Grundlehren Text Editions

Jean-Baptiste Hiriart-Urruty
Claude Lemaréchal

Fundamentals
of Convex Analysis

With 66 Figures

Springer

Jean-Baptiste Hiriart-Urruty

Département de Mathématiques
Université Paul Sabatier
118, route de Narbonne
31062 Toulouse
France
e-mail: jbhu@cict.fr

Claude Lemaréchal

INRIA, Rhône Alpes
ZIRST
655, avenue de l'Europe
38330 Montbonnot
France
e-mail: Claude.Lemarechal@inrialpes.fr

Corrected Second Printing 2004

Preface

This book is an abridged version of our two-volume opus *Convex Analysis and Minimization Algorithms* [18], about which we have received very positive feedback from users, readers, lecturers ever since it was published – by Springer-Verlag in 1993. Its pedagogical qualities were particularly appreciated, in the combination with a rather advanced technical material.

Now [18] has a dual but clearly defined nature:
– an introduction to the basic concepts in convex analysis,
– a study of convex minimization problems (with an emphasis on numerical algorithms),

and insists on their mutual interpenetration. It is our feeling that the above basic introduction is much needed in the scientific community. This is the motivation for the present edition, our intention being to create a tool useful to teach convex analysis. We have thus extracted from [18] its "backbone" devoted to convex analysis, namely Chaps III–VI and X. Apart from some local improvements, the present text is mostly a copy of the corresponding chapters. The main difference is that we have deleted material deemed too advanced for an introduction, or too closely attached to numerical algorithms.

Further, we have included exercises, whose degree of difficulty is suggested by 0, 1 or 2 stars *. Finally, the index has been considerably enriched.

Just as in [18], each chapter is presented as a "lesson", in the sense of our old masters, treating of a given subject in its entirety. After an introduction presenting or recalling elementary material, there are five such lessons:
– A Convex sets (corresponding to Chap. III in [18]),
– B Convex functions (Chap. IV in [18]),
– C Sublinearity and support functions (Chap. V),
– D Subdifferentials in the finite-valued case (VI),
– E Conjugacy (X).

Thus, we do not go beyond conjugacy. In particular, subdifferentiability of extended-valued functions is intentionally left aside. This allows a lighter book, easier to master and to go through. The same reason led us to skip duality which, besides, is more related to optimization. Readers interested by these topics can always read the relevant chapters in [18] (namely Chaps XI and XII).

During the French Revolution, the writer of a bill on public instruction complained: "Le défaut ou la disette de bons ouvrages élémentaires a été, jusqu'à présent, un des plus grands obstacles qui s'opposaient au perfectionnement de l'instruction. La raison de cette disette, c'est que jusqu'à présent les savants d'un mérite éminent ont, presque toujours, *préféré la gloire d'élever l'édifice de la science à la peine d'en éclairer l'entrée.*[1]" Our main motivation here is precisely to "light the entrance" of the monument Convex Analysis. This is therefore not a reference book, to be kept on the shelf by experts who already know the building and can find their way through it; it is far more a book for the purpose of learning and teaching. We call above all on the intuition of the reader, and our approach is very gradual. Nevertheless, we keep constantly in mind the suggestion of A. Einstein: "Everything should be made as simple as possible, but not simpler". Indeed, the content is by no means elementary, and will be hard for a reader not possessing a firm mastery of basic mathematical skill.

We could not completely avoid cross-references between the various chapters; but for many of them, the motivation is to suggest an intellectual link between apparently independent concepts, rather than a technical need for previous results. More than a tree, our approach evokes a spiral, made up of loosely interrelated elements.

Many sections are set in smaller characters. They are by no means reserved to advanced material; rather, they are there to help the reader with illustrative examples and side remarks, that help to understand a delicate point, or prepare some material to come in a subsequent chapter. Roughly speaking, sections in smaller characters can be compared to footnotes, used to avoid interrupting the flow of the development; it can be helpful to skip them during a deeper reading, with pencil and paper. They can often be considered as additional informal exercises, useful to keep the reader alert.

The numbering of sections restarts at 1 in each chapter, and chapter numbers are dropped in a reference to an equation or result from within the same chapter.

Toulouse and Grenoble,
March 2001   J.-B. Hiriart-Urruty, C. LemaréchaL

---

[1] "The lack or scarcity of good, elementary books has been, until now, one of the greatest obstacles in the way of better instruction. The reason for this scarcity is that, until now, scholars of great merit have almost always preferred the glory of constructing the monument of science over the effort of lighting its entrance." D. Guedj: *La Révolution des Savants*, Découvertes, Gallimard Sciences (1988) 130 – 131.

Contents

0. Introduction: Notation, Elementary Results

We start this chapter by listing some basic concepts, which are or should be well-known – but it is good sometimes to return to basics. This gives us the opportunity of making precise the system of notation used in this book. For example, some readers may have forgotten that "i.e." means *id est*, the literal translation of "that is (to say)". If we get closer to mathematics, $S \backslash \{x\}$ denotes the set obtained by depriving a set S of a point $x \in S$. We also mention that, if f is a function, $\overset{-1}{f}(y)$ is the *inverse image* of y, i.e. the set of all points x such that $f(x) = y$. When f is invertible, this set is the singleton $\{f^{-1}(y)\}$.

After these basic recalls, we prove some results on convex functions of one real variable. They are just as basic, but are easily established and will be of some use in this book.

1 Some Facts About Lower and Upper Bounds

1.1 In the *totally ordered* set $\mathbb{R}$, inf E and sup E are respectively the greatest lower bound – the *infimum* – and least upper bound – the *supremum* – of a nonempty subset E, when they exist (as real numbers). Then, they may or may not belong to E; when they do, a more accurate notation is min E and max E. Whenever the relevant infima exist, the following relations are clear enough:

$$\left. \begin{array}{l} \inf(E \cup F) = \min\{\inf E, \inf F\}, \\ F \subset E \implies \inf F \geqslant \inf E, \\ \inf(E \cap F) \geqslant \max\{\inf E, \inf F\}. \end{array} \right\} \tag{1.1}$$

If E is characterized by a certain property P, we use the notation

$$E = \{r \in \mathbb{R} : r \text{ satisfies } P\}.$$

Defining (in $\mathbb{R}$ considered as a real *vector space*) the standard operations on nonempty sets

$$E + F := \{r = e + f : e \in E, f \in F\},$$
$$tE := \{tr : r \in E\} \quad \text{for } t \in \mathbb{R}$$

(the sign ":=" means "equals by definition"), it is also clear that

$$\left.\begin{array}{r}
\inf (E + F) = \inf E + \inf F\,,\\
\inf tE = t \inf E \quad \text{if } t > 0\,,\\
\inf (-E) = -\sup E\,,
\end{array}\right\} \qquad (1.2)$$

whenever the relevant extrema exist.

The word *positive* means "> 0", and *nonpositive* therefore means "$\leqslant 0$"; same conventions with negative and nonnegative. The set of nonnegative numbers is denoted by $\mathbb{R}^+$ and, generally speaking, a substar deprives a set of the point 0. Thus, for example,

$$\mathbb{N}_* = \{1, 2, \ldots\} \quad \text{and} \quad \mathbb{R}_*^+ = \{t \in \mathbb{R} : t > 0\}\,.$$

Squared brackets are used to denote the intervals of $\mathbb{R}$: for example,

$$\mathbb{R} \supset \,]a, b] = \{t \in \mathbb{R} : a < t \leqslant b\}\,.$$

The symbol "$\downarrow$" means convergence from the right, *the limit being excluded*; thus, $t \downarrow 0$ means $t \to 0$ in $\mathbb{R}_*^+$. The words "increasing" and "decreasing" are taken in a broad sense: a sequence (t_k) is increasing when $k > k' \Rightarrow t_k \geqslant t_{k'}$. We use the notation (t_k), or $(t_k)_k$, or $(t_k)_{k \in \mathbb{N}_*}$ for a sequence of elements $t_1, t_2, \ldots$

1.2 Now, to denote a real-valued function f defined on a nonempty set X, we write

$$X \ni x \mapsto f(x) \in \mathbb{R}.$$

The *sublevel-set* of f at level $r \in \mathbb{R}$ is defined by

$$\mathrm{S}_r(f) := \{x \in X : f(x) \leqslant r\}\,.$$

If two functions f and g from X to $\mathbb{R}$ satisfy

$$f(x) \leqslant g(x) \quad \text{for all } x \in X\,,$$

we say that f *minorizes* g (on X), or that g majorizes f.

Computing the number

$$\inf \{f(x) : x \in X\} =: \bar{f} \qquad (1.3)$$

represents a minimization problem posed in X: namely that of finding a so-called *minimizing sequence*, i.e. $(x_k) \subset X$ such that $f(x_k) \to \bar{f}$ when $k \to +\infty$ (note that no structure is assumed on X). In other words, $\bar{f}$ is the largest lower bound $\inf f(X)$ of the subset $f(X) \subset \mathbb{R}$, and will often be called the infimal value, or more simply the *infimum* of f on X. Another notation for (1.3) is $\inf_{x \in X} f(x)$, or also $\inf_X f$. The function f is usually called the objective function, or also infimand. We can also meet supremands, minimands, etc.

From the relations (1.1), (1.2), we deduce (hereafter, $\bar{f}_i$ denotes the infimum of f over X_i for $i = 1, 2$):

$$\inf \{f(x) \ : \ x \in X_1 \cup X_2\} = \min \{\bar{f}_1, \bar{f}_2\},$$
$$X_1 \subset X_2 \implies \bar{f}_1 \geqslant \bar{f}_2,$$
$$\inf \{f(x) \ : \ x \in X_1 \cap X_2\} \geqslant \max \{\bar{f}_1, \bar{f}_2\},$$
$$\inf \{f(x_1) + f(x_2) \ : \ x_1 \in X_1 \text{ and } x_2 \in X_2\} = \bar{f}_1 + \bar{f}_2, \qquad (1.4)$$
$$\inf \{tf(x) \ : \ x \in X\} = t\bar{f}, \quad \text{for } t \geqslant 0,$$
$$\inf \{-f(x) \ : \ x \in X\} = -\sup \{f(x) \ : \ x \in X\},$$

whenever the relevant extrema exist. The last relation is used very often.

The attention of the reader is drawn to (1.4), perhaps the only non-totally trivial among the above relations. Calling $E_1 := f(X_1)$ and $E_2 := f(X_2)$ the *images* of X_1 and X_2 under f, (1.4) represents the sum of the infima $\inf E_1$ and $\inf E_2$. There could just as well be two different infimands, i.e. (1.4) could be written more suggestively

$$\inf \{f(x_1) + g(x_2) \ : \ x_1 \in X_1 \text{ and } x_2 \in X_2\} = \bar{f}_1 + \bar{g}_2$$

(g being another real-valued function). This last relation must not be confused with

$$\inf \{f(x) + g(x) \ : \ x \in X\} \geqslant \bar{f} + \bar{g};$$

here, in the language of (1.4), $X_1 = X_2 = X$, but only the image by f of the diagonal of $X \times X$ is considered.

Another relation requiring some attention is the *decoupling*, or transitivity, of infima: if g sends the Cartesian product $X \times Y$ to $\mathbb{R}$, then

$$\inf \{g(x,y) \ : \ x \in X \text{ and } y \in Y\} =$$
$$= \inf_{x \in X}[\inf_{y \in Y} g(x,y)] = \inf_{y \in Y}[\inf_{x \in X} g(x,y)]. \qquad (1.5)$$

1.3 An *optimal solution* of (1.3) is an $\bar{x} \in X$ such that

$$f(\bar{x}) = \bar{f} \leqslant f(x) \quad \text{for all } x \in X;$$

such an $\bar{x}$ is often called a minimizer, a minimum point, or more simply a *minimum* of f on X. We will also speak of *global minimum*. To say that there exists a minimum is to say that the inf in (1.3) is a min; the infimum $\bar{f} = f(\bar{x})$ can then be called the minimal value. The notation

$$\min \{f(x) \ : \ x \in X\}$$

is the same as (1.3), and says that there does exist a solution; we stress the fact that this notation – as well as (1.3) – represents at the same time a *number* and a *problem to solve*. It is sometimes convenient to denote by

$$\text{Argmin} \{f(x) \ : \ x \in X\}$$

the set of optimal solutions of (1.3), and to use "argmin" if the solution is unique.

It is worth mentioning that the decoupling property (1.5) has a translation in terms of Argmin's. More precisely, the following properties are easy to see:

– If $(\bar{x}, \bar{y})$ minimizes g over $X \times Y$, then $\bar{y}$ minimizes $g(\bar{x}, \cdot)$ over Y and $\bar{x}$ minimizes over X the function

$$\varphi(x) := \inf \{g(x,y) \; : \; y \in Y\}.$$

– Conversely, if $\bar{x}$ minimizes φ over X and if $\bar{y}$ minimizes $g(\bar{x}, \cdot)$ over Y, then $(\bar{x}, \bar{y})$ minimizes g over $X \times Y$.

Needless to say, symmetric properties are established, interchanging the roles of x and y.

1.4 In our context, X is equipped with a topology; actually X is a subset of some finite-dimensional real vector space, call it $\mathbb{R}^n$; the topology is then that induced by a norm. The *interior* and *closure* of X are denoted by int X and cl X respectively; its boundary is bd X.

The concept of limit is assumed familiar. We recall that the *limes inferior* (in the ordered set $\mathbb{R}$) is the smallest cluster point.

Remark 1.1 The standard terminology is *lower limit* ("abbreviated" as lim inf!) This terminology is unfortunate, however: a limit must be a well-defined unique element; otherwise, expressions such as "$f(x)$ has a limit" are ambiguous. □

Thus, to say that $\ell = \liminf_{x \to x^*} f(x)$, with $x^* \in \text{cl}\, X$, means: for all $\varepsilon > 0$,

there is a neighborhood $N(x^*)$ such that $f(x) \geqslant \ell - \varepsilon$ for all $x \in N(x^*)$,

and

in any neighborhood $N(x^*)$, there is $x \in N(x^*)$ such that $f(x) \leqslant \ell + \varepsilon$;

in particular, if $x^* \in X$, we certainly have $\ell \leqslant f(x^*)$.

Let $x^* \in X$. If $f(x^*) \leqslant \liminf_{x \to x^*} f(x)$, then f is said to be *lower semi-continuous* (l.s.c) at x^*; *upper semi-continuity*, which means $f(x^*) \geqslant \limsup f(x)$, is not much used in our context. It is well-known that, if X is a compact set on which f is continuous, then the lower bound $\bar{f}$ exists and (1.3) has a solution. Actually, lower semi-continuity (of f on the whole compact X) suffices: if (x_k) is a minimizing sequence, with some cluster point $x^* \in X$, we have

$$f(x^*) \leqslant \liminf_{k \to \infty} f(x_k) = \lim_{k \to \infty} f(x_k) = \bar{f}.$$

Another observation is: let E be such that cl $E \subset X$; if f is continuous on cl E, then

$$\inf \{f(x) \; : \; x \in E\} = \inf \{f(x) \; : \; x \in \text{cl}\, E\}.$$

This relation is wrong if f is only l.s.c, though: then, only (1.1) gives useful relations.

Related with (1.3), another problem is whether a given minimizing sequence (x_k) converges to an optimal solution when $k \to +\infty$. This problem is really distinct from (1.3): for example, with $X := \mathbb{R}$, $f(0) := 0$, $f(x) := 1/|x|$ for $x \neq 0$, the sequence defined by $x_k = k$ is minimizing but does not converge to the minimum 0 when $k \to +\infty$.

2 The Set of Extended Real Numbers

In convex analysis, there are serious reasons for wanting to give a meaning to (1.3), for arbitrary f and X. For this, two additional elements are appended to $\mathbb{R}$: $+\infty$ and $-\infty$.

If $E \subset \mathbb{R}$ is nonempty but unbounded from above, we set $\sup E = +\infty$; similarly, $\inf E = -\infty$ if E is unbounded from below. Then consider the case of an empty set: to maintain a relation such as (1.1)

$$\inf (E \cup \emptyset)[= \inf E] = \min \{\inf E, \inf \emptyset\} \quad \text{for all } \emptyset \neq E \subset \mathbb{R},$$

we have no choice and we set $\inf \emptyset = +\infty$. Naturally, $\sup \emptyset = -\infty$, and this maintains the relation $\inf (-E) = -\sup E$ in (1.2).

It should be noted that the world of convex analysis is not symmetric, it is *unilateral*. In particular, $+\infty$ and $-\infty$ do not play the same role, and it suffices for our purpose to consider the set $\mathbb{R} \cup \{+\infty\}$. Extending the notation of the intervals of $\mathbb{R}$, this set will also be denoted by $]-\infty, +\infty]$.

To extend the structure of $\mathbb{R}$ to this new set, the natural rules are adopted:

$$
\begin{array}{rl}
\text{order:} & x \leqslant +\infty \text{ for all } x \in \mathbb{R} \cup \{+\infty\}\,; \\
\text{addition:} & (+\infty) + x = x + (+\infty) = +\infty \text{ for all } x \in \mathbb{R} \cup \{+\infty\}\,; \\
\text{multiplication:} & t \cdot (+\infty) = +\infty \text{ for all } 0 < t \in \mathbb{R} \cup \{+\infty\}\,.
\end{array}
$$

Thus, we see that

– the structured set $(\mathbb{R} \cup \{+\infty\}, +)$ is not a group, just because $+\infty$ has no opposite;
– it is a fortiori not a field, a second reason being that we avoid writing $t \times (+\infty)$ for $t \leqslant 0$.

On the other hand, we leave it to the reader to check that the other axioms are preserved (for the order, the addition and the multiplication); so some calculus can at least be done in $\mathbb{R} \cup \{+\infty\}$.

Actually, $\mathbb{R} \cup \{+\infty\}$ is nothing more than an *ordered convex cone*, analogous to the set $\mathbb{R}_*^+$ of positive numbers. In particular, observe the following continuity properties:

$$
\begin{array}{rll}
(x_k, y_k) \to (x, y) \text{ in } [\mathbb{R} \cup \{+\infty\}]^2 & \Longrightarrow & x_k + y_k \to x + y \text{ in } \mathbb{R} \cup \{+\infty\}\,; \\
(t_k, x_k) \to (t, x) \text{ in } \mathbb{R}_*^+ \times (\mathbb{R} \cup \{+\infty\}) & \Longrightarrow & t_k x_k \to tx \text{ in } \mathbb{R} \cup \{+\infty\}\,.
\end{array}
$$

In this book, starting from Chap. B, the minimization problems of §1 – and in particular (1.3) – will be understood as posed in $\mathbb{R} \cup \{+\infty\}$. The advantage of this is to give a systematic meaning to all the relations of §1. On the other hand, the reader should not feel too encumbered by this new set, which takes the place of the familiar set of real numbers where algebra is "easy". First of all, $\mathbb{R} \cup \{+\infty\}$ is relevant only as far as images of functions are concerned: any algebraic manipulations involving no term $f(x)$ is "safe" and requires no special attention. When some $f(x)$ is involved, the following pragmatic attitude can be adopted:

– comparison and addition: no problems in $\mathbb{R} \cup \{+\infty\}$, just as in $\mathbb{R}$;
– subtraction: before subtracting $f(x)$, make sure that $f(x) < +\infty$;
– multiplication: think of a term like $tf(x)$ as the multiplication of the vector $f(x)$ by the scalar t; if $t \leqslant 0$, make sure that $f(x) < +\infty$ (note: in convex analysis, the product of functions $f(x)g(x)$ is rarely used, and multiplication by -1 puts (1.3) in a different world);
– division: same problems as in $\mathbb{R}$, namely avoid division by 0;
– convergence: same problems as in $\mathbb{R}$, namely pay attention to $\infty - \infty$ and $0 \cdot (+\infty)$;
– in general, do not overuse expressions like $tf(x)$ with $t \leqslant 0$, or $r - f(x)$, etc.: they do not fit well with the conical structure of $\mathbb{R} \cup \{+\infty\}$.

3 Linear and Bilinear Algebra

3.0 Let us start with the *model-situation* of $\mathbb{R}^n$, the real n-dimensional vector space of n-uples $x = (\xi^1, \ldots, \xi^n)$. In this space, the vectors $e_1, \ldots, e_n$, where each e_i has coordinates $(0, \ldots, 0, 1, 0, \ldots, 0)$ (the "1" in i^{th} position) form a basis, called the *canonical* basis. The linear mappings from $\mathbb{R}^m$ to $\mathbb{R}^n$ are identified with the $n \times m$ *matrices* which represent them in the canonical bases; *vectors* of $\mathbb{R}^n$ are thus naturally identified with $n \times 1$ matrices.

The space $\mathbb{R}^n$ is equipped with the canonical, or standard, Euclidean structure with the help of the scalar product

$$x = (\xi^1, \ldots, \xi^n), \ y = (\eta^1, \ldots, \eta^n) \quad \mapsto \quad x^\top y := \sum_{i=1}^n \xi^i \eta^i$$

(also denoted by $x \cdot y$). Then we can speak of the Euclidean space $(\mathbb{R}^n, ^\top)$.

3.1 More generally, a *Euclidean space* is a real vector space, say X, of *finite dimension*, say n, equipped with a *scalar product* denoted by $\langle \cdot, \cdot \rangle$. Recall that a scalar (or inner) product is a bilinear symmetric mapping $\langle \cdot, \cdot \rangle$ from $X \times X$ to $\mathbb{R}$, satisfying $\langle x, x \rangle > 0$ for $x \neq 0$.

(a) If a basis $\{b_1, \ldots, b_n\}$ has been chosen in X, along which two vectors x and y have the coordinates $(\xi^1, \ldots, \xi^n)$ and $(\eta^1, \ldots, \eta^n)$, we have

$$\langle x, y \rangle = \sum_{i,j=1}^n \xi^i \eta^j \langle b_i, b_j \rangle .$$

This can be written $\langle x, y \rangle = x^\top Q y$, where Q is a symmetric positive definite $n \times n$ matrix ($S_n(\mathbb{R})$ will denote the set of symmetric matrices). In this situation, to equip X with a scalar product is actually to take a symmetric positive definite matrix.

The simplest matrix Q is the identity matrix I, or I_n, which corresponds to the scalar product

$$\langle x, y \rangle = x^\top y = \sum_{i=1}^{n} \xi^i \eta^i \, ,$$

called the *dot-product*. For this particular product, one has $\langle b_i, b_j \rangle = \delta_{ij}$ (δ_{ij} is the symbol of Kronecker: $\delta_{ij} = 0$ if $i \neq j$, $\delta_{ii} = 1$). The basis $\{b_1, \ldots, b_n\}$ is said to be *orthonormal* for this scalar product; and this scalar product is of course the only one for which the given basis is orthonormal.

Thus, whenever we have a basis in X, we know all the possible ways of equipping X with a Euclidean structure.

(b). Reasoning in the other direction, let us start from a Euclidean space $(X, \langle \cdot, \cdot \rangle)$ of dimension n. It is possible to find a basis $\{b_1, \ldots, b_n\}$ of X, which is orthonormal for the given scalar product (i.e. which satisfies $\langle b_i, b_j \rangle = \delta_{ij}$ for $i, j = 1, \ldots, n$). If two vectors x and y are expressed in terms of this basis, $\langle x, y \rangle$ can be written $x^\top y$.

Use the space $\mathbb{R}^n$ of §3.0 and denote by $\varphi : \mathbb{R}^n \to X$ the unique *linear operator* (isomorphism of vector spaces) satisfying $\varphi(e_i) = b_i$ for $i = 1, \ldots, n$. Then

$$x^\top y = \langle \varphi(x), \varphi(y) \rangle \quad \text{for all } x \text{ and } y \text{ in } \mathbb{R}^n \, ,$$

so the Euclidean structure is also carried over by φ, which is therefore an isomorphism of Euclidean spaces as well. Thus, *any* Euclidean space $(X, \langle \cdot, \cdot \rangle)$ of dimension n is *isomorphic* to $(\mathbb{R}^n, ^\top)$, which explains the importance of this last space. However, given a Euclidean space, an orthonormal basis need not be easy to construct; said otherwise, one must sometimes content oneself with a scalar product imposed by the problem considered.

Example 3.1 Vector spaces of matrices form a rich field of applications for the techniques and results of convex analysis. The set of $p \times q$ matrices forms a vector space of dimension pq, in which a natural scalar product of two matrices M and N is ($\operatorname{tr} A := \sum_{i=1}^{n} A_{ii}$ is the *trace* of the $n \times n$ matrix A)

$$\langle M, N \rangle := \operatorname{tr} M^\top N = \sum_{i=1}^{p} \sum_{j=1}^{q} M_{ij} N_{ij} \, . \qquad \square$$

(c). A subspace V of $(X, \langle \cdot, \cdot \rangle)$ can be equipped with the Euclidean structure defined by

$$V \times V \ni (x, y) \mapsto \langle x, y \rangle \, .$$

Unless otherwise specified, we will generally use this *induced* structure, with the same notation for the scalar product in V and in X.

More importantly, let $(X_1, \langle \cdot, \cdot \rangle_1)$ and $(X_2, \langle \cdot, \cdot \rangle_2)$ be two Euclidean spaces. Their Cartesian product $X = X_1 \times X_2$ can be made Euclidean via the scalar product

$$((x_1, x_2), (y_1, y_2)) = (x, y) \mapsto \langle\!\langle x, y \rangle\!\rangle := \langle x_1, y_1 \rangle_1 + \langle x_2, y_2 \rangle_2 \, .$$

This is not compulsory: cases may occur in which the product-space X has its own Euclidean structure, not possessing this "decomposability" property.

3.2 Let $(X, \langle \cdot, \cdot \rangle)$ and $(Y, \langle\!\langle \cdot, \cdot \rangle\!\rangle)$ be two Euclidean spaces, knowing that we could write just as well $X = \mathbb{R}^n$ and $Y = \mathbb{R}^m$.

(a). If A is a linear operator from X to Y, the *adjoint* of A is the unique operator A^* from Y to X, defined by

$$\langle A^* y, x \rangle = \langle\!\langle y, Ax \rangle\!\rangle \quad \text{for all } (x, y) \in X \times Y.$$

There holds $(A^*)^* = A$. When both X and Y have orthonormal bases (as is the case with canonical bases for the dot-product in the respective spaces), the matrix representing A^* in these bases is the *transpose* of the matrix representing A.

Consider the case $(Y, \langle\!\langle \cdot, \cdot \rangle\!\rangle) = (X, \langle \cdot, \cdot \rangle)$. When A is invertible, so is A^*, and then $(A^*)^{-1} = (A^{-1})^*$. When $A^* = A$, we say that A is self-adjoint, or *symmetric*. If, in addition,

$$\langle Ax, x \rangle > 0 \quad [\text{resp.} \geqslant 0] \quad \text{for all } 0 \neq x \in X,$$

then A is *positive definite* [resp. positive semi-definite] and we write $A \succ 0$ [resp. $A \succ 0$]. When $X = Y$ is equipped with an orthonormal basis, symmetric operators can be characterized in terms of matrices: A is symmetric [resp. symmetric positive (semi)-definite] if and only if the matrix representing A (in the orthonormal basis) is symmetric [resp. symmetric positive (semi)-definite].

(b). When the image-space Y is $\mathbb{R}$, an operator is rather called a *form*. If ℓ is a linear form on $(X, \langle \cdot, \cdot \rangle)$, there exists a unique $s \in X$ such that $\ell(x) = \langle s, x \rangle$ for all $x \in X$. If q is a quadratic form on $(X, \langle \cdot, \cdot \rangle)$, there exists a unique symmetric operator Q such that

$$q(x) := \tfrac{1}{2} \langle Qx, x \rangle \quad \text{for all } x \in X$$

(the coefficient 1/2 is useful to simplify most algebraic manipulations).

Remark 3.2 The correspondence $\ell \leftrightarrows s$ is a triviality in $(\mathbb{R}^n, ^\top)$ (just transpose the $1 \times n$ matrices to vectors) but this is deceiving. Indeed, it is the correspondence $X \leftrightarrows X^*$ between a space and its *dual* that is being considered. For two vectors s and x of X, it is good practice to think of the scalar product $\langle s, x \rangle$ as the action of the first argument s (a slope, representing an element in the dual) on the second argument x; this helps one to understand what one is doing. Likewise, the operator Q associated with a quadratic form sends X to X^*; and an adjoint A^* is from Y^* to X^*.   □

3.3 Two subspaces U and V of $(X, \langle \cdot, \cdot \rangle)$ are mutually *orthogonal* if $\langle u, v \rangle = 0$ for all $u \in U$ and $v \in V$, a relation denoted by $U \perp V$. On the other hand, U and V are *generators* of X if $U + V = X$. For given U, we denote by $U^\perp$ the *orthogonal supplement* of U, i.e. the unique subspace orthogonal to U such that U and $U^\perp$ form a generator of X.

Let $A : X \to Y$ be an arbitrary linear operator, X and Y having arbitrary scalar products. As can easily be seen,

$$\operatorname{Ker} A := \{x \in X : Ax = 0 \in Y\}$$

and

$$\operatorname{Im} A^* := \{x \in X \ : \ x = A^* s \quad \text{for some } s \in Y\}$$

are orthogonal generators of X. In other words,

$$\operatorname{Ker} A = (\operatorname{Im} A^*)^\perp .$$

This is a very important relation; one must learn to use it quasi-mechanically, remembering that $A^{**} = A$ and $U^{\perp\perp} = U$. For example, if A is symmetric, $\operatorname{Im} A$ is the orthogonal supplement of $\operatorname{Ker} A$.

Important examples of linear operators from $(X, \langle \cdot, \cdot \rangle)$ to itself are *orthogonal projections*: if H is a subspace of X, the operator $\mathrm{p}_H : X \to X$ of orthogonal projection onto H is defined by:

$$\begin{vmatrix} \mathrm{p}_H x = 0 & \text{for} & x \in H^\perp , \\ \mathrm{p}_H x = x & \text{for} & x \in H , \\ \mathrm{p}_H \text{ is completed by linearity in between.} \end{vmatrix}$$

This p_H is *symmetric* and *idempotent* (i.e. $\mathrm{p}_H \circ \mathrm{p}_H = \mathrm{p}_H$). Conversely, a linear operator p which is symmetric and idempotent is an orthogonal projection; of course, it is the projection onto the subspace $\operatorname{Im} p$.

3.4 If A is a symmetric linear operator on X, remember that $(\operatorname{Im} A)^\perp = \operatorname{Ker} A$. Then consider the operator $\mathrm{p}_{\operatorname{Im} A}$ of orthogonal projection onto $\operatorname{Im} A$. For given $y \in X$, there is a unique $x = x(y)$ in $\operatorname{Im} A$ such that $Ax = \mathrm{p}_{\operatorname{Im} A} y$; furthermore, the mapping $y \mapsto x(y)$ is linear. This mapping is called the *pseudo-inverse*, or generalized inverse, of A (more specifically, it the pseudo-inverse of Moore and Penrose). We denote it by A^-; other notations are A^+, $A^\#$, etc.

We recall some useful properties of the pseudo-inverse: $\operatorname{Im} A^- = \operatorname{Im} A$; $A^- A = A A^- = \mathrm{p}_{\operatorname{Im} A}$; and if A is positive semi-definite, so is A^-.

4 Differentiation in a Euclidean Space

A Euclidean space $(X, \langle \cdot, \cdot \rangle)$ is a normed vector space (certainly complete) thanks to the norm

$$X \ni x \mapsto \|x\| := \sqrt{\langle x, x \rangle} ,$$

called the *Euclidean* norm associated with $\langle \cdot, \cdot \rangle$. We denote by

$$B(x, r) := \{y \in X \ : \ \|y - x\| \leqslant r\}$$

the *ball* of center $x \in X$ and radius $r > 0$. In particular, $B(0, 1)$ is called the unit ball, whose boundary is the *unit sphere*

$$\widetilde{B}(0, 1) := \{y \in X \ : \ \|y\| = 1\} .$$

The norm and scalar product are related by the fundamental *Cauchy-Schwarz inequality* (no "t", please)

$$|\langle s, x\rangle| \leqslant \|s\|\,\|x\| \quad \text{for all } (s, x) \in X \times X\,.$$

Remember that all norms are equivalent in our finite-dimensional space X: if $\|\!|\cdot|\!\|$ is another norm, there are positive numbers ℓ and L such that

$$\ell\|\!|x|\!\| \leqslant \|x\| \leqslant L\|\!|x|\!\| \quad \text{for all } x \in X\,.$$

However, the Euclidean norm $\|\cdot\|$ plays a special role.

4.1 We will denote by $\varepsilon(\cdot)$ a generic function (from some normed space to another) which tends to 0 when its argument tends to 0. For example, the continuity of f at x can be expressed by
$$f(x + h) = f(x) + \varepsilon(h)\,.$$

When we need to distinguish *speeds* of convergence, multiplicative factors can be used. For example,
$$\|h\|^{\alpha}\varepsilon(h) \quad \text{for } \alpha = 1, 2, \ldots$$

denotes the functions tending to 0 faster than $\|h\|, \|h\|^2, \ldots$ A more handy notation for $\|h\|^{\alpha}\varepsilon(h)$ is $o(\|h\|^{\alpha})$ (pronounce "little oh of ...").

Beware that these are only notations, and algebraic manipulation with them should be done very carefully. For example $\varepsilon(\cdot) \equiv r\varepsilon(\cdot)$ for all $r \neq 0$, hence in particular $\varepsilon(\cdot) - \varepsilon(\cdot) = \varepsilon(\cdot)$! Always keep the definitions in mind; for example, to say that a function $h \mapsto \varphi(h)$ is $o(\|h\|)$ means:

$$\forall \varepsilon > 0 \ \exists \delta > 0 \quad \text{such that} \quad \|h\| \leqslant \delta \implies \|\varphi(h)\| \leqslant \varepsilon\|h\|\,.$$

With this last notation, a function $f : \Omega \to \mathbb{R}$, defined on an open set $\Omega \subset X$, is said to be *differentiable* at $x \in \Omega$ if there exists a linear form ℓ on X such that

$$f(x + h) = f(x) + \ell(h) + o(\|h\|)\,.$$

This linear form ℓ, denoted by $f'(x)$, $\mathrm{D}f(x)$ or $\mathrm{d}f(x)$, is called the *differential* of f at x. According to §3.2(b), it can be represented by a unique element of X; this element is called the *gradient* of f at x, denoted by $\nabla f(x)$, and is therefore defined by
$$f'(x)(h) = \langle \nabla f(x), h\rangle \quad \text{for all } h \in X\,.$$

Example 4.1 Let $H \subset X$ be a subspace, equipped with the Euclidean structure induced by $(X, \langle\cdot, \cdot\rangle)$ as in §3.1(c). If f is differentiable at $x \in H$, its gradient $\nabla f(x)$ is obtained from

$$f(x + h) = f(x) + \langle \nabla f(x), h\rangle + o(\|h\|)\,. \tag{4.1}$$

Then define the function $f_H : H \to \mathbb{R}$ to be the restriction of f to H. This f_H is differentiable at x and its gradient $\nabla f_H(x)$ is the vector (of H!) satisfying

$$f(x + h) = f(x) + \langle \nabla f_H(x), h \rangle + o(\|h\|) \quad \text{for all } h \in H \, .$$

Its computation is simple: in view of the properties of an orthogonal projection,

$$\langle \nabla f(x), h \rangle = \langle p_H \nabla f(x), h \rangle \quad \text{for all } h \in H \, .$$

Plugging this into (4.1), we see that $\nabla f_H(x) = p_H \nabla f(x)$. □

It is important to realize that the representation of $f'(x)$ by $\nabla f(x)$ changes if $\langle \cdot, \cdot \rangle$ is changed. The gradient depends on the scalar product; but the differential does not.

If the space is equipped with an orthonormal basis, along which x has the coordinates $\xi^1, \ldots, \xi^n$, then $f'(x)$ is represented by the *row-matrix*

$$\left[\frac{\partial f}{\partial \xi^1}(x), \ldots, \frac{\partial f}{\partial \xi^n}(x) \right]$$

and $\nabla f(x)$ is the *vector* of $\mathbb{R}^n$ whose coordinates are $(\partial f / \partial \xi^i)(x)$ for $i = 1, \ldots, n$.

4.2 More generally, a function F from $\Omega \subset X$ to some other Euclidean space, say $Y = \mathbb{R}^m$, is differentiable at $x \in \Omega$ if there exists a linear operator L from X to Y yielding the first-order approximation

$$F(x + h) = F(x) + L(h) + o(\|h\|) \, .$$

The differential L of F at x is also called the *Jacobian operator* of F at x, again denoted by $F'(x)$, $\mathrm{D}F(x)$, or also $\mathrm{J}F(x)$. Nothing is really new with respect to the scalar case of §4.1; denoting by $f_1, \ldots, f_m$ the component-functions of F along some basis of Y, F is differentiable at x if and only if each f_j is such and

$$\mathrm{J}F(x)(h) = (f_1'(x)(h), \ldots, f_m'(x)(h)) \quad \text{for all } h \in X \, .$$

The matrix representation of $\mathrm{J}F(x)$ along the bases of X and Y is an $m \times n$ matrix, whose $(i, j)^{\mathrm{th}}$ element is $(\partial f_i / \partial \xi^j)(x)$.

Given a scalar-valued function f, differentiable on Ω, consider the function $y \mapsto f'(y)$, sending Ω to the space of linear forms on X. If this new function is in turn differentiable at x, we obtain the *second-order* differential (of f at x). This defines a *bilinear form* via

$$X \times X \ni (h, k) \mapsto [(f')'(x)(h)](k) =: f''(x)(h, k) \, ,$$

which is also *symmetric*; as such, it induces a quadratic form on X (for which we will use the same notation).

If X is equipped with a scalar product $\langle \cdot, \cdot \rangle$, §3.2(b) tells us that the quadratic form $f''(x)$ defines a symmetric operator: the *Hessian* of f at x, denoted by $\nabla^2 f(x)$, or $\mathrm{H}f(x)$. Just as the gradient, the Hessian depends on the scalar product; and there holds the *second-order approximation* of f at x:

$$f(x + h) = f(x) + \langle \nabla f(x), h \rangle + \tfrac{1}{2} \langle \nabla^2 f(x) h, h \rangle + o(\|h\|^2) \, .$$

With an orthonormal basis and $x = (\xi^1, \ldots, \xi^n)$, $\nabla^2 f(x)$ is represented by a symmetric matrix whose $(i, j)^{\mathrm{th}}$ element is $(\partial^2 f / \partial \xi^i \partial \xi^j)(x)$, called the Hessian matrix of f at x.

5 Set-Valued Analysis

5.1 If S is a nonempty closed subset of $\mathbb{R}^n$, we denote by

$$\mathrm{d}_S(x) := \min_{y \in S} \|y - x\|$$

the *distance* from x to S. Given two nonempty closed sets S_1 and S_2, consider

$$\mathrm{e}_H(S_1/S_2) := \sup \{\mathrm{d}_{S_2}(x) \, : \, x \in S_1\},$$

called the *excess* of S_1 over S_2: geometrically,

$$\mathrm{e}_H(S_1/S_2) \leqslant \varepsilon \quad \text{means} \quad S_1 \subset S_2 + B(0, \varepsilon).$$

The *Hausdorff-distance* Δ_H between S_1 and S_2 is then the symmetrization of the above concept:

$$\Delta_H(S_1, S_2) := \max \{\mathrm{e}_H(S_1/S_2), \mathrm{e}_H(S_2/S_1)\}.$$

One checks immediately that $\Delta_H(S_1, S_2) \in \mathbb{R}^+ \cup \{+\infty\}$, but Δ_H is a finite-valued function when restricted to *bounded* closed sets. Also

$$\Delta_H(S_1, S_2) = 0 \iff S_1 = S_2,$$
$$\Delta_H(S_1, S_2) = \Delta_H(S_2, S_1),$$
$$\Delta_H(S_1, S_3) \leqslant \Delta_H(S_1, S_2) + \Delta_H(S_2, S_3).$$

In other words, Δ_H does define a distance on the family of nonempty compact subsets of $\mathbb{R}^n$.

5.2 A mapping F which, to $x \in X$, associates a subset of $\mathbb{R}^n$, is called a multi-valued, or set-valued mapping, or more simply a *multifunction*; we use the notation

$$X \ni x \longmapsto F(x) \subset \mathbb{R}^n$$

(watch the arrow, longer than in §1.2). The *domain* dom F of F is the set of $x \in X$ such that $F(x) \neq \emptyset$. Its *image* (or range) $F(X)$ and *graph* gr F are the unions of the sets $F(x) \subset \mathbb{R}^n$ and $\{x\} \times F(x) \subset X \times \mathbb{R}^n$ respectively, when x describes X (or, more precisely, dom F). A *selection* of F is a particular function $f : \mathrm{dom}\, F \to \mathbb{R}^n$ with $f(x) \in F(x)$ for all x.

The concept of convergence is here much more tricky than in the single-valued case. First of all, since a limit is going to be a set anyway, the following concept is relevant: the *limes exterior* of $F(x)$ for $x \to x^*$ is the set of all cluster points of all selections (here, $x^* \in \mathrm{cl}\,\mathrm{dom}\, F$). In other words, $y \in \lim \mathrm{ext}_{x \to x^*} F(x)$ means: there exists a sequence $(x_k, y_k)_k$ such that

$$y_k \in F(x_k), \ x_k \to x^* \text{ and } y_k \to y \quad \text{when } k \to +\infty.$$

Note that this does not depend on multi-valuedness: each $F(x)$ might well be a singleton for all x. For example,

$$\lim_{t\downarrow 0}\mathrm{ext}\{\sin(1/t)\} = [-1,+1]\,.$$

The *limes interior* of $F(x)$ for $x \to x^*$ is the set of limits of all convergent selections: $y \in \lim\mathrm{int}_{x\to x^*} F(x)$ means that there exists a function $x \mapsto f(x)$ such that

$$f(x) \in F(x) \text{ for all } x \quad \text{and} \quad f(x) \to y \text{ when } x \to x^*\,.$$

Clearly enough, one always has $\lim\mathrm{int}\, F(x) \subset \lim\mathrm{ext}\, F(x)$; when these two sets are equal, the common set is the *limit* of $F(x)$ when $x \to x^*$.

Remark 5.1 The above concepts are classical but one usually speaks of lim sup and lim inf. The reason is that the lim ext [resp. lim int] is the largest [resp. smallest] cluster set for the order "$\subset$". Such a terminology is however misleading, since this order does not generalize "$\leqslant$".

For example, with $X = \{1, 1/2, \ldots, 1/k, \ldots\}$ (and $x^* = 0$), what are the lim sup and lim inf of the sequence $((-1)^k)$ for $k \to +\infty$? With the classical terminology, the answer is confusing:

– If $(-1)^k$ is considered as a *number* in the ordered set $(\mathbb{R}, \leqslant)$, then

$$\lim\sup\,(-1)^k = 1 \quad \text{and} \quad \lim\inf\,(-1)^k = -1\,.$$

– If $(-1)^k$ is considered as a *singleton* in the ordered set $(2^{\mathbb{R}}, \subset)$, then

$$\lim\sup\,(-1)^k = \{-1,+1\} \quad \text{and} \quad \lim\inf\,(-1)^k = \emptyset\,. \qquad \square$$

Beware that the above *limes* may cover somewhat pathological behaviours. Take for example the multifunction

$$]0,+\infty[\ni t \longmapsto F(t) := \{0, 1/t\} \subset \mathbb{R}.$$

Then $\lim\mathrm{ext}_{t\downarrow 0} F(t) = \{0\}$, a set which does not reflect the intuitive idea that a lim ext should connote. Note: in this example, $\Delta_H[F(t), \{0\}] = \mathrm{e}[F(t)/\{0\}] \to +\infty$ when $t \downarrow 0$. The same pathological behaviour of the Hausdorff distance occurs in the following example:

$$F(t) := [0, 1/t] \quad \text{for } t > 0 \quad \text{and} \quad F(0) = [0,+\infty[\,.$$

Then $F(0) = \lim_{t\downarrow 0} F(t)$ but $\Delta_H[F(t), F(0)] \equiv +\infty\,.$

5.3 The multifunction F is said to be bounded-valued, closed-valued, convex-valued etc. when the sets $F(x)$ are bounded, closed, convex etc. If its graph is a closed set, we say that the multifunction is *closed*. In order to avoid the nasty situations mentioned above, a convenient property is *local boundedness*: we say that the multifunction F is locally bounded near x^* when:

$$\text{For some neighborhood } N \text{ of } x^* \text{ and bounded set } B \subset \mathbb{R}^n, \tag{5.1}$$
$$N \subset \mathrm{dom}\, F \quad \text{and} \quad F(N) \subset B\,.$$

If F is locally bounded near every x^* in a set S, we say that F is locally bounded on S. Then a multifunction F satisfying (5.1) is

– *outer semi-continuous* at x^* when

$$\lim_{x \to x^*} \operatorname{ext} F(x) \subset F(x^*)\,,$$

– *inner semi-continuous* at x^* when

$$F(x^*) \subset \liminf_{x \to x^*} F(x)\,.$$

– *continuous* when it is both outer and inner semi-continuous.

When F is closed-valued, these definitions can be made more handy thanks to (5.1), namely: for all $\varepsilon > 0$, there is a neighborhood $N(x^*)$ such that $x \in N(x^*)$ implies

$$\begin{array}{ll} F(x) \subset F(x^*) + B(0,\varepsilon) & \text{[outer semi-continuity]} \\ F(x^*) \subset F(x) + B(0,\varepsilon)\,. & \text{[inner semi-continuity]} \end{array} \tag{5.2}$$

It is straightforward to check that (5.2) has an equivalent in terms of excesses:

$$\begin{array}{ll} e[F(x)/F(x^*)] \leqslant \varepsilon & \text{[outer semi-continuity]} \\ e[F(x^*)/F(x)] \leqslant \varepsilon\,. & \text{[inner semi-continuity]} \end{array}$$

In words, outer semi-continuity at x^* means: all the points in $F(x)$ are close to $F(x^*)$ if the varying x is close to the fixed x^*. When moving away from x^*, F does not expand suddenly. Inner semi-continuity is the other way round: $F(x)$ does not explode when the varying x reaches x^*. If the mapping is actually single-valued, both definitions coincide with that of a continuous function at x^*. In practice, outer semi-continuity is frequently encountered, while inner semi-continuity is less natural.

