle \geq 0\},$$

which contains C as a proper subset, by Theorem 3.1.7. Since C is a closed cone, we conclude that the corresponding open halfspace

$$K_+ = \{x \in \mathbf{R}^n \mid \langle c, x\rangle > 0\}.$$

contains a vector v that does not belong to C. The opposite vector $-v$, which lies in the opposite open halfspace, does not belong to C, either. Compare figure 4.7.

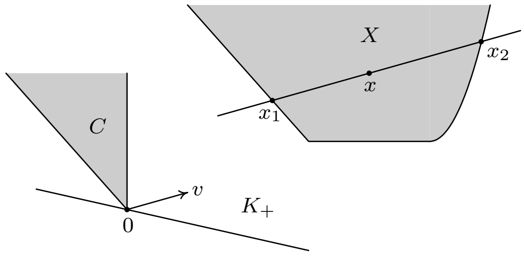

Figure 4.7. An illustration for the proof of Theorem 4.2.1.

We have produced two opposite vectors $\pm v$, both lying outside the recession cone. The two opposite halflines $x + \vec{v}$ and $x - \vec{v}$ from a point $x \in X$ therefore both intersect the complement of X.

The intersection between X and the line through x with direction vector v, which is a closed convex set, is thus either a closed line segment $[x_1, x_2]$ containing x and with endpoints belonging to the boundary of X, or the singleton set $\{x\}$ with x belonging to the boundary of X. In the first case, x is a convex combination of the boundary points x_1 and x_2. So, x lies in the convex hull of the boundary in both cases. This completes the proof. □

It is now possible to give a complete description of line-free closed convex sets in terms of extreme points and recession cones.

Theorem 4.2.2. *A nonempty closed convex set X has extreme points if and only if X is line-free, and if X is line-free, then*

$$X = \operatorname{cvx}(\operatorname{ext} X) + \operatorname{recc} X.$$

Proof. First suppose that the set X is not line-free. Its recessive subspace will then, by definition, contain a nonzero vector y, and this implies that the two points $x \pm y$ lie in X for each point $x \in X$. Therefore, x being the midpoint of the line segment $]x - y, x + y[$, is not an extreme point. This proves that the set $\operatorname{ext} X$ of extreme points is empty.

Next suppose that X is line-free. We claim that $\operatorname{ext} X \neq \emptyset$ and that $X = \operatorname{cvx}(\operatorname{ext} X) + \operatorname{recc} X$, and we will prove this by induction on the dimension of the set X.

Our claim is trivially true for zero-dimensional sets X, i.e. sets consisting of just one point. If $\dim X = 1$, then either X is a halfline $a + \overrightarrow{v}$ with one extreme point a and recession cone equal to $\overrightarrow{v}$, or a line segment $[a, b]$ with two extreme points a, b and recession cone equal to $\{0\}$, and the equality in the theorem is clearly satisfied in both cases.

Now assume that $n = \dim X \geq 2$ and that our claim is true for all line-free closed convex sets X with dimension less than n. By Theorems 4.1.1 and 4.2.1,

$$X = \operatorname{cvx}\bigl(\bigcup F\bigr),$$

where the union is taken over all proper faces F of X. Each proper face F is a nonempty line-free closed convex subset of a supporting hyperplane H and has a dimension which is less than or equal to $n - 1$. Therefore, $\operatorname{ext} F \neq \emptyset$ and

$$F = \operatorname{cvx}(\operatorname{ext} F) + \operatorname{recc} F,$$

by our induction hypothesis.

Since $\operatorname{ext} F \subseteq \operatorname{ext} X$ (by Theorem 4.1.3), it follows that $\operatorname{ext} X \neq \emptyset$. Moreover, $\operatorname{recc} F$ is a subset of $\operatorname{recc} X$, so we have the inclusion

$$F \subseteq \operatorname{cvx}(\operatorname{ext} X) + \operatorname{recc} X$$

for each face F. The union $\bigcup F$ is consequently included in the convex set $\operatorname{cvx}(\operatorname{ext} X) + \operatorname{recc} X$. Hence

$$X = \operatorname{cvx}\bigl(\bigcup F\bigr) \subseteq \operatorname{cvx}(\operatorname{ext} X) + \operatorname{recc} X \subseteq X + \operatorname{recc} X = X,$$

so $X = \operatorname{cvx}(\operatorname{ext} X) + \operatorname{recc} X$, and this completes the induction and the proof of the theorem. $\square$

The recession cone of a compact set is equal to the null cone, and the following result is therefore an immediate corollary of Theorem 4.2.2.

Corollary 4.2.3. *Each nonempty compact convex set has extreme points and is equal to the convex hull of its extreme points.*

We shall now formulate and prove the anaologue of Theorem 4.2.2 for convex cones, and in order to simplify the notation we shall use the following convention: If $\mathcal{A}$ is a family of rays, we let $\operatorname{con}\mathcal{A}$ denote the cone

$$\operatorname{con}\Big(\bigcup_{R\in\mathcal{A}} R\Big),$$

i.e. $\operatorname{con}\mathcal{A}$ is the cone that is generated by the vectors on the rays in the family $\mathcal{A}$. If we choose a nonzero vector a_R on each ray $R \in \mathcal{A}$ and let $A = \{a_R \mid R \in \mathcal{A}\}$, then of course $\operatorname{con}\mathcal{A} = \operatorname{con} A$.

The cone $\operatorname{con}\mathcal{A}$ is clearly finitely generated if $\mathcal{A}$ is a finite family of rays, and we obtain a set of generators by choosing one nonzero vector from each ray.

Theorem 4.2.4. *A closed convex cone X has extreme rays if and only if the cone is proper and not equal to the null cone $\{0\}$. If X is a proper closed convex cone, then*

$$X = \operatorname{con}(\operatorname{exr} X).$$

Proof. First suppose that the cone X is not proper, and let $R = \overrightarrow{x}$ be an arbitrary ray in X. We will prove that R can not be an extreme ray.

Since X is non-proper, there exists a nonzero vector a in the intersection $X \cap (-X)$. First suppose that R is equal to $\overrightarrow{a}$ or to $-\overrightarrow{a}$. Then both R and its opposite ray $-R$ lie in X, and this means that R is not an extreme ray.

Next suppose $R \neq \pm\overrightarrow{a}$. The vectors x and a are then linearly independent, and the two rays $R_1 = \overrightarrow{x+a}$ and $R_2 = \overrightarrow{x-a}$ are consequently distinct and non-opposite rays in the cone X. Since $x = \frac{1}{2}(x+a) + \frac{1}{2}(x-a)$, we conclude that R lies between R_1 and R_2. Thus, R is not an extreme ray in this case either, and this proves that non-proper cones have no extreme rays.

The equality $X = \operatorname{con}(\operatorname{exr} X)$ is trivially true for the null cone, since $\operatorname{exr}\{0\} = \emptyset$ and $\operatorname{con}\emptyset = \{0\}$. To prove that the equality holds for all nontrivial proper closed convex cones and that these cones do have extreme rays, we only have to modify slightly the induction proof for the corresponding part of Theorem 4.2.2.

The start of the induction is of course trivial, since one-dimensional proper cones are rays. So suppose our assertion is true for all cones of dimension less than or equal to $n-1$, and let X be a proper closed n-dimensional convex cone. X is then, in particular, a line-free set, whence $X = \operatorname{cvx}\bigl(\bigcup F\bigr)$, where the union is taken over all proper faces F of the cone. Moreover, since X is a convex cone, $\operatorname{cvx}\bigl(\bigcup F\bigr) \subseteq \operatorname{con}\bigl(\bigcup F\bigr) \subseteq \operatorname{con} X = X$, and we conclude that

$$X = \operatorname{con}\bigl(\bigcup F\bigr). \tag{4.1}$$

We may of course delete the trivial face $F = \{0\}$ from the above union without destroying the identity, and every remaining face F is a proper closed convex cone of dimension less than or equal to $n-1$ with $\operatorname{exr} F \neq \emptyset$ and $F = \operatorname{con}(\operatorname{exr} F)$, by our induction assumption. Since $\operatorname{exr} F \subseteq \operatorname{exr} X$, it now follows that the set $\operatorname{exr} X$ is nonempty and that $F \subseteq \operatorname{con}(\operatorname{exr} X)$.

The union $\bigcup F$ of the faces is thus included in the cone $\operatorname{con}(\operatorname{exr} X)$, so it follows from equation (4.1) that $X \subseteq \operatorname{con}(\operatorname{exr} X)$. Since the converse inclusion is trivial, we have equality $X = \operatorname{con}(\operatorname{exr} X)$, and the induction step is now complete. □

The recession cone of a line-free convex set is a proper cone. The following structure theorem for convex sets is therefore an immediate consequence of Theorems 4.2.2 and 4.2.4.

Theorem 4.2.5. *If X is a nonempty line-free closed convex set, then*

$$X = \operatorname{cvx}(\operatorname{ext} X) + \operatorname{con}(\operatorname{exr}(\operatorname{recc} X)).$$

The study of arbitrary closed convex sets is reduced to the study of line-free such sets by the following theorem, which says that every non-line-free closed convex set is a cylinder with a line-free convex set as its base and with the recessive subspace $\operatorname{lin} X$ as its "axis".

Theorem 4.2.6. *Suppose X is a closed convex set in $\mathbf{R}^n$. The intersection $X \cap (\operatorname{lin} X)^\perp$ is then a line-free closed convex set and*

$$X = \operatorname{lin} X + X \cap (\operatorname{lin} X)^\perp.$$

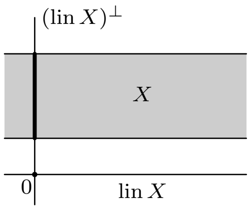

Figure 4.8. Illustration for Theorem 4.2.6.

Proof. Each $x \in \mathbf{R}^n$ has a unique decomposition $x = y + z$ with $y \in \operatorname{lin} X$ and $z \in (\operatorname{lin} X)^\perp$. If $x \in X$, then z lies in X, too, since

$$z = x - y \in X + \operatorname{lin} X = X.$$

This proves the inclusion $X \subseteq \operatorname{lin} X + X \cap (\operatorname{lin} X)^\perp$, and the converse inclusion follows from $\operatorname{lin} X + X \cap (\operatorname{lin} X)^\perp \subseteq \operatorname{lin} X + X = X$. $\square$

Exercises

4.1 Find $\operatorname{ext} X$ and decide whether $X = \operatorname{cvx}(\operatorname{ext} X)$ when

a) $X = \{x \in \mathbf{R}^2_+ \mid x_1 + x_2 \geq 1\}$

b) $X = \big([0,1] \times [0,1[\,\big) \cup \big([0, \tfrac{1}{2}] \times \{1\}\big)$

c) $X = \operatorname{cvx}\big(\{x \in \mathbf{R}^3 \mid (x_1 - 1)^2 + x_2^2 = 1,\ x_3 = 0\} \cup \{(0,0,1), (0,0,-1)\}\big)$.

4.2 Prove that $\operatorname{ext}(\operatorname{cvx} A) \subseteq A$ for each subset A of $\mathbf{R}^n$.

4.3 Let $X = \operatorname{cvx} A$ and suppose the set A is minimal in the following sense: If $B \subseteq A$ och $X = \operatorname{cvx} B$, then $B = A$. Prove that $A = \operatorname{ext} X$.

4.4 Let x_0 be a point in a convex set X. Prove that $x_0 \in \operatorname{ext} X$ if and only if the set $X \setminus \{x_0\}$ is convex.

4.5 Give an example of a compact convex subset of $\mathbf{R}^3$ such that the set of extreme points is not closed.

4.6 A point x_0 in a convex set X is called an *exposed point* if the singleton set $\{x_0\}$ is a face, i.e. if there exists a supporting hyperplane H of X such that $X \cap H = \{x_0\}$.

a) Prove that every exposed point is an extreme point of X.

b) Give an example of a closed convex set in $\mathbf{R}^2$ with an extreme point that is not exposed.

4.7 There is a more general definition of the face concept which runs as follows:

A *face* of a convex set X is a convex subset F of X such that every closed line segment in X with a relative interior point in F lies entirely in F, i.e.

$$(a, b \in X \ \& \ ]a, b[\cap F \neq \emptyset) \implies a, b \in F.$$

Let us call faces according to this definition *general faces* in order to distinguish them from faces according to our old definition, which we call *exposed faces*, provided they are proper, i.e. different from the faces X and $\emptyset$.

The empty set $\emptyset$ and X itself are apparently general faces of X, and all extreme points of X are general faces, too.

Prove that the general faces of a convex set X have the following properties.

a) Each exposed face is a general face.
b) There is a convex set with a general face that is not an exposed face.
c) If F is a general face of X and F' is a general face of F, then F' is a general face of X, but the coresponding result is not true in general for exposed faces.
d) If F is a general face of X and C is an arbitrary convex subset of X such that $F \cap \operatorname{rint} C \neq \emptyset$, then $C \subseteq F$.
e) If F is a general face of X, then $F = X \cap \operatorname{cl} F$. In particular, F is closed if X is closed.
f) If F_1 and F_2 are two general faces of X and $\operatorname{rint} F_1 \cap \operatorname{rint} F_2 \neq \emptyset$, then $F_1 = F_2$.
g) If F is a general face of X and $F \neq X$, then $F \subseteq \operatorname{rbdry} X$.

Chapter 5

Polyhedra

We have already obtained some isolated results on polyhedra, but now is the time to collect these and to complement them in order to get a complete description of this important class of convex sets.

5.1 Extreme points and extreme rays

Polyhedra and extreme points

Each polyhedron X in $\mathbf{R}^n$, except for the entire space, is an intersection of finitely many closed halfspaces and may therefore be written in the form

$$X = \bigcap_{j=1}^{m} K_j,$$

with

$$K_j = \{x \in \mathbf{R}^n \mid \langle c_j, x\rangle \geq b_j\}$$

for suitable nonzero vectors c_j in $\mathbf{R}^n$ and real numbers b_j. Using matrix notation,

$$X = \{x \in \mathbf{R}^n \mid Cx \geq b\},$$

where C is an $m \times n$-matrix with ${c_j}^T$ as rows, and $b = \begin{bmatrix} b_1 & b_2 & \dots & b_m \end{bmatrix}^T$.

Let

$$\begin{aligned} K_j^\circ &= \{x \in \mathbf{R}^n \mid \langle c_j, x\rangle > b_j\} = \operatorname{int} K_j, \quad \text{and} \\ H_j &= \{x \in \mathbf{R}^n \mid \langle c_j, x\rangle = b_j\} = \operatorname{bdry} K_j. \end{aligned}$$

The sets K_j° are open halfspaces, and the H_j are hyperplanes.

If $b = 0$, i.e. if all hyperplanes H_j are linear subspaces, then X is a polyhedral cone.

The polyhedron X is clearly a subset of the closed halfspace K_j, which is bounded by the hyperplane H_j. Let

$$F_j = X \cap H_j.$$

If there is a point in common between the hyperplane H_j and the polyhedron X, without X being entirely contained in H, then H is a supporting hyperplane of X, and the set F_j is a proper face of X. But F_j is a face of X also in the cases when $X \cap H_j = \emptyset$ or $X \subseteq H_j$, due to our convention regarding non-proper faces. Of course, the faces F_j are polyhedra.

All points of a face F_j (proper as well as non-proper) are boundary points of X. Since

$$X = \bigcap_{j=1}^{m} K_j^\circ \cup \bigcup_{j=1}^{m} F_j,$$

and all points in the open set $\bigcap_{j=1}^m K_j^\circ$ are interior points of X, we conclude that

$$\operatorname{int} X = \bigcap_{j=1}^{m} K_j^\circ \qquad \text{and} \qquad \operatorname{bdry} X = \bigcup_{j=1}^{m} F_j.$$

The set $\operatorname{ext} X$ of extreme points of the polyhedron X is a subset of the boundary $\bigcup_{j=1}^m F_j$, and the extreme points are characterized by the following theorem.

Theorem 5.1.1. *A point x_0 in the polyhedron $X = \bigcap_{j=1}^m K_j$ is an extreme point if and only if there exists a subset I of the index set $\{1, 2, \dots, m\}$ such that $\bigcap_{j\in I} H_j = \{x_0\}$.*

Proof. Suppose there exists such an index set I. The intersection

$$F = \bigcap_{j\in I} F_j = X \cap \bigcap_{j\in I} H_j = \{x_0\}$$

is a face of X, by Theorem 4.1.2, and x_0 is obviously an extreme point of F. Therefore, x_0 is also an extreme point of X, by Theorem 4.1.3.

Now suppose, conversely, that there is no such index set I, and let J be an index set that is maximal with respect to the property $x_0 \in \bigcap_{j\in J} H_j$. (Remember that the intersection over an empty index set is equal to the entire space $\mathbf{R}^n$, so $J = \emptyset$ if x_0 is an interior point of X.) The intersection $\bigcap_{j\in J} H_j$ is an affine subspace, which by assumption consists of more than one point and, therefore, contains a line $\{x_0 + tv \mid t \in \mathbf{R}\}$ through x_0. The line is obviously also contained in the larger set $\bigcap_{j\in J} K_j$.

Since x_0 is an interior point of the halfspace K_j for all indices $j \notin J$, we conclude that the points $x_0 + tv$ belong to all these halfspaces for all sufficiently small values of $|t|$. Consequently, there is a number $\delta > 0$ such that the line segment $[x_0 - \delta v, x_0 + \delta v]$ lies in $X = \bigcap_{j \in J} K_j \cap \bigcap_{j \notin J} K_j$, which means that x_0 is not an extreme point. $\square$

The condition $\bigcap_{j \in I} H_j = \{x_0\}$ means that the corresponding system of linear equations

$$\langle c_j, x \rangle = b_j, \quad j \in I,$$

in n unknowns has a unique solution. A necessary condition for this to be true is that the index set I contains at least n elements. And if the system has a unique solution and there are more than n equations, then it is always possible to obtain a quadratic subsystem with a uniqe solution by eliminating suitably selected equations.

Hence, the condition $m \geq n$ is necessary for the polyhedron $X = \bigcap_{j=1}^{m} K_j$ to have at least one extreme point. (This also follows from Theorem 2.7.7, for if $m < n$, then

$$\dim \operatorname{lin} X = \dim\{x \in \mathbf{R}^n \mid Cx = 0\} = n - \operatorname{rank} C \geq n - m > 0,$$

which means that X is not line-free.)

Theorem 5.1.1 gives us the following method for finding all extreme points of the polyhedron X when $m \geq n$:

Solve for each subset J of $\{1, 2, \dots, m\}$ with n elements the corresponding linear system $\langle c_j, x\rangle = b_j$, $j \in J$. If the system has a unique solution x_0, and the solution lies in X, i.e. satisfies the remaining linear inequalities $\langle c_j, x\rangle \geq b_j$, then x_0 is an extreme point of X.

The number of extreme points of X is therefore bounded by $\binom{m}{n}$, which is the number of subsets J of $\{1, 2, \dots, m\}$ with n elements. In particular, we have proved the following theorem.

Theorem 5.1.2. *Polyhedra have finitely many extreme points.*

Polyhedral cones and extreme rays

A polyhedral cone in $\mathbf{R}^n$ is an intersection $X = \bigcap_{j=1}^m K_j$ of conic halfspaces K_j which are bounded by hyperplanes H_j through the origin, and the faces $F_j = X \cap H_j$ are polyhedral cones. Our next theorem is a direct analogue of Theorem 5.1.1.

Theorem 5.1.3. *A point x_0 in the polyhedral cone X generates an extreme ray $R = \overrightarrow{x_0}$ of the cone if and only if $-x_0 \notin X$ and there exists a subset I of the index set $\{1, 2, \dots, m\}$ such that $\bigcap_{j\in I} H_j = \{tx_0 \mid t \in \mathbf{R}\}$.*

Proof. Suppose there exists such an index set I and that $-x_0$ does not belong to the cone X. Then

$$\bigcap_{j\in I} F_j = X \cap \bigcap_{j\in I} H_j = R.$$

By Theorem 4.1.2, this means that R is a face of the cone X. The ray R is an extreme ray of the face R, of course, so it follows from Theorem 4.1.3 that R is an extreme ray of X.

If $-x_0$ belongs to X, then X is not a proper cone, and hence X has no extreme rays according to Theorem 4.2.4.

It remains to show that R is not an extreme ray in the case when $-x_0 \notin X$ and there is no index set I with the property that the intersection $\bigcap_{j\in I} H_j$ is equal to the line through 0 and x_0. So let J be a maximal index set satisfying the condition $x_0 \in \bigcap_{j\in J} H_j$. Due to our assumption, the intersection $\bigcap_{j\in J} H_j$ is then a linear subspace of dimension greater than or equal to two, and therefore it contains a vector v which is linearly independent of x_0. The vectors $x_0 + tv$ and $x_0 - tv$ both belong to $\bigcap_{j\in J} H_j$, and consequently also to $\bigcap_{j\in J} K_j$, for all real numbers t. When $|t|$ is a sufficiently small number, the two vectors also belong to the halfspaces K_j for indices $j \notin J$, because

x_0 is an interior point of K_j for these indices j. Therefore, there exists a positive number δ such that the vectors $x_+ = x_0 + \delta v$ and $x_- = x_0 - \delta v$ both belong to the cone X. The two vectors x_+ and x_- are linearly independent and $x_0 = \frac{1}{2}x_+ + \frac{1}{2}x_-$, so it follows that the ray $R = \overrightarrow{x_0}$ lies between the two rays $\overrightarrow{x_+}$ and $\overrightarrow{x_-}$ in X, and R is therefore not an extreme ray. □

Thus, to find all the extreme rays of the cone

$$X = \bigcap_{j=1}^{m} \{x \in \mathbf{R}^n \mid \langle c_j, x\rangle \geq 0\}$$

we should proceed as follows. First choose an index set J consisting of $n-1$ elements from the set $\{1, 2, \dots, m\}$. This can be done in $\binom{m}{n-1}$ different ways. Then solve the corresponding homogeneous linear system $\langle c_j, x\rangle = 0$, $j \in J$. If the solution set is one-dimensional, than pick a solution x_0. If x_0 satisfies the remaining linear inequalities and $-x_0$ does not, then $R = \overrightarrow{x_0}$ is an extreme ray. If, instead, $-x_0$ satisfies the remaining linear inequalities and x_0 does not, then $-R$ is an extreme ray. Since this is the only way to obtain extreme rays, we conclude that the number of extreme rays is bounded by the number $\binom{m}{n-1}$. In particular, we get the following corollary.

Theorem 5.1.4. *Polyhedral cones have finitely many extreme rays.*

5.2 Polyhedral cones

Theorem 5.2.1. *A cone is polyhedral if and only if it is finitely generated.*

Proof. We first show that every polyhedral cone is finitely generated.

By Theorem 4.2.6, every polyhedral cone X can be written in the form

$$X = \operatorname{lin} X + X \cap (\operatorname{lin} X)^{\perp},$$

and $X \cap (\operatorname{lin})^{\perp}$ is a line-free, i.e. proper, polyhedral cone. Let B be a set consisting of one point from each extreme ray of $X \cap (\operatorname{lin} X)^{\perp}$; then B is a finite set and

$$X \cap (\operatorname{lin} X)^{\perp} = \operatorname{con} B,$$

according to Theorems 5.1.4 and 4.2.4.

Let $e_1, e_2, \dots, e_d$ be a basis for the linear subspace $\operatorname{lin} X$, and put $e_0 = -(e_1 + e_2 + \cdots + e_d)$. The cone $\operatorname{lin} X$ is generated as a cone by the set $A = \{e_0, e_1, \dots, e_d\}$, i.e.

$$\operatorname{lin} X = \operatorname{con} A.$$

Summing up,

$$X = \operatorname{lin} X + X \cap (\operatorname{lin} X)^{\perp} = \operatorname{con} A + \operatorname{con} B = \operatorname{con}(A \cup B),$$

which shows that the cone X is finitely generated by the set $A \cup B$.

Next, suppose that X is a finitely generated cone so that $X = \operatorname{con} A$ for some finite set A. We start by the observation that the dual cone X^+ is polyhedral. Indeed, if $A \neq \emptyset$ then

$$X^+ = A^+ = \{x \in \mathbf{R}^n \mid \langle x, a\rangle \geq 0 \text{ for all } a \in A\} = \bigcap_{a\in A} \{x \in \mathbf{R}^n \mid \langle a, x\rangle \geq 0\}$$

is an intersection of finitely many conical halfspaces, i.e. a polyhedral cone. And if $A = \emptyset$, then $X = \{0\}$ and $X^+ = \mathbf{R}^n$.

The already proven part of the theorem now implies that the dual cone X^+ is finitely generated. But the dual cone of X^+, i.e. the bidual cone X^{++}, is then polyhedral, too. Since the bidual cone X^{++} coincides with the original cone X, by Corollary 3.2.4, we conclude the X is a polyhedral cone. □

We are now able two prove two results that were left unproven in Chapter 2.6; compare Corollary 2.6.9.

Theorem 5.2.2. *(i) The intersection $X \cap Y$ of two finitely generated cones X and Y is a finitely generated cone.*

(ii) The inverse image $T^{-1}(X)$ of a finitely generated cone X under a linear map T is a finitely generated cone.

Proof. The intersection of two conical polyhedra is obviously a conical polyhedron, and the same holds for the inverse image of a conical polyhedron under a linear map. The theorem is therefore a corollary of Theorem 5.2.1. □

5.3 The internal structure of polyhedra

Polyhedra are by definition intersections of finite collections of closed halfspaces, and this can be viewed as an external description of polyhedra. We shall now give an internal description of polyhedra in terms of extreme points and extreme rays, and the following structure theorem is the main result of this chapter.

Theorem 5.3.1. *A nonempty subset X of $\mathbf{R}^n$ is a polyhedron if and only if there exist two finite subsets A and B of $\mathbf{R}^n$ with $A \neq \emptyset$ such that*

$$X = \operatorname{cvx} A + \operatorname{con} B.$$

The cone $\operatorname{con} B$ is then equal to the recession cone $\operatorname{recc} X$ of X. If the polyhedron is line-free, we may choose for A the set $\operatorname{ext} X$ of all extreme points of X, and for B a set consisting of one nonzero point from each extreme ray of the recession cone $\operatorname{recc} X$.

Proof. We first prove that polyhedra have the stated decomposition. So let X be a polyhedron and put $Y = X \cap (\operatorname{lin} X)^{\perp}$. Then, Y is a line-free polyhedron, and

$$X = \operatorname{lin} X + Y = \operatorname{lin} X + \operatorname{recc} Y + \operatorname{cvx}(\operatorname{ext} Y),$$

by Theorems 4.2.6 and 4.2.2. The two polyhedral cones $\operatorname{lin} X$ and $\operatorname{recc} Y$ are, according to Theorem 5.2.1, generated by two finite sets B_1 and B_2, respectively, and their sum is generated by the finite set $B = B_1 \cup B_2$. The set $\operatorname{ext} Y$ is finite, by Theorem 5.1.2, so the representation

$$X = \operatorname{cvx} A + \operatorname{con} B$$

is now obtained by taking $A = \operatorname{ext} Y$.

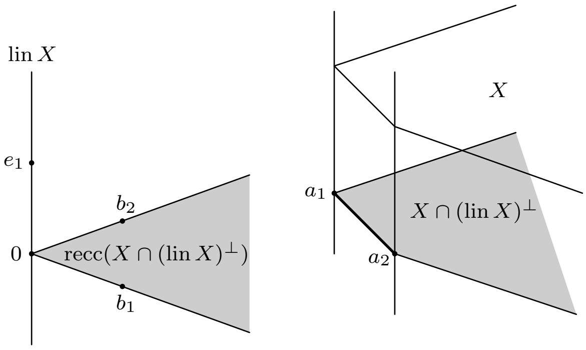

Figure 5.1. An illustration for Theorem 5.3.1. The right part of the figure depicts an unbounded polyhedron X in $\mathbf{R}^3$. Its recessive subspace $\operatorname{lin} X$ is one-dimensional and is generated as a cone by e_1 and -e_1. The intersection $X \cap (\operatorname{lin} X)^\perp$, which is shadowed, is a line-free polyhedron with two extreme points a_1 and a_2. The recession cone $\operatorname{recc}(X \cap (\operatorname{lin} X)^\perp)$ is generated by b_1 and b_2. The representation $X = \operatorname{cvx} A + \operatorname{con} B$ is obtained by taking $A = \{a_1, a_2\}$ and $B = \{e_1, -e_1, b_1, b_2\}$.

The cone $\operatorname{con} B$ is closed and the convex set $\operatorname{cvx} A$ is compact, since the sets A and B are finite. Hence, $\operatorname{con} B = \operatorname{recc} X$ by Corollary 2.7.13.

If X is a line-free polyhedron, then

$$X = \operatorname{cvx}(\operatorname{ext} X) + \operatorname{con}(\operatorname{exr}(\operatorname{recc} X)),$$

by Theorems 4.2.2 and 4.2.4, and this gives us the required representation of X with $A = \operatorname{ext} X$ and with B as a set consisting of one nonzero point from each extreme ray of $\operatorname{recc} X$.

To prove the converse, suppose that $X = \operatorname{cvx} A + \operatorname{con} B$, where $A = \{a_1, \dots, a_p\}$ and $B = \{b_1, \dots, b_q\}$ are finite sets. Consider the cone Y in $\mathbf{R}^n \times \mathbf{R}$ that is generated by the finite set $(A \times \{1\}) \cup (B \times \{0\})$. The cone Y is polyhedral according to Theorem 5.2.1, which means that there is an $m \times (n+1)$-matrix C such that

$$(x, x_{n+1}) \in Y \quad \Leftrightarrow \quad C \begin{bmatrix} x \\ x_{n+1} \end{bmatrix} \geq 0. \tag{5.1}$$

(Here $\begin{bmatrix} x \\ x_{n+1} \end{bmatrix}$ denotes the vector $(x_1, \dots, x_n, x_{n+1})$ written as a column matrix.)

Let C' denote the submatrix of C which consists of all columns but the last, and let c' be the last column of the matrix C. Then

$$C \begin{bmatrix} x \\ x_{n+1} \end{bmatrix} = C'x + x_{n+1}c',$$

which means that the equivalence (5.1) may be written as

$$(x, x_{n+1}) \in Y \quad \Leftrightarrow \quad C'x + x_{n+1}c' \geq 0.$$

By definition, a vector $(x, 1) \in \mathbf{R}^n \times \mathbf{R}$ belongs to the cone Y if and only if there exist nonnegative numbers $\lambda_1, \lambda_2, \dots, \lambda_p$ and $\mu_1, \mu_2, \dots, \mu_q$ such that

$$\begin{cases} x = \lambda_1 a_1 + \lambda_2 a_2 + \cdots + \lambda_p a_p + \mu_1 b_1 + \mu_2 b_2 + \cdots + \mu_q b_q \\ 1 = \lambda_1 + \lambda_2 + \cdots + \lambda_p \end{cases}$$

i.e. if and only if $x \in \operatorname{cvx} A + \operatorname{con} B$. This yields the equivalences

$$x \in X \quad \Leftrightarrow \quad (x, 1) \in Y \quad \Leftrightarrow \quad C'x + c' \geq 0,$$

which means that $X = \{x \in \mathbf{R}^n \mid C'x \geq -c'\}$. Thus, X is a polyhedron. $\square$

5.4 Polyhedron preserving operations

Theorem 5.4.1. *The intersection of finitely many polyhedra in $\mathbf{R}^n$ is a polyhedron.*

Proof. Trivial. □

Theorem 5.4.2. *Suppose X is a polyhedron in $\mathbf{R}^n$ and that $T\colon \mathbf{R}^n \to \mathbf{R}^m$ is an affine map. The image $T(X)$ is then a polyhedron in $\mathbf{R}^m$.*

Proof. The assertion is trivial if the polyhedron is empty, so suppose it is nonempty and write it in the form

$$X = \operatorname{cvx} A + \operatorname{con} B,$$

where $A = \{a_1, \dots, a_p\}$ and $B = \{b_1, \dots, b_q\}$ are finite sets. Each $x \in X$ has then a representation of the form

$$x = \sum_{j=1}^{p} \lambda_j a_j + \sum_{j=1}^{q} \mu_j b_j = \sum_{j=1}^{p} \lambda_j a_i + \sum_{j=1}^{q} \mu_j b_j - (\sum_{j=1}^{q} \mu_j)0$$

with nonnegative coefficients λ_j and μ_j and $\sum_{j=1}^{p} \lambda_j = 1$, i.e. each $x \in X$ is an affine combination of elements in the set $A \cup B \cup \{0\}$. Since T is an affine map,

$$Tx = \sum_{j=1}^{p} \lambda_j Ta_j + \sum_{j=1}^{q} \mu_j Tb_j - (\sum_{j=1}^{q} \mu_j)T0 = \sum_{j=1}^{p} \lambda_j Ta_j + \sum_{j=1}^{q} \mu_j (Tb_j - T0).$$

This shows that the image $T(X)$ is of the form

$$T(X) = \operatorname{cvx} A' + \operatorname{con} B'$$

with $A' = T(A)$ and $B' = -T0 + T(B) = \{Tb_1 - T0, \dots, Tb_q - T0\}$. So the image $T(X)$ is a polyhedron, by Theorem 5.3.1. □

Theorem 5.4.3. *Suppose Y is a polyhedron in $\mathbf{R}^m$ and that $T\colon \mathbf{R}^n \to \mathbf{R}^m$ is an affine map. The inverse image $T^{-1}(Y)$ is then a polyhedron in $\mathbf{R}^n$.*

Proof. First assume that Y is a closed halfspace in $\mathbf{R}^m$ (or the entire space $\mathbf{R}^m$), i.e. that $Y = \{y \in \mathbf{R}^m \mid \langle c, y\rangle \geq b\}$. (The case $Y = \mathbf{R}^m$ is obtained by $c = 0$ and $b = 0$.) The affine map T can be written in the form $Tx = Sx + y_0$, with S as a linear map and y_0 as a vector in $\mathbf{R}^m$. This gives us

$$T^{-1}(Y) = \{x \in \mathbf{R}^n \mid \langle c, Tx\rangle \geq b\} = \{x \in \mathbf{R}^n \mid \langle S^T c, x\rangle \geq b - \langle c, y_0\rangle\}.$$

So $T^{-1}(Y)$ is a closed halfspace in $\mathbf{R}^n$ if $S^T c \neq 0$, the entire space $\mathbf{R}^n$ if $S^T c = 0$ and $b \leq \langle c, y_0\rangle$, and the empty set $\emptyset$ if $S^T c = 0$ and $b > \langle c, y_0\rangle$.

In the general case, $Y = \bigcap_{j=1}^{p} K_j$ is an intersection of finitely many closed halfspaces. Since $S^{-1}(Y) = \bigcap_{j=1}^{p} S^{-1}(K_j)$, the inverse image $S^{-1}(Y)$ is an intersection of closed halfspaces, the empty set, or the entire space $\mathbf{R}^n$. Thus, $S^{-1}(Y)$ is a polyhedron. □

Theorem 5.4.4. *The Cartesian product $X \times Y$ of two polyhedra X and Y is a polyhedron.*

Proof. Suppose X lies in $\mathbf{R}^m$ and Y lies in $\mathbf{R}^n$. The set $X \times \mathbf{R}^n$ is a polyhedron since it is the inverse image of X under the projection $(x, y) \mapsto x$, and $\mathbf{R}^m \times Y$ is a polyhedron for a similar reason. It follows that $X \times Y$ is a polyhedron, because $X \times Y = (X \times \mathbf{R}^n) \cap (\mathbf{R}^m \times Y)$. □

Theorem 5.4.5. *The sum $X + Y$ of two polyhedra in $\mathbf{R}^n$ is a polyhedron.*

Proof. The sum $X + Y$ is equal to the image of $X \times Y$ under the linear map $(x, y) \to x + y$, so the theorem is a consequence of the previous theorem and Theorem 5.4.2. □

5.5 Separation

It is possible to obtain sharper separation results for polyhedra than for general convex sets. Compare the following two theorems with Theorems 3.1.6 and 3.1.5.

Theorem 5.5.1. *If X and Y are two disjoint polyhedra, then there exists a hyperplane that strictly separates the two polyhedra.*

Proof. The difference $X - Y$ of two polyhedra X and Y is a closed set, since it is a polyhedron according to Theorem 5.4.5. So it follows from Theorem 3.1.6 that there exists a hyperplane that strictly separates the two polyhedra, if they are disjoint. □

Theorem 5.5.2. *Let X be a convex set, and let Y be a polyhedron that is disjoint from X. Then there exists a hyperplane that separates X and Y and does not contain X as a subset.*

Proof. We prove the theorem by induction over the dimension n of the surrounding space $\mathbf{R}^n$.

The case $n = 1$ is trivial, so suppose the assertion of the theorem is true when the dimension is $n - 1$, and let X be a convex subset of $\mathbf{R}^n$ that is disjoint from the polyhedron Y. An application of Theorem 3.1.5 gives us a hyperplane H that separates X and Y and, as a consequence, does not contain both sets as subsets. If X is not contained in H, then we are done. So suppose that X is a subset of H. The polyhedron Y then lies in one of the two closed halfspaces defined by the hyperplane H. Let us denote this closed halfspace by H_+, so that $Y \subseteq H_+$, and let H_{++} denote the corresponding open halfspace.

If $Y \subseteq H_{++}$, then Y and H are disjoint polyhedra, and an application of Theorem 5.5.1 gives us a hyperplane that strictly separates Y and H. Of course, this hyperplane also strictly separates Y and X, since X is a subset of H.

This proves the case $Y \subseteq H_{++}$, so it only remains to consider the case when Y is a subset of the closed halfspace H_+ without being a subset of the corresponding open halfspace, i.e. the case

$$Y \subseteq H_+, \quad Y \cap H \neq \emptyset.$$

Due to our induction hypothesis, it is possible to separate the nonempty polyhedron $Y_1 = Y \cap H$ and X inside the $(n-1)$-dimensional hyperplane H using an affine $(n-2)$-dimensional subset L of H which does not contain X as a subset. L divides the hyperplane H into two closed halves L_+ and L_- with L as their common relative boundary, and with X as a subset of L_- and Y_1 as a subset of L_+.

Let us denote the relative interior of L_- by L_{--}, so that $L_{--} = L_- \setminus L$. The assumption that X is not a subset of L implies that $X \cap L_{--} \neq \emptyset$.

Observe that $Y \cap L_- = Y_1 \cap L$. If $Y_1 \cap L = \emptyset$, then there exists a hyperplane that strictly separates the polyhedra Y och L_-, by Theorem 5.5.1, and since $X \subseteq L_-$, we are done in this case, too.

What remains is to treat the case $Y_1 \cap L \neq \emptyset$, and by performing a translation, if necessary, we may assume that the origin lies in $Y_1 \cap L$, which implies that L is a linear subspace. See figure 5.2.

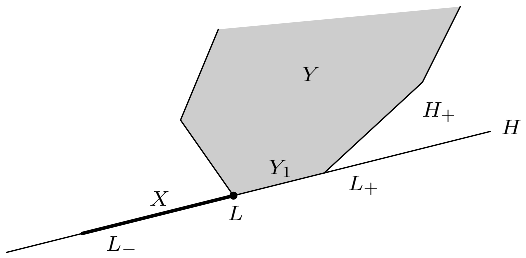

Figure 5.2. Illustration for the proof of Theorem 5.5.2.

Note that the set $H_{++} \cup L_+$ is a cone and that Y is a subset of this cone. Now, consider the cone $\operatorname{con} Y$ generated by the polyhedron Y, and let

$$C = L + \operatorname{con} Y.$$

C is a cone, too, and a subset of the cone $H_{++} \cup L_+$, since Y and L are both subsets of the last mentioned cone. The cone $\operatorname{con} Y$ is polyhedral, because if the polyhedron Y is written as $Y = \operatorname{cvx} A + \operatorname{con} B$ with finite sets A and B, then $\operatorname{con} Y = \operatorname{con}(A \cup B)$ due to the fact that 0 lies in Y. Since the sum of two polyhedral cones is polyhedral, it follows that the cone C is also polyhedral.

The cone C is disjoint from the set L_{--}, since the sets L_{--} and $H_{++} \cup L_+$ are disjoint.

Now write the polyhedral cone C as an intersection $\bigcap K_i$ of finitely many closed halfspaces K_i which are bounded by hyperplanes H_i through the origin. Each halfspace K_i is a cone containing Y as well as L. If a given halfspace K_i contains in addition a point from L_{--}, then it contains the cone generated by that point and L, that is all of L_-. Therefore, since $C = \bigcap K_i$ and $C \cap L_{--} = \emptyset$, we conclude that there exists a halfspace K_i that does not contain any point from L_{--}. In other words, the corresponding boundary hyperplane H_i separates L_- and the cone C and is disjoint from L_{--}. Since $X \subseteq L_-$, $Y \subseteq C$ and $X \cap L_{--} \neq \emptyset$, H_i separates the sets X and Y and does not contain X. This completes the induction step and the proof of the theorem. $\square$

Exercises

5.1 Find the extreme points of the following polyhedra X:

a) $X = \{x \in \mathbf{R}^2 \mid -x_1 + x_2 \le 2,\ x_1 + 2x_2 \ge 2,\ x_2 \ge -1\}$
b) $X = \{x \in \mathbf{R}^2 \mid -x_1 + x_2 \le 2,\ x_1 + 2x_2 \le 2,\ x_2 \ge -1\}$
c) $X = \{x \in \mathbf{R}^3 \mid 2x_1 + x_2 + x_3 \le 4,\ x_1 + 2x_2 + x_3 \le 4,\ x \ge 0\}$
d) $X = \{x \in \mathbf{R}^4 \mid x_1 + x_2 + 3x_3 + x_4 \le 4,\ 2x_2 + 3x_3 \ge 5,\ x \ge 0\}$.

5.2 Find the extreme rays of the cone

$X = \{x \in \mathbf{R}^3 \mid x_1 - x_2 + 2x_3 \ge 0,\ x_1 + 2x_2 - 2x_3 \ge 0,\ x_2 + x_3 \ge 0,\ x_3 \ge 0\}$.

5.3 Find a matrix C such that

$$\operatorname{con}\{(1,-1,1),(-1,0,1),(3,2,1),(-2,-1,0)\} = \{x \in \mathbf{R}^3 \mid Cx \ge 0\}.$$

5.4 Find finite sets A and B such that $X = \operatorname{con} A + \operatorname{cvx} B$ for the following polyhedra:

a) $X = \{x \in \mathbf{R}^2 \mid -x_1 + x_2 \le 2,\ x_1 + 2x_2 \ge 2,\ x_2 \ge -1\}$
b) $X = \{x \in \mathbf{R}^2 \mid -x_1 + x_2 \le 2,\ x_1 + 2x_2 \le 2,\ x_2 \ge -1\}$
c) $X = \{x \in \mathbf{R}^3 \mid 2x_1 + x_2 + x_3 \le 4,\ x_1 + 2x_2 + x_3 \le 4,\ x \ge 0\}$
d) $X = \{x \in \mathbf{R}^4 \mid x_1 + x_2 + 3x_3 + x_4 \le 4,\ 2x_2 + 3x_3 \ge 5,\ x \ge 0\}$.

5.5 Suppose 0 lies in the polyhedron $X = \operatorname{cvx} A + \operatorname{con} B$, where A and B are finite sets. Prove that $\operatorname{con} X = \operatorname{con}(A \cup B)$.

Chapter 6

Convex functions

6.1 Basic definitions

Epigraph and sublevel set

Definition. Let $f\colon X \to \overline{\mathbf{R}}$ be a function with domain $X \subseteq \mathbf{R}^n$ and codomain $\overline{\mathbf{R}}$, i.e. the real numbers extended with ∞. The set

$$\operatorname{epi} f = \{(x,t) \in X \times \mathbf{R} \mid f(x) \leq t\}$$

is called the *epigraph* of the function.

Let α be a real number. The set

$$\operatorname{sublev}_\alpha f = \{x \in X \mid f(x) \leq \alpha\}$$

is called a *sublevel set* of the function, or more precisely, the *α-sublevel set*.

The epigraph is a subset of $\mathbf{R}^{n+1}$, and the word 'epi' means above. So epigraph means above the graph.

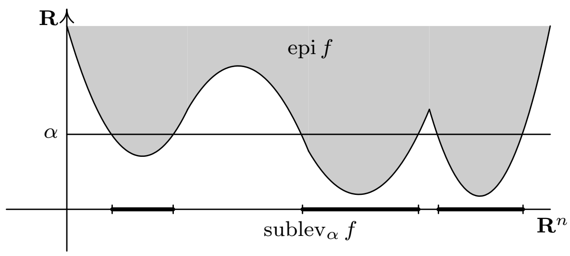

Figure 6.1. Epigraph and a sublevel set

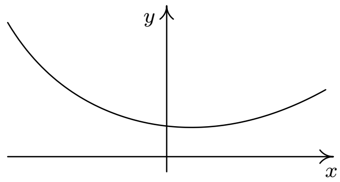

Figure 6.2. The graph of a convex function

We remind the reader of the notation $\operatorname{dom} f$ for the effective domain of f, i.e. the set of points where the function $f\colon X \to \overline{\mathbf{R}}$ is finite. Obviously,

$$\operatorname{dom} f = \{x \in X \mid f(x) < \infty\}$$

is equal to the union of all the sublevel sets of f, and these form an increasing family of sets, i.e.

$$\operatorname{dom} f = \bigcup_{\alpha\in\mathbf{R}} \operatorname{sublev}_\alpha f \qquad \text{and} \qquad \alpha < \beta \Rightarrow \operatorname{sublev}_\alpha f \subseteq \operatorname{sublev}_\beta f.$$

This implies that $\operatorname{dom} f$ is a convex set if all the sublevel sets are convex.

Convex functions

Definition. A function $f\colon X \to \overline{\mathbf{R}}$ is called *convex* if its domain X and epigraph $\operatorname{epi} f$ are convex sets.

A function $f\colon X \to \underline{\mathbf{R}}$ is called *concave* if the function $-f$ is convex.

EXAMPLE 6.1.1. The epigraph of an affine function is a closed halfspace. All affine functions, and in particular all linear functions, are thus convex and concave. □

EXAMPLE 6.1.2. The exponential funcion e^x with $\mathbf{R}$ as domain of definition is a convex function.

To see this, we replace x with $x-a$ in the elementary inequality $e^x \geq x+1$ and obtain the inequality $e^x \geq (x-a)e^a + e^a$, which implies that the epigraph of the exponential function can be expressed as the intersection

$$\bigcap_{a\in\mathbf{R}} \{(x,y) \in \mathbf{R}^2 \mid y \geq (x-a)e^a + e^a\}$$

of a family of closed halfspaces in $\mathbf{R}^2$. The epigraph is thus convex. □

Theorem 6.1.1. *The effective domain* $\operatorname{dom} f$ *and the sublevel sets* $\operatorname{sublev}_\alpha f$ *of a convex function* $f\colon X \to \overline{\mathbf{R}}$ *are convex sets.*

Proof. Suppose that the domain X is a subset of $\mathbf{R}^n$ and consider the projection $P_1\colon \mathbf{R}^n \times \mathbf{R} \to \mathbf{R}^n$ of $\mathbf{R}^n \times \mathbf{R}$ onto its first factor, i.e. $P_1(x,t) = x$. Let furthermore K_α denote the closed halfspace $\{x \in \mathbf{R}^{n+1} \mid x_{n+1} \le \alpha\}$. Then $\operatorname{sublev}_\alpha f = P_1(\operatorname{epi} f \cap K_\alpha)$, for

$$\begin{aligned} f(x) \le \alpha \;&\Leftrightarrow\; \exists t\colon f(x) \le t \le \alpha \;\Leftrightarrow\; \exists t\colon (x,t) \in \operatorname{epi} f \cap K_\alpha \\ &\Leftrightarrow\; x \in P_1(\operatorname{epi} f \cap K_\alpha). \end{aligned}$$

The intersections $\operatorname{epi} f \cap K_\alpha$ are convex sets, and since convexity is preserved by linear maps, it follows that the sublevel sets $\operatorname{sublev}_\alpha f$ are convex. Consequently, their union $\operatorname{dom} f$ is also convex. □

Quasiconvex functions

Many important properties of convex functions are consequences of the mere fact that their sublevel sets are convex. This is the reason for paying special attention to functions with convex sublevel sets and motivates the following definition.

Definition. A function $f\colon X \to \overline{\mathbf{R}}$ is called *quasiconvex* if X and all its sublevel sets $\operatorname{sublev}_\alpha f$ are convex.

A function $f\colon X \to \underline{\mathbf{R}}$ is called *quasiconcave* if $-f$ is quasiconvex.

Convex functions are quasiconvex since their sublevel sets are convex. The converse is not true, because a function f that is defined on some subinterval I of $\mathbf{R}$ is quasiconvex if it is increasing on I, or if it is decreasing on I, or more generally, if there exists a point $c \in I$ such that f is decreasing to the left of c and increasing to the right of c. There are, of course, non-convex functions of this type.

Convex extensions

The effective domain $\operatorname{dom} f$ of a convex (quasiconvex) function $f\colon X \to \overline{\mathbf{R}}$ is convex, and since

$$\begin{aligned} \operatorname{epi} f &= \{(x,t) \in \operatorname{dom} f \times \mathbf{R} \mid f(x) \le t\} \qquad \text{and} \\ \operatorname{sublev}_\alpha f &= \{x \in \operatorname{dom} f \mid f(x) \le \alpha\}, \end{aligned}$$

the restriction $f|_{\operatorname{dom} f}$ of f to $\operatorname{dom} f$ is also a convex (quasiconvex) function, and the restriction has the same epigraph and the same α-sublevel sets as f.

So what is the point of allowing ∞ as a function value of a convex function? We are of course primarily interested in functions with finite values but functions with infinite values arise naturally as suprema or limits of sequences of functions with finite values.

Another benefit of allowing ∞ as a function value of (quasi)convex functions is that we can without restriction assume that they are defined on the entire space $\mathbf{R}^n$. For if $f\colon X \to \overline{\mathbf{R}}$ is a (quasi)convex function defined on a proper subset X of $\mathbf{R}^n$, and if we define the function $\tilde{f}\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ by

$$\tilde{f}(x) = \begin{cases} f(x) & \text{if } x \in X \\ \infty & \text{if } x \notin X, \end{cases}$$

then f and $\tilde{f}$ have the same epigraphs and the same α-sublevel sets. The extension $\tilde{f}$ is therefore also (quasi)convex. Of course, $\operatorname{dom} \tilde{f} = \operatorname{dom} f$.

(Quasi)concave functions have an analogous extension to functions with values in $\underline{\mathbf{R}} = \mathbf{R} \cup \{-\infty\}$.

Alternative characterization of convexity

Theorem 6.1.2. *A function $f\colon X \to \overline{\mathbf{R}}$ with a convex domain of definition X is* *(a) convex if and only if*

$$f(\lambda x + (1-\lambda)y) \le \lambda f(x) + (1-\lambda)f(y) \tag{6.1}$$

for all points $x, y \in X$ and all numbers $\lambda \in]0,1[$;

(b) quasiconvex if and only if

$$f(\lambda x + (1-\lambda)y) \le \max\{f(x), f(y)\} \tag{6.2}$$

for all points $x, y \in X$ and all numbers $\lambda \in]0,1[$.

Proof. (a) Suppose f is convex, i.e. that the epigraph $\operatorname{epi} f$ is convex, and let x and y be two points in $\operatorname{dom} f$. Then the points $(x, f(x))$ and $(y, f(y))$ belong to the epigraph, and the convexity of the epigraph implies that the convex combination

$$\big(\lambda x + (1-\lambda)y, \lambda f(x) + (1-\lambda)f(y)\big)$$

of these two points also belong to the epigraph. This statement is equivalent to the inequality (6.1) being true. If any of the points $x, y \in X$ lies outside $\operatorname{dom} f$, then the inequality is trivially satisfied since the right hand side is equal to ∞ in that case.

To prove the converse, we assume that the inequality (6.1) holds. Let (x, s) and (y, t) be two points in the epigraph, and let $0 < \lambda < 1$. Then $f(x) \le s$ and $f(y) \le t$, by definition, and it therefore follows from the inequality (6.1) that

$$f(\lambda x + (1-\lambda)y) \le \lambda f(x) + (1-\lambda)f(y) \le \lambda s + (1-\lambda)t,$$

so the point $(\lambda x + (1-\lambda)y, \lambda s + (1-\lambda)t)$, i.e. the point $\lambda(x, s) + (1-\lambda)(y, t)$, lies in the epigraph. In orther words, the epigraph is convex.

(b) The proof is analogous and is left to the reader. $\square$

A function $f : X \to \overline{\mathbf{R}}$ is clearly (quasi)convex if and only if the restriction $f|_L$ is (quasi)convex for each line L that intersects X. Each such line has an equation of the form $x = x_0 + tv$, where x_0 is a point in X and v is a vector in $\mathbf{R}^n$, and the corresponding restriction is a one-variable function $g(t) = f(x_0 + tv)$ (with the set $\{t \mid x_0 + tv \in X\}$ as its domain of definition). To decide whether a function is (quasi)convex or not is thus essentially a one-variable problem.

Definition. Let $f\colon X \to \overline{\mathbf{R}}$ be a function defined on a convex cone X. The function is called

- *subadditive* if $f(x+y) \le f(x) + f(y)$ for all $x, y \in X$;
- *positive homogeneous* if $f(\alpha x) = \alpha f(x)$ for all $x \in X$ and all $\alpha \in \mathbf{R}_+$.

Every positive homogeneous, subadditive function is clearly convex. Conversely, every convex, positive homogeneous function f is subadditive, because

$$f(x+y) = 2f(\tfrac{1}{2}x + \tfrac{1}{2}y) \le 2(\tfrac{1}{2}f(x) + \tfrac{1}{2}f(y)) = f(x) + f(y).$$

A *seminorm* on $\mathbf{R}^n$ is a function $f \colon \mathbf{R}^n \to \mathbf{R}$, which is subadditive, positive homogeneous, and *symmetric*, i.e. satisfies the condition

$$f(-x) = f(x) \quad \text{for all } x \in \mathbf{R}^n.$$

The symmetry and homogenouity conditions may of course be merged to the condition

$$f(\alpha x) = |\alpha| f(x) \text{ for all } x \in \mathbf{R}^n \text{ and all } \alpha \in \mathbf{R}.$$

If f is a seminorm, then $f(x) \ge 0$ for all x, since

$$0 = f(0) = f(x - x) \le f(x) + f(-x) = 2f(x).$$

A seminorm f is called a *norm* if $f(x) = 0$ implies $x = 0$. The usual notation for a norm is $\|\cdot\|$.

Seminorms, and in particular norms, are convex functions.

EXAMPLE 6.1.3. The Euclidean norm and the ℓ^1-norm, that were defined in Chapter 1, are special cases of the ℓ^p-norms $\|\cdot\|_p$ on $\mathbf{R}^n$. They are defined for $1 \le p < \infty$ by

$$\|x\|_p = \big(\sum_{i=1}^{n} |x_i|^p\big)^{1/p},$$

and for $p = \infty$ by

$$\|x\|_\infty = \max_{1\le i\le n} |x_i|.$$

The maximum norm $\|\cdot\|_\infty$ is a limiting case, because $\|x\|_p \to \|x\|_\infty$ as $p \to \infty$.

The ℓ^p-norms are obviously positive homogeneous and symmetric and equal to 0 only if $x = 0$. Subadditivity is an immediate consequence of the triangle inequality $|x + y| \le |x| + |y|$ for real numbers when $p = 1$ or $p = \infty$, and of the Cauchy–Schwarz inequality when $p = 2$. For the remaining values of p, subadditivity will be proved in Section 6.4 (Theorem 6.4.3). □

Strict convexity

By strengthening the inequalities in the alternative characterization of convexity, we obtain the following definitions.

Definition. A convex function $f\colon X \to \overline{\mathbf{R}}$ is called *strictly convex* if

$$f(\lambda x + (1-\lambda)y) < \lambda f(x) + (1-\lambda)f(y)$$

for all pairs of distinct points x, $y \in X$ and all $\lambda \in]0, 1[$.

A quasiconvex function f is called *strictly quasiconvex* if inequality (6.2) is strict for all pairs of distinct points x, $y \in X$ and all $\lambda \in]0, 1[$.

A function f is called *strictly concave* (*strictly quasiconcave*) if the function $-f$ is strictly convex (strictly quasiconvex).

EXAMPLE 6.1.4. A quadratic form $q(x) = \langle x, Qx\rangle = \sum_{i,j=1}^{n} q_{ij}x_ix_j$ on $\mathbf{R}^n$ is convex if and only if it is positive semidefinite, and the form is strictly convex if and only if it is positive definite. This follows from the identity

$$\big(\lambda x_i+(1-\lambda)y_i\big)\big(\lambda x_j+(1-\lambda)y_j\big) = \lambda x_ix_j+(1-\lambda)y_iy_j-\lambda(1-\lambda)(x_i-y_i)(x_j-y_j)$$

which after multiplication by q_{ij} and summation yields the equality

$$q(\lambda x + (1-\lambda)y) = \lambda q(x) + (1-\lambda)q(y) - \lambda(1-\lambda)q(x-y).$$

The right hand side is $\le \lambda q(x) + (1-\lambda)q(y)$ for all $0 < \lambda < 1$ if and only if $q(x-y) \ge 0$, which holds for all $x \ne y$ if and only if q is positive semidefinite. Strict inequality requires q to be positive definite. □

Jensen's inequality

The inequalities (6.1) and (6.2) are easily extended to convex combinations of more than two points.

Theorem 6.1.3. *Let f be a function and suppose $x = \lambda_1 x_1 + \lambda_2 x_2 + \cdots + \lambda_m x_m$ is a convex combination of the points $x_1, x_2, \dots, x_m$ in the domain of f.*
(a) If f is convex, then

$$f(x) \le \sum_{j=1}^{m} \lambda_j f(x_j). \qquad \textbf{(Jensen's inequality)} \tag{6.3}$$

If f is strictly convex and $\lambda_j > 0$ for all j, then equality prevails in (6.3) *if and only if $x_1 = x_2 = \cdots = x_m$.*
(b) If f is quasiconvex, then

$$f(x) \le \max_{1\le j\le m} f(x_j). \tag{6.4}$$

If f is strictly quasiconvex and $\lambda_j > 0$ for all j, then equality prevails in (6.4) *if and only if $x_1 = x_2 = \cdots = x_m$.*

Proof. (a) To prove the Jensen inequality we may assume that all coefficients λ_j are positive and that all points x_j lie in $\operatorname{dom} f$, because the right hand side of the inequality is infinite if some point x_j lies outside $\operatorname{dom} f$. Then

$$\bigl(x, \sum_{j=1}^{m} \lambda_j f(x_j)\bigr) = \sum_{j=1}^{m} \lambda_j \bigl(x_j, f(x_j)\bigr),$$

and the right sum, being a convex combination of elements in the epigraph $\operatorname{epi} f$, belongs to $\operatorname{epi} f$. So the left hand side is a point in $\operatorname{epi} f$, and this gives us inequality (6.3).

Now assume that f is strictly convex and that we have equality in Jensen's inequality for the convex combination $x = \sum_{j=1}^{m} \lambda_j x_j$, with positive coefficients λ_j and $m \ge 2$. Let $y = \sum_{j=2}^{m} \lambda_j (1-\lambda_1)^{-1} x_j$. Then $x = \lambda_1 x_1 + (1-\lambda_1) y$, and y is a convex combination of x_2, x_3, …, x_m, so it follows from Jensen's inequality that

$$\begin{aligned}\sum_{j=1}^{m} \lambda_j f(x_j) &= f(x) \le \lambda_1 f(x_1) + (1-\lambda_1) f(y) \\ &\le \lambda_1 f(x_1) + (1-\lambda_1) \sum_{j=2}^{m} \lambda_j (1-\lambda_1)^{-1} f(x_j) = \sum_{j=1}^{m} \lambda_j f(x_j).\end{aligned}$$

Since the left hand side and right hand side of this chain of inequalities and equalities are equal, we conclude that equality holds everywhere. Thus, $f(x) = \lambda_1 f(x_1) + (1 - \lambda_1)f(y)$, and since f is strictly convex, this implies that $x_1 = y = x$.

By symmetri, we also have $x_2 = x, \dots, x_m = x$, and hence $x_1 = x_2 = \dots = x_m$.

(b) Suppose f is quasiconvex, and let $\alpha = \max_{1 \le j \le m} f(x_j)$. If any of the points x_j lies outside $\operatorname{dom} f$, then there is nothing to prove since the right hand side of the inequality (6.4) is infinite. In the opposite case, α is a finite number, and each point x_j belongs to the convex sublevel set $\operatorname{sublev}_\alpha f$, and it follows that so does the point x. This proves inequality (6.4).

The proof of the assertion about equality for strictly quasiconvex functions is analogous with the corresponding proof for strictly convex functions. □

6.2 Operations that preserve convexity

We now describe some ways to construct new convex functions from given convex functions.

Conic combination

Theorem 6.2.1. *Suppose that $f\colon X \to \overline{\mathbf{R}}$ and $g\colon X \to \overline{\mathbf{R}}$ are convex functions and that α and β are nonnegative real numbers. Then $\alpha f + \beta g$ is also a convex function.*

Proof. Follows directly from the alternative characterization of convexity in Theorem 6.1.2. □

The set of convex functions on a given set X is, in other words, a convex cone. So every conic combination $\alpha_1 f_1 + \alpha_2 f_2 + \cdots + \alpha_m f_m$ of convex functions on X is convex.

Note that there is no counterpart of this statement for quasiconvex functions − a sum of quasiconvex functions is not necessarily quasiconvex.

Pointwise limit

Theorem 6.2.2. *Suppose that the functions $f_i\colon X \to \overline{\mathbf{R}}$, $i = 1, 2, 3, \dots,$ are convex and that the limit*

$$f(x) = \lim_{i\to\infty} f_i(x)$$

exists as a finite number or ∞ for each $x \in X$. The limit function $f\colon X \to \overline{\mathbf{R}}$ is then also convex.

Proof. Let x and y be two points in X, and suppose $0 < \lambda < 1$. By passing to the limit in the inequality $f_i(\lambda x + (1-\lambda)y) \le \lambda f_i(x) + (1-\lambda)f_i(y)$ we obtain the following inequality

$$f(\lambda x + (1-\lambda)y) \le \lambda f(x) + (1-\lambda)f(y),$$

which tells us that the limit function f is convex. □

Using Theorem 6.2.2, we may extend the result in Theorem 6.2.1 to infinite sums and integrals. For example, a pointwise convergent infinite sum $f(x) = \sum_{i=1}^{\infty} f_i(x)$ of convex functions is convex.

And if $f(x, y)$ is a function that is convex with respect to the variable x on some set X for each y in a set Y, α is a nonnegative function defined on Y, and the integral $g(x) = \int_Y \alpha(y) f(x, y)\, dy$ exists for all $x \in X$, then g is a convex function on X. This follows from Theorem 6.2.2 by writing the integral as a limit of Riemann sums, or more directly, by integrating the inequalites that characterize the convexity of the functions $f(\cdot, y)$.

Composition with affine maps

Theorem 6.2.3. *Suppose $A\colon V \to \mathbf{R}^n$ is an affine map, that Y is a convex subset of $\mathbf{R}^n$, and that $f\colon Y \to \overline{\mathbf{R}}$ is a convex function. The composition $f \circ A$ is then a convex function on its domain of definition $A^{-1}(Y)$*

Proof. Let $g = f \circ A$. Then, for $x_1, x_2 \in A^{-1}(Y)$ and $0 < \lambda < 1$,

$$\begin{aligned} g(\lambda x_1 + (1-\lambda)x_2) &= f(\lambda Ax_1 + (1-\lambda)Ax_2) \leq \lambda f(Ax_1) + (1-\lambda)f(Ax_2) \\ &= \lambda g(x_1) + (1-\lambda)g(x_2), \end{aligned}$$

which shows that the function g is convex. □

The composition $f \circ A$ of a quasiconvex function f and an affine map A is quasiconvex.

Example 6.2.1. The function $x \mapsto e^{c_1x_1+\cdots+c_nx_n}$ is convex on $\mathbf{R}^n$ since it is a composition of a linear map and the convex exponential function $t \mapsto e^t$. □

Pointwise supremum

Theorem 6.2.4. *Let $f_i\colon X \to \overline{\mathbf{R}}$, $i \in I$, be a family of functions, and define the function $f\colon X \to \overline{\mathbf{R}}$ for $x \in X$ by*

$$f(x) = \sup_{i \in I} f_i(x).$$

Then

(i) f is convex if the functions f_i are all convex;
(ii) f is quasiconvex if the functions f_i are all quasiconvex.

Proof. By the least upper bound definition, $f(x) \leq t$ if and only if $f_i(x) \leq t$ for all $i \in I$, and this implies that

$$\operatorname{epi} f = \bigcap_{i \in I} \operatorname{epi} f_i \quad \text{and} \quad \operatorname{sublev}_t f = \bigcap_{i \in I} \operatorname{sublev}_t f_i$$

for all $t \in \mathbf{R}$. The assertions of the theorem are now immediate consequences of the fact that intersections of convex sets are convex. □

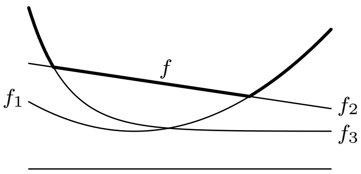

Figure 6.3. $f = \sup f_i$ for a family consisting of three functions.