5.4 Finally, we mention a situation in which limits and continuity of sets have a natural definition: when X is an ordered set, and $x \longmapsto F(x)$ is *nested*. For example, take $X = \mathbb{R}^+$ and let $t \longmapsto F(t)$ satisfy $F(t) \subset F(t')$ whenever $t \geqslant t' > 0$. Then the set

$$\lim_{t \downarrow 0} F(t) := \operatorname{cl} \cup \{F(t)\ :\ t > 0\}$$

coincides with the limit of F defined in §5.2.

6 Recalls on Convex Functions of the Real Variable

We recall that a function f, defined on a nonempty interval I, is said to be *convex on I* when

$$f(\alpha x + (1 - \alpha)x') \leqslant \alpha f(x) + (1 - \alpha)f(x') \tag{6.3}$$

for all pairs of points (x, x') in I and all $\alpha \in\,]0, 1[$.

A fundamental convexity criterion is as follows:

Proposition 6.1 (Criterion of Increasing Slopes) *A function f is convex on an interval I if and only if, for all $x_0 \in I$, the slope-function*

$$x \mapsto \frac{f(x) - f(x_0)}{x - x_0} =: s(x) \tag{6.4}$$

is increasing on $I \backslash \{x_0\}$.

Proof. Draw a picture and work out the necessary algebra.   □

Convex functions enjoy some remarkable regularity properties:

Theorem 6.2 *A function f, convex on an interval I, is continuous on int I.*

Proof. Let $x_0 \in$ int I and take a, b such that $x_0 \in\,]a, b[\, \subset$ int I. For $x \downarrow x_0$ (so $x \in\,]x_0, b[$), write

$$x = \alpha b + (1 - \alpha)x_0 \quad \text{and} \quad x_0 = \beta a + (1 - \beta)x,$$

with $\alpha \downarrow 0$ and $\beta \downarrow 0$. Then apply the relation of definition (6.3) to see that

$$f(x) \leqslant f(x_0) + \alpha[f(b) - f(x_0)] \quad \text{and} \quad \frac{f(x_0)}{1 - \beta} \leqslant \frac{\beta}{1 - \beta}f(a) + f(x)$$

hence, passing to the limit:

$$\limsup f(x) \leqslant f(x_0) \quad \text{and} \quad f(x_0) \leqslant \liminf f(x).$$

Same technique for $x \uparrow x_0$.   □

Theorem 6.3 *A function f convex on an interval I admits a finite left-derivative and a finite right-derivative at each $x_0 \in$ int I:*

$$D_- f(x_0) := \lim_{x \uparrow x_0} \frac{f(x) - f(x_0)}{x - x_0} = \sup_{x < x_0} \frac{f(x) - f(x_0)}{x - x_0} \tag{6.5}$$

$$D_+ f(x_0) := \lim_{x \downarrow x_0} \frac{f(x) - f(x_0)}{x - x_0} = \inf_{x > x_0} \frac{f(x) - f(x_0)}{x - x_0}. \tag{6.6}$$

They satisfy

$$D_- f(x_0) \leqslant D_+ f(x_0). \tag{6.7}$$

Proof. Apply the Criterion 6.1 of increasing slopes: the difference quotient involved in (6.5), (6.6) is just the slope-function s of (6.4). For any two points x, x' in int dom f satisfying $x < x_0 < x'$, $s(x)$ and $s(x')$ are finite numbers satisfying $s(x) \leqslant s(x')$. Furthermore, when $x \uparrow x_0$ [resp. $x \downarrow x_0$], $s(x)$ increases [resp. $s(x')$ decreases], hence they both converge, say as described by the notation (6.5), (6.6); this also proves (6.7).   □

These results have converse statements:

Theorem 6.4 *Let a function f be continuous on an open interval I and possess an increasing right-derivative, or an increasing left-derivative, on I. Then f is convex on I.*

Proof. Assume that f has an increasing right-derivative D_+f. For x, x' in I with $x < x'$ and $u \in]x, x'[$, there holds

$$\frac{f(u) - f(x)}{u - x} \leqslant \sup_{t \in]x,u[} D_+f(t) \leqslant \inf_{t \in]u,x'[} D_+f(t) \leqslant \frac{f(x') - f(u)}{x' - u}$$

(the first and last inequalities come from mean-value theorems – in inequality form – for continuous functions admitting right-derivatives). Then (6.3) is obtained via a multiplication by $x' - x > 0$, knowing that $u = \alpha x + (1 - \alpha)x'$ for some $\alpha \in]0, 1[$. The proof for $D_- f$ is just the same.   □

The case where the two half-derivatives coincide is worth mentioning:

Corollary 6.5 *Let a function f be differentiable with an increasing derivative on an open interval I. Then f is convex on I.*   □

Theorem 6.6 *Let a function f be twice differentiable with a nonnegative second derivative on an open interval I. Then f is convex on I.*

Proof. Immediate from Corollary 6.5.   □

Exercises

1. Given a function $X \ni x \mapsto f(x)$, consider the following minimization problems:

$$(P_1) \quad \begin{cases} \inf f(x), \\ x \in X \end{cases} \qquad \text{and} \qquad (P_2) \quad \begin{cases} \inf_{x,r} r, \\ (x,r) \in X \times \mathbb{R}, \\ f(x) \leqslant r. \end{cases}$$

– Show that they have the same infimal value.
– Show that, if (P_1) has an optimal solution, then the infimum in (P_2) is attained.
– Show that $\bar{x}$ solves (P_1) if and only if there is $\bar{r} \in \mathbb{R}$ such that $(\bar{x}, \bar{r})$ solves (P_2).

2.* Prove the decoupling property (1.5).
     Setting $f_1(x) := \inf g(x, \cdot)$ and taking $(\bar{x}, \bar{y}) \in \text{Argmin}\, g$, show that

$$\bar{x} \in \text{Argmin}\, f_1 \qquad \text{and} \qquad \bar{y} \in \text{Argmin}\, g(\bar{x}, \cdot).$$

Imagine counter-examples to see that converse implications need not hold.

3*. Given a function $g : X \times Y \to \mathbb{R}$, establish the relation

$$\sup_{y \in Y} \left(\inf_{x \in X} g(x, y) \right) \leqslant \inf_{x \in X} \left(\sup_{y \in Y} g(x, y) \right)$$

(start from the inner inequality $\inf g(\cdot, y) \leqslant \sup g(x, \cdot)$, which can be used as a mnemonics: when one minimizes, one obtains something smaller than when one maximizes!)

Setting $f_1(x) := \sup g(x, \cdot)$ and $f_2(y) := \inf g(\cdot, y)$ and taking $\bar{y} \in \text{Argmin } f_2$ (assumed nonempty), compare $\text{Argmin } g(\cdot, \bar{y})$ and $\text{Argmin } f_1$ [hint: use the function $g(x, y) = xy$ to guess what can be proved and what cannot.]

4. Let a, b, c be three real numbers, with $a > 0$. For the function $\mathbb{R} \ni x \mapsto q(x) := ax^2 + 2bx + c$, show the equivalence of the two statements

(i) $q(x) > 0$ for all $x \geqslant 0$,

(ii) $c > 0$ and $b + \sqrt{ac} > 0$.

5. Let $A : \mathbb{R}^n \to \mathbb{R}^n$ be a self-adjoint operator, $b \in \mathbb{R}^n$, and consider the quadratic function

$$\mathbb{R}^n \ni x \mapsto q(x) := \langle Ax, x \rangle - 2\langle b, x \rangle .$$

Show that the three statements

(i) $\inf \{q(x) : x \in \mathbb{R}^n\} > -\infty$,

(ii) $A \succcurlyeq 0$ and $b \in \text{Im } A$,

(iii) the problem $\inf \{q(x) : x \in \mathbb{R}^n\}$ has a solution,

are equivalent. When they hold, characterize the set of minimum points of q, in terms of the pseudo-inverse of A.

6. Let $f :]0, +\infty[\to \mathbb{R}$ be continuous. Show the equivalence of the statements

(i) $f\left(\sqrt{xy}\right) \leqslant \frac{1}{2}[f(x) + f(y)]$ for all $x > 0$ and $y > 0$,

(ii) the function $t \mapsto g(t) := f\left(e^t\right)$ is convex on $\mathbb{R}$.

7. In $S_n(\mathbb{R})$, we denote by $S_n^{++}(\mathbb{R})$ [resp. $S_n^+(\mathbb{R})$] the set of positive [resp. semi]-definite matrices. Characterize the properties $A \in S_n^{++}(\mathbb{R})$ and $A \in S_n^+(\mathbb{R})$ in terms of the eigenvalues of A. What is the boundary of $S_n^+(\mathbb{R})$?

Now $n = 2$ and we use the notation $A = \begin{bmatrix} x & z \\ z & y \end{bmatrix}$ for matrices in $S_2(\mathbb{R})$.

Give necessary and sufficient conditions on x, y, z for which $A \in S_n^+(\mathbb{R})$.

Draw a picture in $\mathbb{R}^3$ to visualize the set $S_2^+(\mathbb{R})$, and its intersection with the matrices of trace 1 in $S_2(\mathbb{R})$.

Show that the boundary of $S_2^+(\mathbb{R})$ is the set of matrices of the form $vv^\top$, with v describing $\mathbb{R}^2$. Give expressions of x, y, z for these matrices, as well as for those of them that have trace 1.

A. Convex Sets

Introduction. Our working space is $\mathbb{R}^n$. We recall that this space has the structure of a real vector space (its elements being called *vectors*), and also of an affine space (a set of *points*); the latter can be identified with the vector-space $\mathbb{R}^n$ whenever an *origin* is specified. It is not always possible, nor even desirable, to distinguish vectors and points.

We equip $\mathbb{R}^n$ with a scalar product $\langle \cdot, \cdot \rangle$, so that it becomes a Euclidean space, and also a complete normed vector space for the norm $\|x\| := \langle x, x \rangle^{1/2}$. If an orthonormal basis is chosen, there is no loss of generality in assuming that $\langle x, y \rangle$ is the usual dot-product $x^\top y$; see §0.3.

The concepts presented in this chapter are of course fundamental, as practically all the subsequent material is based on them (including the study of convex functions). These concepts must therefore be fully mastered, and we will insist particularly on ideas, rather than technicalities.

1 Generalities

1.1 Definition and First Examples

Definition 1.1.1 The set $C \subset \mathbb{R}^n$ is said to be convex if $\alpha x + (1 - \alpha)x'$ is in C whenever x and x' are in C, and $\alpha \in {]}0, 1[$ (or equivalently $\alpha \in [0, 1]$). □

Geometrically, this means that the line-segment

$$[x, x'] := \{\alpha x + (1 - \alpha)x' \ : \ 0 \leqslant \alpha \leqslant 1\}$$

is entirely contained in C whenever its endpoints x and x' are in C. Said otherwise: the set $C - \{c\}$ is a *star-shaped* set whenever $c \in C$ (a star-shaped set is a set containing the segment $[0, x]$ for all its points x). A consequence of the definition is that C is also path-connected, i.e. two arbitrary points in C can be linked by a continuous path.

Examples 1.1.2 (Sets Based on Affinity) Clearly, the convex sets in one dimension are exactly the intervals; let us give some more fundamental examples in several dimensions.

(a) An *affine hyperplane*, or hyperplane, for short, is a set associated with $(s, r) \in \mathbb{R}^n \times \mathbb{R}$ ($s \neq 0$) and defined by

$$H_{s,r} := \{x \in \mathbb{R}^n \ : \ \langle s, x \rangle = r\}.$$

An affine hyperplane is clearly a convex set. Fix s and let r describe $\mathbb{R}$; then the affine hyperplanes $H_{s,r}$ are translations of the same *linear*, or vector, hyperplane $H_{s,0}$. This $H_{s,0}$ is the subspace of vectors that are orthogonal to s and can be denoted by $H_{s,0} = \{s\}^\perp$. Conversely, we say that s is the *normal* to $H_{s,0}$ (up to a multiplicative constant). Affine hyperplanes play a fundamental role in convex analysis; the correspondence between $0 \neq s \in \mathbb{R}^n$ and $H_{s,1}$ is the basis for *duality* in a Euclidean space.

(b) More generally, an affine subspace, or *affine manifold*, is a set V such that the (affine) line $\{\alpha x + (1 - \alpha)x' : \alpha \in \mathbb{R}\}$ is entirely contained in V whenever x and x' are in V (note that a single point is an affine manifold). Again, an affine manifold is clearly convex.

Take $v \in V$; it is easy – but instructive – to show that $V - \{v\}$ is a subspace of $\mathbb{R}^n$, which is independent of the particular v; denote it by V_0. Thus, an affine manifold V is nothing but the translation of some vector space V_0, sometimes called the *direction* (-subspace) of V; we will also say that V_0 and V are *parallel*. One can therefore speak of the *dimension* of an affine manifold V: it is just the dimension of V_0. We summarize in Table 1.1.1 the particular cases of affine manifolds.

Table 1.1.1. Various affine manifolds

Name	Possible definition	Direction	Dimension
point	$\{x\} \ (x \in \mathbb{R}^n)$	$\{0\}$	0
affine line	$\alpha x_1 + (1 - \alpha)x_2$ $\alpha \in \mathbb{R}, \ x_1 \neq x_2$	vector line $\mathbb{R}(x_1 - x_2)$	1
affine hyperplane	$\{x \in \mathbb{R}^n \ : \ \langle s, x \rangle = r\}$ $(s \neq 0, r \in \mathbb{R})$	vector hyperpl. $\{s\}^\perp$	$n - 1$

(c) The *half-spaces* of $\mathbb{R}^n$ are those sets attached to $(s, r) \in \mathbb{R}^n \times \mathbb{R}$, $s \neq 0$, and defined by

$$\{x \in \mathbb{R}^n \ : \ \langle s, x \rangle \leqslant r\} \quad \text{(closed half-space)}$$
$$\{x \in \mathbb{R}^n \ : \ \langle s, x \rangle < r\} \quad \text{(open half-space)};$$

"affine half-space" would be a more accurate terminology. We will often use the notation H^- for closed half-spaces. Naturally, an open [resp. closed] half-space is really an open [resp. closed] set; it is the interior [resp. closure] of the corresponding half-space; and the affine hyperplanes are the boundaries of the half-spaces; all this essentially comes from the continuity of the scalar product $\langle s, \cdot \rangle$.   □

Example 1.1.3 (Simplices) Call $\alpha = (\alpha_1, \ldots, \alpha_k)$ the generic point of the space $\mathbb{R}^k$. The *unit simplex* in $\mathbb{R}^k$ is

$$\Delta_k := \left\{ \alpha \in \mathbb{R}^k \ : \ \sum_{i=1}^k \alpha_i = 1, \ \alpha_i \geqslant 0 \text{ for } i = 1, \ldots, k \right\}.$$

Equipping $\mathbb{R}^k$ with the standard dot-product, $\{e_1, \ldots, e_k\}$ being the canonical basis and $e := (1, \ldots, 1)$ the vector whose coordinates are all 1, we can also write

$$\Delta_k := \left\{ \alpha \in \mathbb{R}^k \; : \; e^\top \alpha = 1, \; e_i^\top \alpha \geqslant 0 \text{ for } i = 1, \ldots, k \right\} . \qquad (1.1.1)$$

Observe the hyperplane and half-spaces appearing in this definition. Unit simplices are convex, compact, and have empty interior – being included in an affine hyperplane. We will often refer to a point $\alpha \in \Delta_k$ as a set of (k) *convex multipliers*.

It is sometimes useful to embed Δ_k in $\mathbb{R}^m$, $m > k$, by appending $m - k$ zeros to the coordinates of $\alpha \in \mathbb{R}^k$, thus obtaining a vector of Δ_m. We mention that a so-called *simplex* of $\mathbb{R}^n$ is the figure formed by $n + 1$ vectors in "nondegenerate positions"; in this sense, the unit simplex of $\mathbb{R}^k$ is a simplex in the affine hyperplane of equation $e^\top \alpha = 1$; see Fig. 1.1.1.
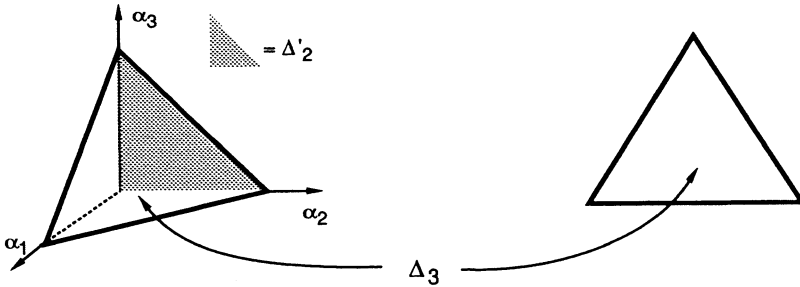
Fig. 1.1.1. Representing a simplex

If we replace $e^\top \alpha = 1$ in (1.1.1) by $e^\top \alpha \leqslant 1$, we obtain another important set, convex, compact, with nonempty interior:

$$\Delta_k' := \left\{ \alpha \in \mathbb{R}^k \; : \; e^\top \alpha \leqslant 1, \; \alpha_i \geqslant 0 \text{ for } i = 1, \ldots, k \right\} .$$

In fact, $\alpha \in \Delta_k'$ means that there is $\alpha_{k+1} \geqslant 0$ such that $(\alpha, \alpha_{k+1}) \in \Delta_{k+1}$. In this sense, the simplex $\Delta_k' \subset \mathbb{R}^k$ can be identified with Δ_{k+1} via a projection operator.

A (unit) simplex is traditionally visualized by a triangle, which can represent Δ_3 or Δ_2'; see Fig. 1.1.1 again. □

Example 1.1.4 (Convex Cones) A *cone* K is a set such that the "open" half-line $\{\alpha x : \alpha > 0\}$ is entirely contained in K whenever $x \in K$. In the usual representation of geometric objects, a cone has an apex; this apex is here at 0 (when it exists: a subspace is a cone but has no apex in this intuitive sense). Also, K is not supposed to contain 0 – this is mainly for notational reasons, to avoid writing $0 \times (+\infty)$ in some situations. A *convex cone* is of course a cone which is convex; an example is the set defined in $\mathbb{R}^n$ by

$$\langle s_j, x \rangle = 0 \text{ for } j = 1, \ldots, m, \quad \langle s_{m+j}, x \rangle \leqslant 0 \text{ for } j = 1, \ldots, p, \qquad (1.1.2)$$

where the s_j's are given in $\mathbb{R}^n$ (once again, observe the hyperplanes and the half-spaces appearing in the above example, observe also that the defining relations must have zero righthand sides).

Convexity of a given set is easier to check if this set is already known to be a cone: in view of Definition 1.1.1, a cone K is convex if and only if $x + y \in K$ whenever x and y lie in K, i.e. $K + K \subset K$. Subspaces are particular convex cones. We leave it as an exercise to show that, to become a subspace, what is missing from a convex cone is just symmetry ($K = -K$).

A very simple cone is the *nonnegative orthant* of $\mathbb{R}^n$

$$\Omega_+ := \left\{ x = (\xi^1, \ldots, \xi^n) \, : \, \xi^i \geqslant 0 \text{ for } i = 1, \ldots, n \right\}.$$

It can also be represented in terms of the canonical basis:

$$\Omega_+ = \left\{ \sum_{i=1}^{n} \alpha_i e_i \, : \, \alpha_i \geqslant 0 \text{ for } i = 1, \ldots, n \right\}$$

or, in the spirit of (1.1.2):

$$\Omega_+ = \left\{ x \in \mathbb{R}^n \, : \, \langle e_i, x \rangle \geqslant 0 \text{ for } i = 1, \ldots, n \right\}.$$

Convex cones will be of fundamental use in the sequel, as they are among the simplest convex sets. Actually, they are important in convex analysis (the "unilateral" realm of inequalities), just as subspaces are important in linear analysis (the "bilateral" realm of equalities).   □

1.2 Convexity-Preserving Operations on Sets

Proposition 1.2.1 *Let $\{C_j\}_{j \in J}$ be an arbitrary family of convex sets. Then their intersection $C := \cap \{C_j : j \in J\}$ is convex.*

Proof. Immediate from the very Definition 1.1.1.   □

Intersecting convex sets is an operation of utmost importance; by contrast, a union of convex sets is usually not convex.

Example 1.2.2 Let $(s_1, r_1), \ldots, (s_m, r_m)$ be m given elements of $\mathbb{R}^n \times \mathbb{R}$ and consider the set
$$\{ x \in \mathbb{R}^n \, : \, \langle s_j, x \rangle \leqslant r_j \text{ for } j = 1, \ldots, m \}. \tag{1.2.1}$$
It is clearly convex, which is confirmed if we view it as an intersection of m half-spaces; see Fig. 1.2.1.

We find it convenient to introduce two notations; $A : \mathbb{R}^n \to \mathbb{R}^m$ is the linear operator which, to $x \in \mathbb{R}^n$, associates the vector with coordinates $\langle s_j, x \rangle$; and in $\mathbb{R}^m$, the notation $a \leqslant b$ means that each coordinate of a is lower than or equal to the corresponding coordinate of b. Then, the set (1.2.1) can be characterized by $Ax \leqslant b$, where $b \in \mathbb{R}^m$ is the vector with coordinates $r_1, \ldots, r_m$.   □
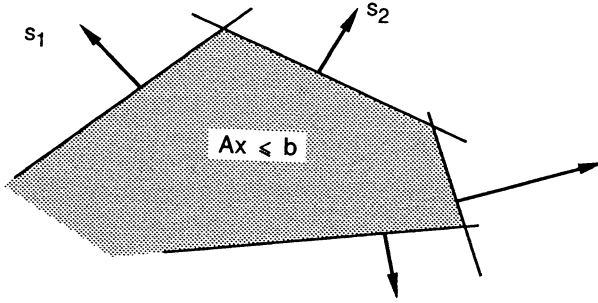
Fig. 1.2.1. An intersection of half-spaces

It is interesting to observe that the above construction applies to the examples of §1.1:
– an affine hyperplane is the intersection of two (closed) half-spaces;
– an affine manifold is the intersection of finitely many affine hyperplanes;
– a unit simplex is the intersection of an affine hyperplane with a closed convex cone;
– a convex cone such as in (1.1.2) is an intersection of a subspace with (homogeneous) half-spaces.

Piecing together these instances of convex sets, we see that they can all be considered as intersections of sufficiently many closed half-spaces. Another observation is that, up to a translation, a hyperplane is the simplest instance of a convex cone – apart from (linear) subspaces. Conclusion: *translations* (the key operations in the affine world), *intersections* and closed *half-spaces* are basic objects in convex analysis.

Convexity is stable under Cartesian product, just as it is under intersection.

Proposition 1.2.3 *For $i = 1, \dots, k$, let $C_i \subset \mathbb{R}^{n_i}$ be convex sets. Then $C_1 \times \cdots \times C_k$ is a convex set of $\mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_k}$.*

Proof. Straightforward. $\square$

The converse is also true; $C_1 \times \cdots \times C_k$ is convex if and only if each C_i is convex, and this results from the next property: stability under affine mappings. We recall that $A : \mathbb{R}^n \to \mathbb{R}^m$ is said to be *affine* when

$$A(\alpha x + (1 - \alpha)x') = \alpha A(x) + (1 - \alpha)A(x')$$

for all x and x' in $\mathbb{R}^n$ and all $\alpha \in \mathbb{R}$. This means that $x \mapsto A(x) - A(0)$ is linear, so an affine mapping can be characterized by a linear mapping A_0 and a point $y_0 := A(0) \in \mathbb{R}^m$:

$$A(x) = A_0 x + y_0 \quad \text{for all } x \in \mathbb{R}^m .$$

It goes without saying that images of affine manifolds under affine mappings are affine manifolds (hence the name!) So is the case as well for convex sets:

Proposition 1.2.4 *Let $A : \mathbb{R}^n \to \mathbb{R}^m$ be an affine mapping and C a convex set of $\mathbb{R}^n$. The image $A(C)$ of C under A is convex in $\mathbb{R}^m$.*

If D is a convex set of $\mathbb{R}^m$, the inverse image $\overset{-1}{A}(D) := \{x \in \mathbb{R}^n : A(x) \in D\}$ is convex in $\mathbb{R}^n$.

Proof. For x and x' in $\mathbb{R}^n$, the image under A of the segment $[x, x']$ is clearly the segment $[A(x), A(x')] \subset \mathbb{R}^m$. This proves the first claim, but also the second: indeed, if x and x' are such that $A(x)$ and $A(x')$ are both in the convex set D, then every point of the segment $[x, x']$ has its image in $[A(x), A(x')] \subset D$. □

Immediate consequences of this last result are:
- the opposite $-C$ of a convex set is convex;
- the sum (called direct sum, or *Minkowski* sum, denoted with the symbol $\oplus$ by some authors)

$$C_1 + C_2 := \{x = x_1 + x_2 \ : \ x_1 \in C_1,\ x_2 \in C_2\}$$

of two convex sets C_1 and C_2 is convex; when $C_2 = \{c_2\}$ is a singleton, we will sometimes use the lighter notation $C_1 + c_2$ for $C_1 + \{c_2\}$;
- more generally, if α_1 and α_2 are two real numbers, the set

$$\alpha_1 C_1 + \alpha_2 C_2 := \{\alpha_1 x_1 + \alpha_2 x_2 \ : \ x_1 \in C_1, x_2 \in C_2\} \tag{1.2.2}$$

is convex: it is the image of the convex set $C_1 \times C_2$ (Proposition 1.2.3) under the linear mapping sending $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$ to $\alpha_1 x_1 + \alpha_2 x_2 \in \mathbb{R}^n$.

We recall here that the sum of two closed sets need not be closed, unless one of the sets is compact; and convexity does not help: with $n = 2$, take for example

$$C_1 := \{(\xi, \eta) : \xi \geqslant 0,\ \eta \geqslant 0,\ \xi\eta \geqslant 1\} \quad \text{and} \quad C_2 := \mathbb{R} \times \{0\}.$$

Example 1.2.5 Let C be convex in $\mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ and take for A the projection from $\mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ onto $\mathbb{R}^{n_1}$, parallel to $\mathbb{R}^{n_2}$: $A(x, y) = x$, see Fig. 1.2.2.

Then realize that the "slice" of C along y

$$C(y) := \{x \in \mathbb{R}^{n_1} \ : \ (x, y) \in C\}$$

and the "shadow" of C over $\mathbb{R}^{n_1}$

$$C_1 := \{x \in \mathbb{R}^{n_1} \ : \ (x, y) \in C \text{ for some } y \in C\}$$

are convex. If, in particular, $C = C_1 \times C_2$ is a product, we obtain the converse to Proposition 1.2.3. □

We finish with a topological operation.

Proposition 1.2.6 *If C is convex, so are its interior* $\operatorname{int} C$ *and its closure* $\operatorname{cl} C$.
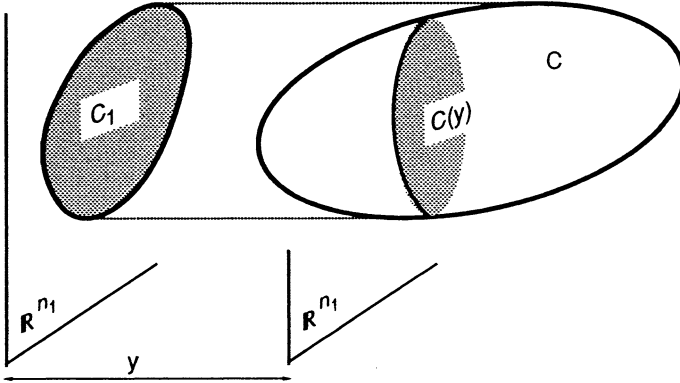
Fig. 1.2.2. Shadow and slice of a convex set

Proof. For given different x and x', and $\alpha \in]0, 1[$, we set $x'' = \alpha x + (1 - \alpha)x' \in]x, x'[$.

Take first x and x' in int C. Choosing $\delta > 0$ such that $B(x', \delta) \subset C$, we show that $B(x'', (1 - \alpha)\delta) \subset C$. As often in convex analysis, it is probably best to draw a picture. The ratio $\|x'' - x\|/\|x' - x\|$ being precisely $1 - \alpha$, Fig. 1.2.3 clearly shows that $B(x'', (1 - \alpha)\delta)$ is just the set $\alpha x + (1 - \alpha)B(x', \delta)$, obtained from segments with endpoints in int C: $x'' \in \text{int } C$.

Now, take x and x' in cl C: we select in C two sequences (x_k) and (x'_k) converging to x and x' respectively. Then, $\alpha x_k + (1 - \alpha)x'_k$ is in C and converges to x'', which is therefore in cl C. $\qquad\square$
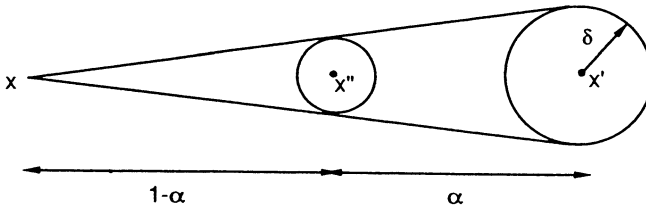
Fig. 1.2.3. Convex sets have convex interiors

The interior of a set is (too) often empty; convexity allows the similar but much more convenient concept of *relative interior*, to be seen below in §2.1. Observe the nonsymmetric character of x and x' in Fig. 1.2.3. It can be exploited to show that the intermediate result $]x, x'[\subset \text{int } C$ remains true even if $x \in \text{cl } C$; a property which will be seen in more detail in §2.1.

1.3 Convex Combinations and Convex Hulls

The operations described in §1.2 took convex sets and made new convex sets with them. The present section is devoted to another operation, which takes a nonconvex set and makes a convex set with it. First, let us recall the following basic facts from linear algebra.

(i) A *linear combination* of elements $x_1, \ldots, x_k$ of $\mathbb{R}^n$ is an element $\sum_{i=1}^{k} \alpha_i x_i$, where the coefficients α_i are arbitrary real numbers.

(ii) A (linear) *subspace* of $\mathbb{R}^n$ is a set containing all its linear combinations; an intersection of subspaces is still a subspace.

(iii) To any nonempty set $S \subset \mathbb{R}^n$, we can therefore associate the intersection of all subspaces containing S. This gives a subspace: the subspace generated by S (or *linear hull* of S), denoted by lin S – other notations are vect S or span S.

(iv) For the $\subset$-relation, lin S is the smallest subspace containing S; it can be constructed directly from S, by collecting all the linear combinations of elements of S.

(v) Finally, $x_1, \ldots, x_k$ are said to be linearly independent if $\sum_{i=1}^{k} \alpha_i x_i = 0$ implies $\alpha_1 = \cdots = \alpha_k = 0$. In $\mathbb{R}^n$, this implies $k \leqslant n$.

Now, let us be slightly more demanding for the coefficients α_i, as follows:

(i') An *affine combination* of elements $x_1, \ldots, x_k$ of $\mathbb{R}^n$ is an element $\sum_{i=1}^{k} \alpha_i x_i$, where the coefficients α_i satisfy $\sum_{i=1}^{k} \alpha_i = 1$.

As explained after Example 1.2.2, "affinity = linearity + translation"; it is therefore not surprising to realize that the development (i) – (v) can be reproduced starting from (i'):

(ii') An affine manifold in $\mathbb{R}^n$ is a set containing all its affine combinations (the equivalence with Example 1.1.2(b) will appear more clearly below in Proposition 1.3.3); it is easy to see that an intersection of affine manifolds is still an affine manifold.

(iii') To any nonempty set $S \subset \mathbb{R}^n$, we can therefore associate the intersection of all affine manifolds containing S. This gives the affine manifold generated by S, denoted aff S: the *affine hull* of S.

(iv') For the $\subset$-relation, aff S is the smallest affine manifold containing S; it can be constructed directly from S, by collecting all the affine combinations of elements of S. To see it, start from $x_0 \in S$, take lin $(S - x_0)$ and come back by adding x_0: the result $x_0 + \text{lin}\,(S - x_0)$ is just aff S.

(v') Finally, the $k + 1$ points $x_0, x_1, \ldots, x_k$ are said affinely independent if the set

$$x_0 + \text{lin}\,\{x_0 - x_0, x_1 - x_0, \ldots, x_k - x_0\} = x_0 + \text{lin}\,\{x_1 - x_0, \ldots, x_k - x_0\}$$

has full dimension, namely k. The above set is just aff $\{x_0, x_1, \ldots, x_k\}$; hence, it does not depend on the index chosen for the translation (here 0). In linear language, the required property is that the k vectors $x_i - x_0$, $i \neq 0$

be linearly independent. Getting rid of the arbitrary index 0, this means that the system of equations

$$\sum_{i=0}^{k} \alpha_i x_i = 0, \quad \sum_{i=0}^{k} \alpha_i = 0 \tag{1.3.1}$$

has the unique solution $\alpha_0 = \alpha_1 = \cdots = \alpha_k = 0$. Considered as elements of $\mathbb{R}^{n+1} = \mathbb{R}^n \times \mathbb{R}$, the vectors $(x_0, 1), (x_1, 1), \ldots, (x_k, 1)$, are linearly independent. In $\mathbb{R}^n$, at most $n + 1$ elements can thus be affinely independent.

If $x_0, x_1, \ldots, x_k$, are affinely independent, $x \in \operatorname{aff} \{x_0, x_1, \ldots, x_k\}$ can be written in a unique way as

$$x = \sum_{i=0}^{k} \alpha_i x_i \quad \text{with} \quad \sum_{i=0}^{k} \alpha_i = 1 .$$

The corresponding coefficients α_i are sometimes called the *barycentric coordinates* of x – even though such a terminology should be reserved to nonnegative α_i's. To say that a set of vectors are affinely dependent is to say that one of them (any one) is an affine combination of the others.

Example 1.3.1 Consider the unit simplex Δ_3 on the left part of Fig. 1.1.1; call $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$, $e_3 = (0, 0, 1)$ the three basis-vectors forming its vertices. The affine hull of $S = \{e_1, e_2\}$ is the affine line passing through e_1 and e_2. For $S = \{e_1, e_2, e_3\}$, it is the affine plane of equation $\alpha_1 + \alpha_2 + \alpha_3 = 1$. The four elements $0, e_1, e_2, e_3$ are affinely independent but the four elements $(1/3, 1/3, 1/3), e_1, e_2, e_3$ are not. □

Passing from (i) to (i') gives a set aff S which is closer to S than lin S, thanks to the extra requirement in (i'). We apply once more the same idea and we pass from affinity to convexity by requiring some more of the α_i's. This gives a new definition, playing the role of (i) and (i'):

Definition 1.3.2 A *convex combination* of elements $x_1, \ldots, x_k$ in $\mathbb{R}^n$ is an element of the form

$$\sum_{i=1}^{k} \alpha_i x_i , \quad \text{where} \quad \sum_{i=1}^{k} \alpha_i = 1 \text{ and } \alpha_i \geqslant 0 \text{ for } i = 1, \ldots, k . \qquad \square$$

A convex combination is therefore a particular affine combination, which in turn is a particular linear combination. Note in passing that all convex combinations of given $x_1, \ldots, x_k$ form a convex set: the image of Δ_k under the linear mapping

$$\mathbb{R}^k \ni (\alpha_1, \ldots, \alpha_k) \mapsto \alpha_1 x_1 + \cdots + \alpha_k x_k \in \mathbb{R}^n .$$

The sets playing the role of linear or affine subspaces of (ii) and (ii') will now be logically called convex, but we have to make sure that this new definition is consistent with Definition 1.1.1.

Proposition 1.3.3 *A set $C \subset \mathbb{R}^n$ is convex if and only if it contains every convex combination of its elements.*

Proof. The condition is sufficient: convex combinations of two elements just make up the segment joining them. To prove necessity, take $x_1, \ldots, x_k$ in C and $\alpha = (\alpha_1, \ldots, \alpha_k) \in \Delta_k$. One at least of the α_i's is positive, say $\alpha_1 > 0$. Then form

$$y_2 := \frac{\alpha_1}{\alpha_1 + \alpha_2} x_1 + \frac{\alpha_2}{\alpha_1 + \alpha_2} x_2 \qquad \left[= \frac{1}{\alpha_1 + \alpha_2} (\alpha_1 x_1 + \alpha_2 x_2) \right]$$

which is in C by Definition 1.1.1 itself. Therefore,

$$y_3 := \frac{\alpha_1 + \alpha_2}{\alpha_1 + \alpha_2 + \alpha_3} y_2 + \frac{\alpha_3}{\alpha_1 + \alpha_2 + \alpha_3} x_3 \qquad \left[= \frac{1}{\sum_{i=1}^3 \alpha_i} \sum_{i=1}^3 \alpha_i x_i \right]$$

is in C for the same reason; and so on until

$$y_k := \frac{\alpha_1 + \cdots + \alpha_{k-1}}{1} y_{k-1} + \frac{\alpha_k}{1} x_k \qquad \left[= \frac{1}{1} \sum_{i=1}^k \alpha_i x_i \right]. \qquad \square$$

The working argument of the above proof is longer to write than to understand. Its basic idea is just *associativity*: a convex combination $x = \sum \alpha_i x_i$ of convex combinations $x_i = \sum \beta_{ij} y_{ij}$ is still a convex combination $x = \sum \sum (\alpha_i \beta_{ij}) y_{ij}$. The same associativity property will be used in the next result.

Because an intersection of convex sets is convex, we can logically define as in (iii), (iii') the *convex hull* co S of a nonempty set S: this is the intersection of all the convex sets containing S.

Proposition 1.3.4 *The convex hull can also be described as the set of all convex combinations:*

$$\text{co}\, S := \cap \{C \ : \ C \text{ is convex and contains } S\}$$
$$= \big\{ x \in \mathbb{R}^n \ : \ \text{for some } k \in \mathbb{N}_*, \text{ there exist } x_1, \ldots, x_k \in S \text{ and} \qquad (1.3.2)$$
$$\alpha = (\alpha_1, \ldots, \alpha_k) \in \Delta_k \ \text{ such that } \sum_{i=1}^k \alpha_i x_i = x \big\}.$$

Proof. Call T the set described in the rightmost side of (1.3.2). Clearly, $T \supset S$. Also, if C is convex and contains S, then it contains all convex combinations of elements in S (Proposition 1.3.3), i.e. $C \supset T$. The proof will therefore be finished if we show that T is convex.

For this, take two points x and y in T, characterized respectively by $(x_1, \alpha_1), \ldots, (x_k, \alpha_k)$ and by $(y_1, \beta_1), \ldots, (y_\ell, \beta_\ell)$; take also $\lambda \in]0, 1[$. Then $\lambda x + (1 - \lambda) y$ is a certain combination of $k + \ell$ elements of S; this combination is convex because its coefficients $\lambda \alpha_i$ and $(1 - \lambda) \beta_j$ are nonnegative, and their sum is

$$\lambda \sum_{i=1}^k \alpha_i + (1 - \lambda) \sum_{j=1}^\ell b_j = \lambda + 1 - \lambda = 1. \qquad \square$$

Example 1.3.5 Take a finite set $\{x_1, \ldots, x_m\}$. To obtain its convex hull, it is not necessary to list all the convex combinations obtained via $\alpha \in \Delta_k$ for all $k = 1, \ldots, m$. In fact, as already seen in Example 1.1.3, $\Delta_k \subset \Delta_m$ if $k \leqslant m$, so we can restrict ourselves to $k = m$. Thus, we see that

$$\text{co}\,\{x_1,\ldots,x_m\} = \left\{ \sum_{j=1}^{m} \alpha_j x_j \;:\; \alpha = (\alpha_1,\ldots,\alpha_m) \in \Delta_m \right\}.$$

Make this example a little more complicated, replacing the collection of points by a collection of convex sets: $S = C_1 \cup \cdots \cup C_m$, where each C_i is convex. A simplification of (1.3.2) can again be exploited here. Indeed, consider a convex combination $\sum_{i=1}^{k} \alpha_i x_i$. It may happen that several of the x_i's belong to the same C_j. To simplify notation, suppose that x_{k-1} and x_k are in C_1; assume also $\alpha_k > 0$. Then set $(\beta_i, y_i) := (\alpha_i, x_i)$, $i = 1,\ldots,k-2$ and

$$\beta_{k-1} := \alpha_{k-1} + \alpha_k, \qquad y_{k-1} := \tfrac{1}{\beta_{k-1}}(\alpha_{k-1}x_{k-1} + \alpha_k x_k) \in C_1\,,$$

so that $\sum_{i=1}^{k} \alpha_i x_i = \sum_{i=1}^{k-1} \beta_i y_i$. Our convex combination (α, x) is useless, in the sense that it can also be found among those with $k-1$ elements. To cut a long story short, associativity of convex combinations yields

$$\text{co}\,S = \left\{ \sum_{i=1}^{m} \alpha_i x_i \;:\; \alpha \in \Delta_m,\ \ x_i \in C_i \text{ for } i = 1,\ldots,m \right\}.$$

From a geometric point of view, the convex hull of $C_1 \cup C_2$ ($m = 2$) is simply constructed by drawing segments, with endpoints in C_1 and C_2; for $C_1 \cup C_2 \cup C_3$, we smear triangles, etc. □

When S is infinite, or has infinitely many convex components, k is a priori unbounded in (1.3.2) and cannot be readily restricted as in the examples above. Yet, a bound on k exists for all S when we consider linear combinations and linear hulls – and consequently in the affine case as well; this is the whole business of *dimension*. In the present case of convex combinations, the same phenomenon is conserved to some extent. For each positive integer k, call S_k the set of all convex combinations of k elements in S: we have $S = S_1 \subset S_2 \subset \cdots \subset S_k \subset \cdots$ The S_k's are not convex but, "at the limit", their union is convex and coincides with $\text{co}\,S$ (Proposition 1.3.4). The theorem below tells us that k does not have to go to $+\infty$: the above sequence actually stops at $S_{n+1} = \text{co}\,S$.

Theorem 1.3.6 (C. Carathéodory) *Any $x \in \text{co}\,S \subset \mathbb{R}^n$ can be represented as a convex combination of $n+1$ elements of S.*

Proof. Take an arbitrary convex combination $x = \sum_{i=1}^{k} \alpha_i x_i$, with $k > n+1$. We claim that one of the x_i's can be assigned a 0-coefficient without changing x. For this, assume that all coefficients α_i are positive (otherwise we are done).

The $k > n+1$ elements x_i are certainly affinely dependent: (1.3.1) tells us that we can find $\beta_1,\ldots,\beta_k$, not all zero, such that $\sum_{i=1}^{k} \beta_i x_i = 0$ and $\sum_{i=1}^{k} \beta_i = 0$. There is at least one positive β_i and we can set $\alpha_i' := \alpha_i - t^*\beta_i$ for $i = 1,\ldots,k$, where

$$t^* := \max\,\{t \geqslant 0 \;:\; \alpha_i - t\beta_i \geqslant 0 \text{ for } i = 1,\ldots,k\} = \min_{\beta_j > 0} \frac{\alpha_j}{\beta_j}\,.$$

These computations are illustrated by Fig. 1.3.1. Clearly enough,

$$\alpha'_i \geqslant 0 \quad \text{for } i = 1, \ldots, k \qquad \begin{array}{l} \text{[automatic if } \beta_i \leqslant 0, \\ \text{by construction of } t^* \text{ if } \beta_i > 0] \end{array}$$

$$\sum_{i=1}^{k} \alpha'_i = \sum_{i=1}^{k} \alpha_i - t^* \sum_{i=1}^{k} \beta_i = 1 \, ;$$

$$\sum_{i=1}^{k} \alpha'_i x_i = x - t^* \sum_{i=1}^{k} \beta_i x_i = x \, ;$$

$$\alpha'_{i_0} = 0 \text{ for some } i_0 \, . \qquad \text{[by construction of } t^*]$$

In other words, we have expressed x as a convex combination of $k - 1$ among the x_i's; our claim is proved.

Now, if $k - 1 = n + 1$, the proof is finished. If not, we can apply the above construction to the convex combination $x = \sum_{i=1}^{k-1} \alpha'_i x_i$ and so on. The process can be continued until there remain only $n + 1$ elements (which may be affinely independent). $\qquad \square$
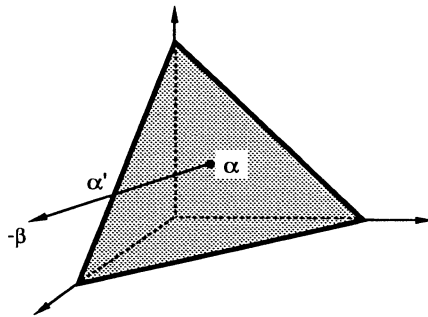
Fig. 1.3.1. Carathéodory's Theorem

The theorem of Carathéodory does not establish the existence of a "*basis*" with $n + 1$ elements, as is the case for linear combinations. Here, the generators x_i may depend on the particular x to be computed. In $\mathbb{R}^2$, think of the corners of a square: any one of these $4 > 2 + 1$ points may be necessary to generate a point in the square; also, the unit disk cannot be generated by finitely many points on the unit circle. By contrast, a subspace of dimension m can be generated by m (carefully selected but) *fixed* generators.

It is not the particular value $n + 1$ which is interesting in the above theorem, but rather the fact that the cardinality of relevant convex combinations is bounded: this is particularly useful when passing to the limit in a sequence of convex combinations. This value $n + 1$ is not of fundamental importance, anyway, and can often be reduced – as in Example 1.3.5: the convex hull of two convex sets in $\mathbb{R}^{100}$ can be generated by 2-combinations; also, the technique of proof shows that it is the dimension of the affine hull of S that counts, not n. Along these lines, we mention without proof a result geometrically very suggestive:

Theorem 1.3.7 (W. Fenchel and L. Bunt) *If $S \subset \mathbb{R}^n$ has no more than n connected components (in particular, if S is connected), then any $x \in \text{co } S$ can be expressed as a convex combination of n elements of S.* $\qquad \square$

This result says in particular that convex and connected one-dimensional sets are the same, namely the intervals. In $\mathbb{R}^2$, the convex hull of a continuous curve can be obtained by joining all pairs of points in it. In $\mathbb{R}^3$, the convex hull of three potatoes is obtained by pasting triangles, etc.

1.4 Closed Convex Sets and Hulls

Closedness is a very important property: most of the convex sets of interest to us in the subsequent chapters will be closed. It is therefore relevant to reproduce the previous section, with the word "closed" added. As far as linearity and affinity are concerned, there is no difference; in words, equalities are not affected when limits are involved. But convexity is another story: when passing from (i), (i') to Definition 1.3.2, inequalities are introduced, together with their accompanying difficulty "$<$ vs. $\leqslant$".