EXAMPLE 6.2.2. A pointwise maximum of finitely many affine functions, i.e. a function of the form

$$f(x) = \max_{1 \le i \le m} (\langle c_i, x \rangle + a_i),$$

is a convex function and is called a convex *piecewise affine* function. $\square$

EXAMPLE 6.2.3. Examples of convex piecewise affine functions f on $\mathbf{R}^n$ are:

(a) The absolute value of the i:th coordinate of a vector

$$f(x) = |x_i| = \max\{x_i, -x_i\}.$$

(b) The maximum norm

$$f(x) = \|x\|_\infty = \max_{1\le i\le n} |x_i|.$$

(c) The sum of the m largest coordinates of a vector

$$f(x) = \max\{x_{i_1} + \cdots + x_{i_m} \mid 1 \le i_1 < i_2 < \cdots < i_m \le n\}. \qquad \square$$

Composition

Theorem 6.2.5. *Suppose that the function $\phi\colon I \to \overline{\mathbf{R}}$ is defined on a real intervall I that contains the range $f(X)$ of the function $f\colon X \to \mathbf{R}$. The composition $\phi \circ f\colon X \to \overline{\mathbf{R}}$ is convex*

(i) if f is convex and ϕ is convex and increasing;

(ii) if f is concave and ϕ is convex and decreasing.

Proof. The inequality

$$\phi\big(f(\lambda x + (1-\lambda)y)\big) \le \phi\big(\lambda f(x) + (1-\lambda) f(y)\big)$$

holds for $x, y \in X$ and $0 < \lambda < 1$ if either f is convex and ϕ is increasing, or f is concave and ϕ is decreasing. If in addition ϕ is convex, then

$$\phi\big(\lambda f(x) + (1-\lambda) f(y)\big) \le \lambda\phi(f(x)) + (1-\lambda)\phi(f(y)),$$

and by combining the two inequalities above, we obtain the inequality that shows that the function $\phi \circ f$ is a convex. $\square$

There is a corresponding result for quasiconvexity: The composition $\phi \circ f$ is quasiconvex if either f is quasiconvex and ϕ is increasing, or f is quasiconcave and ϕ is decreasing.

Example 6.2.4. The function

$$x \mapsto \mathrm{e}^{x_1^2 + x_2^2 + \cdots + x_k^2},$$

where $1 \le k \le n$, is convex on $\mathbf{R}^n$, since the exponential function is convex and increasing, and positive semidefinite quadratic forms are convex. $\square$

Example 6.2.5. The two functions $t \mapsto 1/t$ and $t \mapsto -\ln t$ are convex and decreasing on the interval $]0, \infty[$. So the function $1/g$ is convex and the function $\ln g$ is concave, if g is a concave and positive function. $\square$

Infimum

Theorem 6.2.6. *Let C be a convex subset of $\mathbf{R}^{n+1}$, and let g be the function defined for $x \in \mathbf{R}^n$ by*

$$g(x) = \inf\{t \in \mathbf{R} \mid (x,t) \in C\},$$

with the usual convention $\inf\emptyset = +\infty$. *Suppose there exists a point* x_0 *in the relative interior of the set*

$$X_0 = \{x \in \mathbf{R}^n \mid g(x) < \infty\} = \{x \in \mathbf{R}^n \mid \exists t \in \mathbf{R}\colon (x,t) \in C\}$$

with a finite function value $g(x_0)$. *Then* $g(x) > -\infty$ *for all* $x \in \mathbf{R}^n$, *and* $g\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ *is a convex function with* X_0 *as its effective domain.*

Proof. Let x be an arbitrary point in X_0. To show that $g(x) > -\infty$, i.e. that the set

$$T_x = \{t \in \mathbf{R} \mid (x,t) \in C\}$$

is bounded below, we first choose a point $x_1 \in \operatorname{rint} X_0$ such that x_0 lies on the open line segment $]x, x_1[$, and write $x_0 = \lambda x + (1-\lambda)x_1$ with $0 < \lambda < 1$. We then fix a real number t_1 such that $(x_1, t_1) \in C$, and for $t \in T_x$ define the number t_0 as $t_0 = \lambda t + (1-\lambda)t_1$. The pair (x_0, t_0) is then a convex combination of the points (x,t) and (x_1,t_1) in C, so

$$g(x_0) \le t_0 = \lambda t + (1-\lambda)t_1,$$

by convexity and the definition of g. We conclude that

$$t \ge \frac{1}{\lambda}\big(g(x_0) - (1-\lambda)t_1\big),$$

and this inequality shows that the set T_x is bounded below.

So the function g has $\overline{\mathbf{R}}$ as codomain, and $\operatorname{dom} g = X_0$. Now, let x_1 and x_2 be arbitrary points in X_0, and let λ_1 and λ_2 be two positive numbers with sum 1. To each $\epsilon > 0$ there exist two real numbers t_1 and t_2 such that the two points (x_1,t_1) and (x_2,t_2) lie in C and $t_1 < g(x_1) + \epsilon$ and $t_2 < g(x_2) + \epsilon$. The convex combination $(\lambda_1 x_1 + \lambda_2 x_2, \lambda_1 t_1 + \lambda_2 t_2)$ of the two points lies in C, too, and

$$g(\lambda_1 x_1 + \lambda_2 x_2) \le \lambda_1 t_1 + \lambda_2 t_2 \le \lambda_1 g(x_1) + \lambda_2 g(x_2) + \epsilon.$$

This means that the point $\lambda_1 x_1 + \lambda_2 x_2$ lies in X_0, and by letting ϵ tend to 0 we conclude that $g(\lambda_1 x_1 + \lambda_2 x_2) \le \lambda_1 g(x_1) + \lambda_2 g(x_2)$. Hence, the set X_0 is convex, and the function g is convex. □

We have seen that the pointwise supremum $f(x) = \sup_{i\in I} f_i(x)$ of an arbitrary family of convex functions is convex. So if $f\colon X \times Y \to \overline{\mathbf{R}}$ is a function with the property that the functions $f(\cdot, y)$ are convex on X for each $y \in Y$, and we define the function g on X by $g(x) = \sup_{y\in Y} f(x,y)$, then g is convex, and this is true without any further conditions on the set Y. Our next theorem shows that the corresponding infimum is a convex function, provided f is convex as a function on the product set $X \times Y$.

Theorem 6.2.7. *Suppose $f\colon X \times Y \to \mathbf{R}$ is a convex function, and for each $x \in X$ define*

$$g(x) = \inf_{y \in Y} f(x, y).$$

If there is a point $x_0 \in \operatorname{rint} X$ such that $g(x_0) > -\infty$, then $g(x)$ is a finite number for each $x \in X$, and $g\colon X \to \mathbf{R}$ is a convex function.

Proof. Suppose X is a subset of $\mathbf{R}^n$ and let

$$C = \{(x, t) \in X \times \mathbf{R} \mid \exists y \in Y\colon f(x, y) \le t\}.$$

C is a convex subset of $\mathbf{R}^{n+1}$, because given two points (x_1, t_1) and (x_2, t_2) in C, and two positive numbers λ_1 and λ_2 with sum 1, there exist two points y_1 and y_2 in the convex set Y such that $f(x_i, y_i) \le t_i$ for $i = 1, 2$, and this implies that

$$f(\lambda_1 x_1 + \lambda_2 x_2, \lambda_1 y_1 + \lambda_2 y_2) \le \lambda_1 f(x_1, y_1) + \lambda_2 f(x_2, y_2) \le \lambda_1 t_1 + \lambda_2 t_2,$$

which shows that the convex combination $\lambda_1(x_1, t_1) + \lambda_2(x_2, t_2)$ lies in C. Moreover, $g(x) = \inf\{t \mid (x, t) \in C\}$, so the corollary follows immediately from Theorem 6.2.6. □

Perspective

Definition. Let $f\colon X \to \overline{\mathbf{R}}$ be a function defined on a cone X in $\mathbf{R}^n$. The function $g\colon X \times \mathbf{R}_{++} \to \overline{\mathbf{R}}$, defined by

$$g(x,s) = sf(x/s),$$

is called the *perspective* of f.

Theorem 6.2.8. *The perspective g of a convex function $f\colon X \to \overline{\mathbf{R}}$ with a convex cone X as domain is a convex function.*

Proof. Let (x,s) and (y,t) be two points in $X \times \mathbf{R}_{++}$, and let α, β be two positive numbers with sum 1. Then

$$\begin{aligned} g\big(\alpha(x,s) + \beta(y,t)\big) &= g(\alpha x + \beta y, \alpha s + \beta t) = (\alpha s + \beta t) f\Big(\frac{\alpha x + \beta y}{\alpha s + \beta t}\Big) \\ &= (\alpha s + \beta t) f\Big(\frac{\alpha s}{\alpha s + \beta t} \cdot \frac{x}{s} + \frac{\beta t}{\alpha s + \beta t} \cdot \frac{y}{t}\Big) \\ &\le \alpha s f\Big(\frac{x}{s}\Big) + \beta t f\Big(\frac{y}{t}\Big) = \alpha g(x,s) + \beta g(y,t). \end{aligned}$$ □

Example 6.2.6. By the previous theorem, $f(x) = x_n q(x/x_n)$ is a convex function on $\mathbf{R}^{n-1} \times \mathbf{R}_{++}$ whenever $q(x)$ is a positive semidefinite quadratic form on $\mathbf{R}^{n-1}$. In particular, by choosing the Euclidean norm as quadratic form, we see that the function

$$x \mapsto (x_1^2 + x_2^2 + \cdots + x_{n-1}^2)/x_n$$

is convex on the open halfspace $\mathbf{R}^{n-1} \times \mathbf{R}_{++}$. □

6.3 Maximum and minimum

Minimum points

For an arbitrary function to decide whether a given point is a global minimum point is an intractable problem, but there are good numerical methods for finding local minimum points if we impose some regularity conditions on the function. This is the reason why convexity plays such an important role in optimization theory. A local minimum of a convex function is namely automatically a global minimum.

Let us recall that a point $x_0 \in X$ is a *local minimum point* of the function $f\colon X \to \overline{\mathbf{R}}$ if there exists an open ball $B = B(x_0; r)$ with center at x_0 such that $f(x) \ge f(x_0)$ for all $x \in X \cap B$. The point is a *(global) minimum point* if $f(x) \ge f(x_0)$ for all $x \in X$.

Theorem 6.3.1. *Suppose that the function $f\colon X \to \overline{\mathbf{R}}$ is convex and that $x_0 \in \operatorname{dom} f$ is a local minimum point of f. Then x_0 is a global minimum point. The minimum point is unique if f is strictly convex.*

Proof. Let $x \in X$ be an arbitrary point different from x_0. Since f is a convex function and $\lambda x + (1-\lambda)x_0 \to x_0$ as $\lambda \to 0$, the following inequalities hold for $\lambda > 0$ sufficiently close to 0:

$$f(x_0) \le f(\lambda x + (1-\lambda)x_0) \le \lambda f(x) + (1-\lambda) f(x_0)$$

(with strict inequality in the last place if f is strictly convex). From this follows at once that $f(x) \ge f(x_0)$ (and $f(x) > f(x_0)$, respectively), which proves that x_0 is a global minimum point (and that there are no other minimum points if the convexity is strict) □

Theorem 6.3.2. *The set of minimum points of a quasiconvex function is convex.*

Proof. The assertion is trivial for functions with no minimum point, since the empty set is convex, and for the function which is identically equal to ∞ on X. So, suppose that the quasiconvex function $f\colon X \to \overline{\mathbf{R}}$ has a minimum point $x_0 \in \operatorname{dom} f$. The set of minimum points is then equal to the sublevel set $\{x \in X \mid f(x) \le f(x_0)\}$, which is convex by definition. □

Maximum points

Theorem 6.3.3. *Suppose $X = \operatorname{cvx} A$ and that the function $f\colon X \to \overline{\mathbf{R}}$ is quasiconvex. Then*

$$\sup_{x\in X} f(x) = \sup_{a \in A} f(a).$$

If the function has a maximum, then there is a maximum point in A.

Proof. Let $x \in X$. Since x is a convex combination $x = \sum_{j=1}^m \lambda_j a_j$ of elements $a_j \in A$,

$$f(x) = f(\sum_{j=1}^m \lambda_j a_j) \le \max_{1\le j\le m} f(a_j) \le \sup_{a\in A} f(a),$$

and it follows that

$$\sup_{x\in X} f(x) \le \sup_{a \in A} f(a).$$

The converse inequality being trivial, since A is a subset of X, we conclude that equality holds.

Moreover, if x is a maximum point, then $f(x) \ge \max_{1\le j\le m} f(a_j)$, and combining this with the inequality above, we obtain $f(x) = \max_{1\le j\le m} f(a_j)$, which means that the maximum is certainly attained at some of the points $a_j \in A$. $\square$

Thus, we can find the maximum of a quasiconvex function whose domain is the convex hull of a finite set A, by just comparing finitely many function values. Of course, this may be infeasible if the set A is very large.

Since compact convex sets coincide with the convex hull of their extreme points, we have the following corollary of the previous theorem.

Corollary 6.3.4. *Suppose that X is a compact convex set and that $f\colon X \to \overline{\mathbf{R}}$ is a quasiconvex function. If f has a maximum, then there is a maximum point among the extreme points of X.*

Example 6.3.1. The quadratic form $f(x_1, x_2) = x_1^2 + 2x_1x_2 + 2x_2^2$ is strictly convex, since it positive definite. The maximum of f on the traingle with vertices at the points $(1, 1)$, $(-2, 1)$ and $(0, 2)$ is attained at some of the vertices. The function values at the vertices are 5, 2, and 8, respectively. The maximum value is hence equal to 8, and it is attained at $(0, 2)$. $\square$

A non-constant realvalued convex function can not attain its maximum at an interior point of its domain, because of the following theorem.

Theorem 6.3.5. *A convex function $f\colon X \to \mathbf{R}$ that attains its maximum at a relative interior point of X, is necessarily constant on X.*

Proof. Suppose f has a maximum at the point $a \in \operatorname{rint} X$, and let x be an arbitrary point in X. Since a is a relative interior point, there exists a point $y \in X$ such that a belongs to the open line segment $]x, y[$, i.e. $a = \lambda x + (1-\lambda)y$ for some number λ satisfying $0 < \lambda < 1$. By convexity and since $f(y) \le f(a)$,

$$f(a) = f(\lambda x + (1-\lambda)y) \le \lambda f(x) + (1-\lambda)f(y) \le \lambda f(x) + (1-\lambda)f(a),$$

with $f(x) \ge f(a)$ as conclusion. Since the converse inequality holds trivially, we have $f(x) = f(a)$. The function f is thus equal to $f(a)$ everywhere. $\square$

6.4 Some important inequalities

Many inequalities can be proved by convexity arguments, and we shall give three important examples.

Arithmetic and geometric mean

Definition. Let $\theta_1, \theta_2, \dots, \theta_n$ be given positive numbers with $\sum_{j=1}^n \theta_j = 1$. The *weighted arithmetic mean* A and the *weighted geometric mean* G of n positive numbers $a_1, a_2, \dots, a_n$ with the given numbers $\theta_1, \theta_2, \dots, \theta_n$ as *weights* are defined as

$$A = \sum_{j=1}^n \theta_j a_j \qquad \text{and} \qquad G = \prod_{j=1}^n a_j^{\theta_j}.$$

The usual arithmetic and geometric means are obtained as special cases by taking all weights equal to $1/n$.

We have the following well-known inequality between the arithmetic and the geometric means.

Theorem 6.4.1. *For all positive numbers $a_1, a_2, \dots, a_n$*

$$G \le A$$

with equality if and only if $a_1 = a_2 = \dots = a_n$.

Proof. Let $x_j = \ln a_j$, so that $a_j = e^{x_j} = \exp(x_j)$. The inequality $G \le A$ is now transformed to the inequality

$$\exp\Bigl(\sum_{j=1}^n \theta_j x_j\Bigr) \le \sum_{j=1}^n \theta_j \exp(x_j),$$

which is Jensen's inequality for the strictly convex exponential function, and equality holds if and only if $x_1 = x_2 = \cdots = x_n$, i.e. if and only if $a_1 = a_2 = \cdots = a_n$. $\square$

EXAMPLE 6.4.1. A lot of maximum and minimum problems can be solved by use of the inequality of arithmetic and geometric means. Here follows a general example.

Let f be a function of the form

$$f(x) = \sum_{i=1}^{m} c_i \Big(\prod_{j=1}^{n} x_j^{\alpha_{ij}}\Big), \qquad x \in \mathbf{R}^n$$

where $c_i > 0$ and α_{ij} are real numbers for all i, j.

The function $g(x) = 16x_1 + 2x_2 + x_1^{-1}x_2^{-2}$, corresponding to $n = 2$, $m = 3$, $c = (16, 2, 1)$ and

$$\alpha = [\alpha_{ij}] = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -1 & -2 \end{bmatrix},$$

serves as a typical example of such a function.

Suppose that we want to minimize $f(x)$ over the set $\{x \in \mathbf{R}^n \mid x > 0\}$. This problem can be attacked in the following way. Let $\theta_1, \theta_2, \dots, \theta_m$ be positive numbers with sum equal to 1, and write

$$f(x) = \sum_{i=1}^{m} \theta_i \Big(\frac{c_i}{\theta_i} \prod_{j=1}^{n} x_j^{\alpha_{ij}}\Big).$$

The inequality of arithmetic and geometric means now gives us the following inequality

$$f(x) \geq \prod_{i=1}^{m} \Big(\big(\frac{c_i}{\theta_i}\big)^{\theta_i} \big(\prod_{j=1}^{n} x_j^{\theta_i \alpha_{ij}}\big) \Big) = C(\theta) \cdot \prod_{j=1}^{n} x_j^{\beta_j}, \tag{6.5}$$

with

$$C(\theta) = \prod_{i=1}^{m} \big(\frac{c_i}{\theta_i}\big)^{\theta_i} \qquad \text{and} \qquad \beta_j = \sum_{i=1}^{m} \theta_i \alpha_{ij}.$$

If it is possible to choose the weights $\theta_i > 0$ so that $\sum_{i=1}^{m} \theta_i = 1$ and

$$\beta_j = \sum_{i=1}^{m} \theta_i \alpha_{ij} = 0 \quad \text{for all } j,$$

then inequality (6.5) becomes

$$f(x) \geq C(\theta),$$

and equality occurs if and only if all the products $\dfrac{c_i}{\theta_i}\displaystyle\prod_{j=1}^{n} x_j^{\alpha_{ij}}$ are equal, a condition that makes it possible to determine x. $\square$

Hölder's inequality

Theorem 6.4.2 (Hölder's inequallity). *Suppose* $1 \leq p \leq \infty$ *and let* q *be the dual index defined by the equality*

$$\frac{1}{p} + \frac{1}{q} = 1.$$

Then

$$|\langle x, y\rangle| = \Big|\sum_{j=1}^{n} x_j y_j\Big| \leq \|x\|_p \|y\|_q$$

for all x, $y \in \mathbf{R}^n$. *Moreover, to each* x *there corresponds a* y *with norm* $\|y\|_q = 1$ *such that* $\langle x, y\rangle = \|x\|_p$.

Remark. Observe that $q = 2$ when $p = 2$. Thus, the Cauchy–Schwarz inequality is a special case of Hölder's inequality.

Proof. The case $p = \infty$ follows directly from the triangle inequality for sums:

$$\Big|\sum_{j=1}^{n} x_j y_j\Big| \le \sum_{j=1}^{n} |x_j||y_j| \le \sum_{j=1}^{n} \|x\|_\infty\, |y_j| = \|x\|_\infty \|y\|_1.$$

So assume that $1 \le p < \infty$. Since $\big|\sum_1^n x_j y_j\big| \le \sum_1^n |x_j||y_j|$, and the vector $(|x_1|, \dots, |x_n|)$ has the same ℓ^p-norm as $(x_1, \dots, x_n)$ and the vector $(|y_1|, \dots, |y_n|)$ has the same ℓ^q-norm as $(y_1, \dots, y_n)$, we can without loss of generality assume that the numbers x_j and y_j are positive.

The function $t \mapsto t^p$ is convex on the interval $[0, \infty[$. Hence,

$$\Big(\sum_{j=1}^{n} \lambda_j t_j\Big)^p \le \sum_{j=1}^{n} \lambda_j t_j^p. \tag{6.6}$$

for all positive numbers $t_1, t_2, \dots, t_n$ and all positive numbers $\lambda_1, \lambda_2, \dots, \lambda_n$ with $\sum_1^n \lambda_j = 1$. Now, let us make the particular choice

$$\lambda_j = \frac{y_j^q}{\sum_{j=1}^n y_j^q} \qquad \text{and} \qquad t_j = \frac{x_j y_j}{\lambda_j}.$$

Then

$$\lambda_j t_j = x_j y_j \qquad \text{and} \qquad \lambda_j t_j^p = \frac{x_j^p y_j^p}{y_j^{(p-1)q}} \Big(\sum_{j=1}^{n} y_j^q\Big)^{p-1} = x_j^p \Big(\sum_{j=1}^{n} y_j^q\Big)^{p-1},$$

which inserted in the inequality (6.6) gives

$$\Big(\sum_{j=1}^{n} x_j y_j\Big)^p \le \sum_{j=1}^{p} x_j^p \Big(\sum_{j=1}^{n} y_j^q\Big)^{p-1},$$

and we obtain Hölder's inequality by raising both sides to $1/p$.

It is easy to verify that Hölder's inequality holds with equality and that $\|y\|_q = 1$ if we choose y as follows:

$x = 0:$ All y with norm equal to 1.

$$x \ne 0,\ \ 1 \le p < \infty: \qquad y_j = \begin{cases} \|x\|_p^{-p/q} |x_j|^p / x_j & \text{if } x_j \ne 0, \\ 0 & \text{if } x_j = 0. \end{cases}$$

$$x \ne 0,\ \ p = \infty: \qquad y_j = \begin{cases} |x_j|/x_j & \text{if } j = j_0, \\ 0 & \text{if } j \ne j_0, \end{cases}$$

where j_0 is an index such that $|x_{j_0}| = \|x\|_\infty$. □

Theorem 6.4.3 (Minkowski's inequality)**.** *Suppose $p \geq 1$ and let x and y be arbitrary vectors in $\mathbf{R}^n$. Then*

$$\|x+y\|_p \leq \|x\|_p + \|y\|_p.$$

Proof. Consider the linear forms $x \mapsto f_a(x) = \langle a, x\rangle$ for vectors $a \in \mathbf{R}^n$ satisfying $\|a\|_q = 1$. By Hölder's inequality,

$$f_a(x) \leq \|a\|_q \|x\|_p \leq \|x\|_p,$$

and for each x there exists a vector a with $\|a\|_q = 1$ such that Hölder's inequality holds with equality, i.e. such that $f_a(x) = \|x\|_p$. Thus

$$\|x\|_p = \sup\{f_a(x) \mid \|a\|_q = 1\},$$

and hence, $f(x) = \|x\|_p$ is a convex function by Theorem 6.2.4. Positive homogenouity is obvious, and positive homogeneous convex functions are subadditive, so the proof of Minkowski's inequality is now complete. □

6.5 Solvability of systems of convex inequalities

The solvability of systems of linear inequalities was discussed in Chapter 3. Our next theorem is kind of a generalization of Theorem 3.3.7 and treats the solvability of a system of convex and affine inequalities.

Theorem 6.5.1. *Let $f_i \colon \Omega \to \mathbf{R}$, $i = 1, 2, \dots, m$, be a family of convex functions defined on a convex subset Ω of $\mathbf{R}^n$.*

Let p be an integer in the interval $1 \leq p \leq m$, and suppose if $p < m$ that the functions f_i are restrictions to Ω of affine *functions for $i \geq p+1$ and that the set*

$$\{x \in \operatorname{rint}\Omega \mid f_i(x) \leq 0 \ \textit{for}\ i = p+1, \dots, m\}$$

is nonempty. The following two assertions are then equivalent:

(i) The system

$$\begin{cases} f_i(x) < 0, & i = 1, 2, \dots, p \\ f_i(x) \leq 0, & i = p+1, \dots, m \end{cases}$$

has no solution $x \in \Omega$.

(ii) There exist nonnegative numbers $\lambda_1, \lambda_2, \dots, \lambda_m$, with at least one of the numbers $\lambda_1, \lambda_2, \dots, \lambda_p$ being nonzero, such that

$$\sum_{i=1}^m \lambda_i f_i(x) \geq 0$$

for all $x \in \Omega$.

Remark. The system of inequalities must contain at least one strict inequality, and all inequalities are allowed to be strict (the case $p = m$).

Proof. If the system (i) has a solution x, then the sum in (ii) is obviously negative for the same x, since at least one of its terms is negative and the others are non-positive. Thus, (ii) implies (i).

To prove the converse implication, we assume that the system (i) has no solution and define M to be the set of all $y = (y_1, y_2, \dots, y_m) \in \mathbf{R}^m$ such that the system

$$\begin{cases} f_i(x) < y_i, & i = 1, 2, \dots, p \\ f_i(x) = y_i, & i = p+1, \dots, m \end{cases}$$

has a solution $x \in \Omega$.

The set M is convex, for if y' and y'' are two points in M, $0 \le \lambda \le 1$, and x', x'' are solutions in Ω of the said systems of inequalities and equalities with y' and y'', respectively, as right hand members, then $x = \lambda x' + (1-\lambda)x'' \in \Omega$ will be a solution of the system with $\lambda y' + (1-\lambda)y''$ as its right hand member, due to the convexity and affinity of the functions f_i for $i \le p$ and $i > p$, respectively.

Our assumptions concerning the system (i) imply that $M \cap \mathbf{R}^m_- = \emptyset$. Since $\mathbf{R}^m_-$ is a polyhedron, there exist, by the separation theorem 5.5.2, a nonzero vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_m)$ and a real number α such that the hyperplane $H = \{y \mid \langle \lambda, y \rangle = \alpha\}$ separates M and $\mathbf{R}^m_-$ and does not contain M as subset. We may assume that

$$\lambda_1 y_1 + \lambda_2 y_2 + \dots + \lambda_m y_m \begin{cases} \geq \alpha & \text{for all } y \in M, \\ \leq \alpha & \text{for all } y \in \mathbf{R}^m_-. \end{cases}$$

By first choosing $y = 0$, we see that $\alpha \geq 0$, and by then choosing $y = t\mathbf{e}_i$, where $\mathbf{e}_i$ is the i:th standard basis vector in $\mathbf{R}^m$, and letting t tend to $-\infty$, we conclude that $\lambda_i \geq 0$ for all i.

For each $x \in \Omega$ and $\epsilon > 0$,

$$y = (f_1(x) + \epsilon, \dots, f_p(x) + \epsilon, f_{p+1}(x), \dots, f_m(x))$$

is a point in M. Consequently,

$$\lambda_1(f_1(x) + \epsilon) + \dots + \lambda_p(f_p(x) + \epsilon) + \lambda_{p+1} f_{p+1}(x) + \dots + \lambda_m f_m(x) \geq \alpha \geq 0,$$

and by letting ϵ tend to zero, we obtain the inequality

$$\lambda_1 f_1(x) + \lambda_2 f_2(x) + \dots + \lambda_m f_m(x) \geq 0$$

for all $x \in \Omega$.

If $p = m$, we are done since the vector $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_m)$ is then nonzero, but it remains to prove that some of the coefficients $\lambda_1, \lambda_2, \dots, \lambda_p$ is nonzero when $p < m$. Assume the contrary, i.e. that $\lambda_1 = \lambda_2 = \dots = \lambda_p = 0$, and let

$$h(x) = \sum_{i=p+1}^{m} \lambda_i f_i(x).$$

The function h is affine, and $h(x) \geq 0$ for all $x \in \Omega$. Furthermore, by the assumptions of the theorem, there exists a point x_0 in the relative interior of Ω such that $f_i(x_0) \leq 0$ for all $i \geq p+1$, which implies that $h(x_0) \leq 0$. Thus, $h(x_0) = 0$. This means that the restriction $h|_\Omega$, which is a concave function since h is affine, attains its minimum at a relative interior point, and according to Theorem 6.3.5 (applied to the function $-h|_\Omega$), this implies that the function h is constant and equal to 0 on Ω.

But to each $y \in M$ there corresponds a point $x \in \Omega$ such that $y_i = f_i(x)$ for $i = p+1, \dots m$, and this implies that

$$\langle \lambda, y \rangle = \textstyle\sum_{i=p+1}^{m} \lambda_i f_i(x) = h(x) = 0.$$

We conclude that $\alpha = 0$ and that the hyperplane H contains M, which is a contradiction. Thus, at least one of the coefficients $\lambda_1, \lambda_2, \dots, \lambda_p$ has to be nonzero, and the theorem is proved. □

6.6 Continuity

A real-valued convex function is automatically continuous at all relative interior points of the domain. More precisely, we have the following theorem.

Theorem 6.6.1. *Suppose $f\colon X \to \overline{\mathbf{R}}$ is a convex function and that a is a point in the relative interior of* $\operatorname{dom} f$*. Then there exist a relative open neighborhood U of a in* $\operatorname{dom} f$ *and a constant M such that*

$$|f(x) - f(a)| \leq M\|x - a\|$$

for all $x \in U$. Hence, f is continuous on the relative interior of $\operatorname{dom} f$*.*

Proof. We start by proving a special case of the theorem and then show how to reduce the general case to this special case.

1. So first assume that X is an open subset of $\mathbf{R}^n$, that $\operatorname{dom} f = X$, i.e. that f is a real-valued convex function, that $a = 0$, and that $f(0) = 0$. We will show that if we choose the number $r > 0$ such that the closed hypercube

$$K(r) = \{x \in \mathbf{R}^n \mid \|x\|_\infty \leq r\}$$

is included in X, then there is a constant M such that

$$|f(x)| \leq M\|x\| \tag{6.7}$$

for all x in the closed ball $\overline{B}(0; r) = \{x \in \mathbf{R}^n \mid \|x\| \leq r\}$, where $\|\cdot\|$ is the usual Euclidean norm.

The hypercube $K(r)$ has 2^n extreme points (vertices). Let L denote the largest of the function values of f at these extreme points. Since the convex hull of the extreme points is equal to $K(r)$, it follows from Theorem 6.3.3 that

$$f(x) \leq L$$

for all $x \in K(r)$, and thereby also for all $x \in \overline{B}(0; r)$, because $\overline{B}(0; r)$ is a subset of $K(r)$.

We will now make this inequality sharper. To this end, let x be an arbitrary point in $\overline{B}(0; r)$ different from the center 0. The halfline from 0 through x intersects the boundary of $\overline{B}(0; r)$ at the point

$$y = \frac{r}{\|x\|}x,$$

and since x lies on the line segment $[0, y]$, x is a convex combination of its end points. More precisely, $x = \lambda y + (1 - \lambda)0$ with $\lambda = \|x\|/r$. Therefore, since f is convex,

$$f(x) \leq \lambda f(y) + (1 - \lambda) f(0) = \lambda f(y) \leq \lambda L = \frac{L}{r}\|x\|.$$

The above inequality holds for all $x \in \overline{B}(0;r)$. To prove the same inequality with $f(x)$ replaced by $|f(x)|$, we use the fact that the point $-x$ belongs to $\overline{B}(0;r)$ if x does so, and the equality $0 = \frac{1}{2}x + \frac{1}{2}(-x)$. By convexity,

$$0 = f(0) \le \frac{1}{2}f(x) + \frac{1}{2}f(-x) \le \frac{1}{2}f(x) + \frac{L}{2r}\|x\|,$$

which simplifies to the inequality

$$f(x) \ge -\frac{L}{r}\|-x\| = -\frac{L}{r}\|x\|.$$

This proves that inequality (6.7) holds for $x \in \overline{B}(0;r)$ with $M = L/r$.

2. We now turn to the general case. Let n be the dimension of the set $\operatorname{dom} f$. The affine hull of $\operatorname{dom} f$ is equal to the set $a + V$ for some n-dimensional linear subspace V, and since V is isomorphic to $\mathbf{R}^n$, we can obtain a bijective linear map $T\colon \mathbf{R}^n \to V$ by choosing a coordinate system in V.

The inverse image Y of the relative interior of $\operatorname{dom} f$ under the map $y \mapsto a + Ty$ of $\mathbf{R}^n$ onto $\operatorname{aff}(\operatorname{dom} f)$ is an open convex subset of $\mathbf{R}^n$, and Y contains the point 0. Define the function $g\colon Y \to \mathbf{R}$ by

$$g(y) = f(a + Ty) - f(a).$$

Then, g is a convex function, since g is composed by a convex function and an affine function, and $g(0) = 0$.

For $x = a + Ty \in \operatorname{rint}(\operatorname{dom} f)$ we now have $f(x) - f(a) = g(y)$ and $x - a = Ty$, so in order to prove the general case of our theorem, we have to show that there is a constant M such that $|g(y)| \leq M\|Ty\|$ for all y in some neighborhood of 0. But the map $y \mapsto \|Ty\|$ is a norm on $\mathbf{R}^n$, and since all norms are equivalent, it suffices to show that there is a constant M such that

$$|g(y)| \leq M\|y\|$$

for all y in some neighborhood of 0, and that is exactly what we did in step 1 of the proof. So the theorem is proved. □

The following corollary follows immediately from Theorem 6.6.1, because affine sets have no relative boundary points.

Corollary 6.6.2. *A convex function $f\colon X \to \mathbf{R}$ with an affine subset X as domain is continuous.*

For functions f with a closed interval $I = [a, b]$ as domain, convexity imposes no other restrictions on the function value $f(b)$ than that it has to be greater than or equal to $\lim_{x\to b-} f(x)$. Thus, a convex function need not be continuous at the endpoint b, and a similar remark holds for the left endpoint, of course. For example, a function f, that is identically equal to zero on $I \setminus \{a, b\}$, is convex if $f(a) \geq 0$ and $f(b) \geq 0$. Cf. exercise 7.6.

6.7 The recessive subspace of convex functions

Example 6.7.1. Let $f\colon \mathbf{R}^2 \to \mathbf{R}$ be the convex function

$$f(x_1, x_2) = x_1 + x_2 + \mathrm{e}^{(x_1-x_2)^2}.$$

The restrictions of f to lines with direction given by the vector $v = (1, 1)$ are affine functions, since

$$f(x + tv) = f(x_1 + t, x_2 + t) = x_1 + x_2 + 2t + \mathrm{e}^{(x_1-x_2)^2} = f(x) + 2t.$$

Let $V = \{x \in \mathbf{R}^2 \mid x_1 = x_2\}$ be the linear subspace of $\mathbf{R}^2$ spanned by the vector v, and consider the the orthogonal decomposition $\mathbf{R}^2 = V^\perp + V$. Each $x \in \mathbf{R}^2$ has a corresponding unique decomposition $x = y + z$ with $y \in V^\perp$ and $z \in V$, namely

$$y = \tfrac{1}{2}(x_1 - x_2, x_2 - x_1) \text{ and } z = \tfrac{1}{2}(x_1 + x_2, x_1 + x_2).$$

Moreover, since $z = \frac{1}{2}(x_1 + x_2)v = z_1 v$,

$$f(x) = f(y+z) = f(y) + 2z_1 = f|_{V^\perp}(y) + 2z_1.$$

So there is a corresponding decomposition of f as a sum of the restriction of f to $V^\perp$ and a linear function on V. It is easily verified that the vector $(v, 2) = (1, 1, 2)$ spans the recessive subspace $\operatorname{lin}(\operatorname{epi} f)$, and that V is equal to the image $P_1(\operatorname{lin}(\operatorname{epi} f))$ of $\operatorname{lin}(\operatorname{epi} f)$ under the projection P_1 of $\mathbf{R}^2 \times \mathbf{R}$ onto the first factor $\mathbf{R}^2$. □

The result in the previous example can be generalized, and in order to describe this generalization we need a definition.

Definition. Let $f\colon X \to \overline{\mathbf{R}}$ be a function defined on a subset X of $\mathbf{R}^n$. The linear subspace $V_f = P_1(\operatorname{lin}(\operatorname{epi} f))$, where $P_1\colon \mathbf{R}^n \times \mathbf{R} \to \mathbf{R}^n$ is the projection of $\mathbf{R}^n \times \mathbf{R}$ onto its first factor $\mathbf{R}^n$, is called the *recessive subspace* of the function f.

Theorem 6.7.1. *Let f be a convex function with recessive subspace V_f.*

(i) A vector v belongs to V_f if and only if there is a unique number α_v such that (v, α_v) belongs to the recessive subspace $\operatorname{lin}(\operatorname{epi} f)$ *of the epigraph of the function.*

(ii) The map $g\colon V_f \to \mathbf{R}$, defined by $g(v) = \alpha_v$ for $v \in V_f$, is linear.

(iii) $\operatorname{dom} f = \operatorname{dom} f + V_f$.

(iv) $f(x+v) = f(x) + g(v)$ for all $x \in \operatorname{dom} f$ and all $v \in V_f$.

(v) If the function f is differentiable at $x \in \operatorname{dom} f$ then $g(v) = \langle f'(x), v\rangle$ for all $v \in V_f$.

(vi) Suppose V is a linear subspace, that $h\colon V \to \mathbf{R}$ is a linear map, that $\operatorname{dom} f + V \subseteq \operatorname{dom} f$, *and that $f(x+v) = f(x) + h(v)$ for all $x \in \operatorname{dom} f$ and all $v \in V$. Then, $V \subseteq V_f$.*

Proof. (i) By definition, $v \in V_f$ if and only if there is a real number α_v such that $(v, \alpha_v) \in \operatorname{lin}(\operatorname{epi} f)$. To prove that the number α_v is uniquely determined by $v \in V_f$, we assume that the pair (v, β) also lies in $\operatorname{lin}(\operatorname{epi} f)$.

The point $(x + tv, f(x) + t\alpha_v)$ belongs to the epigraph for each $x \in \operatorname{dom} f$ and each $t \in \mathbf{R}$, i.e.

$$x + tv \in \operatorname{dom} f \quad \text{and} \quad f(x+tv) \le f(x) + t\alpha_v. \tag{6.8}$$

Hence, $(x + tv, f(x+tv))$ is a point in the epigraph, and our assumption $(v, \beta) \in \operatorname{lin}(\operatorname{epi} f)$ now implies that $(x + tv - tv, f(x+tv) - t\beta)$ is a point in i epi f, too. We conclude that

$$f(x) \le f(x+tv) - t\beta \tag{6.9}$$

for all $t \in \mathbf{R}$. By combining the two inequalities (6.8) and (6.9), we obtain the inequality $f(x) \leq f(x) + (\alpha_v - \beta)t$, which holds for all $t \in \mathbf{R}$. This is possible only if $\beta = \alpha_v$, and proves the uniqueness of the number α_v.

(ii) Let, as before, P_1 be the projection of $\mathbf{R}^n \times \mathbf{R}$ onto $\mathbf{R}^n$, and let P_2 be the projection of $\mathbf{R}^n \times \mathbf{R}$ onto the second factor $\mathbf{R}$. The uniqueness result (i) implies that the restriction of P_1 to the linear subspace $\operatorname{lin}(\operatorname{epi} f)$ is a bijective linear map onto V_f. Let Q denote the inverse of this restriction; the map g is then equal to the composition $P_2 \circ Q$ of the two linear maps P_2 and Q, and this implies that g is a linear function.

(iii) The particular choice of $t = 1$ in (6.8) yields the implication

$$x \in \operatorname{dom} f \;\&\; v \in V_f \;\Rightarrow\; x + v \in \operatorname{dom} f,$$

which proves the inclusion $\operatorname{dom} f + V_f \subseteq \operatorname{dom} f$, and the converse inclusion is of course trivial.

(iv) By choosing $t = 1$ in the inequalities (6.8) and(6.9) and using the fact that $\alpha_v = \beta = g(v)$, we obtain the two inequalities $f(x + v) \leq f(x) + g(v)$ and $f(x) \leq f(x + v) - g(v)$, which when combined prove assertion (iv).

(v) Consider the restriction $\phi(t) = f(x + tv)$ of the function f to the line through the point x with direction $v \in V_f$. By (iii), ϕ is defined for all $t \in \mathbf{R}$, and by (iv), $\phi(t) = f(x) + tg(v)$. Hence, $\phi'(0) = g(v)$. But if f is differentiable at x, then we also have $\phi'(0) = \langle f'(x), v\rangle$ according to the chain rule, and this proves our assertion (v).

(vi) Suppose $v \in V$. If (x, s) is an arbitrary point in the epigraph $\operatorname{epi} f$, then $f(x + tv) = f(x) + h(tv) \le s + th(v)$, which means that the point $(x + tv, s + th(v))$ lies in $\operatorname{epi} f$ for every real number t. This proves that $(v, h(v))$ belongs to $\operatorname{lin}(\operatorname{epi} f)$ and, consequently, that v is a vector in V_f. □

By our next theorem, every convex function is the sum of a convex function with a trivial recessive subspace and a linear function.

Theorem 6.7.2. *Suppose that f is a convex function with recessive subspace V_f. Let $\tilde{f}$ denote the restriction of f to $\operatorname{dom} f \cap V_f^{\perp}$, and let $g\colon V_f \to \mathbf{R}$ be the linear function defined in Theorem 6.7.1. The recessive subspace $V_{\tilde{f}}$ of $\tilde{f}$ is then trivial, i.e. equal to $\{0\}$, $\operatorname{dom} f = \operatorname{dom} f \cap V_f^{\perp} + V_f$, and*

$$f(y + z) = \tilde{f}(y) + g(z)$$

for all $y \in \operatorname{dom} f \cap V_f^{\perp}$ and all $z \in V_f$.

Proof. Each $x \in \mathbf{R}^n$ has a unique decomposition $x = y + z$ with $y \in V_f^{\perp}$ and $z \in V_f$, and if $x \in \operatorname{dom} f$ then $y = x - z \in \operatorname{dom} f + V_f = \operatorname{dom} f$, by Theorem 6.7.1, and hence $y \in \operatorname{dom} f \cap V_f^{\perp}$. This proves that $\operatorname{dom} f = \operatorname{dom} f \cap V_f^{\perp} + V_f$.

The equality $f(y + z) = \tilde{f}(y) + g(z)$ now follows from (iv) in Theorem 6.7.1, so it only remains to prove that $V_{\tilde{f}} = \{0\}$. Suppose $v \in V_{\tilde{f}}$, and let x_0 be an arbitrary point in $\operatorname{dom} \tilde{f}$. Then $x_0 + v$ lies in $\operatorname{dom} \tilde{f}$, too, and since $\operatorname{dom} \tilde{f} \subseteq V_f^{\perp}$ and $V_f^{\perp}$ is a linear subspace, we conclude that $v = (x_0 + v) - x_0$ is a vector in $V_f^{\perp}$. This proves the inclusion $V_{\tilde{f}} \subseteq V_f^{\perp}$.

Theorem 6.7.1 gives us two linear functions $g\colon V_f \to \mathbf{R}$ and $\tilde{g}\colon V_{\tilde{f}} \to \mathbf{R}$ such that $f(x + v) = f(x) + g(v)$ for all $x \in \operatorname{dom} f$ and all $v \in V_f$, and $\tilde{f}(y + w) = \tilde{f}(y) + \tilde{g}(w)$ for all $y \in \operatorname{dom} f \cap V_f^{\perp}$ and all $w \in V_{\tilde{f}}$.

Now, let w be an arbitrary vector in $V_{\tilde{f}}$ and x be an arbitrary point in $\operatorname{dom} f$, and write x as $x = y + v$ with $y \in \operatorname{dom} f \cap V_f^{\perp}$ and $v \in V_f$. The point $y + w$ lies in $\operatorname{dom} f \cap V_f^{\perp}$, and we get the following identities:

$$\begin{aligned} f(x + w) &= f(y + v + w) = f(y + w + v) = f(y + w) + g(v) \\ &= \tilde{f}(y + w) + g(v) = \tilde{f}(y) + \tilde{g}(w) + g(v) \\ &= f(y) + g(v) + \tilde{g}(w) = f(x) + \tilde{g}(w). \end{aligned}$$

Therefore, $V_{\tilde{f}} \subseteq V_f$, by Theorem 6.7.1 (v). Hence, $V_{\tilde{f}} \subseteq V_f^{\perp} \cap V_f = \{0\}$, which proves that $V_{\tilde{f}} = \{0\}$. □

6.8 Closed convex functions

Definition. A convex function is called *closed* if it has a closed epigraph.

Theorem 6.8.1. *A convex function $f\colon X \to \overline{\mathbf{R}}$ is closed if and only if all its sublevel sets are closed.*

Proof. Suppose that X is a subset of $\mathbf{R}^n$ and that f is a closed function. Let

$$X_\alpha = \operatorname{sublev}_\alpha f = \{x \in X \mid f(x) \le \alpha\}$$

be an arbitrary nonempty sublevel set of f, and define Y_α to be the set

$$Y_\alpha = \operatorname{epi} f \cap \{(x, x_{n+1}) \in \mathbf{R}^n \times \mathbf{R} \mid x_{n+1} \le \alpha\}.$$

The set Y_α is closed, being the intersection between the closed epigraph $\operatorname{epi} f$ and a closed halfspace, and $X_\alpha = P(Y_\alpha)$, where $P\colon \mathbf{R}^n \times \mathbf{R} \to \mathbf{R}^n$ is the projection $P(x, x_{n+1}) = x$.

Obviously, the recession cone $\operatorname{recc} Y_\alpha$ contains no nonzero vector of the form $v = (0, v_{n+1})$, i.e. no nonzero vector in the null space $\mathcal{N}(P) = \{0\} \times \mathbf{R}$ of the projection P. Hence, $(\operatorname{recc} Y_\alpha) \cap \mathcal{N}(P) = \{0\}$, so it follows from Theorem 2.7.10 that the sublevel set X_α is closed.

To prove the converse, assume that all sublevel sets are closed, and let (x_0, y_0) be a boundary point of $\operatorname{epi} f$. Let $\big((x_k, y_k)\big)_1^\infty$ be a sequence of points in $\operatorname{epi} f$ that converges to (x_0, y_0), and let ϵ be an arbitrary positive number. Then, since $y_k \to y_0$ as $k \to \infty$, $f(x_k) \le y_k \le y_0 + \epsilon$ for all sufficiently large k, so the points x_k belong to the sublevel set $\{x \in X \mid f(x) \le y_0 + \epsilon\}$ for all sufficiently large k. The sublevel set being closed, it follows that the limit point x_0 lies in the same sublevel set, i.e. $x_0 \in X$ and $f(x_0) \le y_0 + \epsilon$, and since $\epsilon > 0$ is arbitrary, we conclude that $f(x_0) \le y_0$. Hence, (x_0, y_0) is a point in $\operatorname{epi} f$. So $\operatorname{epi} f$ contains all its boundary points and is therefore a closed set. □

Corollary 6.8.2. *Continuous convex functions $f\colon X \to \mathbf{R}$ with closed domains X are closed functions.*

Proof. Follows immediately from Theorem 6.8.1, because the sublevel sets of real-valued continuous functions with closed domains are closed sets. □

Theorem 6.8.3. *All nonempty sublevel sets of a closed convex function have the same recession cone and the same recessive subspace. Hence, all sublevel sets are bounded if one of the nonempty sublevel sets is bounded.*

Proof. Let $f\colon X \to \overline{\mathbf{R}}$ be a closed convex function, and suppose that x_0 is a point in the sublevel set $X_\alpha = \{x \in X \mid f(x) \le \alpha\}$. Since X_α and $\operatorname{epi} f$ are closed convex sets and (x_0, α) is a point in $\operatorname{epi} f$, we obtain the following equivalences:

$$\begin{aligned}
v \in \operatorname{recc} X_\alpha &\Leftrightarrow x_0 + tv \in X_\alpha \quad \text{for all } t \in \mathbf{R}_+ \\
&\Leftrightarrow f(x_0 + tv) \le \alpha \quad \text{for all } t \in \mathbf{R}_+ \\
&\Leftrightarrow (x_0 + tv, \alpha) \in \operatorname{epi} f \quad \text{for all } t \in \mathbf{R}_+ \\
&\Leftrightarrow (x_0, \alpha) + t(v, 0) \in \operatorname{epi} f \quad \text{for all } t \in \mathbf{R}_+ \\
&\Leftrightarrow (v, 0) \in \operatorname{recc}(\operatorname{epi} f),
\end{aligned}$$

with the conclusion that the recession cone

$$\operatorname{recc} X_\alpha = \{v \in \mathbf{R}^n \mid (v, 0) \in \operatorname{recc}(\operatorname{epi} f)\}$$

does not depend on α as long as $X_\alpha \ne \emptyset$. Of course, the same is then true for the recessive subspace

$$\operatorname{lin} X_\alpha = \operatorname{recc} X_\alpha \cap (-\operatorname{recc} X_\alpha) = \{v \in \mathbf{R}^n \mid (v, 0) \in \operatorname{lin}(\operatorname{epi} f)\}.$$

The statement concerning bounded sublevel sets follows from the fact that a closed convex set is bounded if and only if its recession cone is equal to the zero cone $\{0\}$. □

Theorem 6.8.4. *A convex function f, which is bounded on an affine subset M, is constant on M.*

Proof. Let $M = a + U$, where U is a linear subspace, and consider the restriction $g = f|_M$ of f to M. The function g is continuous since all points of M are relative interior points, and closed since the domain M is a closed set. Let $\alpha = \sup\{g(x) \mid x \in M\}$; then $\{x \mid g(x) \leq \alpha\} = M$, so by the previous theorem, all nonempty sublevel sets of g has $\operatorname{lin} M$, that is the subspace U, as their recessive subspace.

Let now x_0 be an arbitrary point in M. Since the recessive subspace of the particular sublevel set $\{x \mid g(x) \leq g(x_0)\}$ is equal to U, we conclude that $g(x_0 + u) \leq g(x_0)$ for all $u \in U$. Hence, $g(x) \leq g(x_0)$ for all $x \in M$, which means that x_0 is a maximum point of g. Since $x_0 \in M$ is arbitrary, all points in M are maximum points, and this implies that g is constant on M. □

6.9 The support function

Definition. Let A be a nonempty subset of $\mathbf{R}^n$. The function $S_A \colon \mathbf{R}^n \to \overline{\mathbf{R}}$, defined by

$$S_A(x) = \sup\{\langle y, x\rangle \mid y \in A\}$$

(with the usual convention that $S_A(x) = \infty$ if the function $y \mapsto \langle y, x\rangle$ is unbounded above on A) is called the *support function* of the set A.

Theorem 6.9.1. *(a) The support function S_A is a closed convex function.*

(b) Suppose A and B are nonempty subsets of $\mathbf{R}^n$, that $\alpha > 0$ and that $C \colon \mathbf{R}^n \to \mathbf{R}^m$ is a linear map. Then

(i) $$S_A = S_{\operatorname{cvx} A} = S_{\operatorname{cl}(\operatorname{cvx} A)}$$

(ii) $$S_{\alpha A} = \alpha S_A$$

(iii) $$S_{A+B} = S_A + S_B$$

(iv) $$S_{A\cup B} = \max\{S_A, S_B\}$$

(v) $$S_{C(A)} = S_A \circ C^T.$$

Proof. (a) The support function S_A is closed and convex, because its epigraph

$$\operatorname{epi} S_A = \{(x,t) \mid \langle y, x\rangle \leq t \text{ for all } y \in A\} = \bigcap_{y\in A}\{(x,t) \mid \langle y, x\rangle \leq t\}$$

is closed, being the intersection of a family of closed halfspaces in $\mathbf{R}^n \times \mathbf{R}$.

(b) Since linear forms are convex, it follows from Theorem 6.3.3 that

$$S_A(x) = \sup\{\langle x, y\rangle \mid y \in A\} = \sup\{\langle x, y\rangle \mid y \in \operatorname{cvx} A\} = S_{\operatorname{cvx} A}(x)$$

for all $x \in \mathbf{R}^n$. Moreover, if a function f is continuous on the closure of a set X, then $\sup_{y\in X} f(y) = \sup_{y\in \operatorname{cl} X} f(y)$, and linear forms are of course continuous. Therefore, $S_{\operatorname{cvx} A}(x) = S_{\operatorname{cl}(\operatorname{cvx} A)}(x)$ for all x.

This proves the identity (i), and the remaining identities are obtained as follows:

$$S_{\alpha A}(x) = \sup_{y\in\alpha A} \langle y, x\rangle = \sup_{y\in A} \langle \alpha y, x\rangle = \alpha \sup_{y\in A} \langle y, x\rangle = \alpha S_A(x).$$

$$\begin{aligned} S_{A+B}(x) &= \sup_{y\in A+B} \langle y, x\rangle = \sup_{y_1\in A,\, y_2\in B} \langle y_1 + y_2, x\rangle \\ &= \sup_{y_1\in A,\, y_2\in B} (\langle y_1, x\rangle + \langle y_2, x\rangle) = \sup_{y_1\in A} \langle y_1, x\rangle + \sup_{y_2\in B} \langle y_2, x\rangle \\ &= S_A(x) + S_B(x). \end{aligned}$$

$$\begin{aligned} S_{A\cup B}(x) &= \sup_{y\in(A\cup B)} \langle y, x\rangle = \max\{\sup_{y\in A} \langle y, x\rangle, \sup_{y\in B} \langle y, x\rangle\} \\ &= \max\{S_A(x), S_B(x)\}. \end{aligned}$$

$$S_{C(A)}(x) = \sup_{y\in C(A)} \langle y, x\rangle = \sup_{z\in A} \langle Cz, x\rangle = \sup_{z\in A} \langle z, C^T x\rangle = S_A(C^T x). \qquad \square$$

Example 6.9.1. The support function of a closed interval $[a, b]$ on the real line is given by

$$S_{[a,b]}(x) = S_{\{a,b\}}(x) = \max\{ax, bx\},$$

since $[a, b] = \operatorname{cvx}\{a, b\}$. $\square$

Example 6.9.2. In order to find the support function of the closed unit ball $\overline{B}_p = \{x \in \mathbf{R}^n \mid \|x\|_p \le 1\}$ with respect to the ℓ^p-norm, we use Hölder's inequality, obtaining

$$S_{\overline{B}_p}(x) = \sup\{\langle x, y\rangle \mid \|y\|_p \le 1\} = \|x\|_q,$$

where the relation between p and q is given by the equation $1/p+1/q = 1$. $\square$

Closed convex sets are completely characterized by their support functions, due to the following theorem.

Theorem 6.9.2. *Suppose that X_1 and X_2 are two nonempty closed convex subsets of $\mathbf{R}^n$ with support functions S_{X_1} and S_{X_2}, respectively. Then*

(a) $$X_1 \subseteq X_2 \;\Leftrightarrow\; S_{X_1} \leq S_{X_2}$$

(b) $$X_1 = X_2 \;\Leftrightarrow\; S_{X_1} = S_{X_2}.$$

Proof. Assertion (b) is an immediate consequence of (a), and the implication $X_1 \subseteq X_2 \Rightarrow S_{X_1} \leq S_{X_2}$ is trivial, so it only remains to prove the converse implication, or equivalently, the implication $X_1 \not\subseteq X_2 \Rightarrow S_{X_1} \not\leq S_{X_2}$.

To prove the latter implication we assume that $X_1 \not\subseteq X_2$, i.e. that there exists a point $x_0 \in X_1 \setminus X_2$. The point x_0 is strictly separable from the closed convex set X_2, which means that there exist a vector $c \in \mathbf{R}^n$ and a number b such that $\langle x, c\rangle \leq b$ for all $x \in X_2$ while $\langle x_0, c\rangle > b$. Consequently,

$$S_{X_1}(c) \geq \langle x_0, c\rangle > b \geq \sup\{\langle x, c\rangle \mid x \in X_2\} = S_{X_2}(c),$$

which shows that $S_{X_1} \not\leq S_{X_2}$. □

By combining the previous theorem with property (i) of Theorem 6.9.1, we obtain the following corollary.

Corollary 6.9.3. *Let A and B be two nonempty subsets of $\mathbf{R}^n$. Then,*

$$S_A = S_B \Leftrightarrow \operatorname{cl}(\operatorname{cvx} A) = \operatorname{cl}(\operatorname{cvx} B).$$

6.10 The Minkowski functional

Let X be a convex subset of $\mathbf{R}^n$ with 0 as an interior point of X. Consider the sets tX for $t \geq 0$. This is an increasing family of sets, whose union equals all of $\mathbf{R}^n$, i.e. $0 \leq s < t \Rightarrow sX \subseteq tX$ and $\bigcup_{t\geq 0} tX = \mathbf{R}^n$.

The family is increasing, because using the convexity of the sets tX and the fact that they contain 0, we obtain the following inclusions for $0 \leq s < t$:

$$sX = \frac{s}{t}(tX) + (1 - \frac{s}{t})\, 0 \subseteq \frac{s}{t}(tX) + (1 - \frac{s}{t})(tX) \subseteq tX.$$

That the union equals $\mathbf{R}^n$ only depends on 0 being an interior point of X. For let $\overline{B}(0; r_0)$ be a closed ball centered at 0 and contained in X. An arbitrary point $x \in \mathbf{R}^n$ will then belong to the set $r_0^{-1}\|x\|X$ since $r_0\|x\|^{-1}x$ lies in $\overline{B}(0; r_0)$.

Now fix $x \in \mathbf{R}^n$ and consider the set $\{t \geq 0 \mid x \in tX\}$. This set is an unbounded subinterval of $[0, \infty[$, and it contains the number $r_0^{-1}\|x\|$. We may therefore define a function

$$\phi_X \colon \mathbf{R}^n \to \mathbf{R}_+$$

by letting

$$\phi_X(x) = \inf\{t \geq 0 \mid x \in tX\}.$$

Obviously,

$$\phi_X(x) \leq r_0^{-1}\|x\| \qquad \text{for all } x.$$

Definition. The function $\phi_X \colon \mathbf{R}^n \to \mathbf{R}_+$ is called the *Minkowski functional* of the set X.

Theorem 6.10.1. *The Minkowski functional ϕ_X has the following properties:*

(i) *For all $x, y \in \mathbf{R}^n$ and all $\lambda \in \mathbf{R}_+$,*
  (a) $\phi_X(\lambda x) = \lambda\phi_X(x)$,
  (b) $\phi_X(x + y) \leq \phi_X(x) + \phi_X(y)$.

(ii) *There exists a constant C such that*

$$|\phi_X(x) - \phi_X(y)| \leq C\|x - y\|$$

for all $x, y \in \mathbf{R}^n$.

(iii) $\operatorname{int} X = \{x \in \mathbf{R}^n \mid \phi_X(x) < 1\}$ *and* $\operatorname{cl} X = \{x \in \mathbf{R}^n \mid \phi_X(x) \leq 1\}$.

The Minkowski functional is, in other words, *positive homogeneous*, *subadditive*, and *Lipschitz continuous*. So it is in particular a convex function.

Proof. (i) The equivalence $x \in tX \Leftrightarrow \lambda x \in \lambda tX$, which holds for $\lambda > 0$, together with the fact that $\phi_X(0) = 0$, implies positive homogenouity.

To prove subadditivity we choose, given $\epsilon > 0$, two positive numbers $s < \phi_X(x) + \epsilon$ and $t < \phi_X(y) + \epsilon$ such that $x \in sX$ and $y \in tX$. The point

$$\frac{1}{s+t}(x+y) = \frac{s}{s+t}\,\frac{x}{s} + \frac{t}{s+t}\,\frac{y}{t}$$

is a point in X, by convexity, and it follows that the point $x + y$ belongs to the set $(s + t)X$. This implies that

$$\phi_X(x+y) \le s + t < \phi_X(x) + \phi_X(y) + 2\epsilon,$$

and since this inequality is true for all $\epsilon > 0$, we conclude that

$$\phi_X(x+y) \le \phi_X(x) + \phi_X(y).$$

(ii) We have already noted that the inequality $\phi_X(x) \le C\|x\|$ holds for all x with $C = r_0^{-1}$. By subadditivity,

$$\phi_X(x) = \phi_X(x - y + y) \le \phi_X(x-y) + \phi_X(y),$$

and hence

$$\phi_X(x) - \phi_X(y) \le \phi_X(x-y) \le C\|x-y\|.$$

For symmetry reasons

$$\phi_X(y) - \phi_X(x) \le C\|y-x\| = C\|x-y\|,$$

and hence $|\phi_X(x) - \phi_X(y)| \le C\|x-y\|$.

(iii) The sets $\{x \in \mathbf{R}^n \mid \phi_X(x) < 1\}$ and $\{x \in \mathbf{R}^n \mid \phi_X(x) \le 1\}$ are open and closed, respectively, since ϕ_X is continuous. Therefore, to prove assertion (iii) it suffices, due to the characterization of $\operatorname{int} X$ as the largest open set contained in X and of $\operatorname{cl} X$ as the smallest closed set containing X, to prove the inclusions

$$\operatorname{int} X \subseteq \{x \in \mathbf{R}^n \mid \phi_X(x) < 1\} \subseteq X \subseteq \{x \in \mathbf{R}^n \mid \phi_X(x) \le 1\} \subseteq \operatorname{cl} X.$$

Suppose $x \in \operatorname{int} X$. Since $tx \to x$ as $t \to 1$, the points tx belong to the interior of X for all numbers t that are sufficiently close to 1. Thus, there exists a number $t_0 > 1$ such that $t_0 x \in X$, i.e. such that $x \in t_0^{-1}X$, which means that $\phi_X(x) \le t_0^{-1} < 1$, and this proves the inclusion

$$\operatorname{int} X \subseteq \{x \in \mathbf{R}^n \mid \phi_X(x) < 1\}.$$

The implications $\phi_X(x) < t \Rightarrow x \in tX \Rightarrow \phi_X(x) \leq t$ are direct consequences of the definition of $\phi_X(x)$, and by choosing $t = 1$ we obtain the inclusions

$$\{x \in \mathbf{R}^n \mid \phi_X(x) < 1\} \subseteq X \subseteq \{x \in \mathbf{R}^n \mid \phi_X(x) \leq 1\}.$$

To prove the remaining inclusion it is now enough to prove the inclusion

$$\{x \in \mathbf{R}^n \mid \phi_X(x) = 1\} \subseteq \operatorname{cl} X.$$

So, suppose $\phi_X(x) = 1$. Then there is a sequence $(t_n)_1^\infty$ of numbers > 1 such that $t_n \to 1$ as $n \to \infty$ and $x \in t_n X$ for all n. The points $t_n^{-1}x$ belong to X for all n, and since $t_n^{-1}x \to x$ as $n \to \infty$, x is a point in the closure $\operatorname{cl} X$. □

Exercises

6.1 Find two quasiconvex functions f_1, f_2 with a non-quasiconvex sum $f_1 + f_2$.

6.2 Prove that the following functions $f\colon \mathbf{R}^3 \to \mathbf{R}$ are convex:

a) $f(x) = x_1^2 + 2x_2^2 + 5x_3^2 + 3x_2x_3$

b) $f(x) = 2x_1^2 + x_2^2 + x_3^2 - 2x_1x_2 + 2x_1x_3$

c) $f(x) = e^{x_1 - x_2} + e^{x_2 - x_1} + x_3^2 - 2x_3$.

6.3 For which values of the real number a is the function

$$f(x) = x_1^2 + 2x_2^2 + ax_3^2 - 2x_1x_2 + 2x_1x_3 - 6x_2x_3$$

convex and strictly convex?

6.4 Prove that the function $f(x) = x_1x_2\cdots x_n$ with $\mathbf{R}^n_+$ as domain is quasiconcave, and that the function $g(x) = (x_1x_2\cdots x_n)^{-1}$ with $\mathbf{R}^n_{++}$ as domain is convex.

6.5 Let $x_{[k]}$ denote the k:th biggest coordinate of the point $x = (x_1, x_2, \dots, x_n)$. In other words, $x_{[1]}, x_{[2]}, \dots, x_{[n]}$ are the coordinates of x in decreasing order. Prove for each k that the function $f(x) = \sum_{i=1}^k x_{[i]}$ is convex.

6.6 Suppose $f\colon \mathbf{R}_+ \to \mathbf{R}$ is convex. Prove that

$$f(x_1) + f(x_2) + \cdots + f(x_n) \leq f(x_1 + x_2 + \cdots + x_n) + (n-1)f(0)$$

for all x_1, x_2, …, $x_n \geq 0$. Note the special case $f(0) = 0$!

6.7 The function f is defined on a convex subset of $\mathbf{R}^n$. Suppose that the function $f(x) + \langle c, x\rangle$ is quasiconvex for each $c \in \mathbf{R}^n$. Prove that f is convex.

6.8 We have derived Corollary 6.2.7 from Theorem 6.2.6. Conversely, prove that Theorem 6.2.6 follows easily from Corollary 6.2.7.

6.9 X is a convex set in $\mathbf{R}^n$ with a nonempty interior, and $f\colon X \to \mathbf{R}$ is a continuous function, whose restriction to $\operatorname{int} X$ is convex. Prove that f is convex.

6.10 Suppose that the function $f\colon X \to \overline{\mathbf{R}}$ is convex. Prove that

$$\inf\{f(x) \mid x \in X\} = \inf\{f(x) \mid x \in \operatorname{rint}(\operatorname{dom} f)\}.$$

6.11 Use the method in Example 6.4.1 to determine the minimum of the function

$$g(x_1, x_2) = 16x_1 + 2x_2 + x_1^{-1}x_2^{-2}$$

over the set $x_1 > 0$, $x_2 > 0$.