To construct a convex hull co S, we followed in §1.3 the path (iii), (iii'): we took the intersection of all convex sets containing S. An intersection of closed sets is still closed, so the following definition is also natural:

Definition 1.4.1 The *closed convex hull* of a nonempty set $S \subset \mathbb{R}^n$ is the intersection of all closed convex sets containing S. It will be denoted by $\overline{\text{co}}\, S$. ☐

Another path was also possible to construct co S, namely to take all possible convex combinations: then, we obtained co S again (Proposition 1.3.4); what about closing it? It turns out we can do that as well:

Proposition 1.4.2 The *closed convex hull* $\overline{\text{co}}\, S$ of Definition 1.4.1 is the closure cl (co S) of the convex hull of S.

Proof. Because cl (co S) is a closed convex set containing S, it contains $\overline{\text{co}}\, S$ as well. On the other hand, take a closed convex set C containing S; being convex, C contains co S; being closed, it contains also the closure of co S. Since C was arbitrary, we conclude $\cap C \supset$ cl co S. ☐

From the very definitions, the operation "taking a hull" is monotone: if $S_1 \subset S_2$, then aff $S_1 \subset$ aff S_2, cl $S_1 \subset$ cl S_2, co $S_1 \subset$ co S_2, and of course $\overline{\text{co}}\, S_1 \subset \overline{\text{co}}\, S_2$. A closed convex hull does not distinguish a set from its closure, just as it does not distinguish it from its convex hull: $\overline{\text{co}}\, S = \overline{\text{co}}(\text{cl}\, S) = \overline{\text{co}}(\text{co}\, S)$.

When computing $\overline{\text{co}}$ via Proposition 1.4.2, the closure operation is necessary (co S need not be closed) and must be performed *after* taking the convex hull: the operations do not commute. Consider the example of Fig. 1.4.1: $S = \{(0,0)\} \cup \{(\xi, 1) : \xi \geqslant 0\}$. It is a closed set but co S fails to be closed: it misses the half-line $(\mathbb{R}^+, 0)$. Nevertheless, this phenomenon can occur only when S is unbounded, a result which comes directly from Carathéodory's Theorem:

Theorem 1.4.3 *If S is bounded [resp. compact], then* co S *is bounded [resp. compact].*

Proof. Let $x = \sum_{i=1}^{n+1} \alpha_i x_i \in$ co S. If S is bounded, say by M, we can write

$$\|x\| \leqslant \sum_{i=1}^{n+1} \alpha_i \|x_i\| \leqslant M \sum_{i=1}^{n+1} \alpha_i = M \,.$$
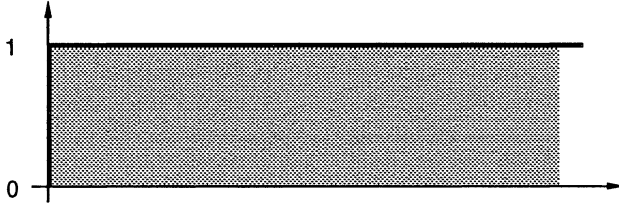
Fig. 1.4.1. A convex hull need not be closed

Now take a sequence $(x^k) \subset \operatorname{co} S$. For each k we can choose $x_1^k, \ldots, x_{n+1}^k$ in S and $\alpha^k \in \Delta_{n+1}$ such that $x^k = \sum_{i=1}^{n+1} \alpha_i^k x_i^k$. Note that Δ_{n+1} is compact. If S is compact, we can extract a subsequence as many times as necessary ($n+2$ times is enough) so that $(\alpha^k)_k$ and each $(x_i^k)_k$ converge: we end up with an index set $K \subset \mathbb{N}$ such that, when $k \to +\infty$,

$$(x_i^k)_{k \in K} \to x_i \in S \quad \text{and} \quad (\alpha^k)_{k \in K} \to \alpha \in \Delta_{n+1}.$$

Passing to the limit for $k \in K$, we see that $(x^k)_{k \in K}$ converges to a point x, which can be expressed as a convex combination of points of S: $x \in \operatorname{co} S$, whose compactness is thus established.   □

Thus, this theorem does allow us to write:

$$S \text{ bounded in } \mathbb{R}^n \quad \Longrightarrow \quad \overline{\operatorname{co}} S = \operatorname{cl} \operatorname{co} S = \operatorname{co} \operatorname{cl} S.$$

Remark 1.4.4 Let us emphasize one point made clear by this and the previous sections: a hull (linear, affine, convex or closed) can be constructed in two ways. In the *inner* way, combinations (linear, affine, convex, or limits) are made with points taken from inside the starting set S. The *outer* way takes sets (linear, affine, convex, or closed) containing S and intersects them.

Even though the first way may seem more direct and natural, it is the second which must often be preferred, at least when closedness is involved. This is especially true when taking the closed convex hull: forming all convex combinations is already a nasty task, which is not even sufficient, as one must close the result afterwards. On the other hand, the external construction of $\overline{\operatorname{co}} S$ is more handy in a set-theoretic framework. We will even see in §4.2(b) that it is not necessary to take in Definition 1.4.1 all closed convex sets containing S: only rather special such sets have to be intersected, namely the closed half-spaces of Example 1.1.2(c).   □

To finish this section, we mention one more hull, often useful. When starting from linear combinations to obtain convex combinations in Definition 1.3.2, we introduced two kinds of constraints on the coefficients: $e^\top \alpha = 1$ and $\alpha_i \geqslant 0$. The first constraint alone yielded affinity; we can take the second alone:

Definition 1.4.5 A *conical combination* of elements $x_1, \ldots, x_k$ is an element of the form $\sum_{i=1}^k \alpha_i x_i$, where the coefficients α_i are nonnegative.

The set of all conical combinations from a given nonempty $S \subset \mathbb{R}^n$ is the *conical hull* of S. It is denoted by $\operatorname{cone} S$.   □

Note that it would be more accurate to speak of convex conical combinations and convex conical hulls. If $\bar{\alpha} := \sum_{i=1}^{k} \alpha_i$ is positive, we can set $\beta_i := \alpha_i / \bar{\alpha}$; then we see that a conical combination of the type $\sum_{i=1}^{k} \alpha_i x_i = \bar{\alpha} \sum_{i=1}^{k} \beta_i x_i$, with $\bar{\alpha} > 0$ and $\beta \in \Delta_k$, is then nothing but a convex combination, multiplied by an arbitrary positive coefficient. We leave it to the reader to realize that

$$\operatorname{cone} S = \mathbb{R}^+ (\operatorname{co} S) = \operatorname{co} (\mathbb{R}^+ S).$$

Thus, $0 \in \operatorname{cone} S$; actually, to form $\operatorname{cone} S$, we intersect all convex cones containing S, and we append 0 to the result. If we close it, we obtain the following definition:

Definition 1.4.6 The *closed conical hull* (or rather closed convex conical hull) of a nonempty set $S \subset \mathbb{R}^n$ is

$$\overline{\operatorname{cone}} \, S := \operatorname{cl} \operatorname{cone} S = \operatorname{cl} \Big\{ \sum_{i=1}^{k} \alpha_i x_i \ : \ \alpha_i \geqslant 0, \ x_i \in S \text{ for } i = 1, \dots, k \Big\}. \qquad \square$$

Theorem 1.4.3 states that the convex hull and closed convex hull of a compact set co-incide, but the property is no longer true for conical hulls: for a counter-example, take the set $\{(\xi, \eta) \in \mathbb{R}^2 \ : \ (\xi - 1)^2 + \eta^2 \leqslant 1\}$. Nevertheless, the result can be recovered with an additional assumption:

Proposition 1.4.7 *Let S be a nonempty compact set such that $0 \notin \operatorname{co} S$. Then*

$$\overline{\operatorname{cone}} \, S = \mathbb{R}^+ (\operatorname{co} S) \quad [= \operatorname{cone} S].$$

Proof. The set $C := \operatorname{co} S$ is compact and does not containing the origin; we prove that $\mathbb{R}^+ C$ is closed. Let $(t_k x_k) \subset \mathbb{R}^+ C$ converge to y; extracting a subsequence if necessary, we may suppose $x_k \to x \in C$; note: $x \neq 0$. We write

$$t_k \frac{x_k}{\|x_k\|} \to \frac{y}{\|x\|},$$

which implies $t_k \to \|y\| / \|x\| =: t \geqslant 0$. Then, $t_k x_k \to tx = y$, which is thus in $\mathbb{R}^+ C$. $\qquad \square$

2 Convex Sets Attached to a Convex Set

2.1 The Relative Interior

Let C be a nonempty convex set in $\mathbb{R}^n$. If $\operatorname{int} C \neq \emptyset$, one easily checks that $\operatorname{aff} C$ is the whole of $\mathbb{R}^n$ (because so is the affine hull of a ball contained in C): we are dealing with a "full dimensional" set. On the other hand, let C be the sheet of paper on which this text is written. Its interior is empty in the surrounding space $\mathbb{R}^3$, but not in the space $\mathbb{R}^2$ of the table on which it is lying; by contrast, note that $\operatorname{cl} C$ is the same in both spaces.

This kind of ambiguity is one of the reasons for introducing the concept of *relative topology*: we recall that a subset A of $\mathbb{R}^n$ can be equipped with the topology relative to A, by defining its "closed balls" $B(x, \delta) \cap A$, for $x \in A$; then A becomes a topological space in its own. In convex analysis, the topology of $\mathbb{R}^n$ is of moderate interest: the topologies relative to *affine manifolds* turn out to be much richer.

Definition 2.1.1 The *relative interior* $\operatorname{ri} C$ (or $\operatorname{relint} C$) of a convex set $C \subset \mathbb{R}^n$ is the interior of C for the topology relative to the affine hull of C. In other words: $x \in \operatorname{ri} C$ if and only if

$$x \in \operatorname{aff} C \quad \text{and} \quad \exists \delta > 0 \text{ such that } (\operatorname{aff} C) \cap B(x, \delta) \subset C.$$

The *dimension* of a convex set C is the dimension of its affine hull, that is to say the dimension of the subspace parallel to $\operatorname{aff} C$. □

Thus, the wording "relative" implicitly means by convention "relative to the affine hull". Of course, note that $\operatorname{ri} C \subset C$. All along this section, and also later in Theorem C.2.2.3, we will see that $\operatorname{aff} C$ is the relevant working topological space. Already now, observe that our sheet of paper above can be moved *ad libitum* in $\mathbb{R}^3$ (but not folded: it would become nonconvex); its affine hull and relative interior move with it, but are otherwise unaltered. Indeed, the relative topological properties of C are the properties of convex sets in $\mathbb{R}^k$, where k is the dimension of C or $\operatorname{aff} C$. Table 2.1.1 gives some examples.

Table 2.1.1. Various relative interiors

C	aff C	dim C	ri C
$\{x\}$	$\{x\}$	0	$\{x\}$
$[x, x']$ $x \neq x'$	affine line generated by x and x'	1	$]x, x'[$
Δ_n	affine manifold of equation $e^\top \alpha = 1$	$n-1$	$\{\alpha \in \Delta_n : \alpha_i > 0\}$
$B(x_0, \delta)$	$\mathbb{R}^n$	n	$\operatorname{int} B(x_0, \delta)$

Remark 2.1.2 The cluster points of a set C are in $\operatorname{aff} C$ (which is closed and contains C), so the relative closure of C is just $\operatorname{cl} C$: a notation $\operatorname{relcl} C$ would be superfluous. On the contrary, the boundary is affected, and we will speak of *relative boundary*: $\operatorname{rbd} C := \operatorname{cl} C \setminus \operatorname{ri} C$. □

A first demonstration of the relevance of our new definition is the following:

Theorem 2.1.3 *If $C \neq \emptyset$, then $\operatorname{ri} C \neq \emptyset$. In fact, $\dim (\operatorname{ri} C) = \dim C$.*

Proof. Let $k := \dim C$. Then C contains $k + 1$ elements affinely independent $x_0, \ldots, x_k$. They generate a simplex $\operatorname{co} \{x_0, \ldots, x_k\} =: \Delta$; see Fig. 2.1.1; $\operatorname{aff} \Delta = \operatorname{aff} C$ because $\Delta \subset C$ and $\dim \Delta = k$. The proof will be finished if we show that Δ has nonempty relative interior.

Take $\bar{x} := \frac{1}{k+1} \sum_{i=0}^k x_i$ (the "center" of Δ) and describe $\operatorname{aff} \Delta$ by points of the form

$$\bar{x} + y = \bar{x} + \sum_{i=0}^k \alpha_i(y) x_i = \sum_{i=0}^k \left[\frac{1}{k+1} + \alpha_i(y) \right] x_i \,,$$

where $\alpha(y) = (\alpha_0(y), \ldots, \alpha_k(y)) \in \mathbb{R}^{k+1}$ solves
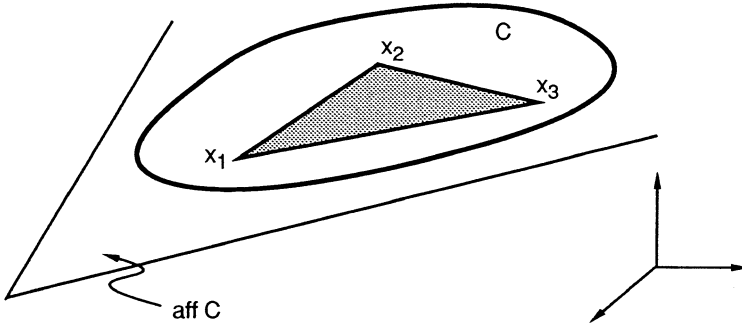
Fig. 2.1.1. A relative interior is nonempty

$$\sum_{i=0}^{k} \alpha_i x_i = y, \quad \sum_{i=0}^{k} \alpha_i = 0.$$

Because this system has a unique solution, the mapping $y \mapsto \alpha(y)$ is (linear and) continuous: we can find $\delta > 0$ such that $\|y\| \leqslant \delta$ implies $|\alpha_i(y)| \leqslant \frac{1}{k+1}$ for $i = 0, \ldots, k$, hence $\bar{x} + y \in \Delta$. In other words, $\bar{x} \in \mathrm{ri}\,\Delta \subset \mathrm{ri}\,C$.

It follows in particular $\dim \mathrm{ri}\,C = \dim \Delta = \dim C$.   □

Remark 2.1.4 We could have gone a little further in our proof, to realize that the relative interior of Δ was

$$\left\{ \sum_{i=0}^{k} \alpha_i x_i \ : \ \sum_{i=0}^{k} \alpha_i = 1, \ \alpha_i > 0 \text{ for } i = 0, \ldots, k \right\}.$$

Indeed, any point in the above set could have played the role of $\bar{x}$ in the proof. Note, incidentally, that the above set is still the relative interior of $\mathrm{co}\,\{x_0, \ldots, x_k\}$, even if the x_i's are not affinely independent.   □

Remark 2.1.5 The attention of the reader is drawn to a detail in the proof of Theorem 2.1.3: $\Delta \subset C$ implied $\mathrm{ri}\,\Delta \subset \mathrm{ri}\,C$ because Δ and C had the same affine hull, hence the same relative topology. Taking the relative interior is not a monotone operation, though: in $\mathbb{R}$, $\{0\} \subset [0, 1]$ but $\{0\} = \mathrm{ri}\,\{0\}$ is not contained in the relative interior $]0, 1[$ of $[0, 1]$.   □

We now turn to a *very useful* technical result; it refines the intermediate result in the proof of Proposition 1.2.6, illustrated by Fig. 1.2.3: when moving "radially" from a point in $\mathrm{ri}\,C$ straight to a point of $\mathrm{cl}\,C$, we stay inside $\mathrm{ri}\,C$.

Lemma 2.1.6 *Let $x \in \mathrm{cl}\,C$ and $x' \in \mathrm{ri}\,C$. Then the half-open segment*

$$]x, x'] = \{\alpha x + (1 - \alpha)x' \ : \ 0 \leqslant \alpha < 1\}$$

is contained in $\mathrm{ri}\,C$.

Proof. Take $x'' = \alpha x + (1 - \alpha)x'$, with $1 > \alpha \geqslant 0$. To avoid writing "$\cap \operatorname{aff} C$" every time, we assume without loss of generality that aff $C = \mathbb{R}^n$.

Since $x \in \operatorname{cl} C$, for all $\varepsilon > 0$, $x \in C + B(0, \varepsilon)$ and we can write

$$
\begin{aligned}
B(x'', \varepsilon) &= \alpha x + (1 - \alpha)x' + B(0, \varepsilon) \\
&\subset \alpha C + (1 - \alpha)x' + (1 + \alpha)B(0, \varepsilon) \\
&= \alpha C + (1 - \alpha)\left\{x' + B\left(0, \tfrac{1+\alpha}{1-\alpha}\varepsilon\right)\right\}.
\end{aligned}
$$

Since $x' \in \operatorname{int} C$, we can choose ε so small that $x' + B\left(0, \tfrac{1+\alpha}{1-\alpha}\varepsilon\right) \subset C$. Then we have

$$
B(x'', \varepsilon) \subset \alpha C + (1 - \alpha)C = C
$$

(where the last equality is just the definition of a convex set). $\qquad\square$

Taking $x \in \operatorname{cl} C$ in the above statement, we see that the relative interior of a convex set is convex; a result refining Proposition 1.2.6.

Remark 2.1.7 We mention an interesting consequence of this result: a half-line issued from $x' \in \operatorname{ri} C$ cannot cut the boundary of C in more than one point; hence, a line meeting $\operatorname{ri} C$ cannot cut $\operatorname{cl} C$ in more than two points: the relative boundary of a convex set is thus a fairly regular object, looking like an "onion skin" (see Fig. 2.1.2). $\qquad\square$
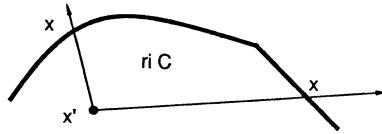
Fig. 2.1.2. The relative boundary of a convex set

Note in particular that $[x, x'] \subset \operatorname{ri} C$ whenever x and x' are in $\operatorname{ri} C$, which confirms that $\operatorname{ri} C$ is convex (cf. Proposition 1.2.6). Actually, $\operatorname{ri} C$, C and $\operatorname{cl} C$ are three convex sets very close together: they are *not distinguished* by the operations "aff", "ri" and "cl".

Proposition 2.1.8 *The three convex sets* $\operatorname{ri} C$, C *and* $\operatorname{cl} C$ *have the same affine hull (and hence the same dimension), the same relative interior and the same closure (and hence the same relative boundary).*

Proof. The case of the affine hull was already seen in Theorem 2.1.3. For the others, the key result is Lemma 2.1.6 (as well as for most other properties involving closures and relative interiors). We illustrate it by restricting our proof to one of the properties, say: $\operatorname{ri} C$ and C have the same closure.

Thus, we have to prove that $\operatorname{cl} C \subset \operatorname{cl}(\operatorname{ri} C)$. Let $x \in \operatorname{cl} C$ and take $x' \in \operatorname{ri} C$ (it is possible by virtue of Theorem 2.1.3). Because $]x, x'] \subset \operatorname{ri} C$ (Lemma 2.1.6), we do have that x is a limit of points in $\operatorname{ri} C$ (and even a "radial" limit); hence x is in the closure of $\operatorname{ri} C$. $\qquad\square$

Remark 2.1.9 This result gives one more argument in favour of our relative topology: if we take a closed convex set C, open it (for the topology of aff C), and close the result, we obtain C again – a very relevant topological property.

Among the consequences of Proposition 2.1.8, we mention the following:

– C and cl C have the same interior – hence the same boundary: in fact, either both interiors are empty (when dim $C = \dim \operatorname{cl} C < n$), or they coincide because the interior equals the relative interior.

– If C_1 and C_2 are two convex sets having the same closure, then they generate the same affine manifold and have the same relative interior. This happens exactly when we have the "sandwich" relation ri $C_1 \subset C_2 \subset \operatorname{cl} C_1$. □

Our relative topology fits rather well with the convexity-preserving operations presented in §1.2. Our first result in these lines concerns intersections and is of paramount importance.

Proposition 2.1.10 *Let the two convex sets C_1 and C_2 satisfy $\operatorname{ri} C_1 \cap \operatorname{ri} C_2 \neq \emptyset$. Then*

$$\operatorname{ri} (C_1 \cap C_2) = \operatorname{ri} C_1 \cap \operatorname{ri} C_2 \tag{2.1.1}$$

$$\operatorname{cl} (C_1 \cap C_2) = \operatorname{cl} C_1 \cap \operatorname{cl} C_2 . \tag{2.1.2}$$

Proof. First we show that $\operatorname{cl} C_1 \cap \operatorname{cl} C_2 \subset \operatorname{cl} (C_1 \cap C_2)$ (the converse inclusion is always true). Given $x \in \operatorname{cl} C_1 \cap \operatorname{cl} C_2$, we pick x' in the nonempty $\operatorname{ri} C_1 \cap \operatorname{ri} C_2$. From Lemma 2.1.6 applied to C_1 and to C_2,

$$]x, x'] \subset \operatorname{ri} C_1 \cap \operatorname{ri} C_2 .$$

Taking the closure of both sides, we conclude

$$x \in \operatorname{cl} (\operatorname{ri} C_1 \cap \operatorname{ri} C_2) \subset \operatorname{cl} (C_1 \cap C_2) ,$$

which proves (2.1.2) because x was arbitrary; the above inclusion is actually an equality.

Now, we have just seen that the two convex sets $\operatorname{ri} C_1 \cap \operatorname{ri} C_2$ and $C_1 \cap C_2$ have the same closure. According to Remark 2.1.9, they have the same relative interior:

$$\operatorname{ri} (C_1 \cap C_2) = \operatorname{ri} (\operatorname{ri} C_1 \cap \operatorname{ri} C_2) \subset \operatorname{ri} C_1 \cap \operatorname{ri} C_2 .$$

It remains to prove the converse inclusion, so let $y \in \operatorname{ri} C_1 \cap \operatorname{ri} C_2$. If we take $x' \in C_1$ [resp. C_2], the segment $[x', y]$ is in aff C_1 [resp. aff C_2] and, by definition of the relative interior, this segment can be stretched beyond y and yet stay in C_1 [resp. C_2] (see Fig. 2.1.3). Take in particular $x' \in \operatorname{ri} (C_1 \cap C_2)$, $x' \neq y$ (if such an x' does not exist, we are done). The above stretching singles out an $x \in C_1 \cap C_2$ such that $y \in]x, x'[$. Then Lemma 2.1.6 applied to $C_1 \cap C_2$ tells us that $y \in \operatorname{ri} (C_1 \cap C_2)$. □

Observe that, if we intersect infinitely many convex sets – instead of two, or a finite number –, the proof of (2.1.2) still works, but certainly not the proof of (2.1.1): the stretching possibility is killed. In (2.1.1), the lefthand side is nonempty (unless $C_1 \cap C_2 = \emptyset$), hence
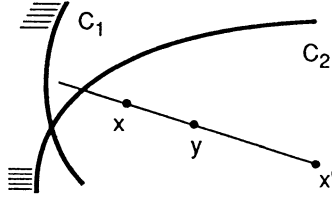
Fig. 2.1.3. The stretching mechanism

nonemptiness of ri $C_1 \cap$ ri C_2 appears as necessary; as for (2.1.2), look again at Remark 2.1.5. Such an assumption is thus very useful; it is usually referred to as a *qualification* assumption, and will be encountered many times in the sequel. Incidentally, it gives another sufficient condition for the monotonicity of the ri-operation (use (2.1.1) with $C_1 \subset C_2$).

We restrict our next statements to the case of the relative interior. Lemma 2.1.6 and Proposition 2.1.8 help in carrying them over to the closure operation.

Proposition 2.1.11 *For $i = 1, \ldots, k$, let $C_i \subset \mathbb{R}^{n_i}$ be convex sets. Then*

$$\text{ri}\,(C_1 \times \cdots \times C_k) = (\text{ri}\,C_1) \times \cdots \times (\text{ri}\,C_k)\,.$$

Proof. It suffices to apply Definition 2.1.1 alone, observing that

$$\text{aff}\,(C_1 \times \cdots \times C_k) = (\text{aff}\,C_1) \times \cdots \times (\text{aff}\,C_k)\,. \qquad \square$$

Proposition 2.1.12 *Let $A : \mathbb{R}^n \to \mathbb{R}^m$ be an affine mapping and C a convex set of $\mathbb{R}^n$. Then*

$$\text{ri}\,[A(C)] = A(\text{ri}\,C)\,. \tag{2.1.3}$$

If D is a convex set of $\mathbb{R}^m$ satisfying $\overset{-1}{A}(\text{ri}\,D) \neq \emptyset$, then

$$\text{ri}\big[\overset{-1}{A}(D)\big] = \overset{-1}{A}(\text{ri}\,D)\,. \tag{2.1.4}$$

Proof. First, note that the continuity of A implies $A(\text{cl}\,S) \subset \text{cl}\,[A(S)]$ for any $S \subset \mathbb{R}^n$. Apply this result to ri C, whose closure is cl C (Proposition 2.1.8), and use the monotonicity of the closure operation:

$$A(C) \subset A(\text{cl}\,C) = A[\text{cl}\,(\text{ri}\,C)] \subset \text{cl}\,[A(\text{ri}\,C)] \subset \text{cl}\,[A(C)]\,;$$

the closed set cl $[A(\text{ri}\,C)]$ is therefore cl $[A(C)]$. Because $A(\text{ri}\,C)$ and $A(C)$ have the same closure, they have the same relative interior (Remark 2.1.9): ri $A(C) = $ ri $[A(\text{ri}\,C)] \subset A(\text{ri}\,C)$.

To prove the converse inclusion, let $w = A(y) \in A(\text{ri}\,C)$, with $y \in \text{ri}\,C$. We choose $z' = A(x') \in \text{ri}\,A(C)$, with $x' \in C$ (we assume $z' \neq w$, hence $x' \neq y$). Using in C the same stretching mechanism as in Fig. 2.1.3, we single out $x \in C$ such that $y \in]x, x'[$, to which corresponds $z = A(x) \in A(C)$. By affinity, $A(y) \in]A(x), A(x')[=]z, z'[$. Thus, z and z' fulfil the conditions of Lemma 2.1.6 applied to the convex set $A(C)$: $w \in \text{ri}\,A(C)$, and (2.1.3) is proved.

The proof of (2.1.4) uses the same technique. $\qquad \square$

As an illustration of the last two results, we see that the relative interior of $\alpha_1 C_1 + \alpha_2 C_2$ is $\alpha_1 \operatorname{ri} C_1 + \alpha_2 \operatorname{ri} C_2$. If we take in particular $\alpha_1 = -\alpha_2 = 1$, we obtain the following theorem:

$$0 \in \operatorname{ri}(C_1 - C_2) \quad \Longleftrightarrow \quad (\operatorname{ri} C_1) \cap (\operatorname{ri} C_2) \neq \emptyset, \tag{2.1.5}$$

which gives one more equivalent form for the qualification condition in Proposition 2.1.10. We will come again to this property on several occasions.

2.2 The Asymptotic Cone

Let x be a point in a closed convex cone K. Draw a picture to see that, for all $d \in K$, the half-line $x + \mathbb{R}^+ d$ is contained in K: $x + td \in K$ for all $t > 0$. Conversely, if $x + \mathbb{R}^+ d \subset K$, i.e. if

$$d \in \frac{K - x}{t} = K - \{\tfrac{1}{t}x\} \quad \text{for all } t > 0,$$

then (K is closed), $d \in K$. In words, a closed convex cone is also the set of directions along which one can go straight to infinity. We now generalize this concept to non-conical sets.

In this section, C will always be a nonempty *closed convex* set. For $x \in C$, let

$$C_\infty(x) := \{d \in \mathbb{R}^n \ : \ x + td \in C \text{ for all } t > 0\}. \tag{2.2.1}$$

Despite the appearances, $C_\infty(x)$ depends only on the behaviour of C "at infinity": in fact, $x + td \in C$ implies that $x + \tau d \in C$ for all $\tau \in [0, t]$ (C is convex). Thus, $C_\infty(x)$ is just the set of directions from which one can go straight from x to infinity, while staying in C. Another formulation is:

$$C_\infty(x) = \bigcap_{t > 0} \frac{C - x}{t}, \tag{2.2.2}$$

which clearly shows that $C_\infty(x)$ is a *closed convex cone*, which of course contains 0. The following property is fundamental.

Proposition 2.2.1 *The closed convex cone $C_\infty(x)$ does not depend on $x \in C$.*

Proof. Take two different points x_1 and x_2 in C; it suffices to prove one inclusion, say $C_\infty(x_1) \subset C_\infty(x_2)$. Let $d \in C_\infty(x_1)$ and $t > 0$, we have to prove $x_2 + td \in C$. With $\varepsilon \in]0, 1[$, consider the point

$$\bar{x}_\varepsilon := x_1 + td + (1 - \varepsilon)(x_2 - x_1) = \varepsilon\left(x_1 + \tfrac{t}{\varepsilon}d\right) + (1 - \varepsilon)x_2.$$

Thus $\bar{x}_\varepsilon \in C$ (use the definitions of $C_\infty(x_1)$ and of a convex set); pass to the limit:

$$x_2 + td = \lim_{\varepsilon \downarrow 0} \bar{x}_\varepsilon \in \operatorname{cl} C = C. \qquad \square$$

It follows that the notation C_∞ is more appropriate:

Definition 2.2.2 (Asymptotic Cone) The *asymptotic cone*, or recession cone of the closed convex set C is the closed convex cone C_∞ defined by (2.2.1) or (2.2.2), in which Proposition 2.2.1 is exploited. $\qquad \square$

Figure 2.2.1 gives three examples in $\mathbb{R}^2$. As for the asymptotic cone of Example 1.2.2, it is the set $\{d \in \mathbb{R}^n : Ad \leqslant 0\}$.

Proposition 2.2.3 *A closed convex set C is compact if and only if $C_\infty = \{0\}$.*

$C =$

$C_\infty =$
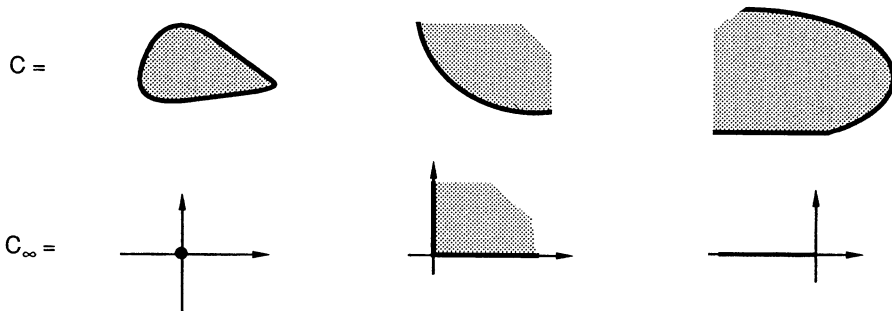
Fig. 2.2.1. Some asymptotic cones

Proof. If C is bounded, it is clear that C_∞ cannot contain any nonzero direction. Conversely, let $(x_k) \subset C$ be such that $\|x_k\| \to +\infty$ (we assume $x_k \neq 0$). The sequence $(d_k := x_k/\|x_k\|)$ is bounded, extract a convergent subsequence: $d = \lim_{k \in K} d_k$ with $K \subset \mathbb{N}$ ($\|d\| = 1$). Now, given $x \in C$ and $t > 0$, take k so large that $\|x_k\| \geqslant t$. Then, we see that

$$x + td = \lim_{k \in K} \left[\left(1 - \tfrac{t}{\|x_k\|}\right)x + \tfrac{t}{\|x_k\|}x_k \right]$$

lies in the closed convex set C, hence $d \in C_\infty$.   □

Remark 2.2.4 Consider the closed convex sets $(C - x)/t$, indexed by $t > 0$. They form a nested decreasing family: for $t_1 < t_2$ and y arbitrary in C,

$$\frac{y - x}{t_2} = \frac{y' - x}{t_1} \quad \text{where} \quad y' := \frac{t_2 - t_1}{t_2}x + \frac{t_1}{t_2}y \in C.$$

Thus, we can write (see §0.5 for the set-limit appearing below)

$$C_\infty = \bigcap_{t>0} \frac{C - x}{t} = \lim_{t \to +\infty} \frac{C - x}{t}, \tag{2.2.3}$$

which interprets C_∞ as a limit of set-valued difference quotients, but with the denominator tending to ∞, instead of the usual 0. This will be seen again later in §5.2.   □

In contrast to the relative interior, the concept of asymptotic cone does not always fit well with usual convexity-preserving operations. We just mention some properties which result directly from the definition of C_∞.

Proposition 2.2.5 – *If $\{C_j\}_{j \in J}$ is a family of closed convex sets having a point in common, then*

$$\left(\cap_{j \in J} C_j\right)_\infty = \cap_{j \in J} (C_j)_\infty.$$

– *If, for $j = 1, \ldots, m$, C_j are closed convex sets in $\mathbb{R}^{n_j}$, then*

$$\left(C_1 \times \cdots \times C_m\right)_\infty = (C_1)_\infty \times \cdots \times (C_m)_\infty.$$

– *Let $A : \mathbb{R}^n \to \mathbb{R}^m$ be a linear operator. If C is closed convex in $\mathbb{R}^n$ and $A(C)$ is closed, then*

$$A(C_\infty) \subset [A(C)]_\infty.$$

– *If D is closed convex in $\mathbb{R}^m$ with nonempty inverse image, then*

$$\left[\overset{-1}{A}(D)\right]_\infty = \overset{-1}{A}(D_\infty).$$ □

Needless to say, convexity does not help to ensure that the image of a closed set under a continuous mapping is closed: take $A(\xi, \eta) = \xi$ (linear) and $C = \{(\xi, \eta) : \eta \geqslant 1/\xi > 0\}$.

2.3 Extreme Points

In this section, C is a nonempty convex set of $\mathbb{R}^n$ and there would be no loss of generality in assuming that it is closed. The reader may make this assumption if he finds it helpful in mastering faster the definitions and properties below; the same remark holds for §2.4.

Definition 2.3.1 (Extreme Point) We say that $x \in C$ is an *extreme point* of C if there are no two different points x_1 and x_2 in C such that $x = 1/2\,(x_1 + x_2)$.   □

Some other ways of expressing the same thing are:
– The representation $x = \alpha x_1 + (1 - \alpha)x_2$ is impossible with two distinct points x_1 and x_2 in C and $\alpha \in\,]0, 1[$: indeed, convexity of C implies that x_1 and x_2 in the definition can be replaced by two other points in the segment $[x_1, x_2]$; this amounts to replacing the number $1/2$ by some other $\alpha \in\,]0, 1[$. In short:

$$x \text{ is an extreme point of } C \text{ if and only if}$$
$$[x = \alpha x_1 + (1 - \alpha)x_2,\ x_i \in C,\ 0 < \alpha < 1] \implies x = x_1 = x_2 .$$

– There is no convex combination $x = \sum_{i=1}^k \alpha_i x_i$ other than $x_1 = \cdots = x_k\ [= x]$.
– Even when deprived of the point x, the set C remains convex.

Examples 2.3.2
– Let C be the unit ball $B(0, 1)$. Multiply by $1/2$ the relation

$$\tfrac{1}{2}\|x_1 + x_2\|^2 = \|x_1\|^2 + \|x_2\|^2 - \tfrac{1}{2}\|x_2 - x_1\|^2 \qquad (2.3.1)$$

to realize that every x of norm 1 is an extreme point of $B(0, 1)$. Likewise, if $Q : \mathbb{R}^n \to \mathbb{R}^n$ is a positive definite symmetric linear operator, any x with $\langle Qx, x\rangle = 1$ is an extreme point of the convex set $\{x \in \mathbb{R}^n : \langle Qx, x\rangle \leqslant 1\}$.
On the other hand, if $\langle Q\cdot, \cdot\rangle^{1/2}$ is replaced by the ℓ_1-norm, the corresponding unit ball has finitely many extreme points.
– If C is a convex cone, a nonzero $x \in C$ has no chance of being an extreme point.
– An affine manifold, a half-space have no extreme points.   □

The *set of extreme points* of C will be denoted by $\operatorname{ext} C$. We mention here that it is a closed set when $n \leqslant 2$; but in general, $\operatorname{ext} C$ has no particular topological or linear properties. Along the lines of the above examples, there is at least one case where there exist extreme points:

Proposition 2.3.3 *If C is compact, then $\operatorname{ext} C \neq \emptyset$.*

Proof. Because C is compact, there is $\bar{x} \in C$ maximizing the continuous function $x \mapsto \|x\|^2$. We claim that $\bar{x}$ is extremal. In fact, suppose that there are x_1 and x_2 in C with $\bar{x} = 1/2\,(x_1 + x_2)$. Then, with $x_1 \neq x_2$ and using (2.3.1), we obtain the contradiction

$$\|\bar{x}\|^2 = \left\|\tfrac{1}{2}(x_1 + x_2)\right\|^2 < \tfrac{1}{2}(\|x_1\|^2 + \|x_2\|^2) \leqslant \tfrac{1}{2}(\|\bar{x}\|^2 + \|\bar{x}\|^2) = \|\bar{x}\|^2 . \quad \square$$

The definitions clearly imply that any extreme point of C is on its boundary, and even on its relative boundary. The essential result on extreme points is the following, which we will prove later in §4.2(c).

Theorem 2.3.4 (H. Minkowski) *Let C be compact, convex in $\mathbb{R}^n$. Then C is the convex hull of its extreme points:* $C = \mathrm{co}\,(\mathrm{ext}\,C)$. $\qquad\qquad \square$

Combined with Carathéodory's Theorem 1.3.6, this result establishes that, if $\dim C = k$, then any element of C is a convex combination of at most $k + 1$ extreme points of C.

Example 2.3.5 Take $C = \mathrm{co}\,\{x_1, \ldots, x_m\}$. All the extreme points of C are present in the list $x_1, \ldots, x_m$; but of course, the x_i's are not all necessarily extremal. Let $\mu \leqslant m$ be the number of extreme points of C, suppose to simplify that these are $x_1, \ldots, x_\mu$. Then $C = \mathrm{co}\,\{x_1, \ldots, x_\mu\}$ and this representation is *minimal*, in the sense that removing one of the generators $x_1, \ldots, x_\mu$ effectively changes C. The case $\mu = n + 1$ corresponds to a simplex in $\mathbb{R}^n$. If $\mu > n + 1$, then for any $x \in C$, there is a representation $x = \sum_{i=1}^{\mu} \alpha_i x_i$ in which at least $\mu - n - 1$ among the $\alpha_i's$ are zero. $\qquad\qquad \square$

A higher-dimensional generalization of extreme points can be defined. Consider again Definition 2.3.1, and replace "the point $x \in C$" by "the convex subset $F \subset C$". Our definition is then generalized as follows: the convex subset $F \subset C$ is *extremal* if there are no two points x_1 and x_2 in $C\backslash F$ such that $1/2\,(x_1 + x_2) \in F$.

Once again, the number $1/2$ has nothing special and can be replaced by any other $\alpha \in \,]0, 1[$. The above statement can be rephrased in reversed logic as: if x_1 and x_2 in C are such that $\alpha x_1 + (1 - \alpha)x_2 \in F$ for some $\alpha \in \,]0, 1[$, then x_1 and x_2 are in F as well. Convexity of F then implies that the whole segment $[x_1, x_2]$ is in F, and we end up with the traditional definition:

Definition 2.3.6 (Faces) A nonempty convex subset $F \subset C$ is a *face* of C if it satisfies the following property: every segment of C, having in its relative interior an element of F, is entirely contained in F. In other words,

$$\left.\begin{array}{r} (x_1, x_2) \in C \times C \quad \text{and} \\ \exists \alpha \in \,]0, 1[\,:\ \alpha x_1 + (1 - \alpha)x_2 \in F \end{array}\right\} \implies [x_1, x_2] \subset F . \qquad (2.3.2)$$
$$\square$$

Being convex, a face has its own affine hull, closure, relative interior and dimension. By construction, extreme points appear as faces that are singletons, i.e. 0-dimensional faces:

$$x \in \text{ext } C \quad \Longleftrightarrow \quad \{x\} \text{ is a face of } C.$$

One-dimensional faces, i.e. segments that are faces of C, are called *edges* of C; and so on until $(k - 1)$-dimensional faces (where $k = \dim C$), called *facets* ... and the only k-dimensional face of C, which is C itself.

A useful property is the "transmission of extremality": if $x \in C' \subset C$ is an extreme point of C, then it is a fortiori an extreme point of the smaller set C'. When C' is a face of C, the converse is also true:

Proposition 2.3.7 *Let F be a face of C. Then any extreme point of F is an extreme point of C.*

Proof. Take $x \in F \subset C$ and assume that x is not an extreme point of C: there are different x_1, x_2 in C and $\alpha \in {]}0, 1{[}$ such that $x = \alpha x_1 + (1 - \alpha)x_2 \in F$. From the very definition (2.3.2) of a face, this implies that x_1 and x_2 are in F: x cannot be an extreme point of F. $\square$

This property can be generalized to: if F' is a face of F, which is itself a face of C, then F' is a face of C. We mention also: the relative interiors of the faces of C form a partition of C. Examine Example 2.3.5 to visualize its faces, their relative interiors and the above partition. The C of Fig. 2.3.1, with its extreme point x, gives a less trivial situation; make a three-dimensional convex set by rotating C around the axis Δ: we obtain a set with no one-dimensional face.
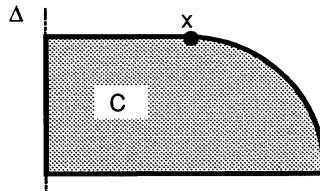
Fig. 2.3.1. A special extreme point

2.4 Exposed Faces

The rationale for extreme points is an inner construction of convex sets, as is particularly illustrated by Theorem 2.3.4 and Example 2.3.5. We mentioned in the important Remark 1.4.4 that a convex set could also be constructed externally, by taking intersections of convex sets containing it (see Proposition 1.3.4: if S is convex, then $S = \text{co } S$). To prepare a deeper analysis, coming in §4.2(b) and §5.2, we need the following fundamental definition, based on Example 1.1.2.

Definition 2.4.1 (Supporting Hyperplane) An affine hyperplane $H_{s,r}$ is said to *support* the set C when C is entirely contained in one of the two closed half-spaces delimited by $H_{s,r}$: say

$$\langle s, y \rangle \leqslant r \quad \text{for all } y \in C. \tag{2.4.1}$$

It is said to support C *at* $x \in C$ when, in addition, $x \in H_{s,r}$: (2.4.1) holds, as well as $\langle s, x \rangle = r$. $\quad\square$

See Fig. 2.4.1 for an illustration. Up to now, it is only a formal definition; existence of some supporting hyperplane will be established later in §4.2(a). Naturally, the inequality-sign could be reversed in (2.4.1): $H_{s,r}$ supports C when $H_{-s,-r}$ supports C. Note also that if $x \in C$ has a hyperplane supporting C, then $x \in \operatorname{bd} C$.
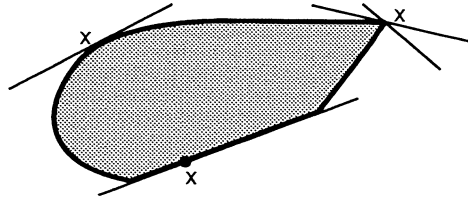
Fig. 2.4.1. Supporting hyperplanes at various points

Definition 2.4.2 (Exposed Faces, Vertices) The set $F \subset C$ is an *exposed face* of C if there is a supporting hyperplane $H_{s,r}$ of C such that $F = C \cap H_{s,r}$.

An *exposed point*, or *vertex*, is a 0-dimensional exposed face, i.e. a point $x \in C$ at which there is a supporting hyperplane $H_{s,r}$ of C such that $H_{s,r} \cap C$ reduces to $\{x\}$. $\quad\square$

See Fig. 2.4.1 again. A supporting hyperplane $H_{s,r}$ may or may not touch C. If it does, the contact-set is an exposed face. If it does at a singleton, this singleton is called an exposed point. As an intersection of convex sets, an exposed face is convex. The next result justifies the wording.

Proposition 2.4.3 *An exposed face is a face.*

Proof. Let F be an exposed face, with its associated support $H_{s,r}$. Take x_1 and x_2 in C:

$$\langle s, x_i \rangle \leqslant r \quad \text{for } i = 1, 2; \tag{2.4.2}$$

take also $\alpha \in \,]0, 1[$ such that $\alpha x_1 + (1 - \alpha) x_2 \in F \subset H_{s,r}$:

$$\langle s, \alpha x_1 + (1 - \alpha) x_2 \rangle = r.$$

Suppose that one of the relations (2.4.2) holds as strict inequality. By convex combination $(0 < \alpha < 1!)$, we obtain the contradiction $\langle s, \alpha x_1 + (1 - \alpha) x_2 \rangle < r$. $\quad\square$

The simple technique used in the above proof appears often in convex analysis: if a convex combination of inequalities holds as an equality, then so does each individual inequality whose coefficient is positive.

Remark 2.4.4 Comparing with Proposition 2.3.7, we see that the property of transmission of extremality applies to exposed faces as well: if x is an extreme point of the exposed face $F \subset C$, then $x \in \text{ext } C$.   □

One could believe (for example from Fig. 2.4.1) that the converse to Proposition 2.4.3 is true. Figure 2.3.1 immediately shows that this intuition is false: the extreme point x is not exposed. Exposed faces form therefore a proper subset of faces. The difference is slight, however: a result of S. Straszewicz (1935) establishes that any extreme point of a closed convex set C is a limit of exposed points in C. In other words,

$$\exp C \subset \text{ext } C \subset \text{cl} \, (\exp C)$$

if $\exp C$ denotes the set of exposed points in C. Comparing with Minkowski's result 2.3.4, we see that $C = \overline{\text{co}}(\exp C)$ for C convex and compact. We also mention that a facet is automatically exposed (the reason is that $n-1$, the dimension of a facet, is also the dimension of the hyperplane involved when exposing faces).