6.12 Find the Minkowski functional of

a) the closed unit ball $\overline{B}(0;1)$ in $\mathbf{R}^n$ with respect to the ℓ^p-norm $\|\cdot\|_p$;

b) the halfspace $\{x \in \mathbf{R}^n \mid x_1 \le 1\}$.

6.13 Let X be a convex set with 0 as interior point and suppose that the set is symmetric with respect to 0, i.e. $x \in X \Rightarrow -x \in X$. Prove that the Minkowski functional ϕ_X is a norm, i.e. that

(i) $\phi_X(x+y) \le \phi_X(x) + \phi_X(y)$;

(ii) $\phi_X(\lambda x) = \lambda\phi_X(x)$ for all $\lambda \in \mathbf{R}$;

(iii) $\phi_X(x) = 0 \Leftrightarrow x = 0$.

Chapter 7

Smooth convex functions

This chapter is devoted to the study of smooth convex functions, i.e. convex functions that are differentiable. A prerequisite for differentiability at a point is that the function is defined and finite in a neighborhood of the point. Hence, it is only meaningful to study differentiability properties at interior points of the domain of the function, and by passing to the restriction of the function to the interior of its domain, we may as well assume from the beginning that the domain of definition is open. That is the reason for assuming all domains to be open and all function values to be finite in this chapter.

7.1 Convex functions on R

Let f be a real-valued function that is defined in a neighborhood of the point $x \in \mathbf{R}$. The one-sided limit

$$f'_+(x) = \lim_{t\to 0+} \frac{f(x+t) - f(x)}{t},$$

if it exists, is called the *right derivative* of f at the point x. The *left derivative* $f'_-(x)$ is similarly defined as the one-sided limit

$$f'_-(x) = \lim_{t\to 0-} \frac{f(x+t) - f(x)}{t}.$$

The function is obviously differentiable at the point x if and only if the right and the left derivatives both exist and are equal, and the derivative $f'(x)$ is in that case equal to their common value.

The left derivative of the function $f\colon I \to \mathbf{R}$ can be expressed as a right derivative of the function $\check{f}$, defined by

$$\check{f}(x) = f(-x) \quad \text{for all } x \in -I,$$

because

$$f'_-(x) = \lim_{t\to 0+} \frac{f(x-t)-f(x)}{-t} = -\lim_{t\to 0+} \frac{\check{f}(-x+t)-\check{f}(-x)}{t}$$

and hence

$$f'_-(x) = -\check{f}'_+(-x).$$

Observe that the function $\check{f}$ is convex if f is convex.

The basic differentiability properties of convex functions are consequences of the following lemma, which has an obvious interpretation in terms of slopes of various chords. Cf. figure 7.1.

Lemma 7.1.1. *Suppose f is a real-valued convex function that is defined on a subinterval of $\mathbf{R}$ containing the points $x_1 < x_2 < x_3$. Then*

$$\frac{f(x_2)-f(x_1)}{x_2-x_1} \le \frac{f(x_3)-f(x_1)}{x_3-x_1} \le \frac{f(x_3)-f(x_2)}{x_3-x_2}.$$

The above inequalities are strict if f is strictly convex.

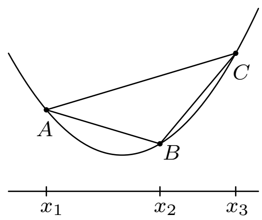

Figure 7.1. A geometric interpretation of Lemma 7.1.1: If k_{PQ} denotes the slope of the chord PQ, then $k_{AB} \le k_{AC} \le k_{BC}$.

Proof. Write $x_2 = \lambda x_3 + (1-\lambda)x_1$; then $\lambda = \dfrac{x_2-x_1}{x_3-x_1}$ is a number in the interval $]0,1[$. By convexity,

$$f(x_2) \le \lambda f(x_3) + (1-\lambda)f(x_1),$$

which simplifies to $f(x_2)-f(x_1) \le \lambda(f(x_3)-f(x_1))$, and this is equivalent to the leftmost of the two inequalities in the lemma.

The rightmost inequality is obtained by applying the already proven inequality to the convex function $\check{f}$. Since $-x_3 < -x_2 < -x_1$,

$$\frac{f(x_2)-f(x_3)}{x_3-x_2} = \frac{\check{f}(-x_2)-\check{f}(-x_3)}{-x_2-(-x_3)} \le \frac{\check{f}(-x_1)-\check{f}(-x_3)}{-x_1-(-x_3)} = \frac{f(x_1)-f(x_3)}{x_3-x_1},$$

and multiplication by -1 gives the desired result.

The above inequalities are strict if f is strictly convex. $\square$

The differentiability properties of convex one-variable functions are given by the following theorem.

Theorem 7.1.2. *Suppose $f : I \to \mathbf{R}$ is a convex function with an open subinterval I of $\mathbf{R}$ as its domain. Then:*

(a) *The function f has right and left derivatives at all points $x \in I$, and $f'_-(x) \leq f'_+(x)$.*

(b) *If $f'_-(x) \leq a \leq f'_+(x)$, then*

$$f(y) \geq f(x) + a(y - x) \qquad \text{for all } y \in I.$$

The above inequality is strict for $y \neq x$, if f is strictly convex.

(c) *If $x < y$, then $f'_+(x) \leq f'_-(y)$, and the inequality is strict if f is strictly convex.*

(d) *The functions $f'_+ : I \to \mathbf{R}$ and $f'_- : I \to \mathbf{R}$ are increasing, and they are strictly increasing if f is strictly convex.*

(e) *The set of points $x \in I$ where the function is not differentiable, is finite or countable.*

Proof. Fix $x \in I$ and let

$$F(t) = \frac{f(x+t) - f(x)}{t}.$$

The domain of F is an open interval J_x with the point 0 removed.

We start by observing that if $s, t, u \in J_x$ and $u < 0 < t < s$, then

$$F(u) \le F(t) \le F(s) \tag{7.1}$$

(and the inequalities are strict if f is strictly convex).

The right inequality $F(t) \le F(s)$ follows directly from the left inequality in Lemma 7.1.1 by choosing $x_1 = x$, $x_2 = x + t$ and $x_3 = x + s$, and the left inequality $F(u) \le F(t)$ follows from the inequality between the extreme ends in the same lemma by instead choosing $x_1 = x + u$, $x_2 = x$ and $x_3 = x + t$.

It follows from inequality (7.1) that the function $F(t)$ is increasing for $t > 0$ (strictly increasing if f is strictly convex) and bounded below by $F(u_0)$, where u_0 is an arbitrary negative number in the domain of F. Hence, the limit

$$f'_+(x) = \lim_{t \to 0+} F(t)$$

exists and

$$F(t) \ge f'_+(x)$$

for all $t > 0$ in the domain of F (with strict inequality if f is strictly convex). By replacing t with $y - x$, we obtain the following implication for $a \le f'_+(x)$:

$$y > x \Rightarrow f(y) - f(x) \ge f'_+(x)(y - x) \ge a(y - x) \tag{7.2}$$

(with strict inequality if f is strictly convex).

The same argument, applied to the function $\check{f}$ and the point $-x$, shows that $\check{f}'_+(-x)$ exists, and that

$$-y > -x \Rightarrow \check{f}(-y) - \check{f}(-x) \ge -a(-y - (-x))$$

if $-a \le \check{f}'_+(-x)$. Since $f'_-(x) = -\check{f}'_+(-x)$, this means that the left derivative $f'_-(x)$ exists and that the implication

$$y < x \Rightarrow f(y) - f(x) \ge a(y - x) \tag{7.3}$$

is true for all constants a satisfying $a \ge f'_-(x)$. The implications (7.2) and (7.3) are both satisfied if $f'_-(x) \le a \le f'_+(x)$, and this proves assertion (b).

Using inequality (7.1) we conclude that $F(-t) \le F(t)$ for all sufficiently small values of t. Hence

$$f'_-(x) = \lim_{t\to 0+} F(-t) \le \lim_{t\to 0+} F(t) = f'_+(x),$$

and this proves assertion (a).

As a special case of assertion (b), we have the two inequalities

$$f(y) - f(x) \ge f'_+(x)(y-x) \quad \text{and} \quad f(x) - f(y) \ge f'_-(y)(x-y),$$

and division by $y - x$ now results in the implication

$$y > x \Rightarrow f'_+(x) \le \frac{f(y)-f(x)}{y-x} \le f'_-(y).$$

(If f is strictly convex, we may replace $\le$ with $<$ at both places.) This proves assertion (c).

By combining (c) with the inequality in (a) we obtain the implication

$$x < y \Rightarrow f'_+(x) \le f'_-(y) \le f'_+(y),$$

which shows that the right derivative f'_+ is increasing. That the left derivative is increasing is proved in a similar way. (And the derivatives are strictly increasing if f is strictly convex.)

To prove the final assertion (e) we define I_x to be the open interval $]f'_-(x), f'_+(x)[$. This interval is empty if the derivative $f'(x)$ exists, and it is nonempty if the derivative does not exist, and intervals I_x and I_y belonging to different points x and y are disjoint because of assertion (c). Now choose, for each point x where the derivative does not exist, a rational number r_x in the interval I_x. Since different intervals are pairwise disjoint, the chosen numbers will be different, and since the set of rational numbers is countable, there are at most countably many points x at which the derivative does not exist. □

Definition. The line $y = f(x_0) + a(x - x_0)$ is called a *supporting line* of the function $f\colon I \to \mathbf{R}$ at the point $x_0 \in I$ if

$$f(x) \ge f(x_0) + a(x - x_0) \tag{7.4}$$

for all $x \in I$.

A supporting line at the point x_0 is a line which passes through the point $(x_0, f(x_0))$ and has the entire function curve $y = f(x)$ above (or on) itself. It is, in other words, a (one-dimensional) supporting hyperplane of the epigraph of f at the point $(x_0, f(x_0))$. The concept will be generalized for functions of several variables in the next chapter.

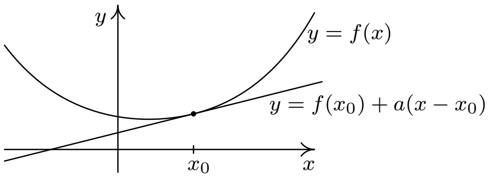

Figure 7.2. A supporting line.

Assertion (b) of the preceding theorem shows that convex functions with open domains have supporting lines at each point, and that the tangent is a supporting line at points where the derivative exists. By our next theorem, the existence of supporting lines is also a sufficient condition for convexity.

Theorem 7.1.3. *Suppose that the function* $f: I \to \mathbf{R}$, *where* I *is an open interval, has a supporting line at each point in* I. *Then,* f *is a convex function.*

Proof. Suppose that $x, y \in I$ and that $0 < \lambda < 1$, and let a be the constant belonging to the point $x_0 = \lambda x + (1-\lambda)y$ in the definition (7.4) of a supporting line. Then we have $f(x) \geq f(x_0) + a(x - x_0)$ and $f(y) \geq f(x_0) + a(y - x_0)$. By multiplying the first inequality by λ and the second inequality by $(1-\lambda)$, and then adding the two resulting inequalities, we obtain

$$\lambda f(x) + (1-\lambda)f(y) \geq f(x_0) + a\big(\lambda x + (1-\lambda)y - x_0\big) = f(x_0).$$

So the function f is convex.

Observe that if the inequality (7.4) is strict for all $x \neq x_0$ and for all $x_0 \in I$, then f is strictly convex. □

For differentiable functions we now obtain the following necessary and sufficient condition for convexity.

Theorem 7.1.4. *A differentiable function $f \colon I \to \mathbf{R}$ is convex if and only if its derivative f' is increasing. And it is strictly convex if and only if the derivative is strictly increasing.*

Proof. Assertion (d) in Theorem 7.1.2 shows that the derivative of a (strictly) convex function is (strictly) increasing.

To prove the converse, we assume that the derivative f' is increasing. By the mean value theorem, if x and x_0 are distinct points in I, there exists a point ξ between x and x_0 such that

$$\frac{f(x) - f(x_0)}{x - x_0} = f'(\xi) \begin{cases} \geq f'(x_0) & \text{if } x > x_0, \\ \leq f'(x_0) & \text{if } x < x_0. \end{cases}$$

Multiplication by $x - x_0$ results, in both cases, in the inequality

$$f(x) - f(x_0) \geq f'(x_0)(x - x_0),$$

which shows that $y = f(x_0) + f'(x_0)(x - x_0)$ is a supporting line of the function f at the point x_0. Therefore, f is convex by Theorem 7.1.3.

The above inequalites are strict if the derivative is strictly increasing, and we conclude that f is strictly convex in that case. □

For two times differentiable functions we obtain the following corollary.

Corollary 7.1.5. *A two times differentiable function $f \colon I \to \mathbf{R}$ is convex if and only if $f''(x) \geq 0$ for all $x \in I$. The function is strictly convex if $f''(x) > 0$ for all $x \in I$.*

Proof. The derivative f' is increasing (strictly increasing) if the second derivative f'' is nonnegative (positive). And the second derivative is nonnegative if the derivative is increasing. □

Remark. A continuous function $f\colon J \to \mathbf{R}$ with a non-open interval J as domain is convex if (and only if) the restriction of f to the interior of J is convex. Hence, if the derivative exists and is increasing in the interior of J, or if the second derivative exists and $f''(x) \geq 0$ for all interior points x of the interval, then f is convex on J. Cf. exercise 7.7.

EXAMPLE 7.1.1. The functions $x \mapsto e^x$, $x \mapsto -\ln x$ and $x \mapsto x^p$, where $p > 1$, are strictly convex on their domains $\mathbf{R}$, $]0, \infty[$ and $[0, \infty[$, respectively, because their first derivatives are strictly increasing functions. □

7.2 Differentiable convex functions

A differentiable one-variable function f is convex if and only if its derivative is an increasing function. In order to generalize this result to functions of several variables it is necessary to express the condition that the derivative is increasing in a generalizable way. To this end, we note that the derivative f' is increasing on an interval if and only if $f'(x+h)h \geq f'(x)h$ for all numbers x and $x+h$ in the interval, and this inequality is also meaningful for functions f of several variables if we interpret $f'(x)h$ as the value of the linear form $Df(x)$ at h. The inequality generalizing that the derivative of a function of several variables is increasing will thus be written $Df(x+h)[h] \geq Df(x)[h]$, or using gradient notation, $\langle f'(x+h), h\rangle \geq \langle f'(x), h\rangle$.

Theorem 7.2.1. *Let X be an open convex subset of $\mathbf{R}^n$, and let $f\colon X \to \mathbf{R}$ be a differentiable function. The following three conditions are equivalent:*

(i) f is a convex function.
(ii) $f(x+v) \geq f(x) + Df(x)[v]$ for all $x, x+v \in X$.
(iii) $Df(x+v)[v] \geq Df(x)[v]$ for all $x, x+v \in X$.

The function f is strictly convex if and only if the inequalities in (ii) and (iii) can be replaced by strict inequalities when $v \neq 0$.

Proof. Let us for given points x and $x+v$ in X consider the restriction $\phi_{x,v}$ of f to the line through x with direction v, i.e. the one-variable function

$$\phi_{x,v}(t) = f(x+tv)$$

with the open interval $I_{x,v} = \{t \in \mathbf{R} \mid x+tv \in X\}$ as domain. The functions $\phi_{x,v}$ are differentiable with derivative $\phi'_{x,v}(t) = Df(x+tv)[v]$, and f is convex if and only if all restrictions $\phi_{x,v}$ are convex.

(i) $\Rightarrow$ (ii) So if f is convex, then $\phi_{x,v}$ is a convex function, and it follows from Theorem 7.1.2 (b) that $\phi_{x,v}(t) \geq \phi_{x,v}(0) + \phi'_{x,v}(0)t$ for all $t \in I_{x,v}$, which means that $f(x+tv) \geq f(x) + Df(x)[v]\,t$ for all t such that $x + tv \in X$. We now obtain the inequality in (ii) by choosing $t = 1$.

(ii) $\Rightarrow$ (iii) We obtain inequality (iii) by adding the two inequalities

$$f(x+v) \geq f(x) + Df(x)[v] \quad \text{and} \quad f(x) \geq f(x+v) + Df(x+v)[-v].$$

(iii) $\Rightarrow$ (i) Suppose (iii) holds, and let $y = x + sv$ and $w = (t-s)v$. If $t > s$, then

$$\begin{aligned}
\phi'_{x,v}(t) - \phi'_{x,v}(s) &= Df(x+tv)[v] - Df(x+sv)[v] \\
&= Df(y+w)[v] - Df(y)[v] \\
&= \frac{1}{t-s}\big(Df(y+w)[w] - Df(y)[w]\big) \geq 0,
\end{aligned}$$

which means that the derivative $\phi'_{x,v}$ is increasing. The functions $\phi_{x,v}$ are thus convex.

This proves the equivalence of assertions (i), (ii) and (iii), and by replacing all inequalities in the proof by strict inequalities, we obtain the corresponding equivalent assertions for strictly convex functions. $\square$

The derivative of a differentiable function is equal to zero at a local minimum point. For convex functions, the converse is also true.

Theorem 7.2.2. *Suppose $f\colon X \to \mathbf{R}$ is a differentiable convex function. Then $\hat{x} \in X$ is a global minimum point if and only if $f'(\hat{x}) = 0$.*

Proof. That the derivative equals zero at a minimum point is a general fact, and the converse is a consequence of property (ii) in the previous theorem, for if $f'(\hat{x}) = 0$, then $f(x) \geq f(\hat{x}) + Df(\hat{x})[x - \hat{x}] = f(\hat{x})$ for all $x \in X$. □

Convexity can also be expressed by a condition on the second derivative, and the natural substitute for the one-variable condition $f''(x) \geq 0$ is that the second derivative should be positive semidefinite.

Theorem 7.2.3. *Let X be an open convex subset of $\mathbf{R}^n$, and suppose that the function $f\colon X \to \mathbf{R}$ is two times differentiable. Then f is convex if and only if the second derivative $f''(x)$ is positive semidefinite for all $x \in X$.*

If $f''(x)$ is positive definite for all $x \in X$, then f is strictly convex.

Proof. The one-variable functions $\phi_{x,v}(t) = f(x + tv)$ are now two times differentiable with second derivative

$$\phi''_{x,v}(t) = D^2 f(x + tv)[v, v] = \langle v, f''(x + tv)v\rangle.$$

Since f is convex if and only if all functions $\phi_{x,v}$ are convex, f is convex if and only if all second derivatives $\phi''_{x,v}$ are nonnegative functions .

If the second derivative $f''(x)$ is positive semidefinite for all $x \in X$, then $\phi''_{x,v}(t) = \langle v, f''(x+tv)v\rangle \geq 0$ for all $x \in X$ and all $v \in \mathbf{R}^n$, which means that the second derivatives $\phi''_{x,v}$ are nonnegative funtions. Conversely, if the second derivatives $\phi''_{x,v}$ are nonnegative, then in particular $\langle v, f''(x)\rangle = \phi''_{x,v}(0) \geq 0$ for all $x \in X$ and all $v \in \mathbf{R}^n$, and we conclude that the second derivative $f''(x)$ is positive semidefinite for all $x \in X$.

If the second derivatives $f''(x)$ are all positive definite, then $\phi''_{x,v}(t) > 0$ for $v \neq 0$, which implies that the functions $\phi_{x,v}$ are strictly convex, and then f is strictly convex, too. □

7.3 Strong convexity

The function surface of a convex functions bends upwards, but there is no lower positive bound on the curvature. By introducing such a bound we obtain the notion of strong convexity.

Definition. Let μ be a positive number. A function $f\colon X \to \overline{\mathbf{R}}$ is called *μ-strongly convex* if the function $f(x) - \frac{1}{2}\mu\|x\|^2$ is convex, and the function f is called *strongly convex* if it is μ-strongly convex for some positive number μ.

Theorem 7.3.1. *A differentiable function $f\colon X \to \mathbf{R}$ with a convex domain is μ-strongly convex if and only if the following two mutually equivalent inequalities are satisfied for all $x, x+v \in X$:*

(i) $$Df(x+v)[v] \geq Df(x)[v] + \mu\|v\|^2$$

(ii) $$f(x+v) \geq f(x) + Df(x)[v] + \tfrac{1}{2}\mu\|v\|^2.$$

Proof. Let $g(x) = f(x) - \frac{1}{2}\mu\|x\|^2$ and note that $g'(x) = f'(x) - \mu x$ and that consequently $Df(x)[v] = Dg(x)[v] + \mu\langle x, v\rangle$.

If f is μ-strongly convex, then g is a convex function, and so it follows from Theorem 7.2.1 that

$$\begin{aligned} Df(x+v)[v] - Df(x)[v] &= Dg(x+v)[v] - Dg(x)[v] + \mu\langle x+v, v\rangle - \mu\langle x, v\rangle \\ &\geq \mu\langle v, v\rangle = \mu\|v\|^2, \end{aligned}$$

i.e. inequality (i) is satisfied.

(i) $\Rightarrow$ (ii): Assume (i) holds, and define the function Φ for $0 \leq t \leq 1$ by

$$\Phi(t) = f(x+tv) - f(x) - Df(x)[v]\,t.$$

Then $\Phi'(t) = Df(x+tv)[v] - Df(x)[v] = \dfrac{1}{t}\big(Df(x+tv)[tv] - Df(x)[tv]\big)$, and it now follows from inequality (i) that

$$\Phi'(t) \geq t^{-1}\mu\|tv\|^2 = \mu\|v\|^2\,t.$$

By integrating the last inequality over the interval $[0,1]$ we obtain

$$\Phi(1) = \Phi(1) - \Phi(0) \geq \tfrac{1}{2}\mu\|v\|^2,$$

which is the same as inequality (ii).

If inequality (ii) holds, then

$$\begin{aligned} g(x+v) &= f(x+v) - \tfrac{1}{2}\mu\|x+v\|^2 \geq f(x) + Df(x)[v] + \tfrac{1}{2}\mu\|v\|^2 - \tfrac{1}{2}\mu\|x+v\|^2 \\ &= g(x) + \tfrac{1}{2}\mu\|x\|^2 + Dg(x)[v] + \mu\langle x, v\rangle + \tfrac{1}{2}\mu\|v\|^2 - \tfrac{1}{2}\mu\|x+v\|^2 \\ &= g(x) + Dg(x)[v]. \end{aligned}$$

The function g is thus convex, by Theorem 7.2.1, and $f(x) = g(x) + \frac{1}{2}\mu\|x\|^2$ is consequently μ-strongly convex. $\square$

Theorem 7.3.2. *A twice differentiable function $f\colon X \to \mathbf{R}$ with a convex domain is μ-strongly convex if and only if*

$$\langle v, f''(x)v\rangle = D^2 f(x)[v,v] \geq \mu\|v\|^2 \tag{7.5}$$

for all $x \in X$ and all $v \in \mathbf{R}^n$.

Remark. If A is a symmetric operator, then

$$\min_{v\neq 0} \frac{\langle v, Av\rangle}{\|v\|^2} = \lambda_{\min},$$

where $\lambda_{\min}$ is the smallest eigenvalue of the operator. Thus, a two times differentiable function f with a convex domain is μ-strongly convex if and only if the eigenvalues of the hessian $f''(x)$ are greater than or equal to μ for each x in the domain.

Proof. Let $\phi_{x,v}(t) = f(x + tv)$. If condition (7.5) holds, then

$$\phi''_{x,v}(t) = D^2 f(x + tv)[v,v] \geq \mu\|v\|^2$$

for all t in the domain of the function. Using Taylor's formula with remainder term, we therefore conclude that

$$\phi_{x,v}(t) = \phi_{x,v}(0) + \phi'_{x,v}(0)t + \tfrac{1}{2}\phi''_{x,v}(\xi)t^2 \geq \phi_{x,v}(0) + \phi'_{x,v}(0)t + \tfrac{1}{2}\mu\|v\|^2\, t^2.$$

For $t = 1$ this amounts to inequality (ii) in Theorem 7.3.1, and hence f is a μ-strongly convex function.

Conversely, if f is μ-strongly convex, then by Theorem 7.3.1 (i)

$$\frac{\phi'_{x,v}(t) - \phi'_{x,v}(0)}{t} = \frac{Df(x+tv)[tv] - Df(x)[tv]}{t^2} \geq \mu\|v\|^2.$$

Taking the limit as $t \to 0$ we obtain

$$D^2 f(x)[v,v] = \phi''_{x,v}(0) \geq \mu\|v\|^2. \qquad \square$$

7.4 Convex functions with Lipschitz continuous derivatives

The rate of convergence of classical iterative algorithms for minimizing functions depends on the variation of the deriviative − the more the derivative varies in a neighborhood of the minimum point, the slower the convergence. The size of the Lipschitz constant is a measure of the variation of the derivative for functions with a Lipschitz continuous derivative. Therefore, we start with a result which for arbitrary functions connects Lipschitz continuity of the first derivative to bounds on the second derivative.

Theorem 7.4.1. *Suppose f is a twice differentiable function and that X is a convex subset of its domain.*

(i) If $\|f''(x)\| \leq L$ for all $x \in X$, then the derivative f' is Lipschitz continuous on X with Lipschitz constant L.

(ii) If the derivative f' is Lipschitz continuous on the set X with constant L, then $\|f''(x)\| \leq L$ for all $x \in \operatorname{int} X$.

Proof. (i) Suppose that $\|f''(x)\| \leq L$ for all $x \in X$, and let x and y be two points in X. Put $v = y - x$, let w be an arbitrary vector with $\|w\| = 1$, and define the function ϕ for $0 \leq t \leq 1$ by

$$\phi(t) = Df(x+tv)[w] = \langle f'(x+tv), w\rangle.$$

Then ϕ is differentiable with derivative

$$\phi'(t) = D^2 f(x+tv)[w,v] = \langle w, f''(x+tv)v\rangle$$

so it follows from the Cauchy-Schwarz inequality that

$$|\phi'(t)| \leq \|w\|\|f''(x+tv)v\| \leq \|f''(x+tv)\|\|v\| \leq L\|v\|,$$

since $x+tv$ is a point in X. By the mean value theorem, $\phi(1)-\phi(0) = \phi'(s)$ for some point $s \in]0,1[$. Consequently,

$$|\langle f'(y) - f'(x), w\rangle| = |\phi(1) - \phi(0)| = |\phi'(s)| \leq L\|y-x\|.$$

Since w is an arbitrary vector of norm 1, we conclude that

$$\|f'(y) - f'(x)\| = \sup_{\|w\|=1} \langle f'(y) - f'(x), w\rangle \leq L\|y-x\|,$$

i.e. the derivative f' is Lipschitz continuous on X with constant L.

(ii) Assume conversely that the first derivative f' is Lipschitz continuous on the set X with constant L. Let x be a point in the interior of X, and let v and w be arbitrary vectors with norm 1. The function

$$\phi(t) = Df(x+tv)[w] = \langle f'(x+tv, w\rangle$$

is then defined and differentiable and the point $x+tv$ lies in X for all t in a neighborhood of 0, and it follows that

$$\begin{aligned}|\phi(t) - \phi(0)| &= |\langle f'(x+tv) - f'(x), w\rangle| \leq \|f'(x+tv) - f'(x)\|\|w\| \\ &\leq L\|tv\| = L|t|.\end{aligned}$$

Division by t and passing to the limit as $t \to 0$ results in the inequality

$$|\langle w, f''(x)v\rangle| = |\phi'(0)| \leq L$$

with the conclustion that

$$\|f''(x)\| = \sup_{\|v\|=1} \|f''(x)v\| = \sup_{\|v\|,\|w\|=1} \langle w, f''(x)v\rangle \leq L. \qquad \square$$

Definition. A differentiable function $f\colon X \to \mathbf{R}$ belongs to the class $\mathcal{S}_{\mu,L}(X)$ if f is μ-strongly convex and the derivative f' is Lipschitz continuous with constant L. The quotient $Q = L/\mu$ is called the *condition number* of the class.

Due to Theorem 7.3.1, a differentiable function f with a convex domain X belongs to the class $\mathcal{S}_{\mu,L}(X)$ if and only if it satisfies the following two inequalities for all $x, x+v \in X$:

$$\langle f'(x+v) - f'(x), v\rangle \geq \mu\|v\|^2 \quad \text{and} \quad \|f'(x+v) - f'(x)\| \leq L\|v\|.$$

If we combine the first of these two inequalities with the Cauchy–Schwarz inequality, we obtain the inequality $\mu\|v\| \leq \|f'(x+v) - f'(x)\|$, so we conclude that $\mu \leq L$ and $Q \geq 1$.

EXAMPLE 7.4.1. Strictly convex quadratic functions

$$f(x) = \tfrac{1}{2}\langle x, Px\rangle + \langle q, x\rangle + r$$

belong to the class $\mathcal{S}_{\lambda_{\min},\lambda_{\max}}(\mathbf{R}^n)$, where $\lambda_{\min}$ and $\lambda_{\max}$ denote the smallest and the largest eigenvalue, respectively, of the positive definite matrix P.

For $f'(x) = Px + q$ and $f''(x) = P$, whence

$$D^2 f(x)[v, v] = \langle v, Pv\rangle \geq \lambda_{\min}\|v\|^2 \qquad \text{and}$$
$$\|f'(x + v) - f'(x)\| = \|Pv\| \leq \|P\|\,\|v\| = \lambda_{\max}\|v\|.$$

The condition number of the quadratic function f is thus equal to the quotient $\lambda_{\max}/\lambda_{\min}$ between the largest and the smallest eigenvalue. □

The sublevel sets $\{x \mid f(x) \leq \alpha\}$ of a strictly convex quadratic function f are ellipsoids for all values of α greater than the minimum value of the function, and the ratio of the longest and the shortest axes of any of these ellipsoids is equal to $\sqrt{\lambda_{\max}/\lambda_{\min}}$, i.e. to the square root of the condition number Q. This ratio is obviously also equal to the ratio of the radii of the smallest ball containing and the largest ball contained in the ellipsoid. As we shall see, something similar applies to all functions in the class $\mathcal{S}_{\mu,L}(\mathbf{R}^n)$.

Theorem 7.4.2. *Let f be a function in the class $\mathcal{S}_{\mu,L}(\mathbf{R}^n)$ with minimum point $\hat{x}$, and let α be a number greater than the minimum value $f(\hat{x})$. Then*

$$B(\hat{x};r) \subseteq \{x \in X \mid f(x) \leq \alpha\} \subseteq B(\hat{x};R),$$

where $r = \sqrt{2L^{-1}(\alpha - f(\hat{x}))}$ and $R = \sqrt{2\mu^{-1}(\alpha - f(\hat{x}))}$.

Remark. Note that $R/r = \sqrt{L/\mu} = \sqrt{Q}$.

Proof. Since $f'(\hat{x}) = 0$ we obtain the following inequalities from Theorems 1.1.2 and 7.3.1 (by replacing a and x respectively with $\hat{x}$ and v with $x - \hat{v}$):

$$f(\hat{x}) + \tfrac{1}{2}\mu\|x - \hat{x}\|^2 \leq f(x) \leq f(\hat{x}) + \tfrac{1}{2}L\|x - \hat{x}\|^2.$$

Hence, $x \in S = \{x \in X \mid f(x) \leq \alpha\}$ implies

$$\tfrac{1}{2}\mu\|x - \hat{x}\|^2 \leq f(x) - f(\hat{x}) \leq \alpha - f(\hat{x}) = \tfrac{1}{2}\mu R^2,$$

which means that $\|x - \hat{x}\| \leq R$ and proves the inclusion $S \subseteq B(\hat{x};R)$.

And if $x \in B(\hat{x};r)$, then $f(x) \leq f(\hat{x}) + \frac{1}{2}Lr^2 = \alpha$, which means that $x \in S$ and proves the inclusion $B(\hat{x};r) \subseteq S$. □

Convex functions on $\mathbf{R}^n$ with Lipschitz continuous derivatives are characterized by the following theorem.

Theorem 7.4.3. *A differentiable function $f\colon \mathbf{R}^n \to \mathbf{R}$ is convex and its derivative is Lipschitz continuous with Lipschitz constant L if and only if the following mutually equivalent inequalities are fulfilled for all $x, v \in \mathbf{R}^n$:*

(i) $$f(x) + Df(x)[v] \leq f(x+v) \leq f(x) + Df(x)[v] + \frac{L}{2}\|v\|^2$$

(ii) $$f(x+v) \geq f(x) + Df(x)[v] + \frac{1}{2L}\|f'(x+v) - f'(x)\|^2$$

(iii) $$Df(x+v)[v] \geq Df(x)[v] + \frac{1}{L}\|f'(x+v) - f'(x)\|^2$$

Proof. That inequality (i) has to be satisfied for convex functions with a Lipschitz continuous derivative is a consequence of Theorems 1.1.2 and 7.2.1.

(i) ⇒ (ii): Let $w = f'(x+v) - f'(x)$, and apply the right inequality in (i) with x replaced by $x + v$ and v replaced by $-L^{-1}w$; this results in the inequality

$$f(x + v - L^{-1}w) \leq f(x+v) - L^{-1}Df(x+v)[w] + \tfrac{1}{2}L^{-1}\|w\|^2.$$

The left inequality in (i) with $v - L^{-1}w$ instead of v yields

$$f(x + v - L^{-1}w) \geq f(x) + Df(x)[v - L^{-1}w].$$

By combining these two new inequalities, we obtain

$$f(x+v) \geq f(x) + Df(x)[v - L^{-1}w] + L^{-1}Df(x+v)[w] - \tfrac{1}{2}L^{-1}\|w\|^2$$

$$\begin{aligned}
&= f(x) + Df(x)[v] + L^{-1}\big(Df(x+v)[w] - Df(x)[w]\big) - \tfrac{1}{2}L^{-1}\|w\|^2 \\
&= f(x) + Df(x)[v] + L^{-1}\langle f'(x+v) - f'(x), w\rangle - \tfrac{1}{2}L^{-1}\|w\|^2 \\
&= f(x) + Df(x)[v] + L^{-1}\langle w, w\rangle - \tfrac{1}{2}L^{-1}\|w\|^2 \\
&= f(x) + Df(x)[v] + \tfrac{1}{2}L^{-1}\|w\|^2,
\end{aligned}$$

and that is inequality (ii).

(ii) $\Rightarrow$ (iii): Add inequality (ii) to the inequality obtained by changing x to $x+v$ and v to $-v$. The result is inequality (iii).

Let us finally assume that inequality (iii) holds. The convexity of f is then a consequence of Theorem 7.2.1, and by combining (iii) with the Cauchy–Schwarz inequality, we obtain the inequality

$$\begin{aligned}
\frac{1}{L}\|f'(x+v) - f'(x)\|^2 \le Df(x+v)[v] - Df(x)[v] &= \langle f'(x+v) - f'(x), v\rangle \\
&\le \|f'(x+v) - f'(x)\| \cdot \|v\|,
\end{aligned}$$

which after division by $\|f'(x+v) - f'(x)\|$ gives us the desired conclusion: the derivative is Lipschitz continuous with Lipschitz constant L. □

Theorem 7.4.4. *If $f \in \mathcal{S}_{\mu,L}(\mathbf{R}^n)$, then*

$$Df(x+v)[v] \ge Df(x)[v] + \frac{\mu L}{\mu + L}\|v\|^2 + \frac{1}{\mu + L}\|f'(x+v) - f'(x)\|^2$$

for all $x, v \in \mathbf{R}^n$.

Proof. Let $g(x) = f(x) - \frac{1}{2}\mu\|x\|^2$; the function g is then convex, and since $Dg(x)[v] = Df(x)[v] - \mu\langle x, v\rangle$, it follows from Theorem 1.1.2 that

$$\begin{aligned}
g(x+v) &= f(x+v) - \tfrac{1}{2}\mu\|x+v\|^2 \\
&\le f(x) + Df(x)[v] + \tfrac{1}{2}L\|v\|^2 - \tfrac{1}{2}\mu\|x+v\|^2 \\
&= g(x) + \tfrac{1}{2}\mu\|x\|^2 + Dg(x)[v] + \mu\langle x, v\rangle + \tfrac{1}{2}L\|v\|^2 - \tfrac{1}{2}\mu\|x+v\|^2 \\
&= g(x) + Dg(x)[v] + \tfrac{1}{2}(L-\mu)\|v\|^2.
\end{aligned}$$

This shows that g satisfies condition (i) in Theorem 7.4.3 with L replaced by $L - \mu$. The derivative g' is consequently Lipschitz continuous with constant $L - \mu$. The same theorem now gives us the inequality

$$Dg(x+v)[v] \ge Dg(x)[v] + \frac{1}{L-\mu}\|g'(x+v) - g'(x)\|^2,$$

which is just a reformulation of the inequality in Theorem 7.4.4. □

Exercises

7.1 Show that the following functions are convex.

a) $f(x_1, x_2) = e^{x_1} + e^{x_2} + x_1x_2, \quad x_1 + x_2 > 0$
b) $f(x_1, x_2) = \sin(x_1 + x_2), \quad -\pi < x_1 + x_2 < 0$
c) $f(x_1, x_2) = -\sqrt{\cos(x_1 + x_2)}, \quad -\frac{\pi}{2} < x_1 + x_2 < \frac{\pi}{2}$.

7.2 Is the function $f(x_1, x_2) = x_1^2/x_2 + x_2^2/x_1$ convex in the first quadrant $x_1 > 0$, $x_2 > 0$?

7.3 Show that the function $f(x) = \sum_{j=1}^{n-1} x_j^2/x_n$ is convex in the halfspace $x_n > 0$.

7.4 Show that the following function is concave on the set $[0, 1[\times[0, 1[\times[0, 1[$:

$$\begin{aligned} f(x_1, x_2, x_3) = \ln(1 - x_1) + \ln(1 - x_2) + \ln(1 - x_3) \\ - (x_1^2 + x_2^2 + x_3^2 + x_1x_2 + x_1x_3 + x_2x_3). \end{aligned}$$

7.5 Let I be an interval and suppose that the function $f : I \to \mathbf{R}$ is convex. Show that f is either increasing on the interval, or decreasing on the interval, or there exists a point $c \in I$ such that f is decreasing to the left of c and increasing to the right of c.

7.6 Suppose $f :]a, b[\to \mathbf{R}$ is a convex function.

a) Prove that the two one-sided limits $\lim_{x\to a+} f(x)$ and $\lim_{x\to b-} f(x)$ exist (as finite numbers or $\pm\infty$).

b) Suppose that the interval is finite, and extend the function to the closed interval $[a, b]$ by defining $f(a) = \alpha$ and $f(b) = \beta$. Prove that the extended function is convex if and only if $\alpha \geq \lim_{x\to a+} f(x)$ and $\beta \geq \lim_{x\to b-} f(x)$.

7.7 Prove that a continuous function $f : [a, b] \to \mathbf{R}$ is convex if and only if its restriction to the open interval $]a, b[$ is convex.

7.8 $\mathcal{F}$ is a family of differentiable functions on $\mathbf{R}^n$ with the following two properties:

(i) $f \in \mathcal{F} \;\Rightarrow\; f + g \in \mathcal{F}$ for all affine functions $g : \mathbf{R}^n \to \mathbf{R}$.
(ii) If $f \in \mathcal{F}$ and $f'(x_0) = 0$, then x_0 is a minimum point of f.

Prove that all functions in $\mathcal{F}$ are convex.

7.9 Suppose that $f : X \to \mathbf{R}$ is a twice differentiable convex function. Prove that its recessive subspace V_f is a subset of $\mathcal{N}(f''(x))$ for each $x \in X$.

7.10 Let $f : X \to \mathbf{R}$ be a differentiable function with a convex domain X. Prove that f is quasiconvex if and only if

$$f(x + v) \leq f(x) \;\Rightarrow\; Df(x)[v] \leq 0$$

for all $x, x + v \in X$.
[Hint: It suffices to prove the assertion for functions on $\mathbf{R}$; the general result then follows by taking restrictions to lines.]

7.11 Let $f: X \to \mathbf{R}$ be a twice differentiable function with a convex domain X. Prove the following assertions:

a) If f is quasiconvex, then

$$Df(x)[v] = 0 \;\Rightarrow\; D^2 f(x)[v,v] \geq 0$$

for all $x \in X$ and all $v \in \mathbf{R}^n$.

b) If

$$Df(x)[v] = 0 \;\Rightarrow\; D^2 f(x)[v,v] > 0$$

for all $x \in X$ and all $v \neq 0$, then f is quasiconvex.
[Hint: It is enough to prove the results for functions defined on $\mathbf{R}$.]

7.12 Prove that the function $\alpha_1 f_1 + \alpha_2 f_2$ is $(\alpha_1\mu_1 + \alpha_2\mu_2)$-strongly convex if f_1 is μ_1-strongly convex, f_2 is μ_2-strongly convex and $\alpha_1, \alpha_2 > 0$.

7.13 Prove that if a differentiable μ-strongly convex function $f: X \to \mathbf{R}$ has a minimum at the point $\hat{x}$, then $\|x - \hat{x}\| \leq \mu^{-1}\|f'(x)\|$ for all $x \in X$.

Chapter 8

The subdifferential

We will now generalize a number of results from the previous chapter to convex functions that are not necessarily differentiable everywhere. However, real-valued convex functions with open domains can not be too irregular − they are, as already noted, continuous, and they have direction derivatives.

8.1 The subdifferential

If f is a differentiable function, then $y = f(a) + \langle f'(a), x - a\rangle$ is the equation of a hyperplane that is tangent to the surface $y = f(x)$ at the point $(a, f(a))$. And if f is also convex, then $f(x) \geq f(a) + \langle f'(a), x - a\rangle$ for all x in the domain of the function (Theorem 7.2.1), so the tangent plane lies below the graph of the function and is a supporting hyperplane of the epigraph.

The epigraph of an arbitrary convex function is a convex set, by definition. Hence, through each boundary point belonging to the epigraph of a convex function there passes a supporting hyperplane. The supporting hyperplanes of a convex one-variable function f, defined on an open interval, are given by Theorem 7.1.2, which says that the line $y = f(x_0) + a(x - x_0)$ supports the epigraph at the point $(x_0, f(x_0))$ if (and only if) $f'_-(x_0) \leq a \leq f'_+(x_0)$.

The existence of supporting hyperplanes characterizes convexity, and this is a reason for a more detailed study of this concept.

Definition. Let $f\colon X \to \overline{\mathbf{R}}$ be a function defined on a subset X of $\mathbf{R}^n$. A vector $c \in \mathbf{R}^n$ is called a *subgradient of f at the point $a \in X$* if the inequality

$$f(x) \geq f(a) + \langle c, x - a\rangle \qquad (8.1)$$

holds for all $x \in X$.

The set of all subgradients of f at a is called the *subdifferential* of f at a and is denoted by $\partial f(a)$.

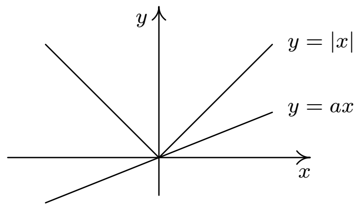

Figure 8.1. The line $y = ax$ is a supporting line of the function $f(x) = |x|$ at the origin if $-1 \leq a \leq 1$.

Remark. The inequality (8.1) is of course satisfied by all points $a \in X$ and all vectors $c \in \mathbf{R}^n$ if x is a point in the set $X \setminus \operatorname{dom} f$. Hence, to verify that c is a subgradient of f at a it suffices to verify that the inequality holds for all $x \in \operatorname{dom} f$.

The inequality (8.1) does not hold for any vector c if a is a point in $X \setminus \operatorname{dom} f$ and x is a point in $\operatorname{dom} f$. Hence, $\partial f(a) = \emptyset$ for all $a \in X \setminus \operatorname{dom} f$, except in the trivial case when $\operatorname{dom} f = \emptyset$, i.e. when f is equal to ∞ on the entire set X. In this case we have $\partial f(a) = \mathbf{R}^n$ for all $a \in X$ since the inequality (8.1) is now trivially satisfied by all $a, x \in X$ and all $c \in \mathbf{R}^n$.

Example 8.1.1. The subdifferentials of the one-variable function $f(x) = |x|$ are

$$\partial f(a) = \begin{cases} \{-1\} & \text{if } a < 0, \\ [-1, 1] & \text{if } a = 0, \\ \{1\} & \text{if } a > 0. \end{cases}$$ □

Theorem 8.1.1. *The subdifferentials of an arbitrary function $f \colon X \to \overline{\mathbf{R}}$ are closed and convex sets.*

Proof. For points $a \in \operatorname{dom} f$,

$$\partial f(a) = \textstyle\bigcap_{x \in \operatorname{dom} f} \{c \in \mathbf{R}^n \mid \langle c, x - a \rangle \leq f(x) - f(a)\}$$

is convex and closed, since it is an intersection of closed halfspaces, and the case $a \in X \setminus \operatorname{dom} f$ is trivial. □

Theorem 8.1.2. *A point $a \in X$ is a global minimum point of the function $f \colon X \to \overline{\mathbf{R}}$ if and only if $0 \in \partial f(a)$.*

Proof. The assertion follows immediately from the subgradient definition. □

Our next theorem tells us that the derivative $f'(a)$ is the only subgradient candidate for functions f that are differentiable at a. Geometrically this means that the tangent plane at a is the only possible supporting hyperplane.

Theorem 8.1.3. *Suppose that the function $f\colon X \to \overline{\mathbf{R}}$ is differentiable at the point $a \in \operatorname{dom} f$. Then either $\partial f(a) = \{f'(a)\}$ or $\partial f(a) = \emptyset$.*

Proof. Suppose $c \in \partial f(a)$. By the differentiability definition,

$$f(a+v) - f(a) = \langle f'(a), v\rangle + r(v)$$

with a remainder term $r(v)$ satisfying the condition

$$\lim_{v\to 0} \frac{r(v)}{\|v\|} = 0,$$

and by the subgradient definition, $f(a+v) - f(a) \geq \langle c, v\rangle$ for all v such that $a + v$ belongs to X. Consequently,

$$\frac{\langle c, v\rangle}{\|v\|} \leq \frac{\langle f'(a), v\rangle + r(v)}{\|v\|} \tag{8.2}$$

for all v with a sufficiently small norm $\|v\|$.

Let $\mathbf{e}_j$ be the j:th unit vector. Then $\langle c, \mathbf{e}_j \rangle = c_j$ and $\langle f'(a), \mathbf{e}_j \rangle = \dfrac{\partial f}{\partial x_j}(a)$, so by choosing $v = t\mathbf{e}_j$ in inequality (8.2), noting that $\|t\mathbf{e}_j\| = |t|$, and letting $t \to 0$ from the right and from the left, respectively, we obtain the following two inequalities

$$c_j \le \frac{\partial f}{\partial x_j}(a) \quad \text{and} \quad -c_j \le -\frac{\partial f}{\partial x_j}(a),$$

which imply that $c_j = \dfrac{\partial f}{\partial x_j}(a)$. Hence, $c = f'(a)$, and this proves the inclusion $\partial f(a) \subseteq \{f'(a)\}$. □

We can now reformulate Theorem 7.2.1 as follows: A differentiable function with a convex domain is convex if and only if it has a subgradient (which is then equal to the derivative) everywhere. Our next theorem generalizes this result.

Theorem 8.1.4. *Let $f \colon X \to \overline{\mathbf{R}}$ be a function with a convex domain X.*

(a) If $\operatorname{dom} f$ is a convex set and $\partial f(x) \neq \emptyset$ for all $x \in \operatorname{dom} f$, then f is a convex function.

(b) If f is a convex function, then $\partial f(x) \neq \emptyset$ for all $x \in \operatorname{rint}(\operatorname{dom} f)$.

Proof. (a) Let x and y be two arbitrary points in $\operatorname{dom} f$ and consider the point $z = \lambda x + (1-\lambda)y$, where $0 < \lambda < 1$. By assumption, f has a subgradient c at the point z. Using the inequality (8.1) at the point $a = z$ twice, one time with x replaced by y, we obtain the inequality

$$\begin{aligned}\lambda f(x) + (1-\lambda)f(y) &\ge \lambda\big(f(z) + \langle c, x - z\rangle\big) + (1-\lambda)\big(f(z) + \langle c, y - z\rangle\big)\\ &= f(z) + \langle c, \lambda x + (1-\lambda)y - z\rangle = f(z) + \langle c, 0\rangle = f(z),\end{aligned}$$

which shows that the restriction of f to $\operatorname{dom} f$ is a convex function, and this implies that f itself is convex.

(b) Conversely, assume that f is a convex function, and let a be a point in $\operatorname{rint}(\operatorname{dom} f)$. We will prove that the subdifferential $\partial f(a)$ is nonempty.

The point $(a, f(a))$ is a relative boundary point of the convex set $\operatorname{epi} f$. Therefore, there exists a supporting hyperplane

$$H = \{(x, x_{n+1}) \in \mathbf{R}^n \times \mathbf{R} \mid \langle c, x - a\rangle + c_{n+1}(x_{n+1} - f(a)) = 0\}$$

of $\operatorname{epi} f$ at the point $(a, f(a))$, and we may choose the normal vector (c, c_{n+1}) in such a way that

$$\langle c, x - a\rangle + c_{n+1}(x_{n+1} - f(a)) \ge 0 \tag{8.3}$$

for all points $(x, x_{n+1}) \in \operatorname{epi} f$. We shall see that this implies that $c_{n+1} > 0$.

By applying inequality (8.3) to the point $(a, f(a)+1)$ in the epigraph, we first obtain the inequality $c_{n+1} \geq 0$.

Now suppose that $c_{n+1} = 0$, and put $L = \operatorname{aff}(\operatorname{dom} f)$. Since $\operatorname{epi} f \subseteq L \times \mathbf{R}$ and the supporting hyperplane $H = \{(x, x_{n+1}) \in \mathbf{R}^n \times \mathbf{R} \mid \langle c, x-a\rangle = 0\}$ by definition does not contain $\operatorname{epi} f$ as a subset, it does not contain $L \times \mathbf{R}$ either. We conclude that there exists a point $y \in L$ such that $\langle c, y-a\rangle \neq 0$. Consider the points $y_\lambda = (1-\lambda)a + \lambda y$ for $\lambda \in \mathbf{R}$; these points lie in the afine set L, and $y_\lambda \to a$ as $\lambda \to 0$. Since a is a point in the relative interior of $\operatorname{dom} f$, the points y_λ lie in $\operatorname{dom} f$ if $|\lambda|$ is sufficiently small, and this implies that the inequality (8.3) can not hold for all points $(y_\lambda, f(y_\lambda))$ in the epigraph, because the expression $\langle c, y_\lambda - a\rangle$ $(= \lambda\langle c, y-a\rangle)$ assumes both positive and negative values depending on the sign of λ.

This is a contradiction and proves that $c_{n+1} > 0$, and by dividing inequality (8.3) by c_{n+1} and letting $d = -(1/c_{n+1})c$, we obtain the inequality

$$x_{n+1} \geq f(a) + \langle d, x-a\rangle$$

for all $(x, x_{n+1}) \in \operatorname{epi} f$. In particular, $f(x) \geq f(a) + \langle d, x-a\rangle$ for all $x \in \operatorname{dom} f$, which means that d is a subgradient of f at a. □

It follows from Theorem 8.1.4 that a real-valued function f with an open convex domain X is convex if and only if $\partial f(x) \neq \emptyset$ for all $x \in X$.

Theorem 8.1.5. *The subdifferential $\partial f(a)$ of a convex function f is a compact nonempty set if a is an interior point of* $\operatorname{dom} f$.

Proof. Suppose a is a point in $\operatorname{int}(\operatorname{dom} f)$. The subdifferential $\partial f(a)$ is closed by Theorem 8.1.1 and nonempty by Theorem 8.1.4, so it only remains to prove that it is a bounded set.

Theorem 6.6.1 yields two positive constants M and δ such that the closed ball $\overline{B}(a;\delta)$ lies in $\operatorname{dom} f$ and

$$|f(x) - f(a)| \leq M\|x-a\| \quad \text{for } x \in \overline{B}(a;\delta).$$

Now suppose that $c \in \partial f(a)$ and that $c \neq 0$. By choosing $x = a + \delta c/\|c\|$ in inequality (8.1), we conclude that

$$\delta\|c\| = \langle c, x-a\rangle \leq f(x) - f(a) \leq M\|x-a\| = \delta M$$

with the bound $\|c\| \leq M$ as a consequence. The subdifferential $\partial f(a)$ is thus included in the closed ball $\overline{B}(0;M)$. □

Theorem 8.1.6. *The sublevel sets of a strongly convex function $f\colon X \to \overline{\mathbf{R}}$ are bounded sets.*

Proof. Suppose that f is μ-strongly convex. Let x_0 be a point in the relative interior of $\operatorname{dom} f$, and let c be a subgradient at the point x_0 of the convex function $g(x) = f(x) - \frac{1}{2}\mu\|x\|^2$. Then, for each x belonging to the sublevel set $S = \{x \in X \mid f(x) \leq \alpha\}$,

$$\begin{aligned}\alpha \geq f(x) = g(x) + \tfrac{1}{2}\mu\|x\|^2 &\geq g(x_0) + \langle c, x - x_0\rangle + \tfrac{1}{2}\mu\|x\|^2 \\ &= f(x_0) - \tfrac{1}{2}\mu\|x_0\|^2 + \langle c, x - x_0\rangle + \tfrac{1}{2}\mu\|x\|^2 \\ &= f(x_0) + \tfrac{1}{2}\mu\big(\|x + \mu^{-1}c\|^2 - \|x_0 + \mu^{-1}c\|^2\big),\end{aligned}$$

which implies that

$$\|x + \mu^{-1}c\|^2 \leq \|x_0 + \mu^{-1}c\|^2 + 2\mu^{-1}(\alpha - f(x_0)).$$

The sublevel set S is thus included in a closed ball with center at the point $-\mu^{-1}c$ and radius $R = \sqrt{\|x_0 + \mu^{-1}c\|^2 + 2\mu^{-1}(\alpha - f(x_0))}$. □

Corollary 8.1.7. *If a continuous and strongly convex function has a nonempty closed sublevel set, then it has a unique minimum point.*

In particular, every strongly convex function $f\colon \mathbf{R}^n \to \mathbf{R}$ has a unique minimum point.

Proof. Let f be a continuous, strongly convex function with a nonempty closed sublevel set S. Then S is compact by the previous theorem, so the restriction of f to S assumes a minimum at some point in S, and this point is obviously a global minimum point of f. The minimum point is unique, because strongly convex functions are strictly convex.

A convex function $f\colon \mathbf{R}^n \to \mathbf{R}$ is automatically continuous, and continuous functions on $\mathbf{R}^n$ are closed. Hence, all sublevel sets of a strongly convex function on $\mathbf{R}^n$ are closed, so it follows from the already proven part of the theorem that there is a unique minimum point. □

8.2 Closed convex functions

In this section, we will use the subdifferential to supplement the results on closed convex functions in chapter 6.8 with some new results. We begin with an alternative characterization of closed convex functions.

Theorem 8.2.1. *A convex function $f\colon X \to \overline{\mathbf{R}}$ is closed if and only if, for all convergent sequences $(x_k)_1^\infty$ of points in* $\operatorname{dom} f$ *with limit x_0,*

$$\varliminf_{k\to\infty} f(x_k) \begin{cases} \geq f(x_0) & \text{if } x_0 \in \operatorname{dom} f, \\ = +\infty & \text{if } x_0 \in \operatorname{cl}(\operatorname{dom} f) \setminus \operatorname{dom} f. \end{cases} \tag{8.4}$$

Proof. Suppose that f is closed, i.e. that $\operatorname{epi} f$ is a closed set, and let $(x_k)_1^\infty$ be a sequence in $\operatorname{dom} f$ which converges to a point $x_0 \in \operatorname{cl}(\operatorname{dom} f)$, and put

$$L = \varliminf_{k\to\infty} f(x_k).$$

Let a be an arbitrary point in the relative interior of $\operatorname{dom} f$ and let c be a subgradient of f at the point a. Then $f(x_k) \geq f(a) + \langle c, x_k - a\rangle$ for all k, and since the right hand side converges (to $f(a) + \langle c, x_0 - a\rangle$) as $k \to \infty$, we conclude that the sequence $(f(x_k))_1^\infty$ is bounded below. Its least limit point, i.e. L, is therefore a real number or $+\infty$.

Inequality (8.4) is trivially satisfied if $L = +\infty$, so assume that L is a finite number, and let $(x_{k_j})_{j=1}^\infty$ be a subsequence of the given sequence with the property that $f(x_{k_j}) \to L$ as $j \to \infty$. The points $(x_{k_j}, f(x_{k_j}))$, which belong to $\operatorname{epi} f$, then converge to the point (x_0, L), and since the epigraph is assumed to be closed, we conclude that the limit point (x_0, L) belongs to the epigraph, i.e. $x_0 \in \operatorname{dom} f$ and $L \geq f(x_0)$.

So if x_0 does not belong to $\operatorname{dom} f$ but to $\operatorname{cl}(\operatorname{dom} f) \setminus \operatorname{dom} f$, then we must have $L = +\infty$. This proves that (8.4) holds.

Conversely, suppose (8.4) holds for all convergent sequences, and let $((x_k.t_k))_1^\infty$ be a sequence of points in $\operatorname{epi} f$ which converges to a point (x_0, t_0). Then, $(x_k)_1^\infty$ converges to x_0 and $(t_k)_1^\infty$ converges to t_0, and since $f(x_k) \le t_k$ for all k, we conclude that

$$\underline{\lim}_{k\to\infty} f(x_k) \le \underline{\lim}_{k\to\infty} t_k = t_0.$$

In particular, $\underline{\lim}_{k\to\infty} f(x_k) < +\infty$, so it follows from inequality (8.4) that $x_0 \in \operatorname{dom} f$ and that $f(x_0) \le t_0$. This means that the limit point (x_0, t_0) belongs to $\operatorname{epi} f$. Hence, $\operatorname{epi} f$ contains all its boundary points and is therefore a closed set, and this means that f is a closed function. □

Corollary 8.2.2. *Suppose that $f\colon X \to \overline{\mathbf{R}}$ is a convex function and that its effective domain* $\operatorname{dom} f$ *is relative open. Then, f is closed if and only if $\lim_{k\to\infty} f(x_k) = +\infty$ for each sequence $(x_k)_1^\infty$ of points in* $\operatorname{dom} f$ *that converges to a relative boundary point of* $\operatorname{dom} f$.

Proof. Since a convex function is continuous at all points in the relative interior of its effective domain, we conclude that $\lim_{k\to\infty} f(x_k) = f(x_0)$ for each sequence $(x_k)_1^\infty$ of points in $\operatorname{dom} f$ that converges to a point x_0 in $\operatorname{dom} f$. Condition (8.4) of the previous theorem is therefore fulfilled if and only if $\lim_{k\to\infty} f(x_k) = +\infty$ for all sequences $(x_k)_1^\infty$ in $\operatorname{dom} f$ that converge to a point in $\operatorname{rbdry}(\operatorname{dom} f)$. □

So a convex function with an affine set as effective domain is closed (and continuous), because affine sets lack relative boundary points.

EXAMPLE 8.2.1. The convex function $f(x) = -\ln x$ with $\mathbf{R}_{++}$ as domain is closed, since $\lim\limits_{x\to 0} f(x) = +\infty$. □

Theorem 8.2.3. *If the function $f\colon X \to \overline{\mathbf{R}}$ is convex and closed, then*

$$f(x) = \lim_{\lambda\to 1^-} f(\lambda x + (1-\lambda)y)$$

for all $x, y \in \operatorname{dom} f$.

Proof. The inequality

$$\overline{\lim}_{\lambda\to 1^-} f(\lambda x + (1-\lambda)y) \le \overline{\lim}_{\lambda\to 1^-} (\lambda f(x) + (1-\lambda)f(y)) = f(x)$$

holds for all convex functions f, and the inequality

$$\underline{\lim}_{\lambda\to 1^-} f(\lambda x + (1-\lambda)y) \ge f(x)$$

holds for all closed convex functions f according to Theorem 8.2.1. □

Theorem 8.2.4. *Suppose that f and g are two closed convex functions, that*

$$\operatorname{rint}(\operatorname{dom} f) = \operatorname{rint}(\operatorname{dom} g)$$

and that

$$f(x) = g(x)$$

for all $x \in \operatorname{rint}(\operatorname{dom} f)$. Then $f = g$.

We remind the reader that the equality $f = g$ should be interpreted as $\operatorname{dom} f = \operatorname{dom} g$ and $f(x) = g(x)$ for all points x in the common effective domain.

Proof. If $\operatorname{rint}(\operatorname{dom} f) = \emptyset$, then $\operatorname{dom} f = \operatorname{dom} g = \emptyset$, and there is nothing to prove, so suppose that x_0 is a point in $\operatorname{rint}(\operatorname{dom} f)$. Then, $\lambda x + (1-\lambda)x_0$ lies in $\operatorname{rint}(\operatorname{dom} f)$, too, for each $x \in \operatorname{dom} f$ and $0 < \lambda < 1$, and it follows from our assumptions and Theorem 8.2.3 that

$$g(x) = \lim_{\lambda \to 1^-} g(\lambda x + (1-\lambda)x_0) = \lim_{\lambda \to 1^-} f(\lambda x + (1-\lambda)x_0) = f(x).$$

Hence, $g(x) = f(x)$ for all $x \in \operatorname{dom} f$, and it follows that $\operatorname{dom} f \subseteq \operatorname{dom} g$. The converse inclusion holds by symmetry, so $\operatorname{dom} f = \operatorname{dom} g$. □

Theorem 8.2.5. *Let $f\colon X \to \overline{\mathbf{R}}$ and $g\colon Y \to \overline{\mathbf{R}}$ be two closed convex functions with $X \cap Y \neq \emptyset$. The sum $f + g\colon X \cap Y \to \overline{\mathbf{R}}$ is then a closed convex function.*

Proof. The theorem follows from the characterization of closedness in Theorem 8.2.1. Let $(x_k)_1^\infty$ be a convergent sequence of points in $\operatorname{dom}(f+g)$ with limit point x_0. If x_0 belongs to $\operatorname{dom}(f+g)$ ($= \operatorname{dom} f \cap \operatorname{dom} g$), then

$$\varliminf_{k\to\infty} (f(x_k) + g(x_k)) \geq \varliminf_{k\to\infty} f(x_k) + \varliminf_{k\to\infty} g(x_k) \geq f(x_0) + g(x_0),$$

and if x_0 does not belong to $\operatorname{dom}(f+g)$, then we use the trivial inclusion

$$\operatorname{cl}(A \cap B) \setminus A \cap B \subseteq (\operatorname{cl} A \setminus A) \cup (\operatorname{cl} B \setminus B),$$

with $A = \operatorname{dom} f$ and $B = \operatorname{dom} g$, to conclude that the sum $f(x_k) + g(x_k)$ tends to $+\infty$, because one of the two sequences $(f(x_k))_1^\infty$ and $(g(x_k))_1^\infty$ tends to $+\infty$ while the other either tends to $+\infty$ or has a finite limes inferior. $\square$

The closure

Definition. Let $f\colon X \to \overline{\mathbf{R}}$ be a function defined on a subset of $\mathbf{R}^n$ and define $(\operatorname{cl} f)(x)$ for $x \in \mathbf{R}^n$ by

$$(\operatorname{cl} f)(x) = \inf\{t \mid (x,t) \in \operatorname{cl}(\operatorname{epi} f)\}.$$

The function $\operatorname{cl} f\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ is called the *closure of f*.

Theorem 8.2.6. *The closure $\operatorname{cl} f$ of a convex function f, whose effective domain is a nonempty subset of $\mathbf{R}^n$, has the following properties:*

(i) *The closure $\operatorname{cl} f\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ is a convex function.*
(ii) $\operatorname{dom} f \subseteq \operatorname{dom}(\operatorname{cl} f) \subseteq \operatorname{cl}(\operatorname{dom} f)$.
(iii) $\operatorname{rint}(\operatorname{dom}(\operatorname{cl} f)) = \operatorname{rint}(\operatorname{dom} f)$.
(iv) $(\operatorname{cl} f)(x) \leq f(x)$ *for all* $x \in \operatorname{dom} f$.
(v) $(\operatorname{cl} f)(x) = f(x)$ *for all* $x \in \operatorname{rint}(\operatorname{dom} f)$.
(vi) $\operatorname{epi}(\operatorname{cl} f) = \operatorname{cl}(\operatorname{epi} f)$.

Proof. (i) Let x_0 be an arbitrary point in $\operatorname{rint}(\operatorname{dom} f)$, and let c be a subgradient of f at the point x_0. Then $f(x) \geq f(x_0) + \langle c, x - x_0\rangle$ for all $x \in \operatorname{dom} f$, which means that the epigraph $\operatorname{epi} f$ is a subset of the closed set $K = \{(x,t) \in \operatorname{cl}(\operatorname{dom} f) \times \mathbf{R} \mid \langle c, x - x_0\rangle + f(x_0) \leq t\}$. It follows that $\operatorname{cl}(\operatorname{epi} f) \subseteq K$, and hence

$$\begin{aligned}(\operatorname{cl} f)(x) &= \inf\{t \mid (x,t) \in \operatorname{cl}(\operatorname{epi} f)\} \geq \inf\{t \mid (x,t) \in K\} \\ &= f(x_0) + \langle c, x - x_0\rangle > -\infty\end{aligned}$$

for all $x \in \mathbf{R}^n$. So $\overline{\mathbf{R}}$ is a codomain of the function $\operatorname{cl} f$, and since $\operatorname{cl}(\operatorname{epi} f)$ is a convex set, it now follows from Theorem 6.2.6 that $\operatorname{cl} f\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ is a convex function.