Enrich Fig. 2.3.1 as follows: take C', obtained from C by a rotation of $30°$ around Δ ; then consider the convex hull of $C \cup C'$, displayed in Fig. 2.4.2. The point P_1 is an extreme point but not a vertex; P_2 is a vertex. The edge E_1 is not exposed; E_2 is an exposed edge. As for the faces F_1 and F_2, they are exposed because they are facets.
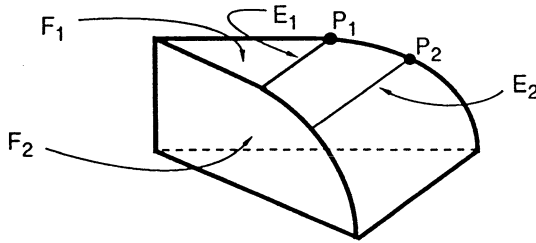
Fig. 2.4.2. Faces and exposed faces

Remark 2.4.5 (Direction Exposing a Face) Let F be an exposed face, and $H_{s,r}$ its associated supporting hyperplane. It results immediately from the definitions that $\langle s, y \rangle \leqslant \langle s, x \rangle$ for all $y \in C$ and all $x \in F$. Another definition of an exposed face can therefore be proposed, as the set of maximizers over C of some linear form: F is an exposed face of C when there is a nonzero $s \in \mathbb{R}^n$ such that

$$F = \left\{ x \in C \, : \, \langle s, x \rangle = \sup_{y \in C} \langle s, y \rangle \right\} . \tag{2.4.3}$$

A relevant notation is thus $F_C(s)$ to designate the exposed face of C associated with $s \in \mathbb{R}^n$; it can also be called the face of C *exposed by* s. For a unified notation, we will consider C itself as exposed by 0: $C = F_C(0)$.   □

Beware that a given s may define no supporting hyperplane at all. Even if it does, it may expose no face (the supremum in (2.4.3) may be not attained). The following result is almost trivial, but very useful: it is "equivalent" to extremize a linear form on a compact set or on its convex hull.

Proposition 2.4.6 *Let C be convex and compact. For $s \in \mathbb{R}^n$, there holds*

$$\max_{x \in C} \langle s, x \rangle = \max_{x \in \text{ext } C} \langle s, x \rangle.$$

Furthermore, the solution-set of the first problem is the convex hull of the solution-set of the second:

$$\text{Argmax}_{x \in C} \langle s, x \rangle = \text{co} \left\{ \text{Argmax}_{x \in \text{ext } C} \langle s, x \rangle \right\}.$$

Proof. Because C is compact, $\langle s, \cdot \rangle$ attains its maximum on $F_C(s)$. The latter set is convex and compact, and as such is the convex hull of its extreme points (Minkowski's Theorem 2.3.4); these extreme points are also extreme in C (Proposition 2.3.7 and Remark 2.4.4). □

3 Projection onto Closed Convex Sets

3.1 The Projection Operator

Denote by p_V the (orthogonal) projection onto a subspace $V \subset \mathbb{R}^n$. The main properties of the operator $x \mapsto p_V(x)$ are to be linear, symmetric, positive semidefinite, idempotent ($p_V \circ p_V = p_V$), nonexpansive ($\|p_V(x)\| \leqslant \|x\|$ for all x); also, it defines a canonical decomposition of $\mathbb{R}^n$ via $x = p_V(x) + p_{V^\perp}(x)$. We will generalize this operator to the case where V is merely convex.

In what follows, C is a *nonempty closed convex* set in $\mathbb{R}^n$. For fixed $x \in \mathbb{R}^n$, we consider the following problem:

$$\inf \left\{ \tfrac{1}{2} \|y - x\|^2 \ : \ y \in C \right\}, \tag{3.1.1}$$

i.e. we are interested in those points (if any) of C that are closest to x for the Euclidean distance. Let $f_x : \mathbb{R}^n \to \mathbb{R}$ be the function defined by

$$\mathbb{R}^n \ni y \mapsto f_x(y) := \tfrac{1}{2} \|y - x\|^2. \tag{3.1.2}$$

For $c \in C$, take the sublevel-set $S := \{y \in \mathbb{R}^n \ : \ f_x(y) \leqslant f_x(c)\}$. Then (3.1.1) is clearly equivalent to $\inf_{y \in C \cap S} f_x(y)$, which has a solution since f_x is continuous and S – hence $C \cap S$ – is compact. We deduce the *existence* of a closest point in C to x; the inf in (3.1.1) is a min.

Note that convexity of C plays no role in the above existence result. Uniqueness, however, depends crucially on convexity: let y_1 and y_2 be two solutions to (3.1.1). Use (2.3.1) with $x_i = y_i - x$ to obtain

$$f_x(y_0) = \tfrac{1}{2}[f_x(y_1) + f_x(y_2)] - \tfrac{1}{8} \|y_2 - y_1\|^2,$$

where $y_0 := 1/2 \, (y_1 + y_2) \in C$; this implies *uniqueness*.

We have thus defined a *projection operator*, namely the mapping $x \mapsto p_C(x)$ which, to each $x \in \mathbb{R}^n$, associates the unique solution $p_C(x)$ of the minimization problem (3.1.1). It is possible to *characterize* $p_C(x)$ differently, as solving a so-called *variational inequality*; and this characterization is the key to all results concerning p_C.

Theorem 3.1.1 *A point $y_x \in C$ is the projection $p_C(x)$ if and only if*

$$\langle x - y_x, y - y_x \rangle \leqslant 0 \quad \text{for all } y \in C \,. \tag{3.1.3}$$

Proof. Call y_x the solution of (3.1.1); take y arbitrary in C, so that $y_x + \alpha(y - y_x) \in C$ for any $\alpha \in \,]0, 1[\,$. Then we can write with the notation (3.1.2)

$$f_x(y_x) \leqslant f_x(y_x + \alpha(y - y_x)) = \tfrac{1}{2}\|y_x - x + \alpha(y - y_x)\|^2 \,.$$

Developing the square, we obtain after simplification

$$0 \leqslant \alpha \langle y_x - x, y - y_x \rangle + \tfrac{1}{2}\alpha^2 \|y - y_x\|^2 \,.$$

Divide by $\alpha \,(> 0)$ and let $\alpha \downarrow 0$ to obtain (3.1.3).

Conversely, suppose that $y_x \in C$ satisfies (3.1.3). If $y_x = x$, then y_x certainly solves (3.1.1). If not, write for arbitrary $y \in C$:

$$0 \geqslant \langle x - y_x, y - y_x \rangle = \langle x - y_x, y - x + x - y_x \rangle =$$
$$= \|x - y_x\|^2 + \langle x - y_x, y - x \rangle \geqslant \|x - y_x\|^2 - \|x - y\|\,\|x - y_x\| \,,$$

where the Cauchy-Schwarz inequality is used. Divide by $\|x - y_x\| > 0$ to see that y_x solves (3.1.1). $\qquad\square$
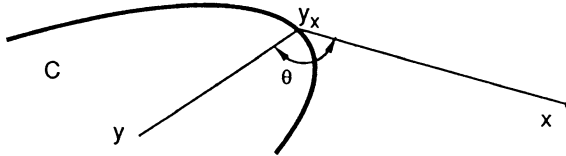
Fig. 3.1.1. The angle-characterization of a projection

Incidentally, this result proves at the same time that the variational inequality (3.1.3) has a unique solution in C. Figure 3.1.1 illustrates the following geometric interpretation: the Cauchy-Schwarz inequality defines the *angle* $\theta \in [0, \pi]$ of two nonzero vectors u and v by

$$\cos \theta := \frac{\langle u, v \rangle}{\|u\|\,\|v\|} \in [-1, +1] \,.$$

Then (3.1.3) expresses the fact that the angle between $y - y_x$ and $x - y_x$ is *obtuse*, for any $y \in C$. Writing (3.1.3) as

$$\langle x - p_C(x), y \rangle \leqslant \langle x - p_C(x), p_C(x) \rangle \quad \text{for all } y \in C \,, \tag{3.1.4}$$

we see that $p_C(x)$ lies in the face of C exposed by $x - p_C(x)$.

Remark 3.1.2 Suppose that C is actually an affine manifold (for example a subspace); then $y_x - y \in C$ whenever $y - y_x \in C$. In this case, (3.1.3) implies that

$$\langle x - y_x, y - y_x \rangle = 0 \quad \text{for all } y \in C \,. \tag{3.1.5}$$

We are back with the classical characterization of the projection onto a subspace, namely that $x - y_x \in C^\perp$ (the subspace orthogonal to C). Passing from (3.1.3) to (3.1.5) shows once more that convex analysis is the unilateral realm of inequalities, in contrast with linear analysis.   □

Some obvious properties of our projection operator are:

– the set $\{x \in \mathbb{R}^n : \mathrm{p}_C(x) = x\}$ of fixed points of p_C is C itself;
– from which it results that $\mathrm{p}_C \circ \mathrm{p}_C = \mathrm{p}_C$, and also that
– p_C is a linear operator if and only if C is a subspace.

More interesting is the following result:

Proposition 3.1.3 *For all $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$, there holds*

$$\|\mathrm{p}_C(x_1) - \mathrm{p}_C(x_2)\|^2 \leqslant \langle \mathrm{p}_C(x_1) - \mathrm{p}_C(x_2), x_1 - x_2 \rangle \,.$$

Proof. Write (3.1.3) with $x = x_1$, $y = \mathrm{p}_C(x_2) \in C$:

$$\langle \mathrm{p}_C(x_2) - \mathrm{p}_C(x_1), x_1 - \mathrm{p}_C(x_1) \rangle \leqslant 0 \,;$$

likewise,

$$\langle \mathrm{p}_C(x_1) - \mathrm{p}_C(x_2), x_2 - \mathrm{p}_C(x_2) \rangle \leqslant 0 \,,$$

and conclude by addition

$$\langle \mathrm{p}_C(x_1) - \mathrm{p}_C(x_2), x_2 - x_1 + \mathrm{p}_C(x_1) - \mathrm{p}_C(x_2) \rangle \leqslant 0 \,. \qquad \square$$

Two immediate consequences are worth noting. One is that

$$0 \leqslant \langle \mathrm{p}_C(x_1) - \mathrm{p}_C(x_2), x_1 - x_2 \rangle \quad \text{for all } (x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n \,,$$

a property expressing that the mapping p_C is, in a way, "monotone increasing". Second, we obtain from the Cauchy-Schwarz inequality:

$$\|\mathrm{p}_C(x_1) - \mathrm{p}_C(x_2)\| \leqslant \|x_1 - x_2\| \,, \tag{3.1.6}$$

i.e. p_C is nonexpansive; in particular, $\|\mathrm{p}_C(x)\| \leqslant \|x\|$ whenever $0 \in C$. However, it is not a contraction: the best Lipschitz constant

$$L := \sup \left\{ \frac{\|\mathrm{p}_C(x_1) - \mathrm{p}_C(x_2)\|}{\|x_1 - x_2\|} \; : \; x_1 \neq x_2, \; x_1 \text{ and } x_2 \text{ out of } C \right\}$$

is equal to 1 (suppose C is a subspace!), unless more is known about the "curvature" of C.

3.2 Projection onto a Closed Convex Cone

As already mentioned in Example 1.1.4, convex cones are important instances of convex sets, somehow intermediate between subspaces and general convex sets. As a result, the projection operator onto a closed convex cone enjoys properties which are finer than those of §3.1, and which come closer to those of the projection onto a subspace.

Definition 3.2.1 (Polar cone) Let K be a convex cone. Its *polar* (called negative polar cone by some authors) is:

$$K^\circ := \{s \in \mathbb{R}^n \; : \; \langle s, x \rangle \leqslant 0 \text{ for all } x \in K\} \; . \qquad \square$$

Let us note some straightforward properties:

– First of all, the polar cone depends on the scalar product: changing $\langle \cdot, \cdot \rangle$ changes K°.
– One easily sees that K° is a closed convex cone (use in particular continuity of the scalar product).
– If K is simply a subspace, then K° is its orthogonal $K^\perp$: polarity generalizes orthogonality in a unilateral way; remember Remark 3.1.2. Incidentally, it will be seen later in §4.2(d) that the polar of K° is nothing but the closure of K.
– Polarity establishes a correspondence in the set of closed convex cones, which is order-reversing: $K' \subset K$ implies $(K')^\circ \supset K^\circ$ (and the converse is true if the relation $K^{\circ\circ} = K$ is admitted for K closed).
– Finally, the only possible element in $K \cap K^\circ$ is 0.

Examples 3.2.2 (see Fig. 3.2.1).

(a). Consider the conical hull of m points $x_1, \ldots, x_m$ in $\mathbb{R}^n$:

$$K = \left\{ \sum_{j=1}^m \alpha_j x_j \; : \; \alpha_j \geqslant 0 \text{ for } j = 1, \ldots, m \right\} .$$

We leave it as an exercise to check the important result:

$$K^\circ = \{s \in \mathbb{R}^n \; : \; \langle s, x_j \rangle \leqslant 0 \text{ for } j = 1, \ldots, m\} .$$

(b). As a particular case, take the usual dot-product for $\langle \cdot, \cdot \rangle$, $\mathbb{R}^n$ being equipped with the canonical basis. Then the polar of the nonnegative orthant

$$\Omega_+ := \{x = (\xi^1, \ldots, \xi^n) \; : \; \xi^i \geqslant 0 \text{ for } i = 1, \ldots, n\}$$

is the nonpositive orthant

$$\Omega_- = (\Omega_+)^\circ = \{s = (\sigma_1, \ldots, \sigma_n) \; : \; \sigma_i \leqslant 0 \text{ for } i = 1, \ldots, n\} .$$

Naturally, such a symmetry is purely due to the fact that the basis vectors are mutually orthogonal.
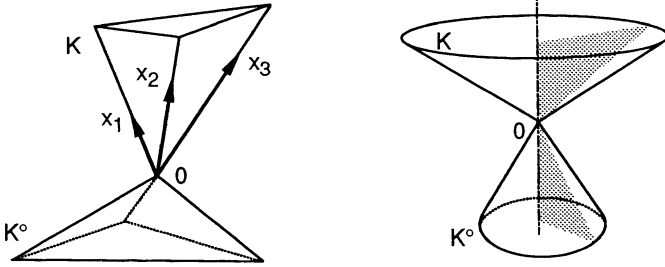
Fig. 3.2.1. Examples of polar cones

(c). Let K be a revolution cone: with $s \in \mathbb{R}^n$ of norm 1 and $\theta \in [0, \pi/2]$,

$$K_s(\theta) := \{x \in \mathbb{R}^n \ : \ \langle s, x \rangle \geqslant \|x\| \cos \theta\}.$$

Then $[K_s(\theta)]^\circ = K_{-s}(\pi/2 - \theta)$.   □

The characterization 3.1.1 takes a special form in the present conical situation.

Proposition 3.2.3 *Let K be a closed convex cone. Then $y_x = \mathrm{p}_K(x)$ if and only if*

$$y_x \in K, \quad x - y_x \in K^\circ, \quad \langle x - y_x, y_x \rangle = 0. \tag{3.2.1}$$

Proof. We know from Theorem 3.1.1 that $y_x = \mathrm{p}_K(x)$ satisfies

$$\langle x - y_x, y - y_x \rangle \leqslant 0 \quad \text{for all } y \in K. \tag{3.2.2}$$

Taking $y = \alpha y_x$, with arbitrary $\alpha \geqslant 0$, this inequality implies $(\alpha - 1)\langle x - y_x, y_x \rangle \leqslant 0$ for all $\alpha \geqslant 0$. Since $\alpha - 1$ can have either sign, this implies $\langle x - y_x, y_x \rangle = 0$ and (3.2.2) becomes

$$\langle y, x - y_x \rangle \leqslant 0 \quad \text{for all } y \in K, \quad \text{i.e.} \quad x - y_x \in K^\circ.$$

Conversely, let y_x satisfy (3.2.1). For arbitrary $y \in K$, use the notation (3.1.2):

$$f_x(y) = \tfrac{1}{2}\|x - y_x + y_x - y\|^2 \geqslant f_x(y_x) + \langle x - y_x, y_x - y \rangle;$$

but (3.2.1) shows that

$$\langle x - y_x, y_x - y \rangle = -\langle x - y_x, y \rangle \geqslant 0,$$

hence $f_x(y) \geqslant f_x(y_x)$: y_x solves (3.1.1).   □

Remark 3.2.4 We already know from (3.1.4) that $\mathrm{p}_K(x)$ lies in the face of K exposed by $x - \mathrm{p}_K(x)$; but (3.2.1) tells us more: by definition of a polar cone, $x - \mathrm{p}_K(x)$ is also in the face of K° exposed by $\mathrm{p}_K(x)$ (a symmetry confirming that $K^{\circ\circ} = K$).

Take for an illustration $K = \Omega_+$ of Example 3.2.2(b): denote by $(\pi^1, \ldots, \pi^n)$ the coordinates of $\mathrm{p}_K(x)$. They are nonnegative because $\mathrm{p}_K(x) \in \Omega_+$; each term $(\xi^i - \pi^i)\pi^i$ is

nonpositive because $x - \mathrm{p}_K(x) \in \Omega_-$. Because their sum is zero, each of these terms is actually zero, i.e.

$$\text{For } i = 1, \ldots, n, \quad \xi^i - \pi^i = 0 \text{ or } \pi^i = 0 \text{ (or both)}.$$

This property can be called a *transversality condition*. Thus, π^i is either ξ^i or 0, for each i; taking the nonnegativity of π into account, we obtain the explicit formula

$$\pi^i = \max\{0, \xi^i\} \quad \text{for } i = 1, \ldots, n.$$

This implies in particular that $\pi^i - \xi^i \geqslant 0$, i.e. $x - \pi \in \Omega_-$. $\Box$

We list some properties which are immediate consequences of the characterization (3.2.1): for all $x \in \mathbb{R}^n$,

$$\mathrm{p}_K(x) = 0 \quad \text{if and only if} \quad x \in K^\circ;$$
$$\mathrm{p}_K(\alpha x) = \alpha \mathrm{p}_K(x) \quad \text{for all } \alpha \geqslant 0;$$
$$\mathrm{p}_K(-x) = -\mathrm{p}_{-K}(x).$$

They somehow generalize the linearity of the projection onto a subspace V. An additional property can be proved, using the obvious relation $(-K)^\circ = -K^\circ$:

$$\mathrm{p}_K(x) + \mathrm{p}_{K^\circ}(x) = x. \tag{3.2.3}$$

It plays the role of $\mathrm{p}_V(x) + \mathrm{p}_{V^\perp}(x) = x$ and connotes the following *decomposition* theorem, generalizing the property $\mathbb{R}^n = V \oplus V^\perp$.

Theorem 3.2.5 (J.-J. Moreau) *Let K be a closed convex cone. For the three elements x, x_1 and x_2 in $\mathbb{R}^n$, the properties below are equivalent:*

(i) $x = x_1 + x_2$ *with* $x_1 \in K$, $x_2 \in K^\circ$ *and* $\langle x_1, x_2 \rangle = 0$;

(ii) $x_1 = \mathrm{p}_K(x)$ *and* $x_2 = \mathrm{p}_{K^\circ}(x)$.

Proof. Straightforward from (3.2.3) and the characterization (3.2.1) of $x_1 = \mathrm{p}_K(x)$. $\Box$

4 Separation and Applications

4.1 Separation Between Convex Sets

Take two disjoint sets S_1 and S_2: $S_1 \cap S_2 = \emptyset$. If, in addition, S_1 and S_2 are convex, some more can be said: it turns out that a simple convex set (namely an affine hyperplane) can be squeezed between S_1 and S_2. This extremely important property follows directly from those of the projection operator onto a convex set.

Theorem 4.1.1 *Let $C \subset \mathbb{R}^n$ be nonempty closed convex, and let $x \notin C$. Then there exists $s \in \mathbb{R}^n$ such that*

$$\langle s, x \rangle > \sup\{\langle s, y \rangle : y \in C\}. \tag{4.1.1}$$

Proof. Set $s := x - p_C(x) \neq 0$. We write (3.1.3) as

$$0 \geqslant \langle s, y - x + s \rangle = \langle s, y \rangle - \langle s, x \rangle + \|s\|^2 .$$

Thus, $\langle s, x \rangle - \|s\|^2 \geqslant \langle s, y \rangle$ for all $y \in C$, and our s is a convenient answer for (4.1.1).   □

Naturally, s could be replaced by $-s$ in (4.1.1) and Theorem 4.1.1 could just be stated as: there exists $s' \in \mathbb{R}^n$ such that $\langle s', x \rangle < \inf_{y \in C} \langle s', y \rangle$. Note that s is certainly nonzero; by positive homogeneity, we may require $\|s\| = 1$.

Geometrically, we know that an $s \neq 0$ defines hyperplanes $H_{s,r}$ as in Example 1.1.2(a), which are translations of each other when r describes $\mathbb{R}$. With s of (4.1.1), pick

$$r = r_s := \tfrac{1}{2}\left(\langle s, x \rangle + \sup_{y \in C} \langle y, s \rangle\right) .$$

Then

$$\langle s, x \rangle - r_s > 0 \quad \text{and} \quad \langle s, y \rangle - r_s < 0 \quad \text{for all } y \in C ,$$

which can be summarized in one sentence: the affine hyperplane H_{s,r_s} *separates* the two convex sets C and $\{x\}$. These two sets are in the opposite (open) half-spaces limited by that hyperplane.

Remark 4.1.2 With relation to this interpretation, Theorem 4.1.1 is often called the *Hahn-Banach* Theorem in geometric form. On the other hand, consider the righthand side of (4.1.1); it suggests a function $\sigma_C : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, called the *support function* of C:

$$\sigma_C(s) := \sup \{\langle s, y \rangle : y \in C\} ,$$

which will be studied thoroughly in Chap. C. If $x \in C$, we have by definition

$$\langle s, x \rangle \leqslant \sigma_C(s) \quad \text{for all } s \in \mathbb{R}^n ;$$

but this actually *characterizes* the elements of C: Theorem 4.1.1 tells us that the converse is true. Therefore the test "$x \in C$?" is equivalent to the test "$\langle \cdot, x \rangle \leqslant \sigma_C$?", which compares the linear function $\langle \cdot, x \rangle$ to the function σ_C. With this interpretation, Theorem 4.1.1 can be formulated in an equivalent analytical way, involving functions instead of hyperplanes; this is called the Hahn-Banach Theorem in *analytical* form.   □

A convenient generalization of Theorem 4.1.1 is the following:

Corollary 4.1.3 (Strict Separation of Convex Sets) *Let C_1, C_2 be two nonempty closed convex sets with $C_1 \cap C_2 = \emptyset$. If C_2 is bounded, there exists $s \in \mathbb{R}^n$ such that*

$$\sup_{y \in C_1} \langle s, y \rangle < \min_{y \in C_2} \langle s, y \rangle . \tag{4.1.2}$$

Proof. The set $C_1 - C_2$ is convex (Proposition 1.2.4) and closed (because C_2 is compact). To say that C_1 and C_2 are disjoint is to say that $0 \notin C_1 - C_2$, so we have by Theorem 4.1.1 an $s \in \mathbb{R}^n$ separating $\{0\}$ from $C_1 - C_2$:

$$\sup \{\langle s, y \rangle : y \in C_1 - C_2\} < \langle s, 0 \rangle = 0 .$$

This means:

$$0 > \sup_{y_1 \in C_1} \langle s, y_1 \rangle + \sup_{y_2 \in C_2} \langle s, -y_2 \rangle$$
$$= \sup_{y_1 \in C_1} \langle s, y_1 \rangle - \inf_{y_2 \in C_2} \langle s, y_2 \rangle.$$

Because C_2 is bounded, the last infimum (is a min and) is finite and can be moved to the lefthand side. $\qquad\square$
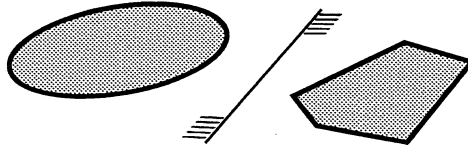
Fig. 4.1.1. Strict separation of two convex sets

Once again, (4.1.2) can be switched over to $\inf_{y \in C_1} \langle s, y \rangle > \max_{y \in C_2} \langle s, y \rangle$. Using the support function of Remark 4.1.2, we can also write (4.1.2) as $\sigma_{C_1}(s) + \sigma_{C_2}(-s) < 0$. Figure 4.1.1 gives the same geometric interpretation as before. Choosing $r = r_s$ strictly between $\sigma_{C_1}(s)$ and $-\sigma_{C_2}(-s)$, we obtain a hyperplane separating C_1 and C_2 strictly: each set is in one of the corresponding open half-spaces.

Fig. 4.1.2. Strict separation needs compactness

When C_1 and C_2 are both unbounded, Corollary 4.1.3 may fail – even though the role of boundedness was apparently minor, but see Fig. 4.1.2. As suggested by this picture, C_1 and C_2 can nevertheless be *weakly* separated, i.e. (4.1.2) can be replaced by a weak inequality. Such a weakening is a bit exaggerated, however: Fig. 4.1.3 shows that (4.1.2) may hold as $\langle s, y_1 \rangle = \langle s, y_2 \rangle$ for all $(y_1, y_2) \in C_1 \times C_2$ if s is orthogonal to aff $(C_1 \cup C_2)$.

For a convenient definition, we need to be more demanding: we say that the two nonempty convex sets C_1 and C_2 are *properly separated* by $s \in \mathbb{R}^n$ when

$$\sup_{y_1 \in C_1} \langle s, y_1 \rangle \leq \inf_{y_2 \in C_2} \langle s, y_2 \rangle \quad \text{and} \quad \inf_{y_1 \in C_1} \langle s, y_1 \rangle < \sup_{y_2 \in C_2} \langle s, y_2 \rangle.$$

This (weak) proper separation property is sometimes just what is needed for technical purposes. It happens to hold under fairly general assumptions on the intersection $C_1 \cap C_2$. We end this section with a possible result, stated without proof.

Theorem 4.1.4 (Proper Separation of Convex Sets) *If the two nonempty convex sets C_1 and C_2 satisfy* $(\text{ri } C_1) \cap (\text{ri } C_2) = \emptyset$, *they can be properly separated.* $\qquad\square$

Observe the qualification assumption coming into play. We have already seen it in Proposition 2.1.10, and we know from (2.1.5) that it is equivalent to $0 \notin \text{ri } (C_1 - C_2)$.
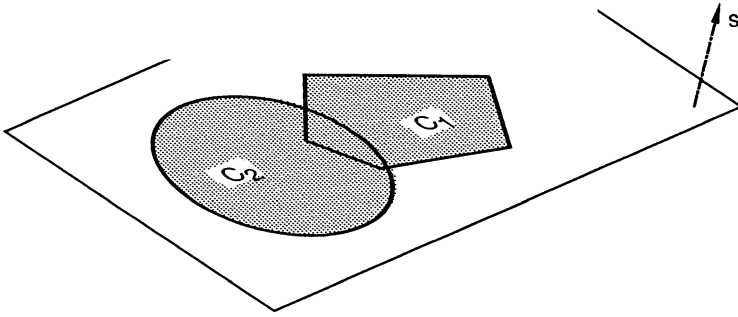
Fig. 4.1.3. An improper separation

4.2 First Consequences of the Separation Properties

The separation properties introduced in §4.1 have many applications. To prove that some set S is contained in a closed convex set C, a possibility is often to argue by contradiction, separating from C a point in $S\backslash C$, and then exploiting the simple structure of the separating hyperplane. Here we review some of these applications, including the proofs announced in the previous sections. Note: our proofs are often fairly short (as is that of Corollary 4.1.3) or geometric. It is a good exercise to develop more elementary proofs, or to support the geometry with detailed calculations.

(a) Existence of Supporting Hyperplanes First of all, we note that a convex set C, not equal to the whole of $\mathbb{R}^n$, does have a supporting hyperplane in the sense of Definition 2.4.1. To see it, use first Proposition 2.1.8: $\operatorname{cl} C \neq \mathbb{R}^n$ (otherwise, we would have the contradiction $C \supset \operatorname{ri} C = \operatorname{ri} \operatorname{cl} C = \operatorname{ri} \mathbb{R}^n = \mathbb{R}^n$). Then take a hyperplane separating $\operatorname{cl} C$ from some $x \notin \operatorname{cl} C$: it is our asserted support of C. Actually, we can prove slightly more:

Lemma 4.2.1 *Let $x \in \operatorname{bd} C$, where $C \neq \emptyset$ is convex in $\mathbb{R}^n$ (naturally $C \neq \mathbb{R}^n$). There exists a hyperplane supporting C at x.*

Proof. Because C, $\operatorname{cl} C$ and their complements have the same boundary (remember Remark 2.1.9), a sequence (x_k) can be found such that

$$x_k \notin \operatorname{cl} C \quad \text{for } k = 1, 2, \dots \quad \text{and} \quad \lim_{k \to +\infty} x_k = x .$$

For each k we have by Theorem 4.1.1 some s_k of norm 1 such that $\langle s_k, x_k - y \rangle > 0$ for all $y \in C \subset \operatorname{cl} C$.

Extract a subsequence if necessary so that $s_k \to s$ (note: $s \neq 0$) and pass to the limit to obtain $\langle s, x - y \rangle \geqslant 0$ for all $y \in C$. This is the required result $\langle s, x \rangle = r \geqslant \langle s, y \rangle$ for all $y \in C$. □

Remark 4.2.2 The above procedure may well end up with a supporting hyperplane containing C: $\langle s, x - y \rangle = 0$ for all $y \in C$, a result of little interest; see also Fig. 4.1.3. This

can happen only when C is a "flat" convex set (dim $C \leqslant n - 1$), in which case our construction should be done in aff C, as illustrated on Fig. 4.2.1. Let us detail such a "relative" construction, to demonstrate a calculation involving affine hulls.
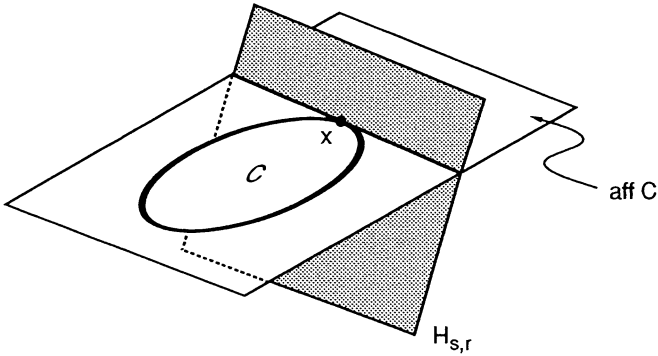
Fig. 4.2.1. Nontrivial supports

Let V be the subspace parallel to aff C, with $U = V^{\perp}$ its orthogonal subspace: by definition, $\langle s, y - x \rangle = 0$ for all $s \in U$ and $y \in C$. Suppose $x \in \text{rbd}\, C$ (the case $x \in \text{ri}\, C$ is hopeless) and translate C to $C_0 := C - \{x\}$. Then C_0 is a convex set in the Euclidean space V and $0 \in \text{rbd}\, C_0$. We take as in 4.2.1 a sequence $(x_k) \subset V \backslash \text{cl}\, C_0$ tending to 0 and a corresponding unitary $s_k \in V$ separating the point x_k from C_0. The limit $s \neq 0$ is in V, separates (not strictly) $\{0\}$ and C_0, i.e. $\{x\}$ and C: we are done.

We will say that $H_{s,r}$ is a *nontrivial* support (at x) if $s \notin U$, i.e. if $s_V \neq 0$, with the decomposition $s = s_V + s_U$. Then C is not contained in $H_{s,r}$: if it were, we would have $r = \langle s, y \rangle = \langle s_V, y \rangle + \langle s_U, x \rangle$ for all $y \in C$. In other words, $\langle s_V, \cdot \rangle$ would be constant on C; by definition of the affine hull and of V, this would mean $s_V \in U$, i.e. the contradiction $s_V = 0$. To finish, note that s_U may be assumed to be 0: if $s_V + s_U$ is a nontrivial support, so is $s_V = s_V + 0$ as well; it corresponds to a hyperplane orthogonal to C. $\square$

(b) Outer Description of Closed Convex Sets Taking the closed convex hull of a set consists in intersecting the closed convex sets containing it. We mentioned in Remark 1.4.4 that convexity allowed the intersection to be restricted to a simple class of closed convex sets: the closed half-spaces. Indeed, let a given set S be contained in some closed half-space; say $S \subset H^-_{s,r} := \{y \in \mathbb{R}^n : \langle s, y \rangle \leqslant r\}$ for some $(s, r) \in \mathbb{R}^n \times \mathbb{R}$ with $s \neq 0$. Then the index-set

$$\begin{aligned} \Sigma_S &:= \{(s, r) \in \mathbb{R}^n \times \mathbb{R} : S \subset H^-_{s,r}\} \\ &= \{(s, r) : \langle s, y \rangle \leqslant r \text{ for all } y \in S\} \end{aligned} \tag{4.2.1}$$

is nonempty. As illustrated by Fig. 4.2.2, we can therefore intersect all the half-spaces indexed in Σ_S. This defines a certain set C^*:

$$S \subset C^* := \cap_{(s,r) \in \Sigma_S} H^-_{s,r} = $$
$$\{z \in \mathbb{R}^n : \langle s, z \rangle \leqslant r \text{ whenever } \langle s, y \rangle \leqslant r \text{ for all } y \in S\}.$$
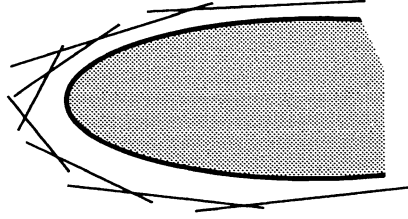
Fig. 4.2.2. Outer construction of a closed convex hull

Theorem 4.2.3 *The closed convex hull of a set $S \neq \emptyset$ is either the whole of $\mathbb{R}^n$ or the set C^* defined above.*

Proof. Suppose $\overline{\text{co}}\, S \neq \mathbb{R}^n$: some nontrivial closed convex set contains S. Then take a hyperplane as described in Lemma 4.2.1 to obtain a closed half-space containing S: the index-set Σ_S is nonempty and C^* is well-defined.

By construction, $C^* \supset \overline{\text{co}}\, S$. Conversely, let $x \notin \overline{\text{co}}\, S$; we can separate $\{x\}$ and $\overline{\text{co}}\, S$: there exists $s_0 \neq 0$ such that $\langle s_0, x \rangle > \sup_{y \in S} \langle s_0, y \rangle =: r_0$. Then $(s_0, r_0) \in \Sigma_S$; but $x \notin H^-_{s_0, r_0}$, hence $x \notin C^*$. $\qquad\square$

The definition of C^*, rather involved, can be slightly simplified: actually, Σ_S is redundant, as it contains much too many r's. Roughly speaking, for given $s \in \mathbb{R}^n$, just take the number
$$r = r_s := \inf \{ r \in \mathbb{R} \; : \; (s, r) \in \Sigma_S \}$$
that is sharp in (4.2.1). Letting s vary, (s, r_s) describes a set Σ^0_S, smaller than Σ_S but just as useful. With this new notation, the expression of $C^* = \overline{\text{co}}\, S$ reduces to
$$\overline{\text{co}}\, S = \left\{ z \in \mathbb{R}^n \; : \; \langle s, z \rangle \leqslant \sup_{y \in S} \langle s, y \rangle \right\}.$$

We find again the support function of Remark 4.1.2 coming into play. Chapter C will follow this development more thoroughly.

Assuming S closed convex in Theorem 4.2.3, we see that a closed convex set can thus be *defined* as the intersection of the closed half-spaces containing it:

Corollary 4.2.4 *The data $(s_j, r_j) \in \mathbb{R}^n \times \mathbb{R}$ for j in an arbitrary index set J is equivalent to the data of a closed convex set C via the relation*
$$C = \{ x \in \mathbb{R}^n \; : \; \langle s_j, x \rangle \leqslant r_j \; \text{for } j \in J \}.$$

Proof. If C is given, define $\{(s_j, r_j)\}_J := \Sigma_C$ as in (4.2.1). If $\{(s_j, r_j)\}_J$ is given, the intersection of the corresponding half-spaces is a closed convex set. $\qquad\square$

Note in the above characterization that $C = \mathbb{R}^n$ with $J = \emptyset$ can be included as an extreme case. Also, the case of a finite J is important:

Definition 4.2.5 (Polyhedral Sets) A *closed convex polyhedron* is an intersection of finitely many half-spaces. Take $(s_1, r_1), \ldots, (s_m, r_m)$ in $\mathbb{R}^n \times \mathbb{R}$, with $s_i \neq 0$ for $i = 1, \ldots, m$; then define

$$P := \{x \in \mathbb{R}^n \ : \ \langle s_j, x \rangle \leqslant r_j \ \text{for} \ j = 1, \ldots, m \} \,,$$

or in matrix notation as in Example 1.2.2, $P = \{x \in \mathbb{R}^n : Ax \leqslant b \}$.

A *closed convex polyhedral cone* is the special case where $b = 0$.   □

(c) Proof of Minkowski's Theorem We turn now to the inner description of a convex set and prove Theorem 2.3.4, asserting that $C = \mathrm{co} \, \mathrm{ext} \, C$ when C is compact convex.

The result is trivially true if $\dim C = 0$, i.e. C is a singleton, with a unique extreme point. Assume for induction that the result is true for compact convex sets of dimension less than k; let C be a compact convex set of dimension k and take $x \in C$. There are two possibilities:

- If $x \in \mathrm{rbd} \, C$, §4.2(a) tells us that there exists a nontrivial hyperplane H supporting C at x. The nonempty compact convex set $C \cap H$ has dimension at most $k - 1$, so $x \in C \cap H$ is a convex combination of extreme points in that set, which is an exposed face of C. Using Remark 2.4.4, these extreme points are also extreme in C.
- If $x \in \mathrm{ri} \, C \, (= C \backslash \mathrm{rbd} \, C)$, take in C a point $x' \neq x$; this is possible for $\dim C > 0$. The affine line generated by x and x' cuts $\mathrm{rbd} \, C$ in at most two points y and z (see Remark 2.1.7, there are really two points because C is compact). From the first part of the proof, y and z are convex combinations of extreme points in C; and so is their convex combination x (associativity of convex combinations).

(d) Bipolar of a Convex Cone The definition of a polar cone was given in §3.2, where some interesting properties were pointed out. Here we can show one more similarity with the concept of orthogonality in linear analysis.

Proposition 4.2.6 *Let K be a convex cone with polar K°; then, the polar $K^{\circ\circ}$ of K° is the closure of K.*

Proof. From the definition itself of a polar cone,

$$x \in K \quad \Longrightarrow \quad \langle s, x \rangle \leqslant 0 \ \text{for all} \ s \in K^\circ \quad \Longrightarrow \quad x \in K^{\circ\circ} \,;$$

thus $K \subset K^{\circ\circ}$, therefore $\mathrm{cl} \, K \subset K^{\circ\circ}$ (a polar cone is always closed).

Conversely, let $x_0 \in K^{\circ\circ}$. If $x_0 \notin \mathrm{cl} \, K$, there is a separating hyperplane (d, α) such that

$$\langle d, x_0 \rangle > \alpha \geqslant \langle d, x \rangle \quad \text{for all} \ x \in K \subset \mathrm{cl} \, K \,.$$

This implies $\alpha \geqslant 0$ (take $x = 0$) and we may assume $\alpha = 0$ (if $x \in K$ has $\langle d, x \rangle > 0$, take $tx \in K$ with $t \to +\infty$). Thus $d \in K^\circ$ and the property $\langle d, x_0 \rangle > 0$ gives the contradiction.   □

Of course, if K is already closed, $K^{\circ\circ} = K$. This has several interesting consequences:

- With relation to (a) above, we observe that every supporting hyperplane of K at $x \in \mathrm{bd} \, K$ also supports K at 0: when dealing with supports to a cone, it is enough to consider *linear* hyperplanes only.

– Besides, consider the index-set Σ_K of (4.2.1): its r-part can be restricted to $\{0\}$; as for its s-part, we see from Definition 3.2.1 that it becomes $K^\circ\backslash\{0\}$. In other words: barring the zero-vector, a closed convex cone is the set of (linear) hyperplanes supporting its polar at 0. This is the outer description of a closed convex cone.

– This has also an impact on the separation Theorem 4.1.1: depending on s, the righthand side in (4.1.1) is either 0 ($s \in K$) or $+\infty$ ($s \notin K$). Once again, linear separating hyperplanes are sufficient and they must all be in the polar cone.

Let us summarize these observations:

Corollary 4.2.7 Let K be a closed convex cone. Then

$$x \in K \quad \Longleftrightarrow \quad \langle s, x \rangle \leqslant 0 \text{ for all } s \in K^\circ.\qquad \square$$

4.3 The Lemma of Minkowski-Farkas

Because of its historical importance, we devote an entire subsection to another consequence of the separation property, known as Farkas' Lemma. Let us first recall a classical result from linear algebra: if A is a matrix with n rows and m columns and $b \in \mathbb{R}^n$, the system $A\alpha = b$ has a solution in $\mathbb{R}^m$ (we say that the system is *consistent*) exactly when

$$b \in \operatorname{Im} A = [\operatorname{Ker} A^\top]^\perp \; ;$$

this can be rewritten $\{b\}^\perp \supset \operatorname{Ker} A^\top$, or

$$\{x \in \mathbb{R}^n \; : \; A^\top x = 0\} \subset \{x \in \mathbb{R}^n \; : \; b^\top x = 0\}.$$

Denoting by $s_1, \ldots, s_m$ the columns of A and using our Euclidean notation, we write the equivalence of these properties as

$$b \in \operatorname{lin}\{s_1, \ldots, s_m\} \quad \text{if and only if}$$
$$\langle b, x \rangle = 0 \quad \text{whenever} \quad \langle s_j, x \rangle = 0 \text{ for } j = 1, \ldots, m.$$

Moving to the unilateral world of convex analysis, we replace linear hulls by conical hulls, and equalities by inequalities. This gives a result dating back to the end of the XIX$^{\text{th}}$ Century, due to J. Farkas and also to H. Minkowski; we state it without proof, as it will be a consequence of Theorem 4.3.4 below.

Lemma 4.3.1 (Farkas I) *Let $b, s_1, \ldots, s_m$ be given in $\mathbb{R}^n$. The set*

$$\{x \in \mathbb{R}^n : \langle s_j, x \rangle \leqslant 0 \quad \text{for } j = 1, \ldots, m\} \qquad (4.3.1)$$

is contained in the set

$$\{x \in \mathbb{R}^n \; : \; \langle b, x \rangle \leqslant 0\} \qquad (4.3.2)$$

if and only if (see Definition 1.4.5 of a conical hull)

$$b \in \operatorname{cone}\{s_1, \ldots, s_m\}. \qquad (4.3.3)$$
$$\square$$

To express the inclusion relation between the sets (4.3.1) and (4.3.2), one also says that the inequality with b is a *consequence* of the joint inequalities with s_j. Another way of expressing (4.3.3) is to say that the system of equations and inequations in α

$$b = \sum_{j=1}^{m} \alpha_j s_j , \quad \alpha_j \geqslant 0 \text{ for } j = 1, \dots, m \tag{4.3.4}$$

has a solution.

Farkas' Lemma is sometimes formulated as an *alternative*, i.e. a set of two statements such that each one is false when the other is true. More precisely, let P and Q be two logical propositions. They are said to form an alternative if one and only one of them is true:

$$P \implies \text{not } Q \quad \text{and} \quad \text{not } P \implies Q$$

or, just as simply:

$$P \iff \text{not } Q \quad [\text{or} \quad Q \iff \text{not } P].$$

This applies to Farkas' Lemma:

Lemma 4.3.2 (Farkas II) *Let $b, s_1, \dots, s_m$ be given in $\mathbb{R}^n$. Then exactly one of the following statements is true.*

$$P := (4.3.4) \text{ has a solution } \alpha \in \mathbb{R}^n .$$

$$Q := \begin{cases} \text{The system of inequations} \\ \langle b, x \rangle > 0 , \quad \langle s_j, x \rangle \leqslant 0 \text{ for } j = 1, \dots, m \\ \text{has a solution } x \in \mathbb{R}^n . \end{cases} \qquad \square$$

Still another formulation is geometric. Call K the convex cone generated by $s_1, \dots, s_m$; as seen in Example 3.2.2, K° is the set (4.3.1). What Farkas' Lemma says is that

$$b \in K \quad [\text{i.e. (4.3.3) holds}] \quad \text{if and only if}$$
$$\langle b, x \rangle \leqslant 0 \quad \text{whenever} \quad x \in K^\circ \quad [\text{i.e. } b \in K^{\circ\circ}].$$

More simply, Farkas' Lemma is: $K^{\circ\circ} = K$; but we know from §4.2(d) that this property holds under the sole condition that K is closed. The proof of Farkas' Lemma therefore reduces to proving the following result:

Lemma 4.3.3 (Farkas III) *Let $s_1, \dots, s_m$ be given in $\mathbb{R}^n$. Then the convex cone*

$$K := \text{cone}\{s_1, \dots, s_m\} = \left\{ \sum_{j=1}^{m} \alpha_j s_j \ : \ \alpha_j \geqslant 0 \text{ for } j = 1, \dots, m \right\}$$

is closed.

Proof. It is quite similar to that of Carathéodory's Theorem 1.3.6. First, the proof is easy if the s_j's are linearly independent: then, the convergence of

$$x^k = \sum_{j=1}^{m} \alpha_j^k s_j \quad \text{for } k \to \infty \tag{4.3.5}$$

is equivalent to the convergence of each $(\alpha_j^k)_k$ to some α_j, which must be nonnegative if each α_j^k in (4.3.5) is nonnegative.

Suppose, on the contrary, that the system $\sum_{j=1}^{m} \beta_j s_j = 0$ has a nonzero solution $\beta \in \mathbb{R}^m$ and assume $\beta_j < 0$ for some j (change β to $-\beta$ if necessary). As in the proof of Theorem 1.3.6, write each $x \in K$ as

$$x = \sum_{j=1}^{m} \alpha_j s_j = \sum_{j=1}^{m} [\alpha_j + t^*(x)\beta_j] s_j = \sum_{j \neq i(x)} \alpha_j' s_j \,,$$

where

$$i(x) \in \operatorname*{Argmin}_{\beta_j < 0} \frac{-\alpha_j}{\beta_j}, \quad t^*(x) := \frac{-\alpha_{i(x)}}{\beta_{i(x)}} \,,$$

so that each $\alpha_j' = \alpha_j + t^*(x)\beta_j$ is nonnegative. Letting x vary in K, we thus construct a decomposition $K = \cup_{i=1}^m K_i$, where K_i is the conical hull of the $m - 1$ generators s_j, $j \neq i$.

Now, if there is some i such that the generators of K_i are linearly dependent, we repeat the argument for a further decomposition of this K_i. After finitely many such operations, we end up with a decomposition of K as a finite union of polyhedral convex cones, each having linearly independent generators. All these cones are therefore closed (first part of the proof), so K is closed as well. □

We are now in a position to state a general version of Farkas' Lemma, with non-homogeneous terms and infinitely many inequalities. Its proof uses in a direct way the separation Theorem 4.1.1.

Theorem 4.3.4 (Generalized Farkas) *Let be given (b, r) and (s_j, ρ_j) in $\mathbb{R}^n \times \mathbb{R}$, where j varies in an (arbitrary) index set J. Suppose that the system of inequalities*

$$\langle s_j, x \rangle \leqslant \rho_j \quad \text{for all } j \in J \tag{4.3.6}$$

has a solution $x \in \mathbb{R}^n$ (the system is consistent). Then the following two properties are equivalent:

(i) *$\langle b, x \rangle \leqslant r$ for all x satisfying (4.3.6);*
(ii) *(b, r) lies in the closed convex conical hull of $S := \{(0, 1)\} \cup \{(s_j, \rho_j)\}_{j \in J}$.*

Proof. [(ii) ⇒ (i)] Let first (b, r) be in $K := \operatorname{cone} S$. In other words, there exists a finite set $\{1, \ldots, m\} \subset J$ and nonnegative $\alpha_0, \alpha_1, \ldots, \alpha_m$ such that (we adopt the convention $\sum_\emptyset = 0$)

$$b = \sum_{j=1}^{m} \alpha_j s_j \quad \text{and} \quad r = \alpha_0 + \sum_{j=1}^{m} \alpha_j \rho_j \, .$$

For each x satisfying (4.3.6) we can write

$$\langle b, x \rangle \leqslant r - \alpha_0 \leqslant r \, . \tag{4.3.7}$$

If, now, (b, r) lies in the closure of K, pass to the limit in (4.3.7) to establish the required conclusion (i) for all (b, r) described by (ii).

$[(i) \Rightarrow (ii)]$ If $(b, r) \notin \operatorname{cl} K$, separate (b, r) from $\operatorname{cl} K$: equipping $\mathbb{R}^n \times \mathbb{R}$ with the scalar product

$$\langle\!\langle (b, r), (d, t) \rangle\!\rangle := \langle b, d \rangle + rt \, ,$$

there exists $(d, -t) \in \mathbb{R}^n \times \mathbb{R}$ such that

$$\sup_{(s, \rho) \in K} [\langle s, d \rangle - \rho t] < \langle b, d \rangle - rt \, . \tag{4.3.8}$$

It follows first that the lefthand supremum is a finite number κ. Then the conical character of K implies $\kappa \leqslant 0$, because $\alpha\kappa \leqslant \kappa$ for all $\alpha > 0$; actually $\kappa = 0$ because $(0, 0) \in K$. In summary, we have singled out $(d, t) \in \mathbb{R}^n \times \mathbb{R}$ such that

	$t \geqslant 0$	[take $(0, 1) \in K$]
$(*)$	$\langle s_j, d \rangle - \rho_j t \leqslant 0$ for all $j \in J$	[take $(s_j, \rho_j) \in K$]
$(**)$	$\langle b, d \rangle - rt > 0 \, .$	[don't forget (4.3.8)]

Now consider two cases:

– If $t > 0$, divide $(*)$ and $(**)$ by t to exhibit the point $x = d/t$ violating (i).

– If $t = 0$, take x_0 satisfying (4.3.6). Observe from $(*)$ that, for all $\alpha > 0$, the point $x(\alpha) = x_0 + \alpha d$ satisfies (4.3.6) as well. Yet, let $\alpha \to +\infty$ in the expression $\langle b, x(\alpha) \rangle = \langle b, x_0 \rangle + \alpha \langle b, d \rangle$ to realize from $(**)$ that $x(\alpha)$ violates (i) if α is large enough.

Thus we have proved in both cases that "not (ii) $\Rightarrow$ not (i)". $\square$

We finish with two comments relating Theorem 4.3.4 with the previous forms of Farkas' Lemma. Take first the homogeneous case, where r and the ρ_j's are all zero. Then the consistency assumption is automatically satisfied (by $x = 0$) and the theorem says:

(i') $[\langle s_j, x \rangle \leqslant 0$ for $j \in J] \Longrightarrow [\langle b, x \rangle \leqslant 0]$

      is equivalent to

(ii') $b \in \overline{\operatorname{cone}} \{ s_j : j \in J \}$.

Second, suppose that $J = \{1, \ldots, m\}$ is a finite set, so the set described by (4.3.6) becomes a closed convex polyhedron, assumed nonempty. A handy matrix notation (assuming the dot-product for $\langle \cdot, \cdot \rangle$), is $A^\top x \leqslant \rho$, if A is the matrix whose columns are the s_j's, and $\rho \in \mathbb{R}^m$ has the coordinates $\rho_1, \ldots, \rho_m$. Then Theorem 4.3.4 writes:

(i") $\{x \in \mathbb{R}^n \; : \; A^\top x \leqslant \rho\} \subset \{x \in \mathbb{R}^n \; : \; b^\top x \leqslant r\}$
is equivalent to

(ii") $\exists \alpha \in \mathbb{R}^m$ such that $\alpha \geqslant 0, \; A\alpha = b, \; \rho^\top \alpha \leqslant r$.

Indeed, it suffices to recall Lemma 4.3.3: the conical hull involved in (ii) of Theorem 4.3.4 is already closed. Beware that the last relation in (ii") is really an inequality.

5 Conical Approximations of Convex Sets

Given a set S and $x \in S$, a fruitful idea is to approximate S near x by a "simpler" set. In classical differential geometry, a "smooth" surface S is approximated by an affine manifold "tangent" to S. This concept is most exploited in the differentiation of a "smooth" function $f : \mathbb{R}^n \to \mathbb{R}$, whose graph has a "tangent" affine hyperplane in $\mathbb{R}^n \times \mathbb{R}$ near $(x, f(x))$:

$$\mathrm{gr}\, f \simeq \{(y, r) \; : \; r - f(x) = \langle \nabla f(x), y - x \rangle\} \, .$$

Because convex sets have no reason to be "smooth", some substitute to affine manifolds must be proposed. We know that affine manifolds are translations of subspaces; say, we approximate S near x by

$$H_S(x) = \{x\} + V_S(x) \simeq S \, ,$$

where $V_S(x)$ is a subspace: the subspace tangent to S at x. It is therefore time to remember §3.2: in the unilateral world of convex analysis, the natural substitutes for subspaces are the *closed convex cones*. Besides, another important object is the set of normals to S at x, i.e. the subspace orthogonal to $V_S(x)$; here, orthogonality will be replaced by *polarity*, as in Moreau's decomposition Theorem 3.2.5.

5.1 Convenient Definitions of Tangent Cones

In order to introduce the convenient objects, we need first to cast a fresh glance at the general concept of tangency. We therefore consider in this subsection an arbitrary *closed* set $S \subset \mathbb{R}^n$.

A direction d is classically called tangent to S at $x \in S$ when it is the derivative at x of some curve drawn on S; it follows that $-d$ is a tangent as well. In our unilateral world, *half-derivatives* are more relevant. Furthermore, sets of discrete type cannot have any tangent direction in the above sense, we will therefore replace *curves* by *sequences*. In a word, our new definition of tangency is as follows:

Definition 5.1.1 (Tangent Directions) Let $S \subset \mathbb{R}^n$ be nonempty. We say that $d \in \mathbb{R}^n$ is a direction *tangent* to S at $x \in S$ when there exists a sequence $(x_k) \subset S$ and a sequence (t_k) such that, when $k \to +\infty$,

$$x_k \to x, \quad t_k \downarrow 0, \quad \frac{x_k - x}{t_k} \to d. \tag{5.1.1}$$

The set of all such directions is called the *tangent cone* (also called the contingent cone, or Bouligand's cone) to S at $x \in S$, denoted by $T_S(x)$.   □

Observe immediately that 0 is always a tangent direction (take $x_k \equiv x$!); also, if d is tangent, so is αd for any $\alpha > 0$ (change t_k to t_k/α!). The terminology "tangent cone" is therefore legal. If $x \in \text{int } S$, $T_S(x)$ is clearly the whole space, so that the only interesting points are those on bd S.

If we set in Definition 5.1.1 $d_k := (x_k - x)/t_k$ [$\to d$], i.e. $x_k = x + t_k d_k$ [$\in S$], we obtain the equivalent formulation:

Proposition 5.1.2 *A direction d is tangent to S at $x \in S$ if and only if*

$$\exists (d_k) \to d, \ \exists (t_k) \downarrow 0 \quad \text{such that} \quad x + t_k d_k \in S \text{ for all } k.   \qquad □$$

A tangent direction thus appears as a set of limits; a limit of tangent directions is therefore a "limit of limits", and is a limit itself:

Proposition 5.1.3 *The tangent cone is closed.*

Proof. Let $(d_\ell) \subset T_S(x)$ be converging to d; for each ℓ take sequences $(x_{\ell,k})_k$ and $(t_{\ell,k})_k$ associated with d_ℓ in the sense of Definition 5.1.1. Fix $\ell > 0$: we can find k_ℓ such that

$$\left\| \frac{x_{\ell,k_\ell} - x}{t_{\ell,k_\ell}} - d_\ell \right\| \leqslant \frac{1}{\ell}.$$

Letting $\ell \to \infty$, we then obtain the sequences $(x_{\ell,k_\ell})_\ell$ and $(t_{\ell,k_\ell})_\ell$ which define d as an element of $T_S(x)$.   □

The examples below confirm that our definition reproduces the classical one when S is "well-behaved", while Fig. 5.1.1 illustrates a case where classical tangency cannot be used.
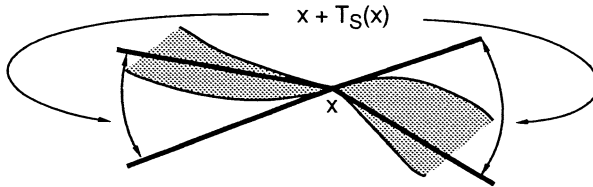
Fig. 5.1.1. Tangency to a "bad" set

Examples 5.1.4 Given m functions $c_1, \ldots, c_m$ continuously differentiable on $\mathbb{R}^n$, consider

$$S := \{x \in \mathbb{R}^n \ : \ c_i(x) = 0 \text{ for } i = 1, \ldots, m\}.$$

Let $x \in S$ be such that the gradients $\nabla c_1(x), \ldots, \nabla c_m(x)$ are linearly independent. Then $T_S(x)$ is the subspace

$$\{d \in \mathbb{R}^n \ : \ \langle \nabla c_i(x), d \rangle = 0 \text{ for } i = 1, \ldots, m\}. \tag{5.1.2}$$

Another example is

$$S := \{x \in \mathbb{R}^n \ : \ c_1(x) \leqslant 0\}.$$

At $x \in S$ such that $c_1(x) = 0$ and $\nabla c_1(x) \neq 0$, $T_S(x)$ is the *half-space*

$$\{d \in \mathbb{R}^n \ : \ \langle \nabla c_1(x), d \rangle \leqslant 0\}. \tag{5.1.3}$$

Both formulae (5.1.2) and (5.1.3) can be proved with the help of the implicit function theorem. This explains the assumptions on the $\nabla c_i(x)$'s; things become more delicate when several inequalities are involved to define S. □

Naturally, the concept of tangency is local, as it depends only on the behaviour of S near x. From its Definition 5.1.1, $T_S(x)$ appears as the set of all possible cluster points of the difference quotients $((y-x)/t)$, with $y \in S$ and $t \downarrow 0$; using set-valued notation (see §0.5):

$$T_S(x) = \lim_{t \downarrow 0} \text{ext} \frac{S - x}{t}. \tag{5.1.4}$$

Another interpretation uses the *distance function* $d_S(x) = \min_{y \in S} \|y - x\|$: $T_S(x)$ can also be viewed as the set of d's such that

$$\liminf_{t \downarrow 0} \frac{d_S(x + td)}{t} = 0. \tag{5.1.5}$$

Knowing that $d_S(x) = 0$ when $x \in S$, the infimand of (5.1.5) can be interpreted as a difference quotient: $[d_S(x + td) - d_S(x)]/t$. Finally, (5.1.5) can be interpreted in a set-formulation: for any $\varepsilon > 0$ and for any $\delta > 0$, there exists $0 < t \leqslant \delta$ such that

$$x + td \in S + B(0, t\varepsilon), \quad \text{i.e.} \quad d \in \frac{S - x}{t} + B(0, \varepsilon).$$

Remark 5.1.5 In (5.1.4), we have taken the tangent cone as a lim ext, which corresponds to a lim inf in (5.1.5). We could have defined another "tangent cone", namely

$$\lim_{t \downarrow 0} \text{int} \frac{S - x}{t}.$$

In this case, (5.1.5) would have been changed to

$$\limsup_{t \downarrow 0} \frac{d_S(x + td)}{t} = 0 \quad [= \lim_{t \downarrow 0} d_S(x + td)/t] \tag{5.1.6}$$

(where the second form relies on the fact that d_S is nonnegative). In a set-formulation as before, we see that (5.1.6) means: for any $\varepsilon > 0$, there exists $\delta > 0$ such that

$$d \in \frac{S - x}{t} + B(0, \varepsilon) \quad \text{for all } 0 < t \leqslant \delta.$$

We will see in §5.2 below that the pair of alternatives (5.1.5) – (5.1.6) is irrelevant for our purpose, because both definitions coincide when S is convex. □

Remark 5.1.6 Still another "tangent cone" would also be possible: one says that d is a *feasible direction* for S at $x \in S$ when there exists $\delta > 0$ such that $x + td \in S$ [i.e. $d \in (S - x)/t$] for all $0 < t \leqslant \delta$.

Once again, we will see that the difference is of little interest: when S is convex, $T_S(x)$ is the closure of the cone of feasible directions thus defined. □

5.2 The Tangent and Normal Cones to a Convex Set

Instead of a general set S, we now consider a *closed convex* set $C \subset \mathbb{R}^n$. In this restricted situation, the tangent cone can be given a more handy expression. The key is to observe that the role of the property $t_k \downarrow 0$ is special in (5.1.5): when both x and $x + t_k d_k$ are in C, then $x + \tau d_k \in C$ for all $\tau \in]0, t_k]$. In particular, $C \subset \{x\} + \mathrm{T}_C(x)$. Indeed, the tangent cone has a *global* character:

Proposition 5.2.1 *The tangent cone to a closed convex set C at $x \in C$ is the closure of the cone generated by $C - \{x\}$:*

$$\mathrm{T}_C(x) = \overline{\mathrm{cone}} \, (C - x) = \mathrm{cl} \, [\mathbb{R}^+ (C - x)]$$
$$= \mathrm{cl} \, \{d \in \mathbb{R}^n \; : \; d = \alpha(y - x), \; y \in C, \; \alpha \geqslant 0\} . \tag{5.2.1}$$

It is therefore convex.

Proof. We have just said that $C - \{x\} \subset \mathrm{T}_C(x)$. Because $\mathrm{T}_C(x)$ is a closed cone (Proposition 5.1.3), it immediately follows that $\mathrm{cl} \, [\mathbb{R}^+ (C - x)] \subset \mathrm{T}_C(x)$. Conversely, for $d \in \mathrm{T}_C(x)$, take (x_k) and (t_k) as in the definition (5.1.1): the point $(x_k - x)/t_k$ is in $\mathbb{R}^+ (C - x)$, hence its limit d is in the closure of this latter set. $\quad\square$

Remark 5.2.2 This new definition is easier to work with – and to master. Furthermore, it strongly recalls Remark 2.2.4: the term in curly brackets in (5.2.1) is just a union,

$$\mathrm{cone} \, (C - x) := \bigcup_{t > 0} \frac{C - x}{t}$$

and, thanks to the monotonicity property of the "difference quotient" $t \mapsto (C - x)/t$, it is also a limit of nested sets:

$$\mathrm{cone} \, (C - x) = \lim_{t \downarrow 0} \frac{C - x}{t}$$

to be compared with the definition (2.2.3) of the asymptotic cone. Having taken a union, or a limit, the closure operation is now necessary, but it was not when we took an intersection. Also, the limit above is unambiguous (it is a union!), and can be understood as the lim ext or the lim int; see Remark 5.1.5. As for Remark 5.1.6, we see that the cone of feasible directions for the convex C at x is just the very last set in brackets in (5.2.1). $\quad\square$

As a closed convex set, $\mathrm{T}_C(x)$ can also be described as an intersection of closed half-spaces – remember §4.2(b). In the present conical situation, the half-spaces can be taken as homogeneous – remember §4.2(d).

Definition 5.2.3 (Normal Cone) The direction $s \in \mathbb{R}^n$ is said to be *normal* to C at $x \in C$ when

$$\langle s, y - x \rangle \leqslant 0 \quad \text{for all } y \in C . \tag{5.2.2}$$

The set of all such directions is called *normal cone* to C at x, denoted by $\mathrm{N}_C(x)$.

$\square$

That $N_C(x)$ is a closed convex cone is clear enough. A normal is a vector s such that the angle between s and $y - x$ is obtuse for all $y \in C$. A consequence of §4.2(a) is that there is a nonzero normal at each x of $\operatorname{bd} C$. Indeed, Theorem 3.1.1 tells us that $v - \operatorname{p}_C(v) \in N_C(\operatorname{p}_C(v))$ for all $v \in \mathbb{R}^n$.

By contrast, $N_C(x) = \{0\}$ for $x \in \operatorname{int} C$. As an example, for a closed half-space $C = H^-_{s,r} = \{y \in \mathbb{R}^n : \langle s, y \rangle \leqslant r\}$, the normals at any point of $H_{s,r}$ are the nonnegative multiples of s.

Proposition 5.2.4 *The normal cone is the polar of the tangent cone.*

Proof. If $\langle s, d \rangle \leqslant 0$ for all $d \in C - x$, the same holds for all $d \in \mathbb{R}^+(C - x)$, as well as for all d in the closure $T_C(x)$ of the latter. Thus, $N_C(x) \subset [T_C(x)]^\circ$.

Conversely, take s arbitrary in $[T_C(x)]^\circ$. The relation $\langle s, d \rangle \leqslant 0$, which holds for all $d \in T_C(x)$, a fortiori holds for all $d \in C - x \subset T_C(x)$; this is just (5.2.2). $\qquad\square$

Knowing that the tangent cone is closed, this result can be combined with Proposition 4.2.6 to obtain a third definition, in addition to Definition 5.1.1 and Proposition 5.2.1:

Corollary 5.2.5 *The tangent cone is the polar of the normal cone:*

$$T_C(x) = \{d \in \mathbb{R}^n : \langle s, d \rangle \leqslant 0 \text{ for all } s \in N_C(x)\}. \qquad\square$$

This describes $T_C(x)$ as an intersection of homogeneous half-spaces and the relationship with §4.2(b) is clear. With the notation thereof, $r_s = 0$ for each s, and the index-set $\Sigma^0_{T_C(x)}$ is nothing more than $N_C(x)$; see again §4.2(d).

It is interesting to note here that normality is again a local concept, even though (5.2.2) does not suggest it. Indeed the normal cone at x to $C \cap B(x, \delta)$ coincides with $N_C(x)$. Also, if C' is "sandwiched", i.e. if $C - \{x\} \subset C' - \{x\} \subset T_C(x)$, then $N_{C'}(x) = N_C(x)$ – and $T_{C'}(x) = T_C(x)$. Let us add that tangent and normal cones to a nonclosed convex set C could be defined if needed: just replace C by $\operatorname{cl} C$ in the definitions.

Another remark is that the tangent and normal cones are "homogeneous" objects, in that they contain 0 as a distinguished element. It is most often the translated version $x + T_C(x)$ that is used and visualized; see Fig. 5.1.1 again.

Let us summarize the concepts introduced so far.
- First, there is the tangent cone, defined via a difference quotient, either set-valued $(C - x)/t$, or using the distance: $\operatorname{d}_C(x + t \cdot)/t$. This cone has several equivalent definitions: (i) as the outer limit of $(C - x)/t$ (i.e. using $\liminf \operatorname{d}_C(x + t \cdot)/t$), (ii) as the inner limit of $(C - x)/t$ (i.e. using $\limsup \operatorname{d}_C(x + t \cdot)/t$), or (iii) as the closed conical hull of $C - x$.
- Then there is the normal cone, which is the polar of the tangent cone.
- Alternatively, one could define first the normal cone, as the cone polar to $C - x$ (or to $\mathbb{R}^+(C - x)$, knowing that polarity does not distinguish a set from its closed conical hull). Then the tangent cone would be defined as the polar to this normal cone.

Needless to say, it is *convexity* of C that makes all these definitions equivalent. If C were not convex, we would obtain just as many different objects.

Examples 5.2.6 *(a).* If $C = K$ is a closed convex cone, $\mathrm{T}_K(0) = K$: the polar K° of a closed convex cone is its normal cone at 0. On the other hand, if $0 \neq x \in K$, then $\mathrm{T}_K(x)$ contains at least one subspace, namely $\mathbb{R}\{x\}$. Actually, we even have

$$\mathrm{N}_K(x) = \{s \in K^\circ \ : \ \langle s, x \rangle = 0\} \quad \text{for } x \neq 0.$$

To see it, observe that $\mathrm{T}_K(x) \supset \{x\}$ means $\mathrm{N}_K(x) \subset \{x\}^\perp$; in other words, the relation "$\langle s, y - x \rangle \leqslant 0$ for all $y \in K$" defining $\mathrm{N}_K(x)$ reduces to "$\langle s, y \rangle \leqslant 0$ $[= \langle s, x \rangle]$ for all $y \in K$".

A cone is a set of vectors defined up to a multiplicative constant, and the value of this constant has often little relevance. As far as the concepts of polarity, tangency, normality are concerned, for example, a closed convex cone T (or N) could equally be replaced by the compact $T \cap B(0,1)$; or also by $\{x \in T : \|x\| = 1\}$, in which the redundancy is totally eliminated.

(b). Take a closed convex polyhedron defined by m *constraints*:

$$C := \{x \in \mathbb{R}^n \ : \ \langle s_j, x \rangle \leqslant r_j \text{ for } j = 1, \ldots, m\} \tag{5.2.3}$$

and define
$$J(x) := \{j = 1, \ldots, m : \langle s_j, x \rangle = r_j\}$$

the index-set of *active constraints* at $x \in C$. Then

$$\mathrm{T}_C(x) = \{d \in \mathbb{R}^n \ : \ \langle s_j, d \rangle \leqslant 0 \text{ for } j \in J(x)\},$$

$$\mathrm{N}_C(x) = \text{cone}\,\{s_j \ : \ j \in J(x)\} = \Big\{\sum_{j \in J(x)} \alpha_j s_j \ : \ \alpha_j \geqslant 0\Big\}.$$

(c). Let C be the unit simplex Δ_n of Example 1.1.3 and $\alpha = (\alpha_1, \ldots, \alpha_n) \in \Delta_n$. If each α_i is positive, i.e. if $\alpha \in \mathrm{ri}\,\Delta_n$, then the tangent cone to Δ_n at α is aff $\Delta_n - \{\alpha\}$, i.e. the linear hyperplane of equation $\sum_{i=1}^n \alpha_i = 0$. Otherwise, with $e := (1, \ldots, 1) \in \mathbb{R}^n$:

$$\mathrm{T}_{\Delta_n}(\alpha) = \{d = (d_1, \ldots, d_n) \ : \ e^\top d = 0, \ d_i \geqslant 0 \text{ if } \alpha_i = 0\}.$$

Using Example (b) above, calling $\{e_1, \ldots, e_n\}$ the canonical basis of $\mathbb{R}^n$ and denoting by $J(\alpha) := \{j : \alpha_j = 0\}$ the active set at α, we obtain the normal cone:

$$\mathrm{N}_{\Delta_n}(\alpha) = \text{cone}\,\big[\{e\} \cup \{-e\} \cup_{j \in J(\alpha)} \{-e_j\}\big]$$
$$= \big\{\textstyle\sum_{j \in \{0\} \cup J(\alpha)} \beta_j e_j \ : \ \beta_j \leqslant 0 \text{ for } j \in J(\alpha)\big\}. \qquad \square$$

5.3 Some Properties of Tangent and Normal Cones

Let us give some properties of the tangent cone, which result directly from Proposition 5.2.1.

- For fixed x, $T_C(x)$ increases if and only if $N_C(x)$ decreases, and these properties happen in particular when C increases.
- The set cone $(C - x)$ and its closure $T_C(x)$ have the same affine hull (actually linear hull!) and the same relative interior. It is not difficult to check that these last sets are (aff $C - x$) and $\mathbb{R}_*^+$ (ri $C - x$) respectively.
- $T_C(x) = $ aff $C - x$ whenever $x \in$ ri C (in particular, $T_C(x) = \mathbb{R}^n$ if $x \in$ int C). As a result, approximating C by $x + T_C(x)$ presents some interest only when $x \in$ rbd C. Along this line, we warn the reader against a too sloppy comparison of the tangent cone with the concept of tangency to a surface: with this intuition in mind, one should rather think of the (relative) *boundary* of C as being approximated by the (relative) boundary of $T_C(x)$.
- The concept of tangent cone fits rather well with the convexity-preserving operations of §1.2. Validating the following calculus rules involves elementary arguments only, and is left as an exercise.

Proposition 5.3.1 *Here, the C's are nonempty closed convex sets.*

(i) *For $x \in C_1 \cap C_2$, there holds*

$$T_{C_1 \cap C_2}(x) \subset T_{C_1}(x) \cap T_{C_2}(x) \quad and \quad N_{C_1 \cap C_2}(x) \supset N_{C_1}(x) + N_{C_2}(x).$$

(ii) *With $C_i \subset \mathbb{R}^{n_i}$, $i = 1, 2$ and $(x_1, x_2) \in C_1 \times C_2$,*

$$T_{C_1 \times C_2}(x_1, x_2) = T_{C_1}(x_1) \times T_{C_2}(x_2),$$

$$N_{C_1 \times C_2}(x_1, x_2) = N_{C_1}(x_1) \times N_{C_2}(x_2).$$

(iii) *With an affine mapping $A(x) = y_0 + A_0 x$ (A_0 linear) and $x \in C$,*

$$T_{A(C)}[A(x)] = \text{cl}\,[A_0 T_C(x)] \quad and \quad N_{A(C)}[A(x)] = \overset{-1}{A_0^*}[N_C(x)].$$

(iv) *In particular (start from (ii), (iii) and proceed as when proving (1.2.2)):*

$$T_{C_1 + C_2}(x_1 + x_2) = \text{cl}\,[T_{C_1}(x_1) + T_{C_2}(x_2)],$$

$$N_{C_1 + C_2}(x_1 + x_2) = N_{C_1}(x_1) \cap N_{C_2}(x_2). \qquad \Box$$

Remark 5.3.2 To obtain equality in (i), an additional qualification assumption is necessary. One was used in Proposition 2.1.10, see also (2.1.5):

$$0 \in \text{ri}\,(C_1 - C_2) \quad \text{or} \quad (\text{ri}\,C_1) \cap (\text{ri}\,C_2) \neq \emptyset \qquad (5.3.1)$$

(the proof of the corresponding statement becomes a bit longer). $\qquad \Box$

Some more properties of tangent and normal cones are worth mentioning, which patch together various notions seen earlier in this chapter.

Proposition 5.3.3 *For $x \in C$ and $s \in \mathbb{R}^n$, the following properties are equivalent:*

(i) $s \in N_C(x)$;

(ii) x is in the *exposed face* $F_C(s)$: $\langle s, x \rangle = \max_{y \in C} \langle s, y \rangle$;

(iii) $x = p_C(x + s)$.

Proof. Nothing really new: everything comes from the definitions of normal cones, supporting hyperplanes, exposed faces, and the characteristic property (3.1.3) of the projection operator. □

This result is illustrated on Fig. 5.3.1 and implies in particular:

$$p_C^{-1}(x) = \{x\} + N_C(x) \quad \text{for all } x \in C .$$

Also,

$$x \neq x' \quad \implies \quad [\{x\} + N_C(x)] \cap [\{x'\} + N_C(x')] = \emptyset$$

(otherwise the projection would not be single-valued).
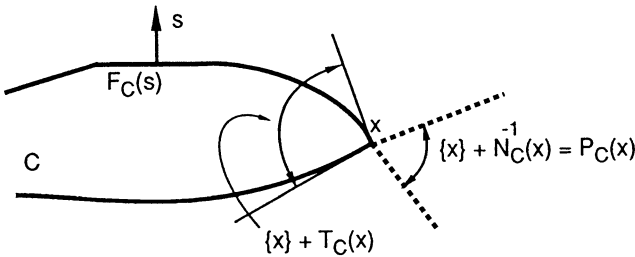
Fig. 5.3.1. Normal cones, projections and exposed faces

Remark 5.3.4 Let us come back again to Fig. 4.2.2. In a first step, fix $x \in C$ and consider only those supporting hyperplanes that pass through x, i.e. those indexed in $N_C(x)$. The corresponding intersection of half-spaces just constructs $T_C(x)$ and ends up with $\{x\} + T_C(x) \supset C$. Note in passing that the closure operation of Proposition 5.2.1 is necessary when $\text{rbd}\, C$ presents some curvature near x.

Then, in a second step, do this operation for all $x \in C$:

$$C \subset \bigcap_{x \in C} [x + T_C(x)] . \tag{5.3.2}$$

A first observation is that x can actually be restricted to the relative boundary of C: for $x \in \text{ri}\, C$, $T_C(x)$ expands to the whole aff $C - x$ and contains all other tangent cones. A second observation is that (5.3.2) actually holds as an equality. In fact, write a point $y \notin C$ as $y = p_C(y) + s$ with $s = y - p_C(y)$; Proposition 5.3.3 tells us that $s \in N_C[p_C(y)]$, hence the nonzero s cannot be in $T_C[p_C(y)]$: we have established $y \notin p_C(y) + T_C[p_C(y)]$, y is not in the righthand side of (5.3.2). In a word:

$$C = \bigcap_{x \in \text{rbd}\, C} [x + T_C(x)] ,$$

which sheds some more light on the outer description of C discussed in §4.2(b). □

Exercises

1. In each of the three cases below, the question is whether $D \subset \mathbb{R}^n$ is convex:
– when $C_1 \subset D \subset C_2$, with C_1 and C_2 convex;
– when the projection of D onto any affine manifold of $\mathbb{R}^n$ is convex;
– when D is closed and satisfies the "mid-point property": $\dfrac{x + y}{2} \in D$ whenever x and y lie in D;
– when $D = \{x : \|x - \bar{x}\| + \|x - y\| \leqslant 1 \text{ for all } y \in S\}$ $(\bar{x} \in \mathbb{R}^n, S \subset \mathbb{R}^n$ arbitrary);
– when $D = \{x : x + S \subset C\}$ (S arbitrary, C convex).

2*. Express the closed convex set $\{(\xi, \eta) \in \mathbb{R}^2 : \xi > 0, \eta > 0, \xi\eta \geqslant 1\}$ as an intersection of half-spaces. What is the set of directions exposing some point in it? Compute the points exposed by such directions (draw a picture).

3**. Give two closed convex cones K_1 and K_2 in $\mathbb{R}^3$ such that $K_1 + K_2$ is not closed (start from two sets in $\mathbb{R}^2$).

4. For C closed convex, establish $C_\infty = \cap_{\varepsilon > 0} \mathrm{cl} \left(\cup_{0 < t \leqslant \varepsilon} tC\right)$.

5**. In the vector space $S_n(\mathbb{R})$, consider the set $\mathcal{E}$ of positive semi-definite matrices $A = [A_{ij}]$ such that $A_{ii} = 1$ for $i = 1, \ldots, n$ (these are called *correlation* matrices).
    If V is a subspace of $\mathbb{R}^n$, use the notation $F_V := \{M \in \mathcal{E}_n : V \subset \mathrm{Ker}\, M\}$.
– Show that F_V is a face of $\mathcal{E}_n$.
– Let F be a face of $\mathcal{E}_n$. Show that $F = F_V$, with $V = \cap_{M \in F} \mathrm{Ker}\, M$.
– Let $M_0 \in \mathcal{E}_n$ be given. Show that the smallest face of $\mathcal{E}_n$ containing M_0 is the set $\{M \in \mathcal{E}_n : \mathrm{Ker}\, M_0 \subset \mathrm{Ker}\, M\}$.
    Show that every face of $\mathcal{E}_n$ is exposed.

6*. Here C and A are two closed sets such that $C \subset A$.
– Show that $p_C \circ p_A = p_C$ if C and A are two subspaces (draw a picture; this is a well-known theorem in elementary geometry).
– Show the same property when C is convex, A being an affine manifold (frequently used with $A = \mathrm{aff}\, C$).
– Show on an example that the property need not hold under mere convexity of C and A.

7*. Show that the polar of $K := \{x = (\xi^1, \ldots, \xi^n) : \xi^1 \leqslant \xi^2 \leqslant \cdots \leqslant \xi^n\}$ (with $n \geqslant 2$) is

$$K^\circ = \left\{y = (\eta^1, \ldots, \eta^n) : \sum_{i=1}^k \eta^i \geqslant 0, \ k = 1, \ldots, n-1 \text{ and } \sum_{i=1}^n \eta^i = 0\right\}.$$

8*. For two integers m and n such that $m \leqslant n$, give all the extreme points of

$$\Pi_m := \left\{\alpha = (\alpha^1, \ldots, \alpha^n) : 0 \leqslant \alpha^i \leqslant 1, \ i = 1, \ldots, n \text{ and } \sum_{i=1}^n \alpha^i = m\right\}$$

(a picture with $n = 3$ will be helpful).

9. Let A_0, A_1,..., A_m be symmetric $n \times n$ matrices. For $x = (\xi^1, \ldots, \xi^n)$, set $A(x) := A_0 + \sum_{i=1}^{m} \xi^i A_i$ and define $C := \{x \in \mathbb{R}^m : A(x) \succcurlyeq 0\}$. Using different techniques, show that C is closed convex.

10. For two nonempty sets A and B in $\mathbb{R}^n$, show that $\operatorname{co}(A + B) = \operatorname{co} A + \operatorname{co} B$.

11*. Given a set S, show that every extreme point of $\operatorname{co} S$ lies in S.

12*. Let C be convex in $\mathbb{R}^n$ and $A : \mathbb{R}^n \to \mathbb{R}^m$ be an injective affine mapping. Show that $A(x)$ is an extreme point of $A(C)$, for each extreme point x in C. Check on an example that the result need not hold without injectivity of A.

13*. Let C be a nonempty closed convex set. Show that $y_x \in C$ is the projection of x onto C if and only if $\langle x - y, y - y_x \rangle \leqslant 0$ for all $y \in C$ (draw a picture).

14*. Let $(C_k)_k$ be a decreasing sequence of closed convex sets, with $C := \cap_k C_k \neq \emptyset$. Show that, for all $x \in \mathbb{R}^n$,

$$\operatorname{p}_{C_k}(x) \to \operatorname{p}_C(x) \text{ and } \operatorname{d}_{C_k}(x) \uparrow \operatorname{d}_C(x) \quad \text{when } k \to +\infty.$$

15*. Let $K := \operatorname{S}_n^+(\mathbb{R})$ denote the (closed convex) cone of positive semi-definite matrices in the space $\operatorname{S}_n(\mathbb{R})$ equipped with the standard scalar product $\langle M, P \rangle := \operatorname{tr} MP$.
– Show that $K^\circ = -K$, a set which can be denoted by $\operatorname{S}_n^-(\mathbb{R})$.
– Characterize the projections $\operatorname{p}_K(M)$ and $\operatorname{p}_{K^\circ}(M)$ of a given matrix M.
– Deduce the Moreau decomposition of M onto K and K° (Thm 3.2.5).

16. If S is an open set, show that $\operatorname{co} S$ is open as well.

17*. Give an example of two closed convex sets C_1 and C_2, such that $C_1 + C_2$ is closed, and $(C_1)_\infty + (C_2)_\infty$ is stricly included in $(C_1 + C_2)_\infty$.

18*. Let $C \subset \mathbb{R}^n$ be nonempty closed convex and $A : \mathbb{R}^n \to \mathbb{R}^m$ be linear. Show that $A(C)$ is closed whenever $C_\infty \cap \operatorname{Ker} A = \{0\}$. As an application, find a sufficient condition for the projection of a closed convex set onto a subspace to be closed.

19. Under which condition is 0 an extreme point of a closed convex cone?
– Show that any x of norm 1 is an extreme point of $B(0, 1)$ (Euclidean unit ball).
– What are the extreme points of $B_1(0, 1)$, the unit ball of the ℓ_1-norm $\|x\|_1 := \sum_{i=1}^{n} |\xi^i|$?
– Same question for $B_\infty(0, 1)$ ($\|x\|_\infty := \max\{|\xi^1|, ..., |\xi^n|\}$). Compare the growth of the number of extreme points, respectively of $B_1(0, 1)$ and $B_\infty(0, 1)$, when $n \to +\infty$.

20*. The diameter of a set S is $\sup\{\|y - x\| : x, y \in S\}$. Show that S and $\operatorname{co} S$ have the same diameter if S is bounded.

21. Show the following relations

$$\text{co}\,(S_1 \cap S_2) \subset (\text{co}\,S_1) \cap (\text{co}\,S_2)\,,$$
$$\text{co}\,(S_1 + S_2) = \text{co}\,S_1 + \text{co}\,S_2\,,$$
$$\text{co}\,(S_1 \cup S_2) = \text{co}\,[(\text{co}\,S_1) \cup (\text{co}\,S_2)]\,;$$

Give an example in which equality does not hold for the intersection.

Show that this equality does hold in the following situation: $S_2 = H$ is a hyperplane and S_1 is contained in H^-, one of the corresponding closed half-spaces.

22. For C convex, show that C is closed if and only if $C \cap L$ is closed, for any affine line L.

23. A closed convex cone K is said to be *acute* when $\langle x, x' \rangle \geqslant 0$ for all x and x' in K.
 – Show that K is acute if and only if $K \subset -K^\circ$.
 – For $p \geqslant 1$, define $K_p := \{(\xi^1, \ldots, \xi^n) : \xi^n \geqslant (\sum_{k=1}^{n-1} |\xi^k|^p)^{1/p}\}$.
   – What is $(K_p)^\circ$?
   – For which values of p is K_p acute?

*24**. For C convex compact, show that $C = \text{co}\,(\text{rbd}\,C)$.

*25**. In the space of $n \times n$ matrices, consider the affine subspace

$$V_n := \{M = [m_{ij}] \; : \; \sum_j m_{ij} = \sum_i m_{ij} = 1\}$$

and the convex compact set (of so-called *bistochastic* matrices of size n)

$$\Sigma_n := \{M \in V_n \; : \; m_{ij} \geqslant 0 \text{ for } i, j = 1, \ldots, n\}\,.$$

 – Show that V_n has dimension $(n-1)^2$ and that aff $\Sigma_n = V_n$.
 – Set $J_n := [\frac{1}{n}]$ (the bistochastic matrix whose all elements are equal; it is a "central" matrix in Σ_n). Check that J_n lies in the relative interior of Σ_n, and that $J_n M = M J_n = J_n$ for all $M \in \Sigma_n$ (J_n "absorbs" the whole of Σ_n).
 – Now we equip the space of matrices with the standard scalar product $\langle M, P \rangle := \text{tr}\,M^\top P$. Denoting by $e \in \mathbb{R}^n$ the vector whose all component are 1, show that $V_n^\perp = \{ue^\top + ev^\top : u \in \mathbb{R}^n, v \in \mathbb{R}^n\}$.
 – Deduce that the orthogonal projection of a matrix M onto V_n is $J_n + K_n M K_n$, where $K_n := I_n - J_n$ (I_n is the identity matrix).

*26**. Let P be a closed convex polyhedron: $P = \{x : Ax = b, x \geqslant 0\}$.
 – For $0 \neq x \in P$, show that x is extremal in P if and only if the columns i of A corresponding to $\xi^i > 0$ are linearly independent.
 – Use this result to prove the following theorem (G. Birkhoff, 1946): the set of bistochastic matrices (see above) is convex compact, and its extreme points are the permutation matrices.

B. Convex Functions

Introduction. The study of convex functions goes together with that of convex sets; accordingly, this chapter and the previous one constitute the first serious steps into the world of convex analysis. This chapter has no pretension to exhaustivity; similarly to Chap. A, it has been kept minimal, containing what is necessary to comprehend the sequel.

1 Basic Definitions and Examples

1.1 The Definitions of a Convex Function

Definition 1.1.1 Let C be a nonempty convex set in $\mathbb{R}^n$. A function $f : C \to \mathbb{R}$ is said to be *convex* on C when, for all pairs $(x, x') \in C \times C$ and all $\alpha \in]0, 1[$, there holds

$$f(\alpha x + (1 - \alpha)x') \leqslant \alpha f(x) + (1 - \alpha)f(x') . \qquad (1.1.1)$$

$\square$

We say that f is *strictly convex* on C when (1.1.1) holds as a strict inequality if $x \neq x'$. An even stronger property is that there exists $c > 0$ such that

$$f(\alpha x + (1 - \alpha)x') \leqslant \alpha f(x) + (1 - \alpha)f(x') - \tfrac{1}{2}c\,\alpha(1 - \alpha)\|x - x'\|^2 \quad (1.1.2)$$

for all $(x, x') \in C \times C$ and all $\alpha \in]0, 1[$. In this case, f is said to be *strongly convex* on C (with *modulus* of strong convexity c). Passing from (1.1.1) to (1.1.2) does not change much the class of functions considered:

Proposition 1.1.2 *The function f is strongly convex on C with modulus c if and only if the function $f - 1/2\,c\|\cdot\|^2$ is convex on C.*

Proof. Use direct calculations in the definition (1.1.1) of convexity applied to the function $f - 1/2\,c\|\cdot\|^2$, namely:

$$f(\alpha x + (1 - \alpha)x') - \tfrac{1}{2}c\|\alpha x + (1 - \alpha)x'\|^2 \leqslant$$
$$\leqslant \alpha f(x) + (1 - \alpha)f(x') - \tfrac{1}{2}c\big[\alpha\|x\|^2 + (1 - \alpha)\|x'\|^2\big] . \qquad \square$$

Although simple, this statement illustrates a useful technique in convex analysis: to prove that a convex function has a certain property, one establishes a related property on a suitable strongly convex perturbation of the given function.

The set C needed in Definition 1.1.1 (which can be the whole space) appears as a sort of domain of definition of f. Of course, it has to be convex so that the lefthand side of (1.1.1) makes sense. In a more modern definition, a convex function f is considered as defined on the whole of $\mathbb{R}^n$, but possibly taking infinite values:

Definition 1.1.3 (The Set Conv $\mathbb{R}^n$) A function $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, not identically $+\infty$, is said to be convex when, for all $(x, x') \in \mathbb{R}^n \times \mathbb{R}^n$ and all $\alpha \in \,]0, 1[$, there holds
$$f(\alpha x + (1 - \alpha)x') \leqslant \alpha f(x) + (1 - \alpha)f(x') \,,$$
considered as an inequality in $\mathbb{R} \cup \{+\infty\}$.

The class of such functions is denoted by $\text{Conv}\,\mathbb{R}^n$.   □

See §0.2 for a short introduction to the set $\mathbb{R} \cup \{+\infty\}$. We mention here that our definition coincides with that of *proper* convexity used by other authors. The distinction is necessary when the value $f(x) = -\infty$ is allowed; but this value is excluded from the very beginning in the present book.

To realize the equivalence between our two definitions, extend an f from Definition 1.1.1 by
$$f(x) := +\infty \quad \text{for} \quad x \notin C \,, \tag{1.1.3}$$
thus obtaining a new f, which is now in $\text{Conv}\,\mathbb{R}^n$. Conversely, consider the following definition (meaningful even for nonconvex f, incidentally):

Definition 1.1.4 (Domain of a Function) The *domain* (or also effective domain) of $f \in \text{Conv}\,\mathbb{R}^n$ is the nonempty set
$$\text{dom}\, f := \{x \in \mathbb{R}^n \,:\, f(x) < +\infty\} \,.   □$$

Clearly enough, an f satisfying (1.1.1), (1.1.3) has a convex domain; given $f \in \text{Conv}\,\mathbb{R}^n$, we can therefore take $C := \text{dom}\, f$ to obtain a convex function in the sense of Definition 1.1.1. Strong convexity is also defined in the spirit of Definition 1.1.3, via (1.1.2) with x and x' varying in $\text{dom}\, f$ or in $\mathbb{R}^n$: it makes no difference. Same remark for strict convexity (checking all these claims is a good exercise to familiarize oneself with computations in $\mathbb{R} \cup \{+\infty\}$).

Now, we recall that the *graph* of an arbitrary function is the set of couples $(x, f(x))$ in $\mathbb{R}^n \times \mathbb{R}$. When moving to the unilateral world of convex analysis, the following is relevant:

Definition 1.1.5 (Epigraph of a Function) Given $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, not identically equal to $+\infty$, the *epigraph* of f is the nonempty set
$$\text{epi}\, f := \{(x, r) \in \mathbb{R}^n \times \mathbb{R} \,:\, r \geqslant f(x)\} \,.$$

Its *strict epigraph* $\text{epi}_s\, f$ is defined likewise, with "$\geqslant$" replaced by "$>$" (beware that the word "strict" here has nothing to do with strict convexity).   □

Defining a *sublevel-sets* (see §0.1.2) by

$$S_r(f) := \{x \in \mathbb{R}^n : f(x) \leqslant r\},$$

we have the equivalent definition

$$(x, r) \in \text{epi}\, f \quad \Longleftrightarrow \quad x \in S_r(f). \tag{1.1.4}$$

The following property is easy to derive, and can be interpreted as giving one more definition of convex functions, which is now of *geometric* nature.

Proposition 1.1.6 *Let* $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ *be not identically equal to* $+\infty$. *The three properties below are equivalent:*

  (i) *f is convex in the sense of Definition* 1.1.3;
 (ii) *its epigraph is a convex set in* $\mathbb{R}^n \times \mathbb{R}$;
(iii) *its strict epigraph is a convex set in* $\mathbb{R}^n \times \mathbb{R}$.

Proof. Left as an exercise.   □

We say that f is *concave* when $-f$ is convex, or equivalently when the *hypograph* of f (revert the inequality in Definition 1.1.5) is a convex set. We will see on examples that either the analytical Definition 1.1.3 or the geometric one coming from 1.1.6 may be more convenient, depending on the situation.