(ii), (iv) and (v) It follows from the inclusion $\operatorname{epi} f \subseteq \operatorname{cl}(\operatorname{epi} f) \subseteq K$ that

$$(\operatorname{cl} f)(x) \begin{cases} \leq \inf\{t \mid (x,t) \in \operatorname{epi} f\} = f(x) < +\infty & \text{if } x \in \operatorname{dom} f, \\ \geq \inf\{t \mid (x,t) \in K\} = \inf \emptyset = +\infty & \text{if } x \notin \operatorname{cl}(\operatorname{dom} f), \end{cases}$$
$$(\operatorname{cl} f)(x_0) \geq \inf\{t \mid (x_0,t) \in K\} = f(x_0).$$

This proves that $\operatorname{dom} f \subseteq \operatorname{dom}(\operatorname{cl} f) \subseteq \operatorname{cl}(\operatorname{dom} f)$, that $(\operatorname{cl} f)(x) \leq f(x)$ for all $x \in \operatorname{dom} f$, and that $(\operatorname{cl} f)(x_0) = f(x_0)$, and since x_0 is an arbitrary point in $\operatorname{rint}(\operatorname{dom} f)$, we conclude that $(\operatorname{cl} f)(x) = f(x)$ for all $x \in \operatorname{rint}(\operatorname{dom} f)$.

(iii) Since $\operatorname{rint}(\operatorname{cl} X) = \operatorname{rint} X$ for arbitrary convex sets X, it follows in particular from (ii) that

$$\operatorname{rint}(\operatorname{dom} f) \subseteq \operatorname{rint}(\operatorname{dom}(\operatorname{cl} f)) \subseteq \operatorname{rint}(\operatorname{cl}(\operatorname{dom} f)) = \operatorname{rint}(\operatorname{dom} f),$$

with the conclusion that $\operatorname{rint}(\operatorname{dom}(\operatorname{cl} f)) = \operatorname{rint}(\operatorname{dom} f)$.

(vi) The implications

$$(x,t) \in \operatorname{cl}(\operatorname{epi} f) \Rightarrow (\operatorname{cl} f)(x) \leq t \Rightarrow (x,t) \in \operatorname{epi}(\operatorname{cl} f)$$

follow immediately from the closure and epigraph definitions. Conversely, suppose that (x,t) is a point in $\operatorname{epi}(\operatorname{cl} f)$, i.e. that $(\operatorname{cl} f)(x) \leq t$, and let $U \times I$ be an open neighborhood of (x,t). The neighborhood I of t contains a number s such that $(x,s) \in \operatorname{cl}(\operatorname{epi} f)$, and since $U \times I$ is also an open neighborhood of (x,s), it follows that $\operatorname{epi} f \cap (U \times I) \neq \emptyset$. This proves that $(x,t) \in \operatorname{cl}(\operatorname{epi} f)$, so we have the implication

$$(x,t) \in \operatorname{epi}(\operatorname{cl} f) \Rightarrow (x,t) \in \operatorname{cl}(\operatorname{epi} f).$$

Thus, $\operatorname{epi}(\operatorname{cl} f) = \operatorname{cl}(\operatorname{epi} f)$. □

Theorem 8.2.7. *If f is a closed convex function, then* $\operatorname{cl} f = f$.

Proof. We have $\operatorname{rint}(\operatorname{dom}(\operatorname{cl} f)) = \operatorname{rint}(\operatorname{dom} f)$ and $(\operatorname{cl} f)(x) = f(x)$ for all $x \in \operatorname{rint}(\operatorname{dom} f)$, by the previous theorem. Therefore it follows from Theorem 8.2.4 that $\operatorname{cl} f = f$. □

8.3 The conjugate function

Definition. Let $f\colon X \to \underline{\overline{\mathbf{R}}}$ be an arbitrary function defined on a subset X of $\mathbf{R}^n$ and define a function f^* on $\mathbf{R}^n$ by

$$f^*(y) = \sup\{\langle y, x\rangle - f(x) \mid x \in X\}$$

for $y \in \mathbf{R}^n$. The function f^* is called the *conjugate function* or the *Fenchel transform* of f.

We use the shorter notation f^{**} for the conjugate function of f^*, i.e. $f^{**} = (f^*)^*$.

The conjugate function f^* of a function $f\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ with a nonempty effective domain is obviously a function $\mathbf{R}^n \to \overline{\mathbf{R}}$, and

$$f^*(y) = \sup\{\langle y, x\rangle - f(x) \mid x \in \operatorname{dom} f\}.$$

There are two trivial cases: If the effective domain of $f\colon X \to \overline{\mathbf{R}}$ is empty, then $f^*(y) = -\infty$ for all $y \in \mathbf{R}^n$, and if $f\colon X \to \overline{\mathbf{R}}$ assumes the value $-\infty$ at some point, then $f^*(y) = +\infty$ for all $y \in \mathbf{R}^n$.

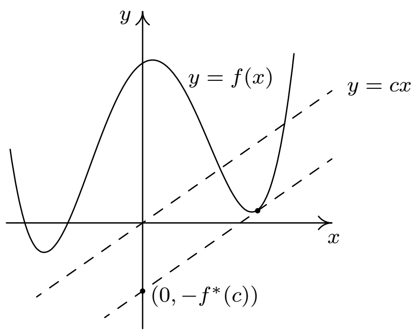

Figure 8.2. A graphical illustration of the conjugate function f^* when f is a one-variable function. The function value $f^*(c)$ is equal to the maximal vertical distance between the line $y = cx$ and the curve $y = f(x)$. If f is differentiable, then $f^*(c) = cx_0 - f(x_0)$ for some point x_0 with $f'(x_0) = c$.

Example 8.3.1. The support functions that were defined in Section 6.9, are conjugate functions. To see this, define for a given subset A of $\mathbf{R}^n$ the function $\chi_A\colon \mathbf{R}^n \to \overline{\mathbf{R}}$ by

$$\chi_A(x) = \begin{cases} 0 & \text{if } x \in A, \\ +\infty & \text{if } x \notin A. \end{cases}$$

The function χ_A is called the *indicator function* of the set A, and it is a convex function if A is a convex set. Obviously,

$$\chi_A^*(y) = \sup\{\langle y, x\rangle \mid x \in A\} = S_A(y)$$

for all $y \in \mathbf{R}^n$, so the support function of A coincides with the conjugate function χ_A^* of the indicator function of A. $\square$

We are primarily interested in conjugate functions of convex functions $f: X \to \overline{\mathbf{R}}$, but we start with some general results.

Theorem 8.3.1. *The conjugate function f^* of a function $f: X \to \overline{\mathbf{R}}$ with a nonempty effective domain is convex and closed.*

Proof. The epigraph $\operatorname{epi} f^*$ consists of all points $(y,t) \in \mathbf{R}^n \times \mathbf{R}$ that satisfy the inequalities $\langle x, y\rangle - t \le f(x)$ for all $x \in \operatorname{dom} f$, which means that it is the intersection of a family of closed halfspaces in $\mathbf{R}^n \times \mathbf{R}$. Hence, $\operatorname{epi} f^*$ is a closed convex set, so the conjugate function f^* is closed and convex. □

Theorem 8.3.2 (Fenchel's inequality)**.** *Let $f: X \to \overline{\mathbf{R}}$ be a function with a nonempty effective domain. Then*

$$\langle x, y\rangle \le f(x) + f^*(y)$$

for all $x \in X$ and all $y \in \mathbf{R}^n$. Moreover, the two sides are equal for a given $x \in \operatorname{dom} f$ if and only if $y \in \partial f(x)$.

Proof. The inequality follows immediately from the definition of $f^*(y)$ as a least upper bound if $x \in \operatorname{dom} f$, and it is trivially true if $x \in X \setminus \operatorname{dom} f$.

Moreover, by the subgradient definition, if $x \in \operatorname{dom} f$ then

$$\begin{aligned} y \in \partial f(x) &\Leftrightarrow f(z) - f(x) \geq \langle y, z - x\rangle \quad \text{for all } z \in \operatorname{dom} f \\ &\Leftrightarrow \langle y, z\rangle - f(z) \leq \langle y, x\rangle - f(x) \quad \text{for all } z \in \operatorname{dom} f \\ &\Leftrightarrow f^*(y) \leq \langle y, x\rangle - f(x) \\ &\Leftrightarrow f(x) + f^*(y) \leq \langle x, y\rangle, \end{aligned}$$

and by combining this with the already proven Fenchel inequality, we obtain the equivalence $y \in \partial f(x) \Leftrightarrow f(x) + f^*(y) = \langle x, y\rangle$. □

By the previous theorem, for all points y in the set $\bigcup\{\partial f(x) \mid x \in \operatorname{dom} f\}$

$$f^*(y) = \langle x_y, y\rangle - f(x_y),$$

where x_y is a point satisfying the condition $y \in \partial f(x_y)$. For differentiable functions f we obtain the points x_y as solutions to the equation $f'(x) = y$. Here follows a concrete example.

EXAMPLE 8.3.2. Let $f\colon]-1, \infty[\to \mathbf{R}$ be the function

$$f(x) = \begin{cases} -x(x+1)^{-1} & \text{if } -1 < x \leq 0, \\ 2x & \text{if } 0 \leq x < 1, \\ (x-2)^2 + 1 & \text{if } 1 \leq x < 2, \\ 2x - 3 & \text{if } x \geq 2. \end{cases}$$

Its graph is shown in the left part of Figure 8.3.

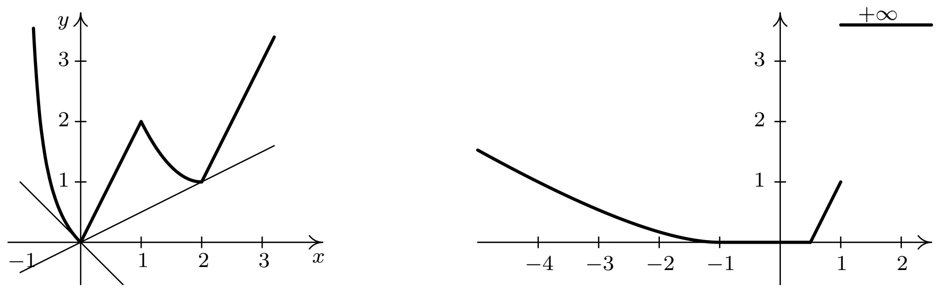

Figure 8.3. To the left the graph of the function $f\colon]-1, \infty[\to \mathbf{R}$, and to the right the graph of the conjugate function $f^*\colon \mathbf{R} \to \overline{\mathbf{R}}$.

A look at the figure shows that the curve $y = f(x)$ lies above all lines that are tangent to the curve at a point (x, y) with $-1 < x < 0$, lies above all lines through the origin with a slope between $f'_{-}(0) = -1$ and the slope of the chord that connects the origin and the point $(2,1)$ on the curve, and lies above all lines through the point $(2,1)$ with a slope between $\frac{1}{2}$ and $f'_{+}(2) = 2$. This means that

$$\begin{aligned}\bigcup\{\partial f(x)\} &= \bigcup_{-1<x<0}\{f'(x)\} \cup \partial f(0) \cup \partial f(2) \cup \bigcup_{x>2}\{f'(x)\}\\ &=]-\infty,-1[\cup [-1,\tfrac{1}{2}] \cup [\tfrac{1}{2},2] \cup \{2\} =]-\infty,2].\end{aligned}$$

The equation $f'(x) = c$ has for $c < -1$ the solution $x = -1 + \sqrt{-1/c}$ in the interval $-1 < x < 0$. Let

$$x_c = \begin{cases} -1+\sqrt{-1/c} & \text{if } c < -1,\\ 0 & \text{if } -1 \le c \le \frac{1}{2},\\ 2 & \text{if } \frac{1}{2} \le c \le 2.\end{cases}$$

Then $c \in \partial f(x_c)$, and it follows from the remark after Theorem 8.3.2 that

$$f^*(c) = cx_c - f(x_c) = \begin{cases} -c-2\sqrt{-c}+1 & \text{if } c < -1,\\ 0 & \text{if } -1 \le c \le \frac{1}{2},\\ 2c-1 & \text{if } \frac{1}{2} \le c \le 2.\end{cases}$$

Since

$$f^*(c) = \sup_{x>-1}\{cx - f(x)\} \ge \sup_{x\ge 2}\{cx - f(x)\} = \sup_{x\ge 2}\{(c-2)x + 3\} = +\infty$$

if $c > 2$, we conclude that $\operatorname{dom} f^* =]-\infty, 2]$. The graph of f^* is shown in the right part of Figure 8.3. □

Theorem 8.3.3. *Let $f\colon X \to \overline{\mathbf{R}}$ be an arbitrary function. Then*

$$f^{**}(x) \le f(x)$$

for all $x \in X$. Furthermore, $f^{**}(x) = f(x)$ if $x \in \operatorname{dom} f$ and $\partial f(x) \ne \emptyset$.

Proof. If $f(x) = +\infty$ for all $x \in X$, then $f^* \equiv -\infty$ and $f^{**} \equiv +\infty$, according to the remarks following the definition of the conjugate function, so the inequality holds with equality for all $x \in X$ in this trivial case.

Suppose, therefore, that $\operatorname{dom} f \ne \emptyset$. Then $\langle x, y\rangle - f^*(y) \le f(x)$ for all $x \in X$ and all $y \in \operatorname{dom} f^*$ because of Fenchel's inequality, and hence $f^{**}(x) = \sup\{\langle x,y\rangle - f^*(y) \mid y \in \operatorname{dom} f^*\} \le f(x)$.

If $\partial f(x) \ne \emptyset$, then Fenchel's inequality holds with equality for $y \in \partial f(x)$. This means that $f(x) = \langle x, y\rangle - f^*(y) \le f^{**}(x)$ and implies that $f(x) = f^{**}(x)$. □

The following corollary follows immediately from Theorem 8.3.3, because convex functions have subgradients at all relative interior points of their effective domains.

Corollary 8.3.4. *If $f\colon X \to \overline{\mathbf{R}}$ is a convex function, then $f^{**}(x) = f(x)$ for all x in the relative interior of* $\operatorname{dom} f$*.*

We will prove that $f^{**} = \operatorname{cl} f$ if f is a convex function, and for this purpose we need the following lemma.

Lemma 8.3.5. *Suppose that f is a convex function and that (x_0, t_0) is a point in $\mathbf{R}^n \times \mathbf{R}$ which does not belong to* $\operatorname{cl}(\operatorname{epi} f)$*. Then there exist a vector $c \in \mathbf{R}^n$ and a real number d such that the "non-vertical" hyperplane*

$$H = \{(x, x_{n+1}) \mid x_{n+1} = \langle c, x\rangle + d\}$$

strictly separates the point (x_0, t_0) from $\operatorname{cl}(\operatorname{epi} f)$*.*

Proof. By the Separation Theorem 3.1.3, there exists a hyperplane

$$H = \{(x, x_{n+1}) \mid c_{n+1}x_{n+1} = \langle c, x\rangle + d\}$$

which strictly separates the point from $\operatorname{cl}(\operatorname{epi} f)$. If $c_{n+1} \neq 0$, we can without loss of generality assume that $c_{n+1} = 1$, and there is nothing more to prove.

So assume that $c_{n+1} = 0$, and choose the signs of c and d so that $\langle c, x_0\rangle + d > 0$ and $\langle c, x\rangle + d < 0$ for all $x \in \operatorname{dom} f$.

Using the subgradient c' at some point in the relative interior of $\operatorname{dom} f$ we obtain an affine function $\langle c', x\rangle + d'$ such that $f(x) \geq \langle c', x\rangle + d'$ for all $x \in \operatorname{dom} f$. This implies that

$$f(x) \geq \langle c', x\rangle + d' + \lambda(\langle c, x\rangle + d) = \langle c' + \lambda c, x\rangle + d' + \lambda d$$

for all $x \in \operatorname{dom} f$ and all positive numbers λ, while

$$\langle c' + \lambda c, x_0\rangle + d' + \lambda d = \langle c', x_0\rangle + d' + \lambda(\langle c, x_0\rangle + d) > t_0$$

for all sufficiently large numbers λ. So the epigraph $\operatorname{epi} f$ lies above the hyperplane

$$H_\lambda = \{(x, x_{n+1}) \mid x_{n+1} = \langle c' + \lambda c, x\rangle + d' + \lambda d\}.$$

and the point (x_0, t_0) lies strictly below the same hyperplane, if the number λ is big enough. By moving the hyperplane H_λ slightly downwards, we obtain a parallel non-vertical hyperplane which strictly separates (x_0, t_0) and $\operatorname{cl}(\operatorname{epi} f)$. □

Lemma 8.3.6. *If $f\colon X \to \overline{\mathbf{R}}$ is a convex function, then*

$$\operatorname{rint}(\operatorname{dom} f^{**}) = \operatorname{rint}(\operatorname{dom} f).$$

Proof. Since $\operatorname{rint}(\operatorname{dom} f) = \operatorname{rint}(\operatorname{cl}(\operatorname{dom} f))$, it suffices to prove the inclusion

$$\operatorname{dom} f \subseteq \operatorname{dom} f^{**} \subseteq \operatorname{cl}(\operatorname{dom} f).$$

The left inclusion follows immediately from the inequality in Theorem 8.3.3. To prove the right inclusion, we assume that $x_0 \notin \operatorname{cl}(\operatorname{dom} f)$ and shall prove that this implies that $x_0 \notin \operatorname{dom} f^{**}$.

It follows from our assumption that the points (x_0, t_0) do not belong to $\operatorname{cl}(\operatorname{epi} f)$ for any number t_0. Thus, given $t_0 \in \mathbf{R}$ there exists, by the previous lemma, a hyperplane $H = \{(x, x_{n+1}) \in \mathbf{R}^n \times \mathbf{R} \mid x_{n+1} = \langle c, x\rangle + d\}$ which strictly separates (x_0, t_0) and $\operatorname{cl}(\operatorname{epi} f)$. Hence, $t_0 < \langle c, x_0\rangle + d$ and $\langle c, x\rangle + d < f(x)$ for all $x \in \operatorname{dom} f$. Consequently,

$$-d \geq \sup\{\langle c, x\rangle - f(x) \mid x \in \operatorname{dom} f\} = f^*(c),$$

and hence

$$t_0 < \langle c, x_0\rangle + d \leq \langle c, x_0\rangle - f^*(c) \leq f^{**}(x_0).$$

Since this holds for all real numbers t_0, it follows that $f^{**}(t_0) = +\infty$, which means that $x_0 \notin \operatorname{dom} f^{**}$. □

Theorem 8.3.7. *If f is a convex function, then $f^{**} = \operatorname{cl} f$.*

Proof. It follows from Lemma 8.3.6 and Theorem 8.2.6 (iii) that

$$\operatorname{rint}(\operatorname{dom} f^{**}) = \operatorname{rint}(\operatorname{dom}(\operatorname{cl} f)),$$

and from Theorem 8.3.4 and Theorem 8.2.6 (v) that

$$f^{**}(x) = (\operatorname{cl} f)(x)$$

for all $x \in \operatorname{rint}(\operatorname{dom} f^{**})$. So the two functions f^{**} and $\operatorname{cl} f$ are equal, according to Theorem 8.2.4, because both of them are closed and convex . □

Corollary 8.3.8. *If f is a closed convex function, then $f^{**} = f$.*

8.4 The direction derivative

Definition. Suppose the function $f\colon X \to \mathbf{R}$ is defined in a neighborhood of x, and let v be an arbitrary vector in $\mathbf{R}^n$. The limit

$$f'(x;v) = \lim_{t\to 0+} \frac{f(x+tv) - f(x)}{t},$$

provided it exists, is called the *direction derivative* of f at the point x in the direction v.

If f is differentiable at x, then obviously $f'(x;v) = Df(x)[v]$.

EXAMPLE 8.4.1. If f is a one-variable function, then

$$f'(x;v) = \begin{cases} f'_+(x)v & \text{if } v > 0, \\ 0 & \text{if } v = 0, \\ f'_-(v)v & \text{if } v < 0. \end{cases}$$

So, the direction derivative is a generalization of left- and right derivatives. □

Theorem 8.4.1. *Let $f\colon X \to \mathbf{R}$ be a convex function with an open domain. The direction derivative $f'(x;v)$ exists for all $x \in X$ and all directions v, and*

$$f(x+v) \geq f(x) + f'(x;v)$$

if $x+v$ lies in X.

Proof. Let $\phi(t) = f(x + tv)$; then $f'(x; v) = \phi'_+(0)$, which exists since convex one-variable functions do have right derivatives at each point by Theorem 7.1.2. Moreover,

$$\phi(t) \geq \phi(0) + \phi'_+(0)\, t$$

for all t in the domain of ϕ, and we obtain the inequality of the theorem by choosing $t = 1$. □

Theorem 8.4.2. *The direction derivative $f'(x; v)$ of a convex function is a positively homogeneous and convex function of the second variable v, i.e.*

$$f'(x; \alpha v) = \alpha f'(x; v) \quad \text{if } \alpha \geq 0$$
$$f'(x; \alpha v + (1 - \alpha)w) \leq \alpha f'(x; v) + (1 - \alpha) f'(x; w) \quad \text{if } 0 \leq \alpha \leq 1.$$

Proof. The homogenouity follows directly from the definition (and holds for arbitrary functions). Moreover, for convex functions f

$$\begin{aligned} f(x + t(\alpha v + (1 - \alpha)w)) - f(x) &= f(\alpha(x + tv) + (1 - \alpha)(x + tw)) - f(x) \\ &\leq \alpha\big(f(x + tv) - f(x)\big) + (1 - \alpha)\big(f(x + tw) - f(x)\big). \end{aligned}$$

The required convexity inequality is now obtained after division by $t > 0$ by passing to the limit as $t \to 0+$. □

Theorem 7.1.2 gives a relation between the subgradient and the direction derivative for convex one-variable functions f – the number c is a subgradient at x if and only if $f'_-(x) \le c \le f'_+(x)$. The subdifferential $\partial f(x)$ is in other words equal to the interval $[f'_-(x), f'_+(x)]$.

We may express this relation using the support function of the subdifferential. Let us recall that the support function S_X of a set X in $\mathbf{R}^n$ is defined as

$$S_X(x) = \sup\{\langle y, x\rangle \mid y \in X\}.$$

For one-variable functions f this means that

$$\begin{aligned} S_{\partial f(x)}(v) = S_{[f'_-(x), f'_+(x)]}(v) = \max\{f'_+(x)v, f'_-(x)v\} &= \begin{cases} f'_+(x)v & \text{if } v > 0,\\ 0 & \text{if } v = 0,\\ f'_-(x)v & \text{if } v < 0 \end{cases}\\ &= f'(x; v). \end{aligned}$$

We will generalize this identity, and to achieve this we need to consider the subgradients of the function $v \mapsto f'(x; v)$. We denote the subdifferential of this function at the point v_0 by $\partial_2 f'(x; v_0)$.

If the function $f\colon X \to \mathbf{R}$ is convex, then so is the function $v \mapsto f'(x; v)$, according to our previous theorem, and the subdifferentials $\partial_2 f'(x; v)$ are thus nonempty sets for all $x \in X$ and all $v \in \mathbf{R}^n$.

Lemma 8.4.3. *Let $f\colon X \to \mathbf{R}$ be a convex function with an open domain X and let x be a point in X. Then:*

(i) $\quad c \in \partial_2 f'(x; 0) \Leftrightarrow f'(x; v) \ge \langle c, v\rangle \quad$ *for all* $v \in \mathbf{R}^n$

(ii) $\quad \partial_2 f'(x; v) \subseteq \partial_2 f'(x; 0) \quad$ *for all* $v \in \mathbf{R}^n$

(iii) $\quad c \in \partial_2 f'(x; v) \Rightarrow f'(x; v) = \langle c, v\rangle$

(iv) $\quad \partial f(x) = \partial_2 f'(x; 0).$

Proof. The equivalence (i) follows directly from the definition of the subgradient and the fact that $f'(x; 0) = 0$.

(ii) and (iii): Suppose $c \in \partial_2 f'(x; v)$ and let $w \in \mathbf{R}^n$ be an arbitrary vector. Then, by homogenouity and the definition of the subgradient, we have the following inequality for $t \ge 0$:

$$tf'(x; w) = f'(x; tw) \ge f'(x; v) + \langle c, tw - v\rangle = f'(x; v) + t\langle c, w\rangle - \langle c, v\rangle,$$

and this is possible for all $t > 0$ only if $f'(x; w) \ge \langle c, w\rangle$. So we conclude from (i) that $c \in \partial_2 f'(x; 0)$, and this proves the inclusion (ii). By choosing $t = 0$

we obtain the inequality $f'(x;v) \leq \langle c, v\rangle$, which implies that $f'(x;v) = \langle c, v\rangle$, and completes the proof of the implication (iii).

(iv) Suppose $c \in \partial_2 f'(x;0)$. By (i) and Theorem 8.4.1,

$$f(y) \geq f(x) + f'(x; y-x) \geq f(x) + \langle c, y-x\rangle$$

for all $y \in X$, which proves that c is a subgradient of f at the point x and gives us the inclusion $\partial_2 f'(x;0) \subseteq \partial f(x)$.

Conversely, suppose $c \in \partial f(x)$. Then $f(x+tv) - f(x) \geq \langle c, tv\rangle = t\langle c, v\rangle$ for all sufficiently small numbers t. Division by $t > 0$ and passing to the limit as $t \to 0+$ results in the inequality $f'(x;v) \geq \langle c, v\rangle$, and it now follows from (i) that $c \in \partial_2 f'(x;0)$. This proves the inclusion $\partial f(x) \subseteq \partial_2 f'(x;0)$, and the proof is now complete. □

Theorem 8.4.4. *Suppose that $f\colon X \to \mathbf{R}$ is a convex function with an open domain. Then*

$$f'(x;v) = S_{\partial f(x)}(v) = \max\{\langle c, v\rangle \mid c \in \partial f(x)\}$$

for all $x \in X$ and all $v \in \mathbf{R}^n$.

Proof. It follows from (i) and (iv) in the preceding lemma that

$$\langle c, v\rangle \leq f'(x;v)$$

for all $c \in \partial f(x)$, and from (ii), (iii) and (iv) in the same lemma that $\langle c, v\rangle$ is equal to $f'(x;v)$ for all subgradients c in the nonempty subset $\partial_2 f'(x;v)$ of $\partial f(x)$. □

8.5 Subdifferentiation rules

Theorem 8.5.1. *Let X be an open convex set, and suppose that $f_i\colon X \to \mathbf{R}$ are convex functions and α_i are nonnegative numbers for $i = 1, 2 \ldots, m$. Define*

$$f = \sum_{i=1}^m \alpha_i f_i.$$

Then

$$\partial f(x) = \sum_{i=1}^m \alpha_i \partial f_i(x).$$

Proof. A sum of compact, convex sets is compact and convex. Therefore, $\sum_{i=1}^m \alpha_i \partial f_i(x)$ is a closed and convex set, just as the set $\partial f(x)$. Hence, by

Theorem 6.9.2 it suffices to prove that the two sets have the same support function. And this follows from Theorems 8.4.4 and 6.9.1, according to which

$$S_{\partial f(x)}(v) = f'(x;v) = \sum_{i=1}^{m} \alpha_i f_i'(x;v) = \sum_{i=1}^{m} \alpha_i S_{\partial f_i}(x) = S_{\sum_{i=1}^{m} \alpha_i \partial f_i(x)}(v). \quad \square$$

Theorem 8.5.2. *Suppose that the functions $f_i\colon X \to \mathbf{R}$ are convex for $i = 1$, $2, \dots, m$, and that their domain X is open, and let*

$$f = \max_{1\le i\le m} f_i.$$

Then

$$\partial f(x) = \operatorname{cvx}\Bigl(\bigcup_{i\in I(x)} \partial f_i(x)\Bigr),$$

for all $x \in X$, where $I(x) = \{i \mid f_i(x) = f(x)\}$.

Proof. The functions f_i are continuous at x and $f_j(x) < f(x)$ for all $j \notin I(x)$. Hence, for all sufficiently small numbers t,

$$f(x+tv) - f(x) = \max_{i\in I(x)} f_i(x+tv) - f(x) = \max_{i\in I(x)} \bigl(f_i(x+tv) - f_i(x)\bigr),$$

and it follows after division by t and passing to the limit that

$$f'(x;v) = \max_{i \in I(x)} f_i'(x;v).$$

We use Theorem 6.9.1 to conclude that

$$\begin{aligned} S_{\partial f(x)}(v) &= f'(x;v) = \max_{i \in I(x)} f_i'(x;v) = \max_{i \in I(x)} S_{\partial f_i(x)}(v) = S_{\bigcup_{i \in I(x)} \partial f_i(x)}(v) \\ &= S_{\operatorname{cvx}(\bigcup_{i \in I(x)} \partial f_i(x))}(v), \end{aligned}$$

and the equality for $\partial f(x)$ is now a consequence of Theorem 6.9.2. □

Our next theorem shows how to compute the subdifferential of a composition with affine functions.

Theorem 8.5.3. *Suppose C is a linear map from $\mathbf{R}^n$ to $\mathbf{R}^m$, that b is a vector in $\mathbf{R}^m$, and that g is a convex function with an open domain in $\mathbf{R}^m$, and let f be the function defined by $f(x) = g(Cx + b)$. Then, for each x in the domain of f,*

$$\partial f(x) = C^T(\partial g(Cx + b)).$$

Proof. The sets $\partial f(x)$ and $C^T(\partial g(Cx + b))$ are convex and compact, so it suffices to show that their support functions are identical. But for each $v \in \mathbf{R}^n$

$$\begin{aligned} f'(x;v) &= \lim_{t \to 0+} \frac{g(C(x + tv) + b) - g(Cx + b)}{t} \\ &= \lim_{t \to 0+} \frac{g(Cx + b + t\,Cv) - g(Cx + b)}{t} = g'(Cx + b; Cv), \end{aligned}$$

so it follows because of Theorem 6.9.1 that

$$S_{\partial f(x)}(v) = f'(x;v) = g'(Cx + b; Cv) = S_{\partial g(Cx+b)}(Cv) = S_{C^T(\partial g(Cx+b))}(v). \ \square$$

The Karush–Kuhn–Tucker theorem

As an application of the subdifferentiation rules we now prove a variant of a theorem by Karush–Kuhn–Tucker on minimization of convex functions with convex constraints. A more thorough treatment of this theme will be given in Part II.

Theorem 8.5.4. *Suppose that the functions f, $g_1, g_2, \dots, g_m$ are convex and defined on an open convex set Ω, and let*

$$X = \{x \in \Omega \mid g_i(x) \le 0 \ \text{for } i = 1, 2, \dots, m.\}$$

Moreover, suppose that there exists a point $\overline{x} \in \Omega$ such that $g_i(\overline{x}) < 0$ for $i = 1, 2, \ldots, m$. (Slater's condition)

Then, $\hat{x} \in X$ is a minimum point of the restriction $f|_X$ if and only if for each $i = 1, 2, \ldots, m$ there exist a subgradient $c_i \in \partial g_i(\hat{x})$ and a scalar $\hat{\lambda}_i \in \mathbf{R}_+$ with the following properties:

(i) $$-\sum_{i=1}^{m} \hat{\lambda}_i c_i \in \partial f(\hat{x}) \qquad \textit{and}$$

(ii) $$\hat{\lambda}_i g_i(\hat{x}) = 0 \quad \textit{for } i = 1, 2, \ldots, m.$$

Remark. If the functions are differentiable, then condition (i) simplifies to

$$\nabla f(\hat{x}) + \sum_{i=1}^{m} \hat{\lambda}_i \nabla g_i(\hat{x}) = 0.$$

Cf. Theorem 11.2.1 in Part II.

Proof. Let $\hat{x}$ be a point in X and consider the convex function

$$h(x) = \max\{f(x) - f(\hat{x}), g_1(x), \ldots, g_m(x)\}$$

with Ω as its domain. Clearly, $h(\hat{x}) = 0$. By defining

$$I(\hat{x}) = \{i \mid g_i(\hat{x}) = 0\},$$

we obtain $I(\hat{x}) = \{i \mid g_i(\hat{x}) = h(\hat{x})\}$, and it follows from Theorem 8.5.2 that

$$\partial h(\hat{x}) = \operatorname{cvx}\bigl(\partial f(\hat{x}) \cup \bigcup\{\partial g_i(\hat{x}) \mid i \in I(\hat{x})\}\bigr).$$

Now assume that $\hat{x}$ is a minimum point of the restriction $f|_X$. Then $h(x) = f(x) - f(\hat{x}) \geq 0$ for all $x \in X$ with equality when $x = \hat{x}$. And if $x \notin X$, then $h(x) > 0$ since $g_i(x) > 0$ for at least one i. Thus, $\hat{x}$ is a global minimum point of h.