Remark 1.1.7 The sublevel-sets of $f \in \text{Conv}\,\mathbb{R}^n$ are convex (possibly empty) subsets of $\mathbb{R}^n$. To construct $S_r(f)$, remember Example A.1.2.5: we cut the epigraph of f by a horizontal blade, forming the intersection $\text{epi}\,f \cap (\mathbb{R}^n \times \{r\})$ of two convex sets; then we project the result down to $\mathbb{R}^n \times \{0\}$ and we change the environment space from $\mathbb{R}^n \times \mathbb{R}$ to $\mathbb{R}^n$. Even though this latter operation changes the topology, it changes neither the closure nor the relative interior.
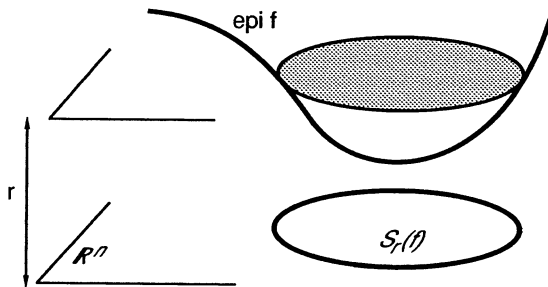
Fig. 1.1.1. Forming a sublevel-set

Conversely, a function whose sublevel-sets are all convex *need not* be convex (see Fig. 1.1.1); such a function is called *quasi-convex*.

Observe that $\text{dom}\,f$ is the union over $r \in \mathbb{R}$ of the sublevel-sets $S_r(f)$, which form a nested family; it is also the projection of $\text{epi}\,f \subset \mathbb{R}^n \times \mathbb{R}$ onto $\mathbb{R}^n$ (so Proposition A.1.2.4 and Example A.1.2.5 confirm its convexity).   □

The basic inequality (1.1.1) can be generalized to convex combinations of more than two points:

Theorem 1.1.8 (Inequality of Jensen) *Let $f \in \text{Conv }\mathbb{R}^n$. Then, for all collections $\{x_1, \ldots, x_k\}$ of points in $\text{dom } f$ and all $\alpha = (\alpha_1, \ldots, \alpha_k)$ in the unit simplex of $\mathbb{R}^k$, there holds (inequality of Jensen in summation form)*

$$f\left(\textstyle\sum_{i=1}^{k} \alpha_i x_i\right) \leqslant \textstyle\sum_{i=1}^{k} \alpha_i f(x_i).$$

Proof. The k points $(x_i, f(x_i)) \in \mathbb{R}^n \times \mathbb{R}$ are clearly in $\text{epi } f$, a convex set. Their convex combination

$$\textstyle\sum_i \alpha_i (x_i, f(x_i)) = \left(\sum_i \alpha_i x_i, \sum_i \alpha_i f(x_i)\right)$$

is also in $\text{epi } f$ (Proposition A.1.3.3). This is just the claimed inequality.   □

Starting from $f \in \text{Conv }\mathbb{R}^n$, we have constructed via Definition 1.1.5 the convex set $\text{epi } f$. Conversely, if $E \subset \mathbb{R}^n \times \mathbb{R}$ is the epigraph of a function in $\text{Conv }\mathbb{R}^n$, this function is directly obtained from $f(x) = \inf_{(x,r) \in E} r$ (recall that $\inf \emptyset = +\infty$; we will see in §1.3(g) what sets are epigraphs of a convex function). In view of this correspondence, the properties of a convex function f are intimately related to those developed in Chap. A, applied to $\text{epi } f$.

We will see later that important functions are those having a closed epigraph. Also, it is clear that $\text{aff epi } f$ contains the vertical lines $\{x\} \times \mathbb{R}$, with $x \in \text{dom } f$. This shows that $\text{epi } f$ cannot be an open set, nor relatively open: take points of the form $(x, f(x) - \varepsilon)$. As a result, $\text{ri epi } f$ cannot be an epigraph, but it is nevertheless of interest to see how this set is constructed:

Proposition 1.1.9 *Let $f \in \text{Conv }\mathbb{R}^n$. The relative interior of $\text{epi } f$ is the union over $x \in \text{ri dom } f$ of the open half-lines with bottom endpoints at $f(x)$:*

$$\text{ri epi } f = \{(x, r) \in \mathbb{R}^n \times \mathbb{R} : x \in \text{ri dom } f, \ r > f(x)\}.$$

Proof. Since $\text{dom } f$ is the image of $\text{epi } f$ under the linear mapping "projection onto $\mathbb{R}^n$", Proposition A.2.1.12 tells us that

$$\text{ri dom } f \text{ is the projection onto } \mathbb{R}^n \text{ of ri epi } f. \tag{1.1.5}$$

Now take x arbitrary in $\text{ri dom } f$. The subset of $\text{ri epi } f$ that is projected onto x is just $(\{x\} \times \mathbb{R}) \cap \text{ri epi } f$, which in turn is $\text{ri}\left[(\{x\} \times \mathbb{R}) \cap \text{epi } f\right]$ (use Proposition A.2.1.10). This latter set is clearly $]f(x), +\infty[$.

In summary, we have proved that, for $x \in \text{ri dom } f$, $(x, r) \in \text{ri epi } f$ if and only if $r > f(x)$. Together with (1.1.5), this proves our claim.   □

Beware that $\text{ri epi } f$ is *not* the strict epigraph of f (watch the side-effect on the relative boundary of $\text{dom } f$).

1.2 Special Convex Functions: Affinity and Closedness

In view of their Definition 1.1.6(ii), convex functions can be classified on the basis of a classification of convex sets in $\mathbb{R}^n \times \mathbb{R}$.

(a) Linear and Affine Functions The epigraph of a linear function is characterized by $s \in \mathbb{R}^n$, and is made of those $(x, r) \in \mathbb{R}^n \times \mathbb{R}$ such that $r \geqslant \langle s, x \rangle$.

Next, we find the epigraphs of *affine functions* f, which are conveniently written in terms of some $x_0 \in \mathbb{R}^n$:

$$\{(x, r) : r \geqslant f(x_0) + \langle s, x - x_0 \rangle\} = \{(x, r) : \langle s, x \rangle - r \leqslant \langle s, x_0 \rangle - f(x_0)\}.$$

In the language of convex sets, the epigraph of an affine function is a closed half-space, characterized by (a constant term and) a vector $(s, -1) \in \mathbb{R}^n \times \mathbb{R}$; the essential property of this vector is to be *non-horizontal*. Affine functions thus play a special role, just as half-spaces did in Chap. A. This explains the interest of the next result; it says a little more than Lemma A.4.2.1, and is actually of paramount importance.

Proposition 1.2.1 *Any $f \in \operatorname{Conv} \mathbb{R}^n$ is minorized by some affine function. More precisely: for any $x_0 \in \operatorname{ri} \operatorname{dom} f$, there is s in the subspace parallel to $\operatorname{aff} \operatorname{dom} f$ such that*

$$f(x) \geqslant f(x_0) + \langle s, x - x_0 \rangle \quad \text{for all } x \in \mathbb{R}^n.$$

In other words, the affine function can be forced to coincide with f at x_0.

Proof. We know that $\operatorname{dom} f$ is the image of $\operatorname{epi} f$ under the linear mapping "projection onto $\mathbb{R}^n$". Look again at the definition of an affine hull (§A.1.3) to realize that $\operatorname{aff} \operatorname{epi} f = (\operatorname{aff} \operatorname{dom} f) \times \mathbb{R}$.

Denote by V the linear subspace parallel to $\operatorname{aff} \operatorname{dom} f$, so that $\operatorname{aff} \operatorname{dom} f = \{x_0\} + V$ with x_0 arbitrary in $\operatorname{dom} f$; then we have

$$\operatorname{aff} \operatorname{epi} f = \{x_0 + V\} \times \mathbb{R}. \tag{1.2.1}$$

We equip $V \times \mathbb{R}$ and $\mathbb{R}^n \times \mathbb{R}$ with the scalar product of product-spaces.

With $x_0 \in \operatorname{ri} \operatorname{dom} f$, Proposition 1.1.9 tells us that $(x_0, f(x_0)) \in \operatorname{rbd} \operatorname{epi} f$ and we can take a nontrivial hyperplane supporting $\operatorname{epi} f$ at $(x_0, f(x_0))$: using Remark A.4.2.2 and (1.2.1), there are $s = s_V \in V$ and $\alpha \in \mathbb{R}$, not both zero, such that

$$\langle s, x \rangle + \alpha r \leqslant \langle s, x_0 \rangle + \alpha f(x_0) \tag{1.2.2}$$

for all (x, r) with $f(x) \leqslant r$. Note: this implies $\alpha \leqslant 0$ (let $r \to +\infty$!)

Because of our choice of s (in V) and x_0 (in $\operatorname{ri} \operatorname{dom} f$), we can take $\delta > 0$ so small that $x_0 + \delta s \in \operatorname{dom} f$, for which (1.2.2) gives

$$\delta \|s\|^2 \leqslant \alpha[f(x_0) - f(x_0 + \delta s)] < +\infty;$$

this shows $\alpha \neq 0$ (otherwise, both s and α would be zero). Without loss of generality, we can assume $\alpha = -1$; then (1.2.2) gives our affine function.   □

Once again, the importance of this result cannot be over-emphasized. With respect to Lemma A.4.2.1, it says that a convex epigraph is supported by a *non-vertical*

hyperplane. Among its consequences, we see that a convex function, having an affine minorant, is bounded from below on every bounded set of $\mathbb{R}^n$.

As already said on several occasions, important special convex sets are the cones, and conical epigraphs will deserve the full Chap. C. So we turn to another class: closed convex sets.

(b) Closed Convex Functions Continuity of functions implies closedness of their graphs. In the unilateral world of convex analysis, one is mainly interested with epigraphs, whose closedness corresponds to *lower* semi-continuity of functions. First, we give some material, independently of any convexity. A function f is lower semi-continuous if, for each $x \in \mathbb{R}^n$,

$$\liminf_{y \to x} f(y) \geqslant f(x). \tag{1.2.3}$$

This relation has to hold in $\mathbb{R} \cup \{+\infty\}$, which complicates things a little; so the following geometric characterizations are useful:

Proposition 1.2.2 *For $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, the following three properties are equivalent:*

  (i) *f is lower semi-continuous on $\mathbb{R}^n$;*
  (ii) *epi f is a closed set in $\mathbb{R}^n \times \mathbb{R}$;*
  (iii) *the sublevel-sets $S_r(f)$ are closed (possibly empty) for all $r \in \mathbb{R}$.*

Proof. [(i) $\Rightarrow$ (ii)] Let $(y_k, r_k)_k$ be a sequence of epi f converging to (x, r) for $k \to +\infty$. Since $f(y_k) \leqslant r_k$ for all k, the l.s.c. relation (1.2.3) readily gives

$$r = \lim r_k \geqslant \liminf f(y_k) \geqslant \liminf_{y \to x} f(y) \geqslant f(x),$$

i.e. $(x, r) \in$ epi f.

[(ii) $\Rightarrow$ (iii)] Construct the sublevel-sets $S_r(f)$ as in Remark 1.1.7: the closed sets epi f and $\mathbb{R}^n \times \{r\}$ have a closed intersection.

[(iii) $\Rightarrow$ (i)] Suppose f is not lower semi-continuous at some x: there is a (sub)sequence (y_k) converging to x such that $f(y_k)$ converges to $\rho < f(x) \leqslant +\infty$. Pick $r \in]\rho, f(x)[$: for k large enough, $f(y_k) \leqslant r < f(x)$; hence $S_r(f)$ contains the tail of (y_k) but not its limit x. Consequently, this $S_r(f)$ is not closed.   □

Beware that, with Definition 1.1.1 in mind, the above statement (i) means more than lower semi-continuity of the restriction of f to C: in (1.2.3), x need not be in dom f. Note also that these concepts and results are independent from convexity. Thus, we are entitled to consider the following definition:

Definition 1.2.3 (Closed Functions) The function $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ is said to be *closed* if it is lower semi-continuous everywhere, or if its epigraph is closed, or if its sublevel-sets are closed.   □

The next step is to take the lower semi-continuous hull of a function f, whose value at $x \in \mathbb{R}^n$ is $\liminf_{y \to x} f(y)$. In view of the proof of Proposition 1.2.2, this operation amounts to closing epi f. When doing so, however, we may slide down to $-\infty$.

Definition 1.2.4 (Closure of a Function) The *closure* (or lower semi-continuous hull) of a function f is the function $\operatorname{cl} f \colon \mathbb{R}^n \to \mathbb{R} \cup \{\pm\infty\}$ defined by:

$$\operatorname{cl} f(x) := \liminf_{y \to x} f(y) \quad \text{for all } x \in \mathbb{R}^n \,, \tag{1.2.4}$$

or equivalently by

$$\operatorname{epi}(\operatorname{cl} f) := \operatorname{cl}(\operatorname{epi} f) \,. \tag{1.2.5}$$
□

An l.s.c. hull may be fairly complicated to compute, though; furthermore, the gap between $f(y)$ and $\operatorname{cl} f(y)$ may be impossible to control when y varies in a neighborhood of a given point x. Now convexity enters into play and makes things substantially easier, without additional assumption on f in the above definition:

– First of all, a convex function is minorized by an affine function (Proposition 1.2.1); closing it cannot introduce the value $-\infty$.
– Second, the issue reduces to the one-dimensional setting, thanks to the following *radial construction* of $\operatorname{cl} f$.

Proposition 1.2.5 *Let $f \in \operatorname{Conv} \mathbb{R}^n$ and $x' \in \operatorname{ri} \operatorname{dom} f$. There holds (in $\mathbb{R} \cup \{+\infty\}$)*

$$\operatorname{cl} f(x) = \lim_{t \downarrow 0} f(x + t(x' - x)) \quad \text{for all } x \in \mathbb{R}^n \,. \tag{1.2.6}$$

Proof. Since $x_t := x + t(x' - x) \to x$ when $t \downarrow 0$, we certainly have

$$(\operatorname{cl} f)(x) \leqslant \liminf_{t \downarrow 0} f(x + t(x' - x)) \,.$$

We will prove the converse inequality by showing that

$$\limsup_{t \downarrow 0} f(x + t(x' - x)) \leqslant r \quad \text{for all } r \geqslant (\operatorname{cl} f)(x)$$

(non-existence of such an r means that $\operatorname{cl} f(x) = +\infty$, the proof is finished).

Thus let $(x, r) \in \operatorname{epi}(\operatorname{cl} f) = \operatorname{cl}(\operatorname{epi} f)$. Pick $r' > f(x')$, hence $(x', r') \in \operatorname{ri} \operatorname{epi} f$ (Proposition 1.1.9). Applying Lemma A.2.1.6 to the convex set epi f, we see that

$$t(x', r') + (1 - t)(x, r) \in \operatorname{ri} \operatorname{epi} f \subset \operatorname{epi} f \quad \text{for all } t \in \,]0, 1] \,.$$

This just means

$$f(x + t(x' - x)) \leqslant tr' + (1 - t)r \quad \text{for all } t \in \,]0, 1]$$

and our required inequality follows by letting $t \downarrow 0$.   □

Another way of expressing the same thing is that, to compute $\operatorname{cl} f$ at some point x, it suffices to consider the restriction of f to a half-line, say $x + \mathbb{R}^+ d$, meeting ri dom f; here, d stands for $x' - x$. The resulting one-dimensional function $\varphi(t) := \operatorname{cl} f(x + td)$ becomes "continuous" from the right at $t = 0$, in the sense that $\varphi(0) = \lim_{t \downarrow 0} \varphi(t)$ – an equality in $\mathbb{R} \cup \{+\infty\}$.

Some simple but important properties come in conjunction with the results of the previous chapters:

Proposition 1.2.6 *For $f \in \operatorname{Conv} \mathbb{R}^n$, there holds*

$$\operatorname{cl} f \in \operatorname{Conv} \mathbb{R}^n \; ; \tag{1.2.7}$$

$$\operatorname{cl} f \text{ and } f \text{ coincide on the relative interior of } \operatorname{dom} f. \tag{1.2.8}$$

Proof. We already know from Proposition A.1.2.6 that epi cl $f = $ cl epi f is a convex set; also cl $f \leqslant f \not\equiv +\infty$; finally, Proposition 1.2.1 guarantees in the relation of definition (1.2.4) that cl $f(x) > -\infty$ for all x: (1.2.7) does hold.

On the other hand, suppose $x \in \operatorname{ri} \operatorname{dom} f$. Then the one-dimensional function $\varphi(t) = f(x + td)$ is continuous at $t = 0$ (Theorem 0.6.2); it follows from Proposition 1.2.5 that cl f coincides with f on ri dom f; besides, cl $f(x)$ is obviously equal to $f(x) = +\infty$ for all $x \notin \operatorname{cl} \operatorname{dom} f$. Altogether, (1.2.8) is true. $\square$

In particular, a finite-valued convex function (dom $f = \mathbb{R}^n$) is lower semi-continuous; actually, Theorem 3.1.2 below will confirm that it is more than that: it is continuous, and even locally Lipschitzian.

Due to their importance, closed convex functions deserve a special notation:

Notation 1.2.7 (The Set $\overline{\operatorname{Conv}} \mathbb{R}^n$) The set of closed convex functions on $\mathbb{R}^n$ is denoted by $\overline{\operatorname{Conv}} \mathbb{R}^n$. $\square$

(c) Outer Construction of Closed Convex Functions The property proved in Proposition 1.2.5 corresponds to a *direct* (or inner) construction of cl f from (1.2.4). Equivalently, cl f can be constructed as the largest l.s.c. (convex) function minorizing f. Correspondingly, the closed (convex) set epi cl f can also be described *externally*, as an intersection of closed (convex) sets. In view of §A.4.2(b), these closed convex sets can be restricted to be closed half-spaces: convexity provides another simplification of the closure operation. Besides, in view of Proposition 1.2.1, these half-spaces can be assumed non-vertical. This is the content of the next result.

Proposition 1.2.8 *The closure of $f \in \operatorname{Conv} \mathbb{R}^n$ is the supremum of all affine functions minorizing f:*

$$\operatorname{cl} f(x) = \sup_{(s,b) \in \mathbb{R}^n \times \mathbb{R}} \left\{ \langle s, x \rangle - b \, : \, \langle s, y \rangle - b \leqslant f(y) \text{ for all } y \in \mathbb{R}^n \right\}. \tag{1.2.9}$$

Proof. A closed half-space containing epi f is characterized by a nonzero vector $(s, \alpha) \in \mathbb{R}^n \times \mathbb{R}$ and a real number b such that

$$\langle s, x \rangle + \alpha r \leqslant b \quad \text{for all } (x, r) \in \operatorname{epi} f \tag{1.2.10}$$

(we equip the graph-space $\mathbb{R}^n \times \mathbb{R}$ with the scalar product of a product-space). Let us denote by $\Sigma \subset \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}$ the index-set of such triples $\sigma = (s, \alpha, b)$, with corresponding half-space

$$H_\sigma^- := \left\{ (x, r) \, : \, \langle s, x \rangle + \alpha r \leqslant b \right\}. \tag{1.2.11}$$

In other words, epi $(\mathrm{cl}\, f) = \mathrm{cl}\,(\mathrm{epi}\, f) = \cap_{\sigma \in \Sigma} H_\sigma^-$.

Because of the particular nature of an epigraph, (1.2.10) implies $\alpha \leqslant 0$ (let $r \to +\infty$) and, by positive homogeneity, the values $\alpha = 0$ and $\alpha = -1$ suffice: Σ can be partitioned in

$$\Sigma_1 := \{(s, -1, b) \,:\, (1.2.10) \text{ holds with } \alpha = -1\}$$

and

$$\Sigma_0 := \{(s, 0, b) \,:\, (1.2.10) \text{ holds with } \alpha = 0\}\,.$$

Indeed, Σ_1 corresponds to affine functions minorizing f (Proposition 1.2.1 tells us that $\Sigma_1 \neq \emptyset$) and Σ_0 to closed half-spaces of $\mathbb{R}^n$ containing dom f (note that $\Sigma_0 = \emptyset$ if dom $f = \mathbb{R}^n$).

We have to prove that, even when $\Sigma_0 \neq \emptyset$, intersecting the half-spaces H_σ^- over Σ or over Σ_1 produces the same set, namely $\mathrm{cl}\,\mathrm{epi}\, f$. For this we take arbitrary $\sigma_0 = (s_0, 0, b_0) \in \Sigma_0$ and $\sigma_1 = (s_1, -1, b_1) \in \Sigma_1$, we set

$$\sigma(t) := (s_1 + ts_0, -1, b_1 + tb_0) \in \Sigma_1 \quad \text{for all } t \geqslant 0\,,$$

and we prove (see Fig. 1.2.1)

$$H_{\sigma_0}^- \cap H_{\sigma_1}^- = \cap_{t \geqslant 0} H_{\sigma(t)}^- =: H^-\,.$$
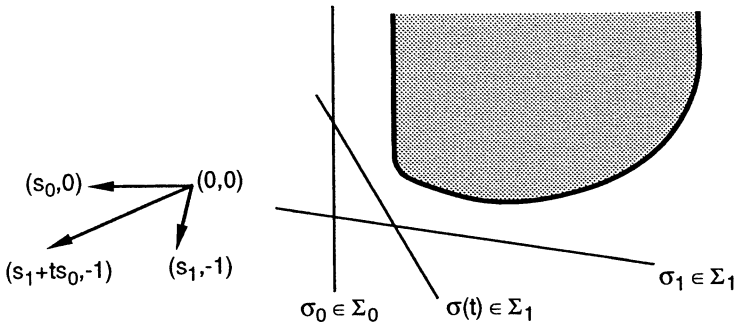
Fig. 1.2.1. Closing a convex epigraph

It results directly from the definition (1.2.11) that an (x, r) in $H_{\sigma_0}^- \cap H_{\sigma_1}^-$ satisfies

$$\langle s_1 + ts_0, x \rangle - (b_1 + tb_0) \leqslant r \quad \text{for all } t \geqslant 0\,, \tag{1.2.12}$$

i.e. $(x, r) \in H^-$. Conversely, take $(x, r) \in H^-$. Set $t = 0$ in (1.2.12) to see that $(x, r) \in H_{\sigma_1}^-$. Also, divide by $t > 0$ and let $t \to +\infty$ to see that $(x, r) \in H_{\sigma_0}^-$. The proof is complete. $\qquad\qquad\square$

1.3 First Examples

(a) Indicator and Support Functions Given a nonempty subset $S \subset \mathbb{R}^n$, the function $\mathrm{i}_S : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ defined by

$$\mathrm{i}_S(x) := \begin{cases} 0 & \text{if } x \in S, \\ +\infty & \text{if not} \end{cases}$$

is called the *indicator* function of S. We mention here that other notations commonly encountered in the literature are δ_S, ψ_S, or even χ_S. Clearly enough, i_S is [closed and] convex if and only if S is [closed and] convex. Indeed, $\mathrm{epi}\,\mathrm{i}_S = S \times \mathbb{R}^+$ by definition.

More generally, if $f \in \mathrm{Conv}\,\mathbb{R}^n$ and if C is a nonempty convex set, the function

$$\varphi(x) := \begin{cases} f(x) & \text{if } x \in C, \\ +\infty & \text{if not} \end{cases}$$

is again convex under one condition: that $\mathrm{dom}\,f$ and C have a nonempty intersection (otherwise φ would be identically $+\infty$). Furthermore, φ is closed when so are f and C. Observe in passing that $\varphi = f + \mathrm{i}_C$.

Attached to a nonempty subset S, another function of interest is the *support* function of S, already encountered in Remark A.4.1.2:

$$\sigma_S(x) := \sup \left\{ \langle s, x \rangle \ : \ s \in S \right\}.$$

It turns out to be closed and convex; this is already suggested by Proposition 1.2.8 and will be confirmed below in §2.1(b). Actually, the importance of this function will motivate an extensive development in Chap. C. Here, we just observe that, for $\alpha > 0$,

$$\sup_{s \in S} \langle s, \alpha x \rangle = \alpha \sup_{s \in S} \langle s, x \rangle,$$

hence $\sigma_S(\alpha x) = \alpha \sigma_S(x)$: the epigraph of a support function is not only closed and convex, but it is a cone in $\mathbb{R}^n \times \mathbb{R}$. Its domain is also a convex cone in $\mathbb{R}^n$:

$$\mathrm{dom}\,\sigma_S = \{a \in \mathbb{R}^n \ : \ \exists r \text{ such that } \langle s, a \rangle \leqslant r \text{ for all } s \in S \}.$$

(b) Piecewise Affine and Polyhedral Functions Let $(s_1, b_1), \ldots, (s_m, b_m)$ be m elements given in $\mathbb{R}^n \times \mathbb{R}$ and consider the function

$$\mathbb{R}^n \ni x \mapsto \check{f}(x) := \max \left\{ \langle s_j, x \rangle - b_j \ : \ j = 1, \ldots, m \right\}. \tag{1.3.1}$$

Such a function is suggestively called *piecewise affine*: $\mathbb{R}^n$ is divided into (at most m) regions in which $\check{f}$ is affine: the j_0^{th} region, possibly empty, is the closed convex polyhedron

$$\{x \in \mathbb{R}^n \ : \ \langle s_{j_0}, x \rangle - b_{j_0} \geqslant \langle s_j, x \rangle - b_j \text{ for } j = 1, \ldots, m \}.$$

This terminology is slightly ambiguous, though: a function whose graph is made up of pieces of affine hyperplanes need not be convex, while (1.3.1) can be seen to produce convex functions only (just as with a support function, convexity and closedness of $\check{f}$ will be confirmed below). It can even be seen that epi $\check{f}$ is a closed convex polyhedron; but again, (1.3.1) cannot describe all polyhedral epigraphs.

A *polyhedral* function will be a function whose epigraph is a closed convex polyhedron. Its most general form is given by Definition A.4.2.5:

$$\operatorname{epi} f = \{(x,r) \in \mathbb{R}^n \times \mathbb{R} \; : \; \langle s_j, x \rangle + \alpha_j r \leqslant b_j \text{ for } j \in J\}\,,$$

where J is a finite set, the $(s, \alpha, b)_j$ being given in $\mathbb{R}^n \times \mathbb{R} \times \mathbb{R}$, $(s_j, \alpha_j) \neq 0$ (and $\mathbb{R}^n \times \mathbb{R}$ is equipped with the scalar product of a product-space). For this set to be an epigraph, each α_j must be nonpositive and, if $\alpha_j < 0$, we may assume without loss of generality $\alpha_j = -1$. Furthermore, we may denote by $\{1, \ldots, m\}$ the subset of J such that $\alpha_j = -1$, and by $\{m+1, \ldots, m+p\}$ the rest (where $\alpha_j = 0$). With these notations, we see that $f(x)$ is given by (1.3.1) whenever x satisfies the set of *constraints*

$$\langle s_j, x \rangle \leqslant b_j \quad \text{for } j = m+1, \ldots, m+p\,;$$

otherwise, $f(x) = +\infty$. Of course, these constraints (usually termed linear, but affine is more correct) define a closed convex polyhedron.

In a word, a polyhedral function is a function which is piecewise affine on its domain, the latter being a closed convex polyhedron. Said otherwise, it is a closed convex function of the form $\check{f} + i_P$, where $\check{f}$ is piecewise affine and i_P is the indicator of a closed convex polyhedron.

(c) Norms and Distances It is a direct consequence of the axioms that a norm is a convex function, finite on the whole space (use Definition 1.1.1). More generally, let C be a nonempty convex set in $\mathbb{R}^n$ and, with an arbitrary norm $\|\cdot\|$, define the *distance function*

$$d_C(x) := \inf\{\|y - x\| \; : \; y \in C\}\,.$$

To establish its convexity, Definition 1.1.1 is again convenient. Take two sequences (y_k) and (y_k') such that, for $k \to +\infty$, $\|y_k - x\|$ and $\|y_k' - x'\|$ tend to $d_C(x)$ and $d_C(x')$ respectively. Then form the sequence $z_k := \alpha y_k + (1 - \alpha)y_k' \in C$ with $\alpha \in \,]0, 1[$; pass to the limit for $k \to +\infty$ in

$$d_C(\alpha x + (1-\alpha)x') \leqslant \|z_k - \alpha x - (1-\alpha)x'\| \leqslant \alpha\|y_k - x\| + (1-\alpha)\|y_k' - x'\|\,.$$

Here again $\operatorname{dom} d_C = \mathbb{R}^n$; the (lower semi-)continuity of d_C follows.

Clearly enough, $d_C = d_{\operatorname{cl} C}$ so, with the help of Proposition A.2.1.8, we see that C, $\operatorname{cl} C$ and $\operatorname{ri} C$ have the same distance function (associated with the same norm $\|\cdot\|$). In particular, d_C is 0 on the whole of $\operatorname{cl} C$; the following variant is slightly more accurate, in that it distinguishes between $\operatorname{int} C$ and $\operatorname{bd} C$:

$$D_C(x) := \begin{cases} d_C(x) & \text{if } x \in C^c\,, \\ -d_{C^c}(x) & \text{if } x \in C\,, \end{cases}$$

where C^c is the complement of C in $\mathbb{R}^n$. Assuming that C and C^c are both nonempty, it is not particularly difficult to prove that D_C is convex, finite everywhere, and that

$$\begin{aligned}
\operatorname{int} C &= \{x \in \mathbb{R}^n \ : \ \mathrm{D}_C(x) < 0\}\,, \\
\operatorname{bd} C &= \{x \in \mathbb{R}^n \ : \ \mathrm{D}_C(x) = 0\}\,, \\
(\operatorname{cl} C)^c &= \{x \in \mathbb{R}^n \ : \ \mathrm{D}_C(x) > 0\}\,.
\end{aligned}$$

(d) Quadratic Forms Let $A : \mathbb{R}^n \to \mathbb{R}^n$ be a symmetric linear operator. Then the quadratic form

$$f(x) := \tfrac{1}{2}\langle Ax, x\rangle$$

is a convex function – with $\operatorname{dom} f = \mathbb{R}^n$ – if and only if A is positive semi-definite, i.e. its eigenvalues are all nonnegative. Call $\lambda_1 \geqslant \cdots \geqslant \lambda_n \geqslant 0$ these eigenvalues; it is well-known that a basis can be formed with the corresponding eigenvectors, and that as a result,

$$\lambda_n \|x\|^2 \leqslant \langle Ax, x\rangle \leqslant \lambda_1 \|x\|^2 \quad \text{for all } x \in \mathbb{R}^n\,.$$

From the first inequality, direct but somewhat tedious calculations yield, with the notation of (1.1.2):

$$f(\alpha x + (1-\alpha)x') \leqslant \alpha f(x) + (1-\alpha)f(x') - \tfrac{1}{2}\lambda_n \alpha(1-\alpha)\|x - x'\|^2\,.$$

Thus, if A is positive definite, f is strongly convex with modulus $\lambda_n > 0$ (while f is not even strictly convex when A is degenerate). A straightforward proof comes also from a general characterization of differentiable strongly convex functions, to be seen below in Theorem 4.1.4 or 4.3.1.

For $r \geqslant 0$, the sublevel-sets of f:

$$\mathrm{S}_r(f) := \left\{x \in \mathbb{R}^n \ : \ \tfrac{1}{2}\langle Ax, x\rangle \leqslant r\right\}$$

are concentric elliptic sets: $\mathrm{S}_{\kappa r}(f) = \sqrt{\kappa}\mathrm{S}_r(f)$. Their common "shape" is given by the eigenvalues of A. They may be degenerate, in that they contain the subspace $\operatorname{Ker} A$ (one should rather speak of elliptic cylinders if $\operatorname{Ker} A \neq \{0\}$). However, $\mathrm{S}_r(f)$ is a neighborhood of the origin for $r > 0$: we have $\mathrm{S}_r(f) \supset B(0, \varepsilon)$ whenever $\lambda_1 \varepsilon^2 \leqslant 2r$.

(e) Sum of Largest Eigenvalues of a Matrix Instead of our working space $\mathbb{R}^n$, consider the vector space $\mathrm{S}_n(\mathbb{R})$ of symmetric $n \times n$ matrices. Denote the eigenvalues of $A \in \mathrm{S}_n(\mathbb{R})$ by $\lambda_1(A) \geqslant \cdots \geqslant \lambda_n(A)$, and consider the sum f_m of the m largest such eigenvalues ($m \leqslant n$ given):

$$\mathrm{S}_n(\mathbb{R}) \ni A \mapsto f_m(A) := \sum_{j=1}^{m} \lambda_j(A)\,.$$

This is a function of A, finite everywhere. Equip $\mathrm{S}_n(\mathbb{R})$ with the standard dot-product of $\mathbb{R}^{n \times n}$:

$$\langle\!\langle A, B \rangle\!\rangle := \operatorname{tr} AB = \sum_{i,j=1}^{n} A_{ij} B_{ij} \, .$$

The function f_m turns out to have the following representation:

$$f_m(A) = \sup \left\{ \langle\!\langle QQ^\top, A \rangle\!\rangle \, : \, Q \in \Omega \right\},$$

where $\Omega := \{Q : Q^\top Q = \mathrm{I}_m\}$ is the set of matrices made up of m orthonormal n-columns. Indeed, Ω is compact and the above supremum is attained at Q formed with the (normalized) eigenvectors associated with $\lambda_1, \ldots, \lambda_m$. Keeping Proposition 1.2.8 in mind, this explains that f_m is convex, as being a supremum of linear functions on $\mathrm{S}_n(\mathbb{R})$.

Naturally, $f_1(A)$ is the largest eigenvalue of A, while $f_n(A)$ is the trace of A, a linear function of A. It follows by taking differences that $f_n - f_m$ (for example the smallest eigenvalue $\lambda_n = f_n - f_{n-1}$) is a concave function on $\mathrm{S}_n(\mathbb{R})$.

(f) Volume of Ellipsoids Still in the space of symmetric matrices $\mathrm{S}_n(\mathbb{R})$, define the function

$$A \mapsto f(A) := \begin{cases} \log\left(\det A^{-1}\right) & \text{if } A \text{ is positive definite,} \\ +\infty & \text{if not.} \end{cases}$$

It will be seen in §3.1 that the concave finite-valued function $\lambda_n(\cdot)$ is continuous. The domain of f, which is the set of $A \in \mathrm{S}_n(\mathbb{R})$ such that $\lambda_n(A) > 0$, is therefore open, and even an open convex cone. It turns out that f is convex. To see it, start from the inequality

$$\det\left[\alpha A + (1 - \alpha)A'\right] \geqslant (\det A)^\alpha (\det A')^{1-\alpha} \, ,$$

valid for all symmetric positive definite matrices A and A' (and $\alpha \in \,]0, 1[\,$); take the inverse of each side; remember that the inverse of the determinant is the determinant of the inverse; finally, pass to the logarithms.

Geometrically, consider again an elliptic set

$$E_A := \{x \in \mathbb{R}^n \, : \, x^\top A x \leqslant 1\}$$

where A is a symmetric positive definite matrix. Up to a positive multiplicative constant (which is the volume of the unit ball E_{I_n}), the volume of E_A is precisely $\sqrt{\det A^{-1}}$.

Because $\operatorname{dom} f$ is open, $\operatorname{ri} \operatorname{dom} f = \operatorname{int} \operatorname{dom} f = \operatorname{dom} f$, which establishes the lower semi-continuity of f on its domain. Furthermore, suppose $A_k \to A$ with A not positive definite; by continuity of the concave function $\lambda_n(\cdot)$, A is positive semi-definite and the smallest eigenvalue of A_k tends to 0: $f(A_k) \to +\infty$. The function f is closed.

(g) Epigraphical Hull and Lower-Bound Function of a Convex Set Given a nonempty convex set $C \subset \mathbb{R}^n \times \mathbb{R}$, an interesting question is: when is C the epigraph of some function $f \in \operatorname{Conv} \mathbb{R}^n$? Let us forget for the moment the convexity

issue, which is not really relevant. First, the condition $f(x) > -\infty$ for all x means that C contains no vertical downward half-line:

$$\{r \in \mathbb{R} \ : \ (x, r) \in C\} \text{ is minorized for all } x \in \mathbb{R}^n \ . \tag{1.3.2}$$

A second condition is also obvious: C must be unbounded from above, more precisely

$$(x, r) \in C \quad \Longrightarrow \quad (x, r') \in C \text{ for all } r' > r \ . \tag{1.3.3}$$

The story does not end here, though: C must have a "closed bottom", i.e.

$$[(x, r') \in C \text{ and } r' \downarrow r] \quad \Longrightarrow \quad (x, r) \in C \ . \tag{1.3.4}$$

This time, we are done: a nonempty set C satisfying (1.3.2) – (1.3.4) is indeed an epigraph (of a convex function if C is convex). Alternatively, if C satisfying (1.3.2), (1.3.3) has its bottom open, i.e.

$$(x, r) \in C \quad \Longrightarrow \quad (x, r - \varepsilon) \in C \quad \text{for some } \varepsilon = \varepsilon(x, r) > 0 \ ,$$

then C is a strict epigraph. To cut a long story short: a [strict] epigraph is a union of closed [open] upward half-lines – knowing that we always rule out the value $-\infty$.

The next interesting point is to make an epigraph with a given set: the *epigraphical hull* of $C \subset \mathbb{R}^n \times \mathbb{R}$ is the smallest epigraph containing C. Its construction involves only rather trivial operations in the ordered set $\mathbb{R}$:

  (i) force (1.3.3) by stuffing in everything above C: for each $(x, r) \in C$, add to C all (x, r') with $r' > r$;
 (ii) force (1.3.4) by closing the bottom of C: put (x, r) in C whenever $(x, r') \in C$ with $r' \to r$.

These operations (i), (ii) amount to constructing a function:

$$x \mapsto \ell_C(x) := \inf \{r \in \mathbb{R} \ : \ (x, r) \in C\} \ , \tag{1.3.5}$$

the *lower-bound function* of C; clearly enough, epi ℓ_C is the epigraphical hull of C. We have that $\ell_C(x) > -\infty$ for all x if (and only if) C satisfies (1.3.2).

The construction of an epigraphical hull is illustrated on Fig. 1.3.1, in which the point A and the curve Γ are not in C; nevertheless, there holds (epi$_s$ is the strict epigraph)

$$\text{epi}_s \ \ell_C \subset C + \{0\} \times \mathbb{R}^+ \subset \text{epi} \ \ell_C \subset \text{cl} \, (C + \{0\} \times \mathbb{R}^+) \ . \tag{1.3.6}$$

Theorem 1.3.1 *Let C be a nonempty subset of $\mathbb{R}^n \times \mathbb{R}$ satisfying* (1.3.2), *and let its lower-bound function ℓ_C be defined by* (1.3.5).

  (i) *If C is convex, then $\ell_C \in \text{Conv} \, \mathbb{R}^n$;*
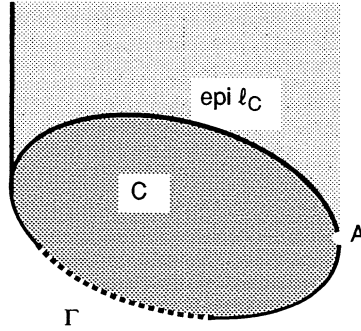 (ii) *If C is closed convex, then $\ell_C \in \overline{\text{Conv}} \, \mathbb{R}^n$.*

Fig. 1.3.1. The lower-bound function

Proof. We use the analytical definition (1.1.1). Take arbitrary $\varepsilon > 0$, $\alpha \in \,]0,1[$ and $(x_i, r_i) \in C$ such that $r_i \leqslant \ell_C(x_i) + \varepsilon$ for $i = 1, 2$.

When C is convex, $(\alpha x_1 + (1 - \alpha)x_2, \alpha r_1 + (1 - \alpha)r_2) \in C$, hence

$$\ell_C(\alpha x_1 + (1 - \alpha)x_2) \leqslant \alpha r_1 + (1 - \alpha)r_2 \leqslant \alpha \ell_C(x_1) + (1 - \alpha)\ell_C(x_2) + \varepsilon.$$

The convexity of ℓ_C follows, since $\varepsilon > 0$ was arbitrary; (i) is proved.

Now take a sequence $(x_k, \rho_k)_k \subset$ epi ℓ_C converging to (x, ρ); we have to prove $\ell_C(x) \leqslant \rho$ (cf. Proposition 1.2.2). By definition of $\ell_C(x_k)$, we can select, for each positive integer k, a real number r_k such that $(x_k, r_k) \in C$ and

$$\ell_C(x_k) \leqslant r_k \leqslant \ell_C(x_k) + \tfrac{1}{k} \leqslant \rho_k + \tfrac{1}{k}. \tag{1.3.7}$$

We deduce first that (r_k) is bounded from above. Also, when ℓ_C is convex, Proposition 1.2.1 implies the existence of an affine function minorizing ℓ_C: (r_k) is bounded from below.

Extracting a subsequence if necessary, we may assume $r_k \to r$. When C is closed, $(x, r) \in C$, hence $\ell_C(x) \leqslant r$; but pass to the limit in (1.3.7) to see that $r \leqslant \rho$; the proof is complete. $\qquad\square$

2 Functional Operations Preserving Convexity

It is natural to build up new convex functions from simpler ones, via operations preserving convexity, or even yielding it. This approach goes together with that of §A.1.2: convex epigraphs can be made up from simpler epigraphs. Here again, proving convexity of the new function will rely either on the analytical definition or on the geometric one, whichever is simpler.

2.1 Operations Preserving Closedness

(a) Positive Combinations of Functions

Proposition 2.1.1 *Let $f_1, \ldots, f_m$ be in $\mathrm{Conv}\,\mathbb{R}^n$ [resp. in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$], let $t_1, \ldots, t_m$ be positive numbers, and assume that there is a point where all the f_j's are finite. Then the function $f := \sum_{j=1}^m t_j f_j$ is in $\mathrm{Conv}\,\mathbb{R}^n$ [resp. in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$].*

Proof. The convexity of f is readily proved from the relation of definition (1.1.1). As for its closedness, start from

$$\liminf_{y \to x} t_j f_j(y) = t_j \liminf_{y \to x} f_j(y) \geq t_j f_j(x)$$

(valid for $t_j > 0$ and f_j closed); then note that the lim inf of a sum is not smaller than the sum of lim inf's.   □

As an example, let $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$ and $C \subset \mathbb{R}^n$ be closed convex, with dom $f \cap C \neq \emptyset$. Then the function $f + i_C$ of Example 1.3(a) is in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$.

(b) Supremum of Convex Functions

Proposition 2.1.2 *Let $\{f_j\}_{j \in J}$ be an arbitrary family of convex [resp. closed convex] functions. If there exists x_0 such that $\sup_J f_j(x_0) < +\infty$, then their pointwise supremum $f := \sup\{f_j : j \in J\}$ is in $\mathrm{Conv}\,\mathbb{R}^n$ [resp. in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$].*

Proof. The key property is that a supremum of functions corresponds to an intersection of epigraphs: epi $f = \cap_{j \in J}$ epi f_j, which conserves convexity and closedness. The only needed restriction is nonemptiness of this intersection.   □

In a way, this result was already announced by Proposition 1.2.8. It has also been used again and again in the examples of §1.3.

Example 2.1.3 (Conjugate Function) Let $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ be a function not identically $+\infty$, minorized by an affine function (i.e., for some $(s_0, b) \in \mathbb{R}^n \times \mathbb{R}$, $f \geq \langle s_0, \cdot \rangle - b$ on $\mathbb{R}^n$). Then the function f^* defined by

$$\mathbb{R}^n \ni s \mapsto f^*(s) := \sup\{\langle s, x \rangle - f(x) : x \in \mathrm{dom}\, f\}$$

is called the *conjugate function* of f, to be studied thoroughly in Chap. E. Observe that $f^*(s_0) \leq b$ and $f^*(s) > -\infty$ for all s because dom $f \neq \emptyset$. Thus, $f^* \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$; this is true without any further assumption on f, in particular its convexity or closedness are totally irrelevant here.   □

(c) Pre-Composition with an Affine Mapping

Proposition 2.1.4 *Let $f \in \mathrm{Conv}\,\mathbb{R}^n$ [resp. $\overline{\mathrm{Conv}}\,\mathbb{R}^n$] and let A be an affine mapping from $\mathbb{R}^m$ to $\mathbb{R}^n$ such that $\mathrm{Im}\, A \cap \mathrm{dom}\, f \neq \emptyset$. Then the function*

$$f \circ A : \quad \mathbb{R}^m \ni x \mapsto (f \circ A)(x) = f(A(x))$$

is in $\mathrm{Conv}\,\mathbb{R}^m$ [resp. $\overline{\mathrm{Conv}}\,\mathbb{R}^m$].

Proof. Clearly $(f \circ A)(x) > -\infty$ for all x; besides, there exists by assumption $y = A(x) \in \mathbb{R}^n$ such that $f(y) < +\infty$. To check convexity, it suffices to plug the relation

$$A(\alpha x + (1 - \alpha)x') = \alpha A(x) + (1 - \alpha)A(x')$$

into the analytical definition (1.1.1) of convexity. As for closedness, it comes readily from the continuity of A when f is itself closed.   □

Example 2.1.5 With f (closed) convex on $\mathbb{R}^n$, take $x_0 \in \operatorname{dom} f$, $d \in \mathbb{R}^n$ and define

$$A : \mathbb{R} \ni t \mapsto A(t) := x_0 + td;$$

this A is affine, its linear part is $t \mapsto A_0 t := td$. The resulting $f \circ A$ appears as (a parametrization of) the restriction of f along the line $x_0 + \mathbb{R}d$, which meets $\operatorname{dom} f$ (at x_0).

This operation is often used in applications. Even from a theoretical point of view, the one-dimensional traces of f are important, in that f itself inherits many of their properties; Proposition 1.2.5 gives an instance of this phenomenon.   □

Remark 2.1.6 With relation to this operation on $f \in \operatorname{Conv} \mathbb{R}^n$ [resp. $\overline{\operatorname{Conv}} \, \mathbb{R}^n$], call V the subspace parallel to aff $\operatorname{dom} f$. Then, fix $x_0 \in \operatorname{dom} f$ and define the convex function $f_0 \in \operatorname{Conv} V$ [resp. $\overline{\operatorname{Conv}} \, V$] by $f_0(y) := f(x_0 + y)$ for all $y \in V$.

This new function is obtained from f by a simple translation, *composed with a restriction* (from $\mathbb{R}^n$ to V). As a result, $\operatorname{dom} f_0$ is now full-dimensional (in V), the relative topology relevant for f_0 is the standard topology of V. This trick is often useful and explains why "flat" domains, instead of full-dimensional, create little difficulties.   □

(d) Post-Composition with an Increasing Convex Function

Proposition 2.1.7 *Let* $f \in \operatorname{Conv} \mathbb{R}^n$ *[resp.* $\overline{\operatorname{Conv}} \, \mathbb{R}^n$ *] and let* $g \in \operatorname{Conv} \mathbb{R}$ *[resp.* $\overline{\operatorname{Conv}} \, \mathbb{R}$*] be increasing. Assume that there is* $x_0 \in \mathbb{R}^n$ *such that* $f(x_0) \in \operatorname{dom} g$, *and set* $g(+\infty) := +\infty$. *Then the composite function* $g \circ f : x \mapsto g(f(x))$ *is in* $\operatorname{Conv} \mathbb{R}^n$ *[resp. in* $\overline{\operatorname{Conv}} \, \mathbb{R}^n$ *].*

Proof. It suffices to check the inequalities of definition: (1.1.1) for convexity, (1.2.3) for closedness.   □

The exponential $g(t) := \exp t$ is convex increasing, its domain is the whole line, so $\exp f(x)$ is a [closed] convex function of $x \in \mathbb{R}^n$ whenever f is [closed] convex. A function $f : \mathbb{R}^n \to \,]0, +\infty]$ is called *logarithmically convex* when $\log f \in \operatorname{Conv} \mathbb{R}^n$ (we set again $\log(+\infty) = +\infty$). Because $f = \exp \log f$, a logarithmically convex function is convex.

As another application, the square of an arbitrary nonnegative convex function (for example a norm) is convex: post-compose it by the function $g(t) = (\max \{0, t\})^2$.

2.2 Dilations and Perspectives of a Function

For a convex function f and $u > 0$, the function

$$f_u : \mathbb{R}^n \ni x \mapsto f_u(x) := uf(x/u)$$

is again convex. This comes from Propositions 2.1.1 and 2.1.4 but can also be seen geometrically: since $f_u(x)/u = f(x/u)$, the epigraphs and sublevel-sets are related by

$$\operatorname{epi} f_u = u \operatorname{epi} f, \quad \operatorname{epi}_s f_u = u \operatorname{epi}_s f, \quad S_r(f_u) = u \, S_{r/u}(f),$$

which express that f_u is a "dilated version" of f.

More interesting, however, is to study f_u as a function of *both* variables x and u, i.e. to consider the set of all dilations of f. We therefore define the *perspective* of f as the function from $\mathbb{R} \times \mathbb{R}^n$ to $\mathbb{R} \cup \{+\infty\}$ given by
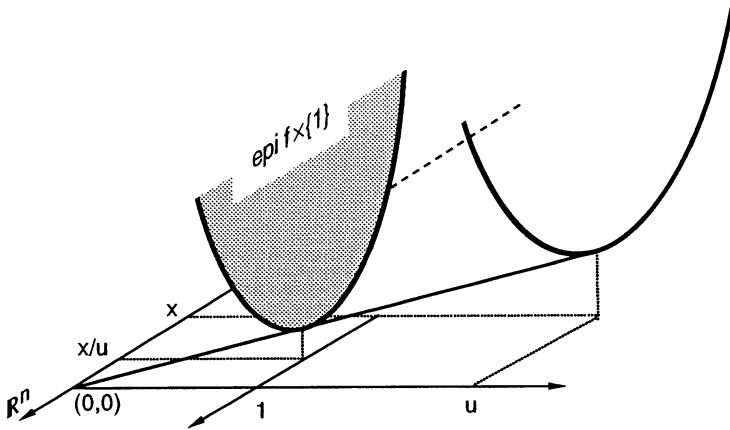
$$\tilde{f}(u, x) := \begin{cases} uf(x/u) & \text{if } u > 0, \\ +\infty & \text{if not}. \end{cases}$$

Proposition 2.2.1 *If $f \in \text{Conv } \mathbb{R}^n$, its perspective $\tilde{f}$ is in $\text{Conv } \mathbb{R}^{n+1}$.*

Proof. Here also, it is better to look at $\tilde{f}$ with "geometric glasses":

$$\begin{aligned} \text{epi } \tilde{f} &= \{(u, x, r) \in \mathbb{R}_*^+ \times \mathbb{R}^n \times \mathbb{R} : f(x/u) \leqslant r/u\} \\ &= \{u(1, x', r') : u > 0, (x', r') \in \text{epi } f\} \\ &= \cup_{u>0}\{u(\{1\} \times \text{epi } f)\} = \mathbb{R}_*^+ (\{1\} \times \text{epi } f) \end{aligned}$$

and epi $\tilde{f}$ is therefore a convex cone.   □

Fig. 2.2.1. The perspective of a convex function

Figure 2.2.1 illustrates the construction of epi $\tilde{f}$, as given in the above proof. Embed epi f into $\mathbb{R} \times \mathbb{R}^n \times \mathbb{R}$, where the first $\mathbb{R}$ represents the extra variable u; translate it horizontally by one unit; finally, take the positive multiples of the result. Observe that, following the same technique, we obtain

$$\text{dom } \tilde{f} = \mathbb{R}_*^+ (\{1\} \times \text{dom } f). \tag{2.2.1}$$

Another observation is that, by construction, epi $\tilde{f}$ [resp. dom $\tilde{f}$] does not contain the origin of $\mathbb{R} \times \mathbb{R}^n \times \mathbb{R}$ [resp. $\mathbb{R} \times \mathbb{R}^n$].

Convexity of a perspective-function is an important property, which we will use later in the following way. For fixed $x_0 \in \text{dom } f$, the function $d \mapsto f(x_0 + d) - f(x_0)$ is obviously convex, so its perspective

$$r(u, d) := u[f(x_0 + d/u) - f(x_0)] \quad \text{(for } u > 0) \tag{2.2.2}$$

is also convex with respect to the couple $(u, d) \in \mathbb{R}_*^+ \times \mathbb{R}^n$. Up to the simple change of variable $u \mapsto t = 1/u$, we recognize a difference quotient.

The next natural question is the closedness of a perspective-function: admitting that f itself is closed, troubles can still be expected at $u = 0$, where we have brutally set $\tilde{f}(0, \cdot) = +\infty$ (possibly not the best idea ...) A relatively simple calculation of cl $\tilde{f}$ is in fact given by Proposition 1.2.5:

Proposition 2.2.2 Let $f \in \overline{\mathrm{Conv}} \, \mathbb{R}^n$ and let $x' \in \mathrm{ri} \, \mathrm{dom} \, f$. Then the closure cl $\tilde{f}$ of its perspective is given as follows:

$$(\mathrm{cl} \, \tilde{f})(u, x) = \begin{cases} u f(x/u) & \text{if } u > 0, \\ \lim_{\alpha \downarrow 0} \alpha f(x' - x + x/\alpha) & \text{if } u = 0, \\ +\infty & \text{if } u < 0. \end{cases}$$

Proof. Suppose first $u < 0$. For any x, it is clear that (u, x) is outside cl dom $\tilde{f}$ and, in view of (1.2.8), cl $\tilde{f}(u, x) = +\infty$.

Now let $u \geqslant 0$. Using (2.2.1), the assumption on x' and the results of §A.2.1, we see that $(1, x') \in \mathrm{ri} \, \mathrm{dom} \, \tilde{f}$, so Proposition 1.2.5 allows us to write

$$(\mathrm{cl} \, \tilde{f})(u, x) = \lim_{\alpha \downarrow 0} \tilde{f} \big((u, x) + \alpha[(1, x') - (u, x)] \big)$$

$$= \lim_{\alpha \downarrow 0} [u + \alpha(1 - u)] \, f\big(\tfrac{x + \alpha(x' - x)}{u + \alpha(1 - u)} \big).$$

If $u = 1$, this reads cl $\tilde{f}(1, x) = \mathrm{cl} \, f(x) = f(x)$ (because f is closed); if $u = 0$, we just obtain our claimed relation. □

Remark 2.2.3 Observe that the behaviour of $\tilde{f}(u, \cdot)$ for $u \downarrow 0$ just depends on the behaviour of f at infinity. If $x = 0$, we have

$$\mathrm{cl} \, \tilde{f}(0, 0) = \lim_{\alpha \downarrow 0} \alpha f(x') = 0 \quad [f(x') < +\infty!].$$

For $x \neq 0$, suppose for example that dom f is bounded; then $f(x' - x + x/\alpha) = +\infty$ for α small enough and cl $\tilde{f}(0, x) = +\infty$. On the other hand, when dom f is unbounded, cl $\tilde{f}(0, \cdot)$ may assume finite values if, at infinity, f does not increase too fast.

For another illustration, we apply here Proposition 2.2.2 to the perspective-function r of (2.2.2). Assuming $x_0 \in \mathrm{ri} \, \mathrm{dom} \, f$, we can take $d' = 0$ – which is in the relative interior of the function $d \mapsto f(x_0 + d) - f(x_0)$ – to obtain

$$(\mathrm{cl} \, r)(0, d) = \lim_{\tau \to +\infty} \frac{f(x_0 - d + \tau d) - f(x_0)}{\tau}.$$

Because $(\tau - 1)/\tau \to 1$ for $\tau \to +\infty$, the last limit can also be written (in $\mathbb{R} \cup \{+\infty\}$)

$$(\mathrm{cl} \, r)(0, d) = \lim_{t \to +\infty} \frac{f(x_0 + td) - f(x_0)}{t}.$$

We will return to all this in §3.2 below. □

2.3 Infimal Convolution

Starting from two functions f_1 and f_2, form the set epi $f_1 +$ epi $f_2 \subset \mathbb{R}^n \times \mathbb{R}$:

$$C := \{(x_1 + x_2, r_1 + r_2) \; : \; r_j \geqslant f_j(x_j) \text{ for } j = 1, 2\}.$$

Under a suitable minorization property, this C has a lower-bound function ℓ_C as in (1.3.5):

$$\ell_C(x) = \inf \{r_1 + r_2 \; : \; r_j \geqslant f_j(x_j) \text{ for } j = 1, 2, \; x_1 + x_2 = x\}.$$

In the above minimization problem, the variables are r_1, r_2, x_1, x_2, but the r_j's can be eliminated; in fact, ℓ_C can be defined as follows.

Definition 2.3.1 Let f_1 and f_2 be two functions from $\mathbb{R}^n$ to $\mathbb{R} \cup \{+\infty\}$. Their *infimal convolution* is the function from $\mathbb{R}^n$ to $\mathbb{R} \cup \{\pm\infty\}$ defined by

$$\begin{aligned}(f_1 \,\vartriangledown\, f_2)(x) &:= \inf \{f_1(x_1) + f_2(x_2) \; : \; x_1 + x_2 = x\} \\ &= \inf_{y \in \mathbb{R}^n}[f_1(y) + f_2(x - y)].\end{aligned} \tag{2.3.1}$$

We will also call "infimal convolution" the *operation* expressed by (2.3.1). It is called *exact* at $x = \bar{x}_1 + \bar{x}_2$ when the infimum is attained at $(\bar{x}_1, \bar{x}_2)$, not necessarily unique. □

This operation is admittedly complex but important and will be encountered on many occasions. Let us observe right here that it corresponds to the (admittedly simple) addition of epigraphs – barring some technicalities, see Remark 2.3.4 below. It is a good exercise to visualize the infimal convolution of an arbitrary convex f_1 and

– $f_2(x) = \mathrm{i}_{\{0\}}(x) + r$ (shift epi f_1 vertically by r);
– $f_2(x) = \mathrm{i}_{\{x_0\}}(x)$ (horizontal shift);
– $f_2(x) = \mathrm{i}_{B(0,r)}(x)$ (horizontal smear);
– in one dimension: $f_2(x) = sx - r$ (it is gr f_2 that wins);
– $f_2(x) = 1 - \sqrt{1 - \|x\|^2}$ for $x \in B(0, 1)$ (the "ball-pen function" of Fig. 2.3.1); translate the bottom of the ball-pen (the origin of $\mathbb{R}^n \times \mathbb{R}$) to each point in gr f_1;
– $f_2(x) = 1/2 \, \|x\|^2$ (similar operation, called the *Moreau-Yosida regularization* of f).
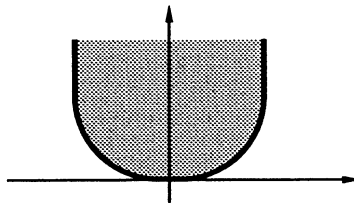
Fig. 2.3.1. The ball-pen function

Remark 2.3.2 The classical (integral) convolution between two functions F_1 and F_2 is

$$(F_1 * F_2)(x) := \int_{\mathbb{R}^n} F_1(y) F_2(x - y) dy \quad \text{for all } x \in \mathbb{R}^n \,.$$

For nonnegative functions, we can consider the "convolution of order p" ($p > 0$):

$$(F_1 *_p F_2)(x) := \left\{ \int [F_1(y) F_2(x - y)]^p dy \right\}^{1/p} \,.$$

When $p \to +\infty$, this integral converges to $\sup_y F_1(y) F_2(x - y)$ (a fact easy to accept). Now take $F_i := \exp(-f_i)$, $i = 1, 2$; we have

$$(F_1 *_\infty F_2)(x) = \sup_y e^{-f_1(y) - f_2(x-y)} = e^{-\inf_y [f_1(y) + f_2(x-y)]} \,.$$

Thus, the infimal convolution appears as a "convolution of infinite order", combined with an exponentiation.   □

To exclude the undesired value $-\infty$ from the range of an inf-convolution, an additional assumption is obviously needed: in one dimension, the infimal convolution of the functions x and $-x$ is identically $-\infty$. Our next result proposes a convenient such assumption.

Proposition 2.3.3 *Let the functions f_1 and f_2 be in $\operatorname{Conv} \mathbb{R}^n$. Suppose that they have a common affine minorant: for some $(s, b) \in \mathbb{R}^n \times \mathbb{R}$,*

$$f_j(x) \geqslant \langle s, x \rangle - b \quad \text{for } j = 1, 2 \text{ and all } x \in \mathbb{R}^n \,.$$

Then their infimal convolution is also in $\operatorname{Conv} \mathbb{R}^n$.

Proof. For arbitrary $x \in \mathbb{R}^n$ and x_1, x_2 such that $x_1 + x_2 = x$, we have by assumption

$$f_1(x_1) + f_2(x_2) \geqslant \langle s, x \rangle - 2b > -\infty \,,$$

and this inequality extends to the infimal value $(f_1 \,\dot{+}\, f_2)(x)$.

On the other hand, it suffices to choose particular values $x_j \in \operatorname{dom} f_j$, $j = 1, 2$, to obtain the point $x_1 + x_2 \in \operatorname{dom}(f_1 \,\dot{+}\, f_2)$. Finally, the convexity of $f_1 \,\dot{+}\, f_2$ results from the convexity of a lower-bound function, as seen in §1.3(g).   □

Remark 2.3.4 To prove that an inf-convolution of convex functions is convex, one can also show the following relation between strict epigraphs:

$$\operatorname{epi}_s (f_1 \,\dot{+}\, f_2) = \operatorname{epi}_s f_1 + \operatorname{epi}_s f_2 \,. \tag{2.3.2}$$

In fact, $(x, r) \in \operatorname{epi}_s (f_1 \,\dot{+}\, f_2)$ if and only if there is $\varepsilon > 0$ such that

$$f_1(x_1) + f_2(x_2) = r + \varepsilon \quad \text{for some } x_1 \text{ and } x_2 \text{ adding up to } x \,.$$

This is equivalent to $f_j(x_j) < r_j$ for some (x_1, r_1) and (x_2, r_2) adding up to (x, r) (set $r_j := f_j(x_j) + \varepsilon/2$ for $j = 1, 2$, to show the "$\Rightarrow$" direction). This last property holds if and only if $(x, r) \in \operatorname{epi}_s f_1 + \operatorname{epi}_s f_2$.

This proof explains why the infimal convolution is sometimes called the (strict) *epigraphic addition*.   □

Similarly to (2.3.2), we have by construction

$$\mathrm{dom}\,(f_1 \,\Diamond\, f_2) = \mathrm{dom}\,f_1 + \mathrm{dom}\,f_2 \,.$$

Let us mention some immediate properties of the infimal convolution:

$$f_1 \,\Diamond\, f_2 = f_2 \,\Diamond\, f_1 \quad \text{(commutativity)} \tag{2.3.3}$$

$$(f_1 \,\Diamond\, f_2) \,\Diamond\, f_3 = f_1 \,\Diamond\, (f_2 \,\Diamond\, f_3) \quad \text{(associativity)} \tag{2.3.4}$$

$$f \,\Diamond\, i_{\{0\}} = f \quad \text{(existence of a neutral element in } \overline{\mathrm{Conv}}\,\mathbb{R}^n) \tag{2.3.5}$$

$$f_1 \leqslant f_2 \quad \Longrightarrow \quad f_1 \,\Diamond\, g \leqslant f_2 \,\Diamond\, g \quad (\Diamond \text{ preserves the order}).$$

With relation to (2.3.3), (2.3.4), more than two functions can of course be inf-convolved:

$$(f_1 \,\Diamond\, \cdots \,\Diamond\, f_m)(x) = \inf \left\{ \sum_{j=1}^{m} f_j(x_j) \; : \; \sum_{j=1}^{m} x_j = x \right\}.$$

Remark 2.3.5 If C_1 and C_2 are nonempty convex sets in $\mathbb{R}^n$, then

$$i_{C_1} \,\Diamond\, i_{C_2} = i_{C_1+C_2} \,.$$

This is due to the additional nature of the inf-convolution, and can also be checked directly; but it leads us to an important observation: since the sum of two closed sets may not be closed, an infimal convolution *need not be closed*, even if it is constructed from two closed functions and if it is exact everywhere. □

Example 2.3.6 Let C be a nonempty convex subset of $\mathbb{R}^n$ and $\|\cdot\|$ an arbitrary norm. Then

$$i_C \,\Diamond\, \|\cdot\| = d_C \,,$$

which confirms the convexity of the distance function d_C. It also shows that inf-convolving two non-closed functions (C need not be closed) may result in a closed function. □

Example 2.3.7 Let f be an arbitrary convex function minorized by some affine function with slope s. Taking an affine function $g = \langle s, \cdot \rangle - b$, we obtain $f \,\Diamond\, g = g - c$, where c is a constant: $c = \sup_y [\langle s, y \rangle - f(y)]$ (already encountered in Example 2.1.3 $c = f^*(s)$, the value at s of the conjugate of f).

Assuming f bounded below, let in particular g be constant: $-g = \bar{f} := \inf_y f(y)$. Then $f \,\Diamond\, (-\bar{f}) = 0$. Do not believe, however, that the infimal convolution provides $\overline{\mathrm{Conv}}\,\mathbb{R}^n$ with the structure of a commutative group: in view of (2.3.5), the 0-function is not the neutral element! □

Example 2.3.8 We have seen (Proposition 1.2.1) that a convex function is indeed minorized by some affine function. The dilated versions $f_u = u f(\cdot/u)$ of a given convex function f are minorized by some affine function with a slope *independent* of $u > 0$, and can be inf-convolved by each other. We obtain $f_u \,\Diamond\, f_{u'} = f_{u+u'}$; the quickest way to prove this formula is probably to use (2.3.2), knowing that $\mathrm{epi}_s\, f_u = u\,\mathrm{epi}_s\, f$.

In particular, inf-convolving m times a function with itself gives a sort of mean-value formula:

$$\tfrac{1}{m}(f \,\Diamond\, \cdots \,\Diamond\, f)(x) = f\left(\tfrac{1}{m}x\right).$$

Observe how a perspective-function gives a meaning to a non-integer number of self-inf-convolutions. □

Example 2.3.9 Consider two quadratic forms

$$f_j(x) = \tfrac{1}{2}\langle A_j x, x \rangle \quad \text{for } j = 1, 2 \,,$$

with A_1 and A_2 symmetric positive definite. Expressing their infimal convolution as

$$\tfrac{1}{2} \inf_y \left[\langle A_1 y, y \rangle + \langle A_2 (x - y), x - y \rangle \right] \,,$$

the minimum can be explicitly worked out, to give $(f_1 \,\dot{\vee}\, f_2)(x) = 1/2\langle A_{12} x, x \rangle$, where $A_{12} := \left(A_1^{-1} + A_2^{-1} \right)^{-1}$.

This formula has an interesting physical interpretation: consider an electrical circuit made up of two generalized resistors A_1 and A_2 connected in parallel. A given current-vector $i \in \mathbb{R}^n$ is distributed among the two branches ($i = i_1 + i_2$), in such a way that the dissipated power $\langle A_1 i_1, i_1 \rangle + \langle A_2 i_2, i_2 \rangle$ is minimal (this is Maxwell's variational principle); see Fig. 2.3.2. In other words, if $i = \bar{\imath}_1 + \bar{\imath}_2$ is the real current distribution, we must have

$$\langle A_1 \bar{\imath}_1, \bar{\imath}_1 \rangle + \langle A_2 \bar{\imath}_2, \bar{\imath}_2 \rangle = \inf_{i_1 + i_2 = i} \left(\langle A_1 i_1, i_1 \rangle + \langle A_2 i_2, i_2 \rangle \right) \,.$$

The unique distribution $(\bar{\imath}_1, \bar{\imath}_2)$ is thus characterized by the formulae

$$A_1 \bar{\imath}_1 = A_2 \bar{\imath}_2 = A_{12} i \,, \tag{2.3.6}$$

from which it follows that $\langle A_1 \bar{\imath}_1, \bar{\imath}_1 \rangle + \langle A_2 \bar{\imath}_2, \bar{\imath}_2 \rangle = \langle A_{12} i, i \rangle$.
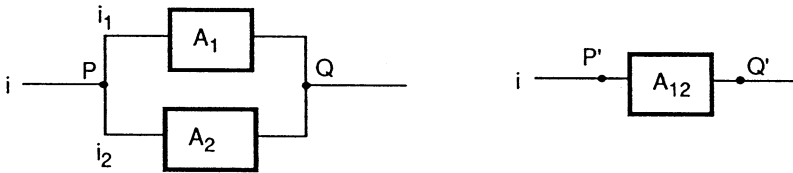
Fig. 2.3.2. Equivalent resistors

Thus, A_{12} plays the role of a generalized resistor equivalent to A_1 and A_2 connected in parallel; when $n = 1$, we get the more familiar relation $1/r = 1/r_1 + 1/r_2$ between ordinary resistances r_1 and r_2. Note an interpretation of the optimality (or equilibrium) condition (2.3.6). The voltage between P and Q [resp. P' and Q'] on Fig. 2.3.2, namely $A_1 \bar{\imath}_1 = A_2 \bar{\imath}_2$ [resp. $A_{12} i$], is independent of the path chosen: either through A_1, or through A_2, or by construction through A_{12}.

The above example of two convex quadratic functions can be extended to general functions, and it gives an economic interpretation of the infimal convolution: let $f_1(x)$ [resp. $f_2(x)$] be the cost of producing x by some production unit U_1 [resp. U_2]. If we want to distribute optimally the production of a given x between U_1 and U_2, we have to solve the minimization problem (2.3.1). $\qquad\square$

2.4 Image of a Function Under a Linear Mapping

Consider a constrained optimization problem, formally written as

$$\inf_{u \in U} \{ \varphi(u) \ : \ c(u) \leqslant x \}, \tag{2.4.1}$$

where the optimization variable is u, the righthand side x being considered as a parameter taken in some ordered set X. The optimal value in such a problem is then a function of x, characterized by the triple (U, φ, c), and taking its values in $\mathbb{R} \cup \pm\{\infty\}$. This is an important function, usually called the *value* function, or *marginal* function, or *perturbation* function, or *primal* function, etc.

Several variants of (2.4.1) are possible: we may encounter equality constraints, some constraints may be included in the objective via an indicator function, etc. A convenient unified formulation is the following:

Definition 2.4.1 (Image Function) Let $A : \mathbb{R}^m \to \mathbb{R}^n$ be a linear operator and let $g : \mathbb{R}^m \to \mathbb{R} \cup \{+\infty\}$. The *image* of g under A is the function $Ag : \mathbb{R}^n \to \mathbb{R} \cup \pm\infty$ defined by

$$(Ag)(x) := \inf \{ g(y) \ : \ Ay = x \} \tag{2.4.2}$$

(here as always, $\inf \emptyset = +\infty$).   □

The terminology comes from the case of an indicator function: when $g = i_C$, with C nonempty in $\mathbb{R}^m$, (2.4.2) writes

$$(Ag)(x) = \begin{cases} 0 & \text{if } x = Ay \text{ for some } y \in C, \\ +\infty & \text{otherwise}. \end{cases}$$

In other words, $Ag = i_{A(C)}$ is the indicator function of the image of C under A (and we know from Proposition A.1.2.4 that this image is convex when C is convex).

Even if U and X in (2.4.1) are Euclidean spaces, we seem to limit the generality when passing to (2.4.2), since only linear constraints are considered. Actually, (2.4.1) can be put in the form (2.4.2): with $X = \mathbb{R}^n$ and $y = (u, v) \in U \times X = \mathbb{R}^m$, define $Ay := v$ and $g(y) := \varphi(u) + i_C(y)$, where

$$C := \{ y = (u, v) \in \mathbb{R}^m \ : \ c(u) \leqslant v \}. \tag{2.4.3}$$

Note that conversely, (2.4.2) can be put in the form (2.4.1) via an analogous trick turning its equality constraints into inequalities.

Theorem 2.4.2 *Let g of Definition 2.4.1 be in $\operatorname{Conv} \mathbb{R}^m$. Assume also that, for all $x \in \mathbb{R}^n$, g is bounded below on the inverse image $\overset{-1}{A}(x) = \{ y \in \mathbb{R}^m \ : \ Ay = x \}$. Then $Ag \in \operatorname{Conv} \mathbb{R}^n$.*

Proof. By assumption, Ag is nowhere $-\infty$; also, $(Ag)(x) < +\infty$ whenever $x = Ay$, with $y \in \operatorname{dom} g$. Now consider the extended operator

$$A' : \ \mathbb{R}^m \times \mathbb{R} \ni (y, r) \mapsto A'(y, r) := (Ay, r) \in \mathbb{R}^m \times \mathbb{R}.$$

The set $A'(\operatorname{epi} g) =: C$ is convex in $\mathbb{R}^n \times \mathbb{R}$, let us compute its lower-bound function (1.3.5): for given $x \in \mathbb{R}^n$,

$$\inf_r\{r \ : \ (x,r) \in C\} = \inf_{y,r}\{r \ : \ Ay = x \text{ and } g(y) \leqslant r\}$$
$$= \inf_y\{g(y) \ : \ Ay = x\} = (Ag)(x)\,,$$

and this proves the convexity of $Ag = \ell_C$. $\qquad\qquad\qquad\qquad\qquad\qquad\square$

Usually, $\overset{-1}{A}(x)$ contains several points – it is an affine manifold of $\mathbb{R}^n$ – and $Ag(x)$ selects one giving the least value of g (admitting that (2.4.2) has a solution). If A is invertible, $Ag = g \circ A^{-1}$; more generally, the above proof discloses the following interpretation: epi (Ag) is the epigraphical hull of the inverse image $\overset{-1}{A'}(\text{epi } g)$ (a convex set in $\mathbb{R}^n \times \mathbb{R}$).

Corollary 2.4.3 *Let* (2.4.1) *have the following form:* $U = \mathbb{R}^p$; $\varphi \in \text{Conv } \mathbb{R}^p$; $X = \mathbb{R}^n$ *is equipped with the canonical basis; the mapping* c *has its components* $c_j \in \text{Conv } \mathbb{R}^p$ *for* $j = 1, \dots, n$. *Suppose also that the optimal value is* $> -\infty$ *for all* $x \in \mathbb{R}^n$, *and that*

$$\text{dom } \varphi \cap \text{dom } c_1 \cap \cdots \cap \text{dom } c_n \neq \emptyset\,. \qquad (2.4.4)$$

Then the value function

$$v_{\varphi,c}(x) := \inf\{\varphi(u) \ : \ c_j(u) \leqslant x_j, \text{ for } j = 1, \dots, n\}$$

lies in $\text{Conv } \mathbb{R}^n$.

Proof. Note first that we have assumed $v_{\varphi,c}(x) > -\infty$ for all x. Take u_0 in the set (2.4.4) and set $M := \max_j c_j(u_0)$; then take $x_0 := (M, \dots, M) \in \mathbb{R}^n$, so that $v_{\varphi,c}(x_0) \leqslant \varphi(u_0) < +\infty$. Knowing that $v_{\varphi,c}$ is an image-function, we just have to prove the convexity of the set (2.4.3); but this in turn comes immediately from the convexity of each c_j. $\qquad\qquad\qquad\qquad\qquad\qquad\square$

Taking the image of a convex function under a linear mapping can be used as a mould to describe a number of other operations – (2.4.1) is indeed one of them. An example is the infimal convolution of §2.3: with f_1 and f_2 in $\text{Conv } \mathbb{R}^n$, define $g \in \text{Conv }(\mathbb{R}^n \times \mathbb{R}^n)$ by

$$g(x_1, x_2) := f_1(x_1) + f_2(x_2)$$

and $A : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}^n$ by

$$A(x_1, x_2) := x_1 + x_2\,.$$

Then we have $Ag = f_1 \uplus f_2$ and (2.3.1) is put in the form (2.4.2). Incidentally, this shows that an image of a closed function need not be closed.

The following example of an image function is frequently encountered:

**Definition 2.4.4 (Marginal Function)** Let $g \in \text{Conv }(\mathbb{R}^n \times \mathbb{R}^m)$. Then

$$\mathbb{R}^n \ni x \mapsto \gamma(x) := \inf\{g(x,y) \ : \ y \in \mathbb{R}^m\}$$

is the *marginal function* of g. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Corollary 2.4.5 *With the above notation, suppose that g is bounded below on the set $\{x\} \times \mathbb{R}^m$, for all $x \in \mathbb{R}^n$. Then the marginal function γ lies in $\operatorname{Conv} \mathbb{R}^n$.*

Proof. The marginal function γ is the image of g under the the linear operator A projecting each $(x, y) \in \mathbb{R}^n \times \mathbb{R}^m$ onto $x \in \mathbb{R}^n$: $A(x, y) = x$.   □

Geometrically, a marginal function is given by Fig. 2.4.1, which explains why convexity is preserved: the strict epigraph of γ is the projection onto $\mathbb{R}^n \times \mathbb{R}$ of the strict epigraph of g ($\subset \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}$). Therefore, $\operatorname{epi}_s \gamma$ is also the image of a convex set under a linear mapping; see again Example A.1.2.5.
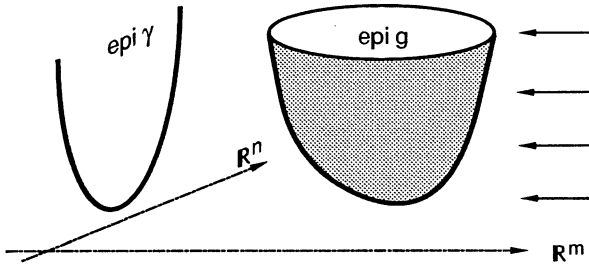
Fig. 2.4.1. The shadow of a convex epigraph

Remark 2.3.5 mentioned that an infimal convolution need not be closed. The above interpretation suggests that a marginal function need not be closed either. A counter-example for $(x, y) \in \mathbb{R} \times \mathbb{R}$ is the (closed) indicator function $g(x, y) = 0$ if $x > 0$, $y > 0$, $xy \geqslant 1$, $g(x, y) = +\infty$ otherwise. Its marginal function $\inf_y g(x, y)$ is the (non-closed) indicator of $]0, +\infty[$. Needless to say, an image-function is not closed in general.

As seen in §2.1(b), supremization preserves convexity. Here, if $g(\cdot, y)$ were concave for each y, γ would therefore be concave: the convexity of γ is a little surprising. It is of course the convexity of g with respect to the *couple* of variables x and y that is crucial.

2.5 Convex Hull and Closed Convex Hull of a Function

Given a (nonconvex) function g, a natural idea coming from §A.1.3 is to take the convex hull $\operatorname{co} \operatorname{epi} g$ of its epigraph. This gives a convex set, which is not an epigraph, but which can be made so by "closing its bottom" via its lower-bound function (1.3.5). As seen in §A.1.3, there are several ways of constructing a convex hull; the next result exploits them, and uses the unit simplex Δ_k.

Proposition 2.5.1 *Let $g : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, not identically $+\infty$, be minorized by some affine function: for some $(s, b) \in \mathbb{R}^n \times \mathbb{R}$,*

$$g(x) \geqslant \langle s, x \rangle - b \quad \text{for all } x \in \mathbb{R}^n \, . \tag{2.5.1}$$

Then, the following three functions f_1, f_2 and f_3 are convex and coincide on $\mathbb{R}^n$:

$$f_1(x) := \inf\{r : (x,r) \in \text{co epi } g\},$$
$$f_2(x) := \sup\{h(x) : h \in \text{Conv } \mathbb{R}^n,\ h \leqslant g\}.$$

$$f_3(x) := \inf\left\{\textstyle\sum_{j=1}^k \alpha_j g(x_j) : k = 1, 2, \ldots \right.$$
$$\left. \alpha \in \Delta_k,\ x_j \in \text{dom } g,\ \textstyle\sum_{j=1}^k \alpha_j x_j = x\right\}. \tag{2.5.2}$$

Proof. We denote by Γ the family of convex functions minorizing g. By assumption, $\Gamma \neq \emptyset$; then the convexity of f_1 results from §1.3(g).

$[f_2 \leqslant f_1]$ Consider the epigraph of any $h \in \Gamma$: its lower-bound function $\ell_{\text{epi } h}$ is h itself; besides, it contains epi g, and co (epi g) as well (see Proposition A.1.3.4). In a word, there holds $h = \ell_{\text{epi } h} \leqslant \ell_{\text{co epi } g} = f_1$ and we conclude $f_2 \leqslant f_1$ since h was arbitrary in Γ.

$[f_3 \leqslant f_2]$ We have to prove $f_3 \in \Gamma$, and the result will follow by definition of f_2; clearly $f_3 \leqslant g$ (take $\alpha \in \Delta_1$!), so it suffices to establish $f_3 \in \text{Conv } \mathbb{R}^n$. First, with (s, b) of (2.5.1) and all x, $\{x_j\}$ and $\{\alpha_j\}$ as described by (2.5.2),

$$\sum_{j=1}^k \alpha_j g(x_j) \geqslant \sum_{j=1}^k \alpha_j(\langle s, x_j\rangle - b) = \langle s, x\rangle - b;$$

hence f_3 is minorized by the affine function $\langle s, \cdot\rangle - b$. Now, take two points (x, r) and (x', r') in the strict epigraph of f_3. By definition of f_3, there are k, $\{\alpha_j\}$, $\{x_j\}$ as described in (2.5.2), and likewise k', $\{\alpha'_j\}$, $\{x'_j\}$, such that $\sum_{j=1}^k \alpha_j g(x_j) < r$ and likewise $\sum_{j=1}^{k'} \alpha'_j g(x'_j) < r'$.

For arbitrary $t \in]0, 1[$, we obtain by convex combination

$$\sum_{j=1}^k t\alpha_j g(x_j) + \sum_{j=1}^{k'} (1-t)\alpha'_j g(x'_j) < tr + (1-t)r'.$$

Observe that

$$\sum_{j=1}^k t\alpha_j x_j + \sum_{j=1}^{k'} (1-t)\alpha'_j x'_j = tx + (1-t)x',$$

i.e. we have in the lefthand side a convex decomposition of $tx + (1-t)x'$ in $k + k'$ elements; therefore, by definition of f_3:

$$f_3(tx + (1-t)x') \leqslant \sum_{j=1}^k t\alpha_j g(x_j) + \sum_{j=1}^{k'} (1-t)\alpha'_j g(x'_j)$$

and we have proved that $\text{epi}_s f_3$ is a convex set: f_3 is convex.

$[f_1 \leqslant f_3]$ Let $x \in \mathbb{R}^n$ and take an arbitrary convex decomposition $x = \sum_{j=1}^{k} \alpha_j x_j$, with α_j and x_j as described in (2.5.2). Since $(x_j, g(x_j)) \in \operatorname{epi} g$ for $j = 1, \ldots, k$,

$$\left(x, \sum_{j=1}^{k} \alpha_j g(x_j) \right) \in \operatorname{co} \operatorname{epi} g$$

and this implies $f_1(x) \leqslant \sum_{j=1}^{k} \alpha_j g(x_j)$ by definition of f_1. Because the decomposition of x was arbitrary within (2.5.2), this implies $f_1(x) \leqslant f_3(x)$. $\square$

Note in (2.5.2) the role of the convention $\inf \emptyset = +\infty$, in case x has no decomposition – which means that $x \notin \operatorname{co} \operatorname{dom} g$. The restrictions $x_j \in \operatorname{dom} g$ could be equally relaxed (an $x_j \notin \operatorname{dom} g$ certainly does not help making the infimum); notationally, α should then be taken in ri Δ_k, so as to avoid the annoying multiplication $0 \times (+\infty)$. Beware that epi $(\operatorname{co} g)$ is *not exactly* the convex hull co $(\operatorname{epi} g)$: we need to close the bottom of this latter set, as in §1.3(g)(ii) – an operation which affects only the relative boundary of co epi g, though. Note also that Carathéodory's Theorem yields an upper bound on k for (2.5.2), namely $k \leqslant (n+1) + 1 = n + 2$.

Instead of co epi g, we can take the closed convex hull $\overline{\operatorname{co}}$ epi $g = \operatorname{cl} \operatorname{co} \operatorname{epi} g$ (see §A.1.4). We obtain a closed set, with in particular a closed bottom: it is already an epigraph, the epigraph of a closed convex function. The corresponding operation that yielded f_1, f_2, f_3 is therefore now simpler. Furthermore, we know from Proposition 1.2.8 that all closed convex functions are redundant to define the function corresponding to f_2: affine functions are enough. We leave it as an exercise to prove the following result:

Proposition 2.5.2 *Let g satisfy the hypotheses of Proposition 2.5.1. Then the three functions below*

$$\bar{f}_1(x) := \inf \{ r : (x, r) \in \overline{\operatorname{co}} \operatorname{epi} g \},$$
$$\bar{f}_2(x) := \sup \{ h(x) : h \in \operatorname{\overline{Conv}} \mathbb{R}^n, \, h \leqslant g \},$$
$$\bar{f}_3(x) := \sup \{ \langle s, x \rangle - b : \langle s, y \rangle - b \leqslant g(y) \text{ for all } y \in \mathbb{R}^n \}$$

are closed, convex, and coincide on $\mathbb{R}^n$ with the closure of the function constructed in Proposition 2.5.1. $\square$

In view of the relationship between the operations studied in this Section 2.5 and the convexification of epi g, the following notation is justified, even if it is not quite accurate.

Definition 2.5.3 (Convex Hulls of a Function) Let $g : \mathbb{R}^n \to \mathbb{R}^n \cup \{+\infty\}$, not identically $+\infty$, be minorized by an affine function. The common function $f_1 = f_2 = f_3$ of Proposition 2.5.1 is called the *convex hull* of g, denoted by co g. The *closed convex hull* of g is any of the functions described by Proposition 2.5.2; it is denoted by $\overline{\operatorname{co}} \, g$ or cl co g. $\square$

If $\{g_j\}_{j\in J}$ is an arbitrary family of functions, all minorized by the same affine function, the epigraph of the [closed] convex hull of the function $\inf_{j\in J} g_j$ is obtained from $\cup_{j\in J}$ epi g_j. An important case is when the g_j's are *convex*; then, exploiting Example A.1.3.5, the formula giving f_3 simplifies: several x_j's corresponding to the same g_i can be compressed to a single convex combination.

Proposition 2.5.4 *Let* $g_1, \ldots, g_m$ *be in* Conv $\mathbb{R}^n$, *all minorized by the same affine function. Then the convex hull of their infimum is the function*

$$\mathbb{R}^n \ni x \mapsto [\mathrm{co}\,(\min_j g_j)](x) =$$
$$\inf \left\{ \sum_{j=1}^m \alpha_j g_j(x_j) \,:\, \alpha \in \Delta_m,\, x_j \in \mathrm{dom}\, g_j,\, \sum_{j=1}^m \alpha_j x_j = x \right\}. \qquad (2.5.3)$$

Proof. Apply Example A.1.3.5 to the convex sets $C_j = \mathrm{epi}\, g_j$. $\qquad\qquad\square$

The above statement was made in the simple situation of finitely many g_j's, but the representation (2.5.3) can be extended to an arbitrary family of convex functions g_j: it suffices to consider in the infimand all the representations of x as convex combinations of finitely many elements $x_j \in \mathrm{dom}\, g_j$.

Along these lines, note that an arbitrary function $g : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ can be seen as an infimum of convex functions: considering $\mathrm{dom}\, g$ as an index-set, we can write

$$g(x) = \inf \left\{ g(x_j) + i_{\{x_j\}}(x) \,:\, x_j \in \mathrm{dom}\, g \right\},$$

where each $g(x_j)$ denotes a (finite) constant function.
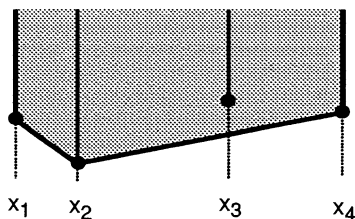
Fig. 2.5.1. A convex hull of needles

Example 2.5.5 Let $(x_1, b_1), \ldots, (x_m, b_m)$ be given in $\mathbb{R}^n \times \mathbb{R}$ and define for $j = 1, \ldots, m$

$$g_j(x) = \begin{cases} b_j & \text{if } x = x_j\,, \\ +\infty & \text{if not}. \end{cases}$$

Then $f := \mathrm{co}\,(\min g_j) = \overline{\mathrm{co}}\,(\min g_j)$ is the polyhedral function with the epigraph illustrated on Fig. 2.5.1, and analytically given by

$$f(x) = \begin{cases} \min \left\{ \sum_{j=1}^m \alpha_j b_j \,:\, \alpha \in \Delta_m,\, \sum_{j=1}^m \alpha_j x_j = x \right\} & \text{if } x \in \mathrm{co}\,\{x_1, \ldots, x_m\}\,, \\ +\infty & \text{if not}. \end{cases}$$

Calling $b \in \mathbb{R}^m$ the vector whose components are the b_i's and A the matrix whose columns are the x_i's, the above minimization problem in α can be written – at least when $x \in \operatorname{co}\{x_1, \ldots, x_m\}$:

$$f(x) = \min\{b^\top \alpha \ : \ \alpha \in \Delta_m, \ A\alpha = x\}. \qquad \qquad \square$$

To conclude this Section 2, Table 2.5.1 summarizes the main operations on functions and epigraphs that we have encountered.

Table 2.5.1. Main operations yielding convexity

Operations on functions: $f =$	Operations on sets: $\operatorname{epi} f$ or $\operatorname{epi}_s f =$	Closedness
$\sum_{j=1}^m t_j f_j$	nothing interesting	preserved
$\sup_{j \in J} f_j$	$\cap_{j \in J} \operatorname{epi} f_j$	preserved
$g \circ A$ (A affine)	$\overset{-1}{A'}(\operatorname{epi} g)$	preserved
$ug(x/u)$	$\mathbb{R}_*^+ (\{1\} \times \operatorname{epi} g)$	must be forced
$f_1 \Updownarrow f_2$	$\operatorname{epi}_s f_1 + \operatorname{epi}_s f_2$	destroyed
Ag (A linear)	epigr. hull of $A'(\operatorname{epi} g)$	destroyed
$\inf_y g(\cdot, y)$	$\operatorname{Proj}_{\mathbb{R}^n \times \mathbb{R}} \operatorname{epi}_s g$	destroyed
$\operatorname{co} g$	epigr. hull of $\operatorname{co} \operatorname{epi} g$	can be forced

3 Local and Global Behaviour of a Convex Function

3.1 Continuity Properties

Convex functions turn out to enjoy remarkable continuity properties: they are locally Lipschitzian on the relative interior of their domain. On the relative boundary of that domain, however, all kinds of continuity may disappear.

We start with a technical lemma.

Lemma 3.1.1 *Let* $f \in \operatorname{Conv} \mathbb{R}^n$ *and suppose there are* x_0, δ, m *and* M *such that*

$$m \leqslant f(x) \leqslant M \quad \text{for all } x \in B(x_0, 2\delta).$$

Then f *is Lipschitzian on* $B(x_0, \delta)$; *more precisely: for all* y *and* y' *in* $B(x_0, \delta)$,

$$|f(y) - f(y')| \leqslant \frac{M - m}{\delta} \|y - y'\|. \qquad (3.1.1)$$

Proof. Look at Fig. 3.1.1: with two different y and y' in $B(x_0, \delta)$, take

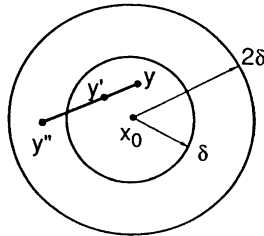$$y'' := y' + \delta \frac{y' - y}{\|y' - y\|} \in B(x_0, 2\delta);$$

by construction, y' lies on the segment $[y, y'']$, namely

$$y' = \frac{\|y' - y\|}{\delta + \|y' - y\|} y'' + \frac{\delta}{\delta + \|y' - y\|} y .$$

Applying the convexity of f and using the postulated bounds, we obtain

$$f(y') - f(y) \leqslant \frac{\|y' - y\|}{\delta + \|y' - y\|} [f(y'') - f(y)] \leqslant \frac{1}{\delta} \|y' - y\| (M - m) .$$

Then, it suffices to exchange y and y' to prove (3.1.1).   □

Fig. 3.1.1. Moving in a neighborhood of x_0

This implies the local Lipschitz continuity of a convex function, as announced.

Theorem 3.1.2 *With $f \in \operatorname{Conv} \mathbb{R}^n$, let S be a convex compact subset of* ri dom f. *Then there exists $L = L(S) \geqslant 0$ such that*

$$|f(x) - f(x')| \leqslant L\|x - x'\| \quad \text{for all } x \text{ and } x' \text{ in } S . \tag{3.1.2}$$

Proof. [*Preliminaries*] First of all, our statement ignores x-values outside the affine hull of the convex set dom f. Instead of $\mathbb{R}^n$, it can be formulated in $\mathbb{R}^d$, where d is the dimension of dom f; alternatively, we may assume ri dom $f = $ int dom f, which will simplify the writing.