Conversely, if $\hat{x}$ is a global minimum point of h, then $h(x) \geq 0$ for all $x \in \Omega$. In particular, for $x \in X$ this means that $h(x) = f(x) - f(\hat{x}) \geq 0$, and hence $\hat{x}$ is a mimimum point of the restriction $f|_X$, too.

Using Theorem 8.1.2 we therefore obtain the following equivalences:

$$\begin{aligned}
\hat{x} \text{ is a minimum point of } f|_X \quad &\Leftrightarrow \quad \hat{x} \text{ is a minimum point of } h \\
&\Leftrightarrow \quad 0 \in \partial h(\hat{x}) \\
&\Leftrightarrow \quad 0 \in \operatorname{cvx}\bigl(\partial f(\hat{x}) \cup \bigcup\{\partial g_i(\hat{x}) \mid i \in I(\hat{x})\}\bigr) \\
&\Leftrightarrow \quad 0 = \lambda_0 c_0 + \sum_{i \in I(\hat{x})} \lambda_i c_i \\
&\Leftrightarrow \quad \lambda_0 c_0 = -\sum_{i \in I(\hat{x})} \lambda_i c_i, \qquad (8.5)
\end{aligned}$$

where $c_0 \in \partial f(\hat{x})$, $c_i \in \partial g_i(\hat{x})$ for $i \in I(\hat{x})$, and the scalars λ_i are nonnegative numbers with sum equal to 1.

We now claim that $\lambda_0 > 0$. To prove this, assume the contrary. Then $\sum_{i \in I(\hat{x})} \lambda_i c_i = 0$, and it follows that

$$\sum_{i \in I(\hat{x})} \lambda_i g_i(\overline{x}) \geq \sum_{i \in I(\hat{x})} \lambda_i \big(g_i(\hat{x}) + \langle c_i, \overline{x} - \hat{x} \rangle\big) = \langle \sum_{i \in I(\hat{x})} \lambda_i c_i, \overline{x} - \hat{x} \rangle = 0,$$

which is a contradiction, since $g_i(\overline{x}) < 0$ for all i and $\lambda_i > 0$ for some $i \in I(\hat{x})$.

We may therefore divide the equality in (8.5) by λ_0, and conditions (i) and (ii) in our theorem are now fulfilled if we define $\hat{\lambda}_i = \lambda_i / \lambda_0$ for $i \in I(\hat{x})$, and $\hat{\lambda}_i = 0$ for $i \notin I(\hat{x})$, and choose arbitrary subgradients $c_i \in \partial g_i(\hat{x})$ for $i \notin I(\hat{x})$. $\square$

Exercises

8.1 Suppose $f\colon \mathbf{R}^n \to \mathbf{R}$ is a strongly convex function. Prove that

$$\lim_{\|x\|\to\infty} f(x) = \infty$$

8.2 Find $\partial f(-1,1)$ for the function $f(x_1,x_2) = \max(|x_1|,|x_2|)$.

8.3 Determine the subdifferential $\partial f(0)$ at the origin for the following functions $f\colon \mathbf{R}^n \to \mathbf{R}$:

a) $f(x) = \|x\|_2$ b) $f(x) = \|x\|_\infty$ c) $f(x) = \|x\|_1$.

8.4 Determine the conjugate functions of the following functions:

a) $f(x) = ax + b$, $\operatorname{dom} f = \mathbf{R}$ b) $f(x) = -\ln x$, $\operatorname{dom} f = \mathbf{R}_{++}$
c) $f(x) = e^x$, $\operatorname{dom} f = \mathbf{R}$ d) $f(x) = x \ln x$, $\operatorname{dom} f = \mathbf{R}_{++}$
e) $f(x) = 1/x$, $\operatorname{dom} f = \mathbf{R}_{++}$.

8.5 Use the relation between the support function S_A and the indicator function χ_A and the fact that $S_A = S_{\mathrm{cl}(\mathrm{cvx}\,A)}$ to prove Corollary 6.9.3, i.e. that

$$\mathrm{cl}(\mathrm{cvx}\,A) = \mathrm{cl}(\mathrm{cvx}\,B) \iff S_A = S_B.$$

Bibliografical and historical notices

Basic references in convex analysis are the books by Rockafellar [1] from 1970 and Hiriart-Urutty–Lemarechal [1] from 1993. Almost all results in Part I and in Chapters 9–10 in Part II of this series can be found in one form or another in Rockafellar's book, which also contains a historical overview with references to the original works in the field. A more accessible book on the same subject is Webster [1]. Among textbooks in convexity with an emphasis on polyhedra, one should mention Stoer–Witzgall [1] and the more combinatorially oriented Grünbaum [1].

The general convexity theory was founded around the turn of the century 1900 by Hermann Minkowski [1, 2] as a byproduct of his number theoretic studies. Among other things, Minkowski introduced the concepts of separation and extreme point, and he showed that every compact convex set is equal to the convex hull of its extreme points and that every polyhedron is finitely generated, i.e. one direction of Theorem 5.3.1) − the converse was noted later by Weyl [1].

The concept of dual cone was introduced by Steinitz [1], who also showed basic results about the recession cone.

The theory of linear inequalities is surprisingly young − a special case of Theorem 3.3.7 (Exercise 3.11a) was proved by Gordan [1], the algebraic version of Farkas's lemma, i.e. Corollary 3.3.3, can be found in Farkas [1], and a closely related result (Exercise 3.11b) is given by Stiemke [1]. The first systematic treatment of the theory is given by Weyl [1] and Motzkin [1]. Significant contributions have also been provided by Tucker [1]. The proof in Chapter 3 of Farkas's lemma has a geometrical character; an alternative algebraic induction proof of the lemma has been given by Kuhn [1].

Extreme points and faces are treated in detail in Klee [1,2].

Jensen [1] studied convex functions of one real variable and showed that convex functions with $\mathbf{R}$ as domain are continuous and have one-sided derivatives everywhere. Jensen's inequality, however, was shown earlier for functions with positive second derivative by Hölder [1].

The conjugate function was introduced by Fenchel [1], and a modern treatment of the theory of convex cones, sets and functions can be found in Fenchel [2], which among other things contains original results about the closure of convex functions and about the subdifferential.

References

Farkas, J.

[1] Theorie der einfachen Ungleichungen, *J. reine angew. Math.* 124 (1902), 1–27.

Fenchel, W.

[1] On conjugate convex functions. *Canad. J. Math.* 1 (1949), 73-77.

[2] *Convex Cones, Sets and Functions.* Lecture Notes, Princeton University, 1951.

Gordan, P.

[1] Über die Auflösung linearer Gleichungen mit reellen Coefficienten, *Math. Ann.* 6 (1873), 23–28.

Grünbaum, B.

[1] *Convex Polytopes.* Interscience publishers, New York, 1967.

Hiriart-Urruty, J.-B. & Lemaréchal, C.

[1] *Convex Analysis and Minimization Algorithms.* Springer, 1993.

Hölder, O.

[1] Über einen Mittelwertsatz, *Nachr. Ges. Wiss. Göttingen*, 38–47, 1989.

Jensen, J.L.W.V.

[1] Sur les fonctions convexes et les inégalités entre les valeur moyennes, *Acta Math.* 30 (1906), 175–193.

Klee, V.

[1] Extremal structure of convex sets, *Arch. Math.* 8 (1957), 234–240.

[2] Extremal structure of convex sets, II, *Math. Z.* 69 (1958), 90–104.

Klee, V. & Minty, G.J.

[1] How Good is the Simplex Algorithm? Pages 159–175 in Shisha, O. (ed.), *Inequalities, III,* Academic Press, 1972.

Koopmans, T.C., ed.

[1] *Activity Analysis of Production and Allocation.* John Wiley & Sons, 1951.

Kuhn, H.W.

[1] Solvability and Consistency for Linear Equations and Inequalities, *Amer. Math. Monthly* 63 (1956), 217–232.

Minkowski, H.

[1] *Geometrie der Zahlen*. Teubner, Leipzig, 1910.

[2] *Gesammelte Abhandlungen von Hermann Minkowski, Vol. 1, 2.* Teubner, Leibzig, 1911

Motzkin, T.

[1] *Beiträge zur Theorie der linearen Ungleichungen.* Azviel, Jerusalem, 1936.

Rockafellar, R.T.

[1] *Convex Analysis.* Princeton Univ. Press., 1970

Steinitz, E.

[1] Bedingt konvergente Reihen und konvexe Systeme, I, II, III, *J. Reine Angew. Math.* 143 (1913), 128–175; 144 (1914), 1–40; 146 (1916), 1–52.

Stiemke, E.

[1] Über positive Lösungen homogener linearer Gleichungen, *Math. Ann.* 76 (1915), 340–342.

Stoer, J. & Witzgall, C.

[1] *Convexity and Optimization in Finite Dimensions I.* Springer-Verlag, 1970.

Tucker, A.W.

[1] Dual Systems of Homogeneous Linear Relations. Pages 3–18 in Kuhn, H.W. & Tucker, A.W. (eds.), *Linear Inequalities and Related Systems*, Princeton Univ. Press, 1956.

Webster, R.

[1] *Convexity.* Oxford University Press, 1954.

Weyl, H.

[1] Elementare Theorie der konvexen Polyeder, *Comment. Math. Helv.* 7 (1935), 290–306.

Answers and solutions to the exercises

Chapter 2

2.2 a) $\{x \in \mathbf{R}^2 \mid 0 \le x_1 + x_2 \le 1,\ x_1, x_2 \ge 0\}$
b) $\{x \in \mathbf{R}^2 \mid \|x\| \le 1\}$
c) $\mathbf{R}^2_{++} \cup \{(0,0)\}$

2.3 E.g. $\{(0,1)\} \cup (\mathbf{R} \times \{0\})$ in $\mathbf{R}^2$.

2.4 $\{x \in \mathbf{R}^3_{++} \mid x_3^2 \le x_1x_2\}$

2.5 Use the triangle inequality

$$\big(\textstyle\sum_1^n (x_j + y_j)^2\big)^{1/2} \le \big(\sum_1^n x_j^2\big)^{1/2} + \big(\sum_1^n y_j^2\big)^{1/2}$$

to show that the set is closed under addition of vectors. Or use the perspective map; see example 2.3.4.

2.6 Follows from the fact that $-\mathbf{e}_k$ is a conic combination of the vectors $\mathbf{e}_0, \mathbf{e}_1, \dots, \mathbf{e}_n$.

2.7 Let X be the halfspace $\{x \in \mathbf{R}^n \mid \langle c, x\rangle \ge 0\}$. Each vector $x \in X$ is a conic combination of c and the vector $y = x - \langle c, x\rangle \|c\|^{-2} c$, and y lies in the $(n-1)$-dimensional subspace $Y = \{x \in \mathbf{R}^n \mid \langle c, x\rangle = 0\}$, which is generated by n vectors as a cone according to the previous exercise. Hence, x is a conic combination of these n vectors and c.

2.8 The intersection between the cone X and the unit circle is a closed circular arc with endpoints x and y, say. The length of the arc is either less than π, equal to π, or equal to 2π. The cone X is proper and generated by the two vectors x and y in the first case. It is equal to a halfspace in the second case and equal to $\mathbf{R}^2$ in the third case, and it is generated by three vectors in both these cases.

2.9 Use exercise 2.8.

2.10 a) $\operatorname{recc} X = \{x \in \mathbf{R}^2 \mid x_1 \ge x_2 \ge 0\}$, $\operatorname{lin} X = \{(0,0)\}$
b) $\operatorname{recc} X = \operatorname{lin} X = \{(0,0)\}$

2.10 c) $\operatorname{recc} X = \{x \in \mathbf{R}^3 \mid 2x_1 + x_2 + x_3 \le 0,\ x_1 + 2x_2 + x_3 \le 0\}$,
$\operatorname{lin} X = \{(t, t, -3t) \mid t \in \mathbf{R}\}$

d) $\operatorname{recc} X = \{x \in \mathbf{R}^3 \mid x_1 \ge |x_2|\}$,
$\operatorname{lin} X = \{x \in \mathbf{R}^3 \mid x_1 = x_2 = 0\}$.

2.12 b) (i) $c(X) = \{x \in \mathbf{R}^2 \mid 0 \le \frac{1}{3}x_1 \le x_2 \le \frac{1}{2}x_1\} = \operatorname{cl}(c(X))$

(ii) $c(X) = \{x \in \mathbf{R}^2 \mid 0 < x_2 \le \frac{1}{2}x_1\} \cup \{(0,0)\}$,
$\operatorname{cl}(c(X) = \{x \in \mathbf{R}^2 \mid 0 \le x_2 \le \frac{1}{2}x_1\}$,

(iii) $c(X) = \{x \in \mathbf{R}^3 \mid x_1x_3 \ge x_2^2,\ x_3 > 0\} \cup \{(0,0,0)\}$,
$\operatorname{cl}(c(X)) = c(X) \cup \{(x_1, 0, 0) \mid x_1 \ge 0\}$.

c) $c(X) = \{(x, x_{n+1}) \in \mathbf{R}^n \times \mathbf{R} \mid \|x\| \le x_{n+1}\}$.

2.14 Let $z_n = x_n + y_n$, $n = 1, 2, \dots$ be a convergent sequence of points in $X + Y$ with $x_n \in X$ and $y_n \in Y$ for all n and limit z_0. The sequence $(y_n)_1^\infty$ ha a convergent subsequence $(y_{n_k})_{k=1}^\infty$ with limit $y_0 \in Y$, since Y is compact. The corresponding subsequence $(z_{n_k} - y_{n_k})_{k=1}^\infty$ of points in X converges to $z_0 - y_0$, and the limit point belongs to X since X is a closed set. Hence, $z_0 = (z_0 - y_0) + y_0$ lies in $X + Y$, and this means that $X + Y$ is a closed set.

Chapter 3

3.1 E.g. $\{x \in \mathbf{R}^2 \mid x_2 \le 0\}$ and $\{x \in \mathbf{R}^2 \mid x_2 \ge e^{x_1}\}$.

3.2 Follows from Theorem 3.1.3 for closed sets and from Theorem 3.1.5 for open sets.

3.4 a) $\mathbf{R}_+ \times \mathbf{R}$ b) $\{0\} \times \mathbf{R}$ c) $\{0\} \times \mathbf{R}_+$ d) $\mathbf{R}_+ \times \mathbf{R}_+$
e) $\{x \in \mathbf{R}^2 \mid x_2 \ge x_1 \ge 0\}$

3.6 a) $X = X^{++} = \{x \in \mathbf{R}^2 \mid x_1 + x_2 \ge 0,\ x_2 \ge 0\}$,
$X^+ = \{x \in \mathbf{R}^2 \mid x_2 \ge x_1 \ge 0\}$

b) $X = X^{++} = \mathbf{R}^2$, $X^+ = \{(0,0)\}$

c) $X = \mathbf{R}^2_{++} \cup \{(0,0)\}$, $X^+ = X^{++} = \mathbf{R}^2_+$

3.7 (i) $\Rightarrow$ (iii): Since $-a_j \notin \operatorname{con} A$, there is, for each j, a vector c_j such that $-\langle c_j, a_j\rangle < 0$ and $\langle c_j, x\rangle \ge 0$ for all $x \in \operatorname{con} A$, which implies that $\langle c_j, a_j\rangle > 0$ and $\langle c_j, a_k\rangle \ge 0$ for all k. It follows that $c = c_1 + c_2 + \cdots + c_m$ works.

(iii) $\Rightarrow$ (ii): Suppose that $\langle c, a_j\rangle > 0$ for all j. Then $\sum_1^m \lambda_j a_j = 0$ implies $0 = \langle c, 0\rangle = \sum_1^m \lambda_j \langle c, a_j\rangle$, so if $\lambda_j \ge 0$ for all j then $\lambda_j \langle c, a_j\rangle = 0$ for all j, with the conclusion that $\lambda_j = 0$ for all j.

(ii) $\Rightarrow$ (i): If there is a vector x such that $x = \sum_1^m \lambda_j a_j$ and $-x = \sum_1^m \mu_j a_j$ with nonnegative scalars λ_j, μ_j, then by addition we obtan the equality $\sum_1^m(\lambda_j + \mu_j)a_j = 0$ with the conclusions $\lambda_j + \mu_j = 0$, $\lambda_j = \mu_j = 0$ and $x = 0$.

3.8 No solution.

3.10 Solvable for $\alpha \leq -2$, $-1 < \alpha < 1$ and $\alpha > 1$.

3.11 a) The systems (S) and (S*) are equivalent to the systems

$$\begin{cases} Ax \geq 0 \\ -Ax \geq 0 \\ Ex \geq 0 \\ \mathbf{1}^T x > 0 \end{cases} \quad \text{and} \quad \begin{cases} A^T(y' - y'') + Ez + \mathbf{1}t = 0 \\ y', y'', z \geq 0, \; t > 0, \end{cases}$$

respectively (with y corresponding to $y'' - y'$). The assertion therefore follows from Theorem 3.3.7.

b) The systems (S) and (S*) are equivalent to the systems

$$\begin{cases} Ax \geq 0 \\ -Ax \geq 0 \\ Ex > 0 \end{cases} \quad \text{and} \quad \begin{cases} A^T(y' - y'') + Ez = 0 \\ y', y'', z \geq 0, \; z \neq 0, \end{cases}$$

respectively. Now apply Theorem 3.3.7.

3.12 By Theorem 3.3.7, the system is solvable if and only if the dual system

$$\begin{cases} A^T(y'-y'')+z+u=0 \\ A(w+u)=0 \\ y',y'',z,w,u\geq 0,\ u\neq 0 \end{cases}$$

has no solution. It follows from the two equations of the dual system that:

$$\begin{aligned} 0 &= -(w+u)^T A^T = -(w+u)^T A^T(y'-y'') = (w+u)^T(z+u) \\ &= w^Tz + w^Tu + u^Tz + u^Tu, \end{aligned}$$

and all the four terms in the last sum are nonnegative. We conclude that $u^Tu=0$, and hence $u=0$. So the dual system has no solution.

Chapter 4

4.1 a) $\operatorname{ext} X = \{(1,0),(0,1)\}$ b) $\operatorname{ext} X = \{(0,0),(1,0),(0,1),(\frac{1}{2},1)\}$
c) $\operatorname{ext} X = \{(0,0,1),(0,0,-1)\} \cup \{(x_1,x_2,0) \mid (x_1-1)^2+x_2^2=1\} \setminus \{(0,0,0)\}$

4.2 Suppose $x \in \operatorname{cvx} A \setminus A$; then $x = \lambda a + (1-\lambda)y$ where $a \in A$, $y \in \operatorname{cvx} A$ and $0<\lambda<1$. It follows that $x \notin \operatorname{ext}(\operatorname{cvx} A)$.

4.3 We have $\operatorname{ext} X \subseteq A$, according to the previous exercise. Suppose that $a \in A \setminus \operatorname{ext} X$. Then $a = \lambda x_1 + (1-\lambda)x_2$, where $x_1,x_2 \in X$, $x_1 \neq x_2$ and $0<\lambda<1$. We have $x_i = \mu_i a + (1-\mu_i)y_i$, where $0 \leq \mu_i < 1$ and $y_i \in \operatorname{cvx}(A \setminus \{a\})$. It now follows from the equality

$$a = (1-\lambda\mu_1-(1-\lambda)\mu_2)^{-1}(\lambda(1-\mu_1)y_1+(1-\lambda)(1-\mu_2)y_2),$$

that a lies in $\operatorname{cvx}(A \setminus \{a\})$.Therefore, $\operatorname{cvx}(A\setminus\{a\}) = \operatorname{cvx} A = X$, which contradicts the minimality of A. Hence, $\operatorname{ext} X = A$.

4.4 The set $X \setminus \{x_0\}$ is convex if and only if $]a,b[\subseteq X \setminus \{x_0\}$ for all $a,b \in X \setminus \{x_0\}$, i.e. if and only if $x_0 \notin]a,b[$ for all $a,b \in X\setminus\{x_0\}$, i.e. if and only if $x_0 \in \operatorname{ext} X$.

4.5 E.g. the set in exercise 4.1 c).

4.6 a) Follows directly from Theorem 4.1.3.
b) The extreme point $(1,0)$ of $\{x \in \mathbf{R}^2 \mid x_2 \leq \sqrt{1-x_1^2},\ |x_1| \leq 1\}$ is not exposed.

4.7 b) A zero-dimensional general face is an extreme point, and a zero-dimensional exposed face is an exposed point. Hence, exercise 4.6 b) contains an example of a general face which is not an exposed face.

c) Suppose that $a, b \in X$ and that the open line segment $]a, b[$ intersects F'. Since $F' \subseteq F$, the same line segment also intersects F, so it follows that $a, b \in F$. But since F' is a general face of F, it follows that $a, b \in F'$. So F' is indeed a general face of X.
The set X in exercise 4.6 b) has $F = \{1\}\times]-\infty, 0]$ as an exposed face, and $F' = \{(1, 0)\}$ is an exposed face of F but not of X.

d) Fix a point $x_0 \in F \cap \operatorname{rint} C$. To each $x \in C$ there is a point $y \in C$ such that x_0 lies on the open line segment $]x, y[$, and it now follows from the definition of a general face that $x \in F$.

e) Use the result in d) on the set $C = X \cap \operatorname{cl} F$. Since $\operatorname{rint} C$ contains $\operatorname{rint} F$ as a subset, $F \cap \operatorname{rint} C \neq \emptyset$, so it follows that $C \subseteq F$. The converse inclusion is of course trivial.

f) Use the result in d) with $F = F_1$ och $C = F_2$, which gives us the inclusion $F_2 \subseteq F_1$. The converse inclusion is obtained analogously.

g) If F is a general face and $F \cap \operatorname{rint} X \neq \emptyset$, then $X \subseteq F$ by d) above. For faces $F \neq X$ we therefore have $F \cap \operatorname{rint} X = \emptyset$, which means that $F \subseteq \operatorname{rbdry} X$.

Chapter 5

5.1 a) $(-\frac{2}{3}, \frac{4}{3})$, and $(4, -1)$ b) $(-\frac{2}{3}, \frac{4}{3})$, $(4, -1)$, and$(-3, -1)$
c) $(0, 0, 0)$, $(2, 0, 0)$, $(0, 2, 0)$, $(0, 0, 4)$, and $(\frac{4}{3}, \frac{4}{3}, 0)$
d) $(0, 4, 0, 0)$, $(0, \frac{5}{2}, 0, 0)$, $(\frac{3}{2}, \frac{5}{2}, 0, 0)$, $(0, 1, 1, 0)$, and $(0, \frac{5}{2}, 0, \frac{3}{2})$

5.2 The extreme rays are generated by $(-2, 4, 3)$, $(1, 1, 0)$, $(4, -1, 1)$, and $(1, 0, 0)$.

5.3 $C = \begin{bmatrix} 1 & -2 & 1 \\ -1 & 2 & 3 \\ -3 & 2 & 5 \end{bmatrix}$

5.4 a) $A = \{(1, 0), (0, 1)\}$, $B = \{(-\frac{2}{3}, \frac{4}{3}), (4, -1)\}$
b) $A = \emptyset$, $B = \{(-\frac{2}{3}, \frac{4}{3}), (4, -1), (-3, -1)\}$
c) $A = \{(1, 1, -3), (-1, -1, 3), (4, -7, -1), (-7, 4, -1)\}$,
$B = \{(0, 0, 0), (2, 0, 0), (0, 2, 0), (0, 0, 4), (\frac{4}{3}, \frac{4}{3}, 0)\}$
d) $A = \emptyset$,
$B = \{(0, 4, 0, 0), (0, \frac{5}{2}, 0, 0), (\frac{3}{2}, \frac{5}{2}, 0, 0), (0, 1, 1, 0), (0, \frac{5}{2}, 0, \frac{3}{2})\}$.

5.5 The inclusion $X = \operatorname{cvx} A + \operatorname{con} B \subseteq \operatorname{con} A + \operatorname{con} B = \operatorname{con}(A \cup B)$ implies that $\operatorname{con} X \subseteq \operatorname{con}(A \cup B)$. Obviously, $A \subseteq \operatorname{cvx} A \subseteq X$. Since $\operatorname{cvx} A$ is a compact set, $\operatorname{recc} X = \operatorname{con} B$, so using the assumption $0 \in X$, we obtain the inclusion $B \subseteq \operatorname{con} B \subseteq X$. Thus, $A \cup B \subseteq X$, and it follows that $\operatorname{con}(A \cup B) \subseteq \operatorname{con} X$.

Chapter 6

6.1 E.g. $f_1(x) = x - |x|$ and $f_2(x) = -x - |x|$.

6.3 $a \geq 5$ and $a > 5$, respectively.

6.4 Use the result of exercise 2.1.

6.5 Follows from $f(x) = \max(x_{i_1} + x_{i_2} + \cdots + x_{i_k})$, where the maximum is taken over all subsets $\{i_1, i_2, \ldots, i_k\}$ of $\{1, 2, \ldots, n\}$ consisting of k elements.

6.6 The inequality is trivial if $x_1 + x_2 + \cdots + x_n = 0$, and it is obtained by adding the n inequalities

$$f(x_i) \leq \frac{x_i}{x_1 + \cdots + x_n} f(x_1 + \cdots + x_n) + \big(1 - \frac{x_i}{x_1 + \cdots + x_n}\big) f(0)$$

if $x_1 + \cdots + x_n > 0$.

6.7 Choose

$$c = \frac{f(x_2) - f(x_1)}{\|x_2 - x_1\|^2}\,(x_1 - x_2),$$

to obtain $f(x_1) + \langle c, x_1\rangle = f(x_2) + \langle c, x_2\rangle$. By quasiconvexity,

$$f(\lambda x_1 + (1-\lambda)x_2) + \langle c, \lambda x_1 + (1-\lambda)x_2\rangle \le f(x_1) + \langle c, x_1\rangle,$$

which simplifies to

$$f(\lambda x_1 + (1-\lambda)x_2) \le \lambda f(x_1) + (1-\lambda) f(x_2).$$

6.8 Let $f\colon \mathbf{R}^n \times \mathbf{R} \to \overline{\mathbf{R}}$ be the function defined by

$$f(x,t) = \begin{cases} t & \text{if } (x,t) \in C \\ +\infty & \text{if } (x,t) \notin C. \end{cases}$$

Then $\inf\{t \in \mathbf{R} \mid (x,t) \in C\} = \inf\{f(x,t) \mid t \in \mathbf{R}\}$, and Theorem 6.2.6 now follows from Corollary 6.2.7.

6.9 Choose, given $x, y \in X$, sequences $(x_k)_1^\infty$, $(y_k)_1^\infty$ of points $x_k, y_k \in \operatorname{int} X$ such that $x_k \to x$ and $y_k \to y$ as $k \to \infty$. Since the points $\lambda x_k + (1-\lambda) y_k$ belong to $\operatorname{int} X$,

$$f(\lambda x_k + (1-\lambda)y_k) \le \lambda f(x_k) + (1-\lambda) f(y_k),$$

and since f is continuous on X, we now obtain the desired inequality $f(\lambda x + (1-\lambda)y) \le \lambda f(x) + (1-\lambda) f(y)$ by passing to the limit.

6.10 Let $m = \inf\{f(x) \mid x \in \operatorname{rint}(\operatorname{dom} f)\}$ and fix a relative interior point x_0 of $\operatorname{dom} f$. If $x \in \operatorname{dom} f$ is arbitrary and $0 < \lambda < 1$, then $\lambda x + (1-\lambda)x_0$ is a relative interior point of $\operatorname{dom} f$, and it follows that

$$m \le f(\lambda x + (1-\lambda)x_0) \le \lambda f(x) + (1-\lambda) f(x_0).$$

The inequality $f(x) \ge m$ now follows by letting $\lambda \to 1$.

6.11 Minimum 8 at $x = (\frac{1}{8}, 2)$.

6.12 a) $\|x\|_p$ b) $\max(x_1, 0)$.

Chapter 7

7.2 Yes.

7.5 Let J be a subinterval of I. If $f'_+(x) \ge 0$ for all $x \in J$, then

$$f(y) - f(x) \ge f'_+(x)(y-x) \ge 0$$

for all $y > x$ in the interval J, i.e. f is increasing on J. If instead $f'_+(x) \le 0$ for all $x \in J$, then $f(y) - f(x) \ge f'_+(x)(y-x) \ge 0$ for all $y < x$, i.e. f is decreasing on J.

Since the right derivative $f'_+(x)$ is increasing on I, there are three different cases to consider. Either $f'_+(x) \ge 0$ for all $x \in I$, and f is then increasing on I, or $f'_+(x) \leq 0$ for all $x \in I$, and f is then decreasing on I, or there is a point $c \in I$ such that $f'_+(x) \leq 0$ to the left of c and $f'_+(x) > 0$ to the right of c, and f is in this case decreasing to the left of c and increasing to the right of c.

7.6 a) The existence of the limits is a consequence of the results of the previous exercise.

b) Consider the epigraph of the extended function.

7.7 Follows directly from exercise 7.6 b).

7.8 Suppose that $f \in \mathcal{F}$. Let $x_0 \in \mathbf{R}^n$ be an arbitrary point, and consider the function $g(x) = f(x) - \langle f'(x_0), x - x_0\rangle$. The function g belongs to $\mathcal{F}$ and $g'(x_0) = 0$. It follows that $g(x) \geq g(x_0)$ for all x, which means that $f(x) \geq f(x_0) + \langle f'(x_0), x - x_0\rangle$ for all x. Hence, f is convex by Theorem 7.2.1.

7.9 $\phi(t) = f(x + tv) = f(x) + t\langle f'(x), v\rangle$ for $v \in V_f$ by Theorem 6.7.1. Differentiate two times to obtain $D^2 f(x)[v, v] = \phi''(0) = 0$, with the conlusion that $f''(x)v = 0$.

7.13 By combining Theorem 7.3.1 (i) with x replaced by $\hat{x}$ and $v = x-\hat{x}$ with the Cauchy-Schwarz inequality, we obtain the inequality $\mu\|x - \hat{x}\|^2 \leq \langle f'(x), x - \hat{x}\rangle \leq \|f'(x)\|\|x - \hat{x}\|$.

Chapter 8

8.1 Suppose that f is μ-strongly convex, where $\mu > 0$, and let c be a subgradient at 0 of the convex function $g(x) = f(x) - \frac{1}{2}\mu\|x\|^2$. Then $f(x) \geq f(0) + \langle c, x\rangle + \frac{1}{2}\mu\|x\|^2$ for all x, and the right-hand side tends to ∞ as $\|x\| \to \infty$. Alternatively, one could use Theorem 8.1.6.

8.2 The line segment $[-\mathbf{e}_1, \mathbf{e}_2]$, where $\mathbf{e}_1 = (1, 0)$ and $\mathbf{e}_2 = (0, 1)$.

8.3 a) $\overline{B}_2(0; 1) = \{x \mid \|x\|_2 \leq 1\}$ b) $\overline{B}_1(0; 1) = \{x \mid \|x\|_1 \leq 1\}$

c) $\overline{B}_\infty(0; 1) = \{x \mid \|x\|_\infty \leq 1\}$.

8.4 a) $\operatorname{dom} f^* = \{a\},\ f^*(a) = b$

b) $\operatorname{dom} f^* = \{x \mid x < 0\},\ f^*(x) = -1 - \ln(-x)$

c) $\operatorname{dom} f^* = \mathbf{R}_+,\ f^*(x) = x \ln x - x,\ f^*(0) = 0$

d) $\operatorname{dom} f^* = \mathbf{R},\ f^*(x) = e^{x-1}$

e) $\operatorname{dom} f^* = \mathbf{R}_-,\ f^*(x) = -2\sqrt{-x}$.

Index

ENDNOTES

Chapter 1

†The intersection of an empty family of sets is usually defined as the entire space, and using this convention the polyhedron $\mathbf{R}^n$ can also be viewed as an intersection of halfspaces.

‡The terminology is not universal. A proper cone is usually called a *salient cone*, while the term *proper cone* is sometimes reserved for cones that are closed, have a nonempty interior and do not contain any lines through the origin.

Chapter 3

†The second condition is usually not included in the definition of separation, but we have included it in order to force a hyperplane H that separates two subsets of a hyperplane H' to be different from H'.

Chapter 4

†For if $X = \{x_0\}$, then $\operatorname{rint} X = \{x_0\}$, $\operatorname{rbdry} X = \emptyset$ and $\operatorname{ext} X = \{x_0\}$.

‡There is an alternative and more general definition of the face concept, see exercise 4.7. Our proper faces are called *exposed faces* by Rockafellar in his standard treatise *Convex Analysis*. Every exposed face is also a face according to the alternative definition, but the two definitions are not equivalent, because there are convex sets with faces that are not exposed.

To see Part II, download:
Linear and Convex Optimization
Convexity and Optimization – Part II