Make this assumption and let $x_0 \in S$. We will prove that there are $\delta = \delta(x_0) > 0$ and $L = L(x_0, \delta)$ such that the ball $B(x_0, \delta)$ is included in int dom f and

$$|f(y) - f(y')| \leqslant L\|y - y'\| \quad \text{for all } y \text{ and } y' \text{ in } B(x_0, \delta) . \tag{3.1.3}$$

If this holds for all $x_0 \in S$, the corresponding balls $B(x_0, \delta)$ will provide a covering of the compact set S, from which we will extract a finite covering $(x_1, \delta_1, L_1), \ldots,$ (x_k, δ_k, L_k). With these balls, we will divide an arbitrary segment $[x, x']$ of the convex set S into finitely many subsegments, of endpoints $y_0 := x, \ldots, y_i, \ldots, y_\ell := x'$. Ordering properly the y_i's, we will have $\|x - x'\| = \sum_{i=1}^{\ell} \|y_i - y_{i-1}\|$; futhermore, f will be Lipschitzian on each $[y_{i-1}, y_i]$ with the common constant $L := \max\{L_1, \ldots, L_k\}$. The required Lipschitz property (3.1.2) will follow.

[Main Step] To establish (3.1.3), we use Lemma 3.1.1, which requires boundedness of f in the neighborhood of x_0. For this, we construct as in the proof of Theorem A.2.1.3 (see Fig. A.2.1.1) a simplex $\Delta = \mathrm{co}\,\{v_0, \ldots, v_n\}r \subset \mathrm{dom}\, f$ having x_0 in its interior: we can take $\delta > 0$ such that $B(x_0, 2\delta) \subset \Delta$.

Then any $y \in B(x_0, 2\delta)$ can be written: $y = \sum_{i=0}^n \alpha_i v_i$ with $\alpha \in \Delta_{n+1}$, so that the convexity of f gives

$$f(y) \leqslant \sum_{i=0}^n \alpha_i f(v_i) \leqslant \max\{f(v_0), \ldots, f(v_n)\} =: M\,.$$

On the other hand, Proposition 1.2.1 tells us that f is bounded from below, say by m, on this very same $B(x_0, 2\delta)$. Our claim is proved: we have singled out $\delta > 0$ such that $m \leqslant f(y) \leqslant M$ for all $y \in B(x_0, 2\delta)$. □

Note that the key-argument in the main step above is to find a (relative) neighborhood of $x \in \mathrm{ri}\,\mathrm{dom}\, f$, which is convex and which has a finite number of extreme points, all lying in $\mathrm{dom}\, f$. The simplex Δ is such a neighborhood, with a minimal number of extreme points

Remark 3.1.3 It follows in particular that f is continuous relatively to the relative interior of its domain, i.e.: for $x_0 \in \mathrm{ri}\,\mathrm{dom}\, f$ and $x \in \mathrm{ri}\,\mathrm{dom}\, f$ converging to x_0, we have that $f(x) \to f(x_0)$.

An equivalent formulation of Theorem 3.1.2 is: f is locally Lipschitzian on the relative interior of its domain, i.e. for all $x_0 \in \mathrm{ri}\,\mathrm{dom}\, f$, there are $L(x_0)$ and $\delta(x_0)$ such that

$$|f(x) - f(x')| \leqslant L(x_0)\|x - x'\| \quad \text{for all } x \text{ and } x' \text{ in the set}$$
$$S(x_0) := B(x_0, \delta(x_0)) \cap \mathrm{aff}\,\mathrm{dom}\, f \subset \mathrm{ri}\,\mathrm{dom}\, f\,.$$

In fact, the bulk of our proof is just concerned with this last statement. Of course, when x_0 gets closer to the relative boundary of $\mathrm{dom}\, f$, the size $\delta(x_0)$ of the allowed neighborhood shrinks to 0; but also, the local Lipschitz constant $L(x_0)$ may grow unboundedly ($\mathrm{gr}\, f$ may become steeper and steeper). □

Because of the phenomenon mentioned in the above remark, we cannot put $\mathrm{ri}\,\mathrm{dom}\, f$ instead of S in Theorem 3.1.2: a convex function need not be Lipschitzian on the relative interior of its domain.

Let us sum up the continuity properties of a convex function.

– First of all it is aff $\mathrm{dom}\, f$, and not $\mathbb{R}^n$, that is the relevant embedding (affine) space: there is no point in studying the behaviour of f when moving out of this space. Continuity, and even Lipschitz continuity, holds when x remains "well inside" $\mathrm{ri}\,\mathrm{dom}\, f$.

– When x approaches $\mathrm{rbd}\,\mathrm{dom}\, f$, continuity may break down: f may go to infinity, or jump discontinuously to some finite value, etc. Still, irregular behaviour of f is limited by Proposition 1.2.5.

– Closing epi f if necessary, we obtain a lower semi-continuous function at a reasonable price (specified by Proposition 1.2.5).

– It remains to ask whether f can be assumed upper semi-continuous (on rbd dom f, and relative to dom f). It turns out that this property automatically holds for univariate functions. With several variables, the answer is no in general, though: a counter-example is

$$\mathbb{R}^2 \ni x = (\xi, \eta) \mapsto f(x) = \sup_{\alpha, \beta} \left\{ \xi\alpha + \eta\beta \ : \ \tfrac{1}{2}\alpha^2 \leqslant \beta \right\} .$$

We see that $f(0) = 0$, and we know from Proposition 2.1.2 that $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$. In fact, the optimal (α, β) (if any) satisfies $1/2\,\alpha^2 = \beta$, so that

$$f(\xi, \eta) = \sup_{\alpha} \left(\tfrac{1}{2}\eta\alpha^2 + \xi\alpha \right) = \begin{cases} 0 & \text{if } \xi = \eta = 0 \,, \\ -\dfrac{\xi^2}{2\eta} & \text{if } \eta < 0 \,, \\ +\infty & \text{otherwise.} \end{cases} \qquad (3.1.4)$$

Thus, when x tends to 0 following the path $\eta = -1/2\,\xi^2$, then $f(x) \equiv 1 > 0 = f(0)$.

To conclude this subsection, we give a rather powerful convergence result: convex functions converging pointwise to some (convex) function f do converge *uniformly* on each compact set contained in the relative interior of dom f. For the sake of simplicity, we limit ourselves here to the case of finite-valued functions. For the general case, just specify that the compact set S in the next statement must be in ri dom f, and adapt the proof accordingly.

Theorem 3.1.4 *Let the convex functions $f_k : \mathbb{R}^n \to \mathbb{R}$ converge pointwise for $k \to +\infty$ to $f : \mathbb{R}^n \to \mathbb{R}$. Then f is convex and, for each compact set S, the convergence of f_k to f is uniform on S.*

Proof. Convexity of f is trivial: pass to the limit in the definition (1.1.1) itself. For uniformity, we want to use Lemma 3.1.1, so we need to bound f_k on S independently of k; thus, let $r > 0$ be such that $S \subset B(0, r)$.

[*Step 1*] First the function $g := \sup_k f_k$ is convex, and $g(x) < +\infty$ for all x because the convergent sequence $(f_k(x))_k$ is certainly bounded. Hence, g is continuous and therefore bounded, say by M, on the compact set $B(0, 2r)$:

$$f_k(x) \leqslant g(x) \leqslant M \quad \text{for all } k \text{ and all } x \in B(0, 2r) \,.$$

Second, the convergent sequence $(f_k(0))_k$ is bounded from below:

$$\mu \leqslant f_k(0) \quad \text{for all } k \,.$$

Then, for $x \in B(0, 2r)$ and all k, use convexity on $[-x, x] \subset B(0, 2r)$:

$$2\mu \leqslant 2f_k(0) \leqslant f_k(x) + f_k(-x) \leqslant f_k(x) + M \,,$$

i.e. the f_k's are bounded from below, independently of k. Thus, we are within the conditions of Lemma 3.1.1: there is some L (independent of k) such that

$$|f_k(y) - f_k(y')| \leqslant L\|y - y'\| \quad \text{for all } k \text{ and all } y, y' \text{ in } B(0, r) . \qquad (3.1.5)$$

Naturally, the same Lipschitz property is transmitted to the limiting function f.

[*Step 2*] Now fix $\varepsilon > 0$. Cover S by the balls $B(x, \varepsilon)$ for x describing S, and extract a finite covering $S \subset B(x_1, \varepsilon) \cup \cdots \cup B(x_m, \varepsilon)$. With x arbitrary in S, take an x_i such that $x \in B(x_i, \varepsilon)$. There is $k_{i,\varepsilon}$ such that, for all $k \geqslant k_{i,\varepsilon}$,

$$|f_k(x) - f(x)| \leqslant |f_k(x) - f_k(x_i)| + |f_k(x_i) - f(x_i)| + |f(x_i) - f(x)| \leqslant (2L+1)\varepsilon$$

where we have also used (3.1.5), knowing that x and x_i are in $S \subset B(0, r)$. The above inequality is then valid uniformly in x, providing that

$$k \geqslant \max\{k_{1,\varepsilon}, \ldots, k_{m,\varepsilon}\} =: k_\varepsilon . \qquad \square$$

3.2 Behaviour at Infinity

Having studied the behaviour of $f(x)$ when x approaches rbd dom f, it remains to consider the case of unbounded x. An important issue is the behaviour of $f(x_0 + td)$ when $t \to +\infty$ (x_0 and d being fixed). Once again, we make this study by viewing epi f just as a convex set (unbounded), and we call for §A.2.2.

Thus we assume $f \in \overline{\text{Conv}} \, \mathbb{R}^n$, which allows us to consider the asymptotic cone (epi $f)_\infty$ of the closed convex set epi f. It is a closed convex cone of $\mathbb{R}^n \times \mathbb{R}$, which clearly contains the half-line $\{0\} \times \mathbb{R}^+$. According to its Definition A.2.2.2,

$$(\text{epi } f)_\infty = \{(d, \rho) \in \mathbb{R}^n \times \mathbb{R} : (x_0, r_0) + t(d, \rho) \in \text{epi } f \text{ for all } t > 0\}, \quad (3.2.1)$$

where (x_0, r_0) is an arbitrary element of epi f. This can be written

$$(\text{epi } f)_\infty = \{(d, \rho) : \text{epi } f + t(d, \rho) \subset \text{epi } f \text{ for all } t > 0\}$$

and, since we already know that (epi $f)_\infty$ is a convex cone:

$$(\text{epi } f)_\infty = \{(d, \rho) : \text{epi } f + (d, \rho) \subset \text{epi } f\} .$$

This object turns out to define a new function:

Proposition 3.2.1 *For $f \in \overline{\text{Conv}} \, \mathbb{R}^n$, the asymptotic cone of* epi f *is the epigraph of the function* $f'_\infty \in \overline{\text{Conv}} \, \mathbb{R}^n$ *defined by*

$$d \mapsto f'_\infty(d) := \sup_{t>0} \frac{f(x_0 + td) - f(x_0)}{t} = \lim_{t \to +\infty} \frac{f(x_0 + td) - f(x_0)}{t} , \quad (3.2.2)$$

where x_0 is arbitrary in dom f.

Proof. Since $(x_0, f(x_0)) \in$ epi f, (3.2.1) tells us that $(d, \rho) \in (\text{epi } f)_\infty$ if and only if $f(x_0 + td) \leqslant f(x_0) + t\rho$ for all $t > 0$, which means

$$\sup_{t>0} \frac{f(x_0 + td) - f(x_0)}{t} \leqslant \rho. \qquad (3.2.3)$$

In other words, $(\mathrm{epi}\, f)_\infty$ is the epigraph of the function whose value at d is the lefthand side of (3.2.3); and this is true no matter how x_0 has been chosen in $\mathrm{dom}\, f$. The rest follows from the fact that the difference quotient in (3.2.3) is closed convex in d, and increasing in t (the function $t \mapsto f(x_0 + td)$ is convex and enjoys the property of increasing slopes, namely Proposition 0.6.1). $\square$

It goes without saying that the expressions appearing in (3.2.2) are independent of x_0: f'_∞ is really a function of d only. By construction, this function is positively homogeneous: $f'_\infty(\alpha d) = \alpha f'_\infty(d)$ for all $\alpha > 0$. Our notation suggests that it is something like a "slope at infinity" in the direction d.

Definition 3.2.2 (Asymptotic function) The function f'_∞ of Proposition 3.2.1 is called the *asymptotic function*, or recession function, of f. $\square$

Consider for example the indicator i_C of a closed convex set C. By definition of the asymptotic cone, we see that $\mathrm{i}_C(x_0 + td) = 0$ for all $t > 0$ if and only if $d \in C_\infty$; we obtain

$$(\mathrm{i}_C)'_\infty = \mathrm{i}_{(C_\infty)}.$$

The next example is more interesting and extends Remark 2.2.3:

Example 3.2.3 Let $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$. Take $x_0 \in \mathrm{dom}\, f$ and consider the convex function $d \mapsto f(x_0 + d) - f(x_0)$, whose domain contains 0, and whose perspective-function is r of (2.2.2). The closure of r can be computed with the help of Proposition 2.2.2: with $x_0 + d'$ arbitrary in $\mathrm{ri}\, \mathrm{dom}\, f$,

$$(\mathrm{cl}\, r)(0, d) = \lim_{\alpha \downarrow 0} \alpha[f(x_0 + d' - d + d/\alpha) - f(x_0)].$$

Note that the term $f(x_0) < +\infty$ can be suppressed, or replaced by $f(x_0 + d')$ (because $\alpha \downarrow 0$); moreover, as in Remark 2.2.3, the above limit is exactly

$$\lim_{t \to +\infty} \frac{f(x_0 + d' + td)}{t} = \lim_{t \to +\infty} \frac{f(x_0 + d' + td) - f(x_0 + d')}{t} = f'_\infty(d).$$

In summary, the function defined by

$$\mathbb{R} \times \mathbb{R}^n \ni (u, d) \mapsto \begin{cases} u[f(x_0 + d/u) - f(x_0)] & \text{if } u > 0, \\ f'_\infty(d) & \text{if } u = 0, \\ +\infty & \text{elsewhere} \end{cases}$$

is in $\overline{\mathrm{Conv}}(\mathbb{R} \times \mathbb{R}^n)$; and only its "$u > 0$-part" depends on the reference point $x_0 \in \mathrm{dom}\, f$. $\square$

Our next result assesses the importance of the asymptotic function.

Proposition 3.2.4 *Let* $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$. *All the nonempty sublevel-sets of* f *have the same asymptotic cone, which is the sublevel-set of* f'_∞ *at the level* 0:

$$\forall r \in \mathbb{R} \text{ with } \mathrm{S}_r(f) \neq \emptyset, \quad [\mathrm{S}_r(f)]_\infty = \{d \in \mathbb{R}^n : f'_\infty(d) \leqslant 0\}.$$

In particular, the following statements are equivalent:

(i) *There is r for which* $S_r(f)$ *is nonempty and compact;*
(ii) *all the sublevel-sets of f are compact;*
(iii) $f'_\infty(d) > 0$ *for all nonzero* $d \in \mathbb{R}^n$.

Proof. By definition (A.2.2.1), a direction d is in the asymptotic cone of the nonempty sublevel-set $S_r(f)$ if and only if

$$x \in S_r(f) \implies [x + td \in S_r(f) \text{ for all } t > 0],$$

which can also be written – see (1.1.4):

$$(x, r) \in \text{epi } f \implies (x + td, r + t \times 0) \in \text{epi } f \text{ for all } t > 0;$$

and this in turn just means that $(d, 0) \in (\text{epi } f)_\infty = \text{epi } f'_\infty$. We have proved the first part of the theorem.

A particular case is when the sublevel-set $S_0(f'_\infty)$ is reduced to the singleton $\{0\}$, which exactly means (iii). This is therefore equivalent to $[S_r(f)]_\infty = \{0\}$ for all $r \in \mathbb{R}$ with $S_r(f) \neq \emptyset$, which means that $S_r(f)$ is compact (Proposition A.2.2.3). The equivalence between (i), (ii) and (iii) is proved.   □

Needless to say, the convexity of f is essential to ensure that *all* its nonempty sublevel-sets have the *same* asymptotic cone. In Remark 1.1.7, we have seen (closed) quasi-convex functions: their sublevel-sets are all convex, and as such they have asymptotic cones, which normally depend on the level.

Definition 3.2.5 (Coercivity) The functions $f \in \overline{\text{Conv}} \, \mathbb{R}^n$ satisfying (i), (ii) or (iii) are called *0-coercive*. Equivalently, the 0-coercive functions are those that "increase at infinity":

$$f(x) \to +\infty \quad \text{whenever} \quad \|x\| \to +\infty,$$

and closed convex 0-coercive functions achieve their minimum over $\mathbb{R}^n$.

An important particular case is when $f'_\infty(d) = +\infty$ for all $d \neq 0$, i.e. when $f'_\infty = i_{\{0\}}$. It can be seen that this means precisely

$$\frac{f(x)}{\|x\|} \to +\infty \quad \text{whenever} \quad \|x\| \to +\infty$$

(to establish this equivalence, extract a cluster point of $(x_k/\|x_k\|)_k$ and use the lower semi-continuity of f'_∞). In words: at infinity, f increases to infinity faster than any affine function; such functions are called *1-coercive*, or sometimes just coercive.   □

Suppose for example that f is quadratic:

$$f(x) = \tfrac{1}{2}\langle Qx, x \rangle + \langle b, x \rangle + c,$$

with Q a positive semi-definite symmetric operator, $b \in \mathbb{R}^n$ and $c \in \mathbb{R}$. Then it is easy to compute

$$f'_\infty(d) = \begin{cases} \langle b, d \rangle & \text{if } d \in \text{Ker } Q, \\ +\infty & \text{if not.} \end{cases}$$

In this particular case, the different sorts of coercivity coincide:

$$f \text{ is 0-coercive} \quad \Longleftrightarrow \quad f \text{ is 1-coercive} \quad \Longleftrightarrow \quad Q \text{ is positive definite.}$$

The word "coercive" alone comes from the study of bilinear forms: for our more general framework of non-quadratic functions, it becomes ambiguous, hence our distinction.

Proposition 3.2.6 *A function* $f \in \overline{\text{Conv}}\, \mathbb{R}^n$ *is Lipschitzian on the whole of* $\mathbb{R}^n$ *if and only if* f'_∞ *is finite on the whole of* $\mathbb{R}^n$. *The best Lipschitz constant for* f *is then*

$$\sup \{f'_\infty(d) \,:\, \|d\| = 1\}. \tag{3.2.4}$$

Proof. When the (convex) function f'_∞ is finite-valued, it is continuous (§3.1) and therefore bounded on the compact unit sphere:

$$\sup \{f'_\infty(d) \,:\, \|d\| = 1\} =: L < +\infty,$$

which implies by positive homogeneity

$$f'_\infty(d) \leqslant L\|d\| \quad \text{for all } d \in \mathbb{R}^n.$$

Now use the definition (3.2.2) of f'_∞:

$$f(x + d) - f(x) \leqslant L\|d\| \quad \text{for all } x \in \text{dom } f \text{ and } d \in \mathbb{R}^n;$$

thus, dom f is the whole space ($f(x + d) < +\infty$ for all d) and we do obtain that L is a global Lipschitz constant for f.

Conversely, let f have a global Lipschitz constant L. Pick $x_0 \in \text{dom } f$ and plug the inequality
$$f(x_0 + td) - f(x_0) \leqslant Lt\|d\| \quad \text{for all } t > 0 \text{ and } d \in \mathbb{R}^n$$
into the definition (3.2.2) of f'_∞ to obtain $f'_\infty(d) \leqslant L\|d\|$ for all $d \in \mathbb{R}^n$.

It follows that f'_∞ is finite everywhere, and the value (3.2.4) does not exceed L. □

Concerning (3.2.4), it is worth mentioning that d can run in the unit ball B (instead of the unit sphere $\widetilde{B}$), and/or the supremand can be replaced by its absolute value; these two replacements are made possible thanks to convexity.

Remark 3.2.7 We mention the following classification, of interest for minimization theory: let $f \in \overline{\text{Conv}}\, \mathbb{R}^n$.

– If $f'_\infty(d) < 0$ for some d, then f is unbounded from below; more precisely: for all $x_0 \in \text{dom } f$, $f(x_0 + td) \downarrow -\infty$ when $t \to +\infty$.

– The condition $f'_\infty(d) > 0$ for all $d \neq 0$ is necessary and sufficient for f to have a nonempty bounded (hence compact) set of minimum points.

– If $f'_\infty \geqslant 0$, with $f'_\infty(d) = 0$ for some $d \neq 0$, existence of a minimum cannot be guaranteed (but if x_0 is minimal, so is the half-line $x_0 + \mathbb{R}^+ d$).

Observe that, if the continuous function $d \mapsto f'_\infty(d)$ is positive for all $d \neq 0$, then it is minorized by some $m > 0$ on the unit sphere $\widetilde{B}$ and this m also minorizes the speed at which f increases at infinity. □

To close this section, we mention some calculus rules on the asymptotic function. They come directly either from the analytic definition (3.2.2), or from the geometric definition epi $f'_\infty = (\text{epi } f)_\infty$ combined with Proposition A.2.2.5.

Proposition 3.2.8

– *Let $f_1, \ldots, f_m$ be m functions of* $\overline{\mathrm{Conv}}\,\mathbb{R}^n$, *and* $t_1, \ldots, t_m$ *be positive numbers. Assume that there is* x_0 *at which each* f_j *is finite. Then,*

$$\text{for } f := \sum_{j=1}^m t_j f_j, \quad \text{we have } f'_\infty = \sum_{j=1}^m t_j (f_j)'_\infty.$$

– *Let* $\{f_j\}_{j \in J}$ *be a family of functions in* $\overline{\mathrm{Conv}}\,\mathbb{R}^n$. *Assume that there is* x_0 *at which* $\sup_{j \in J} f_j(x_0) < +\infty$. *Then,*

$$\text{for } f := \sup_{j \in J} f_j, \quad \text{we have } f'_\infty = \sup_{j \in J} (f_j)'_\infty.$$

– *Let* $A : \mathbb{R}^n \to \mathbb{R}^m$ *be affine with linear part* A_0, *and let* $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^m$. *Assume that* $A(\mathbb{R}^n) \cap \mathrm{dom}\, f \neq \emptyset$. *Then* $(f \circ A)'_\infty = f'_\infty \circ A_0$. $\square$

We mention the corresponding formula for an image-function (2.4.2): it would be $(Ag)'_\infty(d) = \inf \{g'_\infty(z) : Az = d\}$, which can be written symbolically $(Ag)'_\infty = A(g'_\infty)$. However, this formula cannot hold without an additional assumption, albeit to guarantee $Ag \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$. One such assumption is: $g'_\infty(z) > 0$ for all $z \in \mathrm{Ker}\,A \backslash \{0\}$ (appropriate coercivity is added where necessary, so that the infimum in the definition of $(Ag)(x)$ is always attained). Proving this result is not simple.

4 First- and Second-Order Differentiation

Let $C \subset \mathbb{R}^n$ be nonempty and convex. For a function f defined on C ($f(x) < +\infty$ for all $x \in C$), we study here the following questions:

– When f is convex and differentiable on C, what can be said of the gradient ∇f?
– When f is differentiable on C, can we characterize its convexity in terms of ∇f?
– When f is convex on C, what can be said of its first and second differentiability?

We start with the first two questions.

4.1 Differentiable Convex Functions

First we assume that f is differentiable on C. Given $x_0 \in C$, the sentence "f is differentiable at x_0" is meaningful only if f is at least defined in a neighborhood of x_0. Then, it is normal to assume that C is contained in an open set Ω on which f is differentiable.

Theorem 4.1.1 *Let f be a function differentiable on an open set $\Omega \subset \mathbb{R}^n$, and let C be a convex subset of Ω. Then*

(i) *f is convex on C if and only if*

$$f(x) \geqslant f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle \quad \text{for all } (x_0, x) \in C \times C; \quad (4.1.1)$$

(ii) f is strictly convex on C if and only if strict inequality holds in (4.1.1) whenever $x \neq x_0$;

(iii) f is strongly convex with modulus c on C if and only if, for all $(x_0, x) \in C \times C$,

$$f(x) \geqslant f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle + \tfrac{1}{2}c\|x - x_0\|^2 . \qquad (4.1.2)$$

Proof. [(i)] Let f be convex on C: for arbitrary $(x_0, x) \in C \times C$ and $\alpha \in \,]0, 1[$, we have from the definition (1.1.1) of convexity

$$f(\alpha x + (1 - \alpha)x_0) - f(x_0) \leqslant \alpha[f(x) - f(x_0)] .$$

Divide by α and let $\alpha \downarrow 0$: observing that $\alpha x + (1 - \alpha)x_0 = x_0 + \alpha(x - x_0)$, the lefthand side tends to $\langle \nabla f(x_0), x - x_0 \rangle$ and (4.1.1) is established.

Conversely, take x_1 and x_2 in C, $\alpha \in \,]0, 1[$ and set $x_0 := \alpha x_1 + (1 - \alpha)x_2 \in C$. By assumption,

$$f(x_i) \geqslant f(x_0) + \langle \nabla f(x_0), x_i - x_0 \rangle \quad \text{for } i = 1, 2 \qquad (4.1.3)$$

and we obtain by convex combination

$$\alpha f(x_1) + (1 - \alpha)f(x_2) \geqslant f(x_0) + \langle \nabla f(x_0), \alpha x_1 + (1 - \alpha)x_2 - x_0 \rangle$$

which, after simplification, is just the relation of definition (1.1.1).

[(ii)] If f is strictly convex, we have for $x_0 \neq x$ in C and $\alpha \in \,]0, 1[$,

$$f(x_0 + \alpha(x - x_0)) - f(x_0) < \alpha[f(x) - f(x_0)] ;$$

but f is in particular convex and we can use (i):

$$\langle \nabla f(x_0), \alpha(x - x_0) \rangle \leqslant f(x_0 + \alpha(x - x_0)) - f(x_0) ,$$

so the required strict inequality follows.

For the converse, proceed as for (i), starting from strict inequalities in (4.1.3).

[(iii)] Using Proposition 1.1.2, just apply (i) to the function $f - 1/2\,c\|\cdot\|^2$, which is of course differentiable. $\qquad \square$

Thus, a differentiable function is convex when its graph lies above its tangent hyperplanes: for each x_0, f is minorized by its first-order approximation $f(x_0) + \langle \nabla f(x_0), \cdot - x_0 \rangle$ (which coincides with f at x_0). It is strictly convex when the coincidence set is reduced to the singleton $(x_0, f(x_0))$. Finally, f is strongly convex when it is minorized by the quadratic convex function

$$x \mapsto f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle + \tfrac{1}{2}c\|x - x_0\|^2 ,$$

whose gradient at x_0 is also $\nabla f(x_0)$. These tangency properties are illustrated on Fig. 4.1.1.

Remark 4.1.2 Inequality (4.1.1) is fundamental. In case of convexity, the remainder term r in the development

$$f(x) = f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle + r(x_0, x)$$

must be well-behaved; for example, it is nonnegative for all x and x_0; also, $r(x_0, \cdot)$ is convex. $\qquad \square$
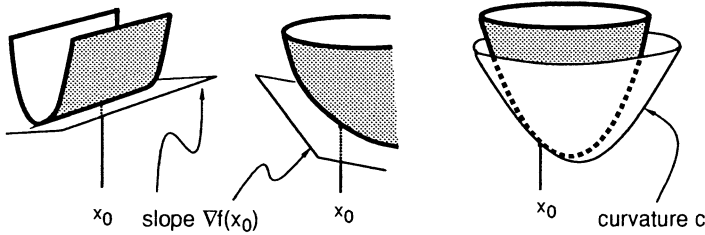
Fig. 4.1.1. Affine and quadratic minorizations

Both f- and ∇f-values appear in the relations dealt with in Theorem 4.1.1; we now proceed to give additional relations, involving ∇f only. We know that a differentiable univariate function is convex if and only if its derivative is monotone increasing (on the interval where the function is studied; see §0.6). Here, we need a generalization of the wording "monotone increasing" to our multidimensional situation. There are several possibilities, one is particularly well-suited to convexity:

Definition 4.1.3 Let $C \subset \mathbb{R}^n$ be convex. The mapping $F : C \to \mathbb{R}^n$ is said to be *monotone* [resp. strictly monotone, resp. strongly monotone with modulus $c > 0$] on C when, for all x and x' in C,

$$\langle F(x) - F(x'), x - x' \rangle \geqslant 0$$
$$[\text{resp. } \langle F(x) - F(x'), x - x' \rangle > 0 \quad \text{whenever} \quad x \neq x',$$
$$\text{resp. } \langle F(x) - F(x'), x - x' \rangle \geqslant c\|x - x'\|^2 \,].$$
□

In the univariate case, the present monotonicity thus corresponds to F being *increasing*. When particularized to a gradient mapping $F = \nabla f$, our definition characterizes the convexity of the underlying potential function f:

Theorem 4.1.4 *Let f be a function differentiable on an open set $\Omega \subset \mathbb{R}^n$, and let C be a convex subset of Ω. Then, f is convex [resp. strictly convex, resp. strongly convex with modulus c] on C if and only if its gradient ∇f is monotone [resp. strictly monotone, resp. strongly monotone with modulus c] on C.*

Proof. We combine the "convex ⇔ monotone" and "strongly convex ⇔ strongly monotone" cases by accepting the value $c = 0$ in the relevant relations such as (4.1.2).

Thus, let f be [strongly] convex on C: in view of Theorem 4.1.1, we can write for arbitrary x_0 and x in C:

$$f(x) \geqslant f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle + \tfrac{1}{2}c\|x - x_0\|^2$$
$$f(x_0) \geqslant f(x) + \langle \nabla f(x), x_0 - x \rangle + \tfrac{1}{2}c\|x_0 - x\|^2 \,,$$

and mere addition shows that ∇f is [strongly] monotone.

Conversely, let (x_0, x_1) be a pair of elements in C. Consider the univariate function $t \mapsto \varphi(t) := f(x_t)$, where $x_t := x_0 + t(x_1 - x_0)$; for t in an open interval

containing $[0,1]$, $x_t \in \Omega$ and φ is well-defined and differentiable; its derivative at t is $\varphi'(t) = \langle \nabla f(x_t), x_1 - x_0 \rangle$. Thus, we have for all $0 \leqslant t' < t \leqslant 1$

$$\begin{aligned}\varphi'(t) - \varphi'(t') &= \langle \nabla f(x_t) - \nabla f(x_{t'}), \ x_1 - x_0 \rangle \\ &= \tfrac{1}{t-t'}\langle \nabla f(x_t) - \nabla f(x_{t'}), x_t - x_{t'} \rangle\end{aligned} \qquad (4.1.4)$$

and the monotonicity relation for ∇f shows that φ' is increasing, φ is therefore convex (Corollary 0.6.5).

For strong convexity, set $t' = 0$ in (4.1.4) and use the strong monotonicity relation for ∇f:

$$\varphi'(t) - \varphi'(0) \geqslant \tfrac{1}{t}c\|x_t - x_0\|^2 = tc\|x_1 - x_0\|^2 . \qquad (4.1.5)$$

Because the differentiable convex function φ is the integral of its derivative, we can write

$$\varphi(1) - \varphi(0) - \varphi'(0) = \int_0^1 [\varphi'(t) - \varphi'(0)]\mathrm{d}t \geqslant \tfrac{1}{2}c\|x_1 - x_0\|^2$$

which, by definition of φ , is just (4.1.2) (the coefficient $1/2$ is $\int_0^1 t \, \mathrm{d}t$!).

The same technique proves the "strictly monotone $\Leftrightarrow$ strictly convex" case; then, (4.1.5) becomes a strict inequality – with $c = 0$ – and remains so after integration. $\qquad\square$

The attention of the reader is drawn on the coefficient c – and not $1/2\,c$ – in the definition 4.1.3 of strong monotonicity. Actually, a sensible rule is: "Use $1/2$ when dealing with a square"; here, the scalar product $\langle \Delta F, \Delta x \rangle$ is homogeneous to a square. Alternatively, remember in Proposition 1.1.2 that the gradient of $1/2\,c\|\cdot\|^2$ at x is cx.

We mention the following example: let $f(x) := \tfrac{1}{2}\langle Ax, x \rangle + \langle b, x \rangle$ be a quadratic convex function: A is symmetric, call $\lambda_n \geqslant 0$ its smallest eigenvalue. Observe that $\nabla f(x) = Ax + b$ and that

$$\langle Ax - Ax', x - x' \rangle = \langle A(x - x'), x - x' \rangle \geqslant \lambda_n\|x - x'\|^2 .$$

Thus ∇f is monotone [strongly with modulus λ_n]. The [strong] convexity of f, in the sense of (1.1.2), has been already alluded to in §1.3(d); but (4.1.2) is easier to establish here: simply write

$$\begin{aligned}f(x) - f(x_0) - \langle \nabla f(x_0), x - x_0 \rangle &= \tfrac{1}{2}\langle Ax, x \rangle - \tfrac{1}{2}\langle Ax_0, x_0 \rangle - \langle Ax_0, x - x_0 \rangle \\ &= \tfrac{1}{2}\langle A(x - x_0), x - x_0 \rangle \geqslant \tfrac{1}{2}\lambda_n\|x - x_0\|^2 .\end{aligned}$$

Note that for this particular class of convex functions, strong and strict convexity are equivalent to each other, and to the positive definiteness of A.

Remark 4.1.5 Do not infer from Theorem 4.1.4 the statement "a monotone mapping is the gradient of a convex function", which is wrong. To be so, the mapping in question must first be a gradient, an issue which we do not study here. We just mention the following property: if Ω is convex and $F : \Omega \to \mathbb{R}^n$ is differentiable, then F is a gradient if and only if its Jacobian operator is symmetric (in 2 or 3 dimensions, curl $F \equiv 0$). $\qquad\square$

4.2 Nondifferentiable Convex Functions

A convex function need not be differentiable over the whole interior of its domain; nevertheless, it is so at "many points" in this set. Before making this sentence mathematically precise, we note the following nice property of convex functions.

Proposition 4.2.1 *For $f \in \mathrm{Conv}\, \mathbb{R}^n$ and $x \in \mathrm{int\, dom}\, f$, the three statements below are equivalent:*

(i) *The function*

$$\mathbb{R}^n \ni d \mapsto \lim_{t \downarrow 0} \frac{f(x + td) - f(x)}{t}$$

  is linear in d;

(ii) *for some basis of $\mathbb{R}^n$ in which $x = (\xi^1, \ldots, \xi^n)$, the partial derivatives $\dfrac{\partial f}{\partial \xi^i}(x)$ exist at x, for $i = 1, \ldots, n$;*

(iii) *f is differentiable at x.*

Proof. First of all, remember from Theorem 0.6.3 that the one-dimensional function $t \mapsto f(x + td)$ has half-derivatives at 0: the limits considered in (i) exist for all d. We will denote by $\{b_1, \ldots, b_n\}$ the basis postulated in (ii), so that $x = \sum_{i=1}^n \xi^i b_i$.

Denote by $d \mapsto \ell(d)$ the function defined in (i); taking $d = \pm b_i$, realize that, when (i) holds,

$$\lim_{\tau \uparrow 0} \frac{f(x + \tau b_i) - f(x)}{-\tau} = \ell(-b_i) = -\ell(b_i) = -\lim_{t \downarrow 0} \frac{f(x + t b_i) - f(x)}{t}.$$

This means that the two half-derivatives at $t = 0$ of the function $t \mapsto f(x + t b_i)$ coincide: the partial derivative of f at x along b_i exists, (ii) holds. That (iii) implies (i) is clear: when f is differentiable at x,

$$\lim_{t \downarrow 0} \frac{f(x + td) - f(x)}{t} = \langle \nabla f(x), d \rangle.$$

We do not really complete the proof here, because everything follows in a straightforward way from subsequent chapters. More precisely, [(ii) $\Rightarrow$ (i)] is Proposition C.1.1.6, which states that the function ℓ is linear on the space generated by the b_i's, whenever it its linear along each b_i. Finally [(i) $\Rightarrow$ (iii)] results from Lemma D.2.1.1 and the proof goes as follows. One of the possible definitions of (iii) is:

$$\lim_{t \downarrow 0, d' \to d} \frac{f(x + td') - f(x)}{t} \quad \text{is linear in } d.$$

Because f is locally Lipschitzian, the above limit exists whenever it exists for fixed $d' = d$ – i.e. the expression in (i).   □

Remark 4.2.2 The above result reveals three interesting properties enjoyed by convex functions:

– First, consider the restriction $\varphi_d(t) := f(x + td)$ of f along a line $x + \mathbb{R}d$. As soon as there are n independent directions, say $d_1, \ldots, d_n$, such that each φ_{d_i} has a derivative at $t = 0$, then the same property holds for all the other possible directions $d \in \mathbb{R}^n$.

– Second, this "radial" differentiability property of φ_d for all d (or for many enough d) suffices to guarantee the "global" (i.e. Fréchet) differentiability of f at x; a property which does not hold in general. It depends crucially on the Lipschitz property of f.

– One can also show that, if f is convex and differentiable in a neighborhood of x, then ∇f is continuous at x. Hence, if Ω is an open convex set, the following equivalence holds true for the convex f:

$$f \text{ differentiable on } \Omega \quad \Longleftrightarrow \quad f \in C^1(\Omega).$$

This rather surprising property will be confirmed by Theorem D.6.2.4.   □

The largest set on which a function can be differentiable is the interior of its domain. Actually, a result due to H. Rademacher (1919) says that a function which is locally Lipschitzian on an open set Ω is differentiable almost everywhere in Ω. This applies to convex functions, which are locally Lipschitzian (Theorem 3.1.2); we state without proof the corresponding result:

Theorem 4.2.3 *Let $f \in \text{Conv}\,\mathbb{R}^n$. The subset of $\text{int dom}\, f$ where f fails to be differentiable is of zero (Lebesgue) measure.*   □

4.3 Second-Order Differentiation

The most useful criterion to recognize a convex function uses second derivatives. For this matter, the best idea is to reduce the question to the one-dimensional case: a function is convex if and only if its restrictions to the segments $[x, x']$ are also convex. These segments can in turn be parametrized via an origin x and a direction d: convexity of f amounts to the convexity of $t \mapsto f(x + td)$. Then, it suffices to apply calculus rules to compute the second derivative of this last function. Our first result mimics Theorem 4.1.4.

Theorem 4.3.1 *Let f be twice differentiable on an open convex set $\Omega \subset \mathbb{R}^n$. Then*

(i) *f is convex on Ω if and only if $\nabla^2 f(x_0)$ is positive semi-definite for all $x_0 \in \Omega$;*

(ii) *if $\nabla^2 f(x_0)$ is positive definite for all $x_0 \in \Omega$, then f is strictly convex on Ω;*

(iii) *f is strongly convex with modulus c on Ω if and only if the smallest eigenvalue of $\nabla^2 f(\cdot)$ is minorized by c on Ω: for all $x_0 \in \Omega$ and all $d \in \mathbb{R}^n$,*

$$\langle \nabla^2 f(x_0)d, d \rangle \geqslant c\|d\|^2.$$

Proof. For given $x_0 \in \Omega$, $d \in \mathbb{R}^n$ and $t \in \mathbb{R}$ such that $x_0 + td \in \Omega$, we will set

$$x_t := x_0 + td \quad \text{and} \quad \varphi(t) := f(x_t) = f(x + td),$$

so that $\varphi''(t) = \langle \nabla^2 f(x_t)d, d \rangle$.

[*(i)*] Assume f is convex on Ω; let (x_0, d) be arbitrary in $\Omega \times \mathbb{R}^n$, with $d \neq 0$: φ is then convex on the open interval $I := \{t \in \mathbb{R} : x_0 + td \in \Omega\}$. It follows

$$0 \leqslant \varphi''(t) = \langle \nabla^2 f(x_t)d, d \rangle \quad \text{for all } t \in I \ni 0$$

and $\nabla^2 f(x_0)$ is positive semi-definite.

Conversely, take an arbitrary $[x_0, x_1] \subset \Omega$, set $d := x_1 - x_0$ and assume $\nabla^2 f(x_t)$ positive semi-definite, i.e. $\varphi''(t) \geqslant 0$, for $t \in [0, 1]$. Then Theorem 0.6.6 tells us that φ is convex on [0,1], i.e. f is convex on $[x_0, x_1]$. The result follows since x_0 and x_1 were arbitrary in Ω.

[(ii)] To establish the strict convexity of f on Ω, we prove that ∇f is strictly mono-tone on Ω: Theorem 4.1.4 will apply. As above, take an arbitrary $[x_0, x_1] \subset \Omega$, $x_1 \neq x_0$, $d := x_1 - x_0$, and apply the mean-value theorem to the function φ', differentiable on [0,1]: for some $\tau \in {]0, 1[}$,

$$\varphi'(1) - \varphi'(0) = \varphi''(\tau) = \langle \nabla^2 f(x_\tau)d, d \rangle > 0$$

and the result follows since

$$\varphi'(1) - \varphi'(0) = \langle \nabla f(x_1) - \nabla f(x_0), x_1 - x_0 \rangle .$$

[(iii)] Using Proposition 1.1.2, apply (i) to the function $f - 1/2\, c\| \cdot \|^2$, whose Hessian operator is $\nabla^2 f - cI_n$ and has the eigenvalues $\lambda - c$, with λ describing the eigenvalues of $\nabla^2 f$. $\qquad\square$

Some differences have appeared with respect to §4.1:

– The sufficiency condition in (ii) is not necessary, even for univariate functions: think of $f(x) = 1/4\, x^4$.
– Theorem 4.1.1 stated that the affine (first-order) approximation of f around x_0 was actually a global minorization – more or less "comfortable". Here, we cannot say that the quadratic second-order approximation (of f around x_0)

$$x \mapsto f(x_0) + \langle \nabla f(x_0), x - x_0 \rangle + \tfrac{1}{2}\langle \nabla^2 f(x_0)(x - x_0), x - x_0 \rangle$$

minorizes f: think of $f(x) = 1/2\, x^2 - 1/4\, x^4$, which is convex for $|x|^2 \leqslant 1/3$.
– The present statements do not characterize convexity on a convex subset $C \subset \Omega$: C must be *open*. The reason is that §4.1 was dealing with the image (through f or ∇f) of *pairs* of points in C (x_0 and x, or x and x'). Here, $\nabla^2 f$ looks at f in the neighborhood of a *single point*, say x_0. Thus, a statement like "f is convex on $C \subset \Omega$ if and only if $\nabla^2 f(\cdot)$ is positive semi-definite on C" may be wrong when C is not open: $f(\xi, \eta) := \xi^2 - \eta^2$ is convex on $C = \mathbb{R} \times \{0\}$ but its Hessian is nowhere positive semi-definite.

Remark 4.3.2 Despite the last comment above, the convexity criterion using sec-ond derivatives is still the most powerful, even if positive (semi-)definiteness is not always easy to check. To recognize a convex function on a non-open set C, the best chance is therefore to use the Hessian operator on $\Omega = \operatorname{int} C$, hope-fully nonempty, and then to try and conclude by passing to the limit: the property $C \subset \operatorname{cl}(\operatorname{int} C) = \operatorname{cl} C$ is useful for that. $\qquad\square$

Example 4.3.3 To illustrate Theorem 4.3.1, consider the function

$$\Omega := \left\{ x = (\xi^1, \ldots, \xi^n) \ : \ \xi^i > 0 \text{ for } i = 1, \ldots, n \right\},$$

$$f : \Omega \ni x \mapsto f(x) := -(\xi^1 \xi^2 \cdots \xi^n)^{1/n}.$$

Direct computations give its second derivatives, i.e. its Hessian operator associated with the dot-product of $\mathbb{R}^n$:

$$\frac{\partial^2 f}{\partial \xi^i \partial \xi^j}(x) = \frac{f(x)}{n^2 \xi^i \xi^j}(1 - n\delta_{ij})$$

where δ_{ij} is Kronecker's symbol. We obtain, with $d = (d^1, \ldots, d^n) \in \mathbb{R}^n$:

$$[\nabla^2 f(x) d]^\top d = \frac{f(x)}{n^2} \left[\left(\sum_{i=1}^n \frac{d^i}{\xi^i} \right)^2 - n \sum_{i=1}^n \left(\frac{d^i}{\xi^i} \right)^2 \right].$$

The ℓ_1- and ℓ_2-norms are related on $\mathbb{R}^n$ by the inequality $\| \cdot \|_1 \leqslant \sqrt{n} \| \cdot \|_2$ (take a vector of the type $(\pm 1, \ldots, \pm 1)$ and use the Cauchy-Schwarz inequality); because f is negative on Ω, the above expression is therefore nonnegative: f is convex. Observe in passing that we obtain an equality if d and x are collinear: our f is positively homogeneous.

Observe here that f can be extended to cl Ω by posing $f(x) = 0$ if some ξ^i is zero. Convexity is preserved, and this illustrates Remark 4.3.2.   □

Remark 4.3.4 (Flat Domains) In all this Section 4, dom f was implicitly assumed full-dimensional, in order to have a nonempty int dom f. When such is not the case, some kind of differentiation can still be performed. In fact, exploit Remark 2.1.6 and make a change of variable, introducing the function $y \mapsto f_0(y) := f(x_0 + y)$; here x_0 is fixed in dom f, y varies in the subspace V parallel to aff dom f. Now, $f_0 \in \text{Conv } V$ and dom f_0 is full-dimensional in V. Equipping V with the induced scalar product $\langle \cdot, \cdot \rangle$ and the induced Lebesgue measure, the main results above can be reproduced. More precisely: almost everywhere in int dom f_0, i.e. for almost all $x_0 + y \in \text{ri dom } f$, there is a vector $s \in V$ (the gradient of f_0 at y) which gives the first-order approximation of f around x_0

$$\forall h \in V, \quad f(x_0 + y + h) = f(x_0 + y) + \langle s, h \rangle + o(\|h\|)$$

(the remainder term $o(\|h\|)$ being nonnegative); this s could be called the "relative gradient" of f at $x := x_0 + y$; it exists at $x_0 + y$ if and only if the function $t \mapsto f(x_0 + y + td)$ has a derivative at $t = 0$ for all $d \in V$.   □

Exercises

1. Consider the function $\mathbb{R} \ni x \mapsto f(x) := \frac{1}{1+x^2}$. Determine the largest intervals on which f is convex.

2. With $g \in \overline{\text{Conv}} \, \mathbb{R}$, define the function $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ by $f(x) := g(\|x\|)$. Show that $f \in \overline{\text{Conv}} \, \mathbb{R}^n$.

Show that the function $\mathbb{R}^2 \ni x \mapsto f(x) := \frac{1}{1+\|x\|^2}$ is "nowhere convex", i.e. on no set $C \subset \mathbb{R}^2$ (study the eigenvalues of its Hessian).

3. Let $f : [0, +\infty[\to \mathbb{R}$ be continuous, twice differentiable on $]0, +\infty[$ and satisfy $f(0) = 0$, $f'' > 0$ on $]0, +\infty[$.
– Show that $f'(x) > \frac{f(x)}{x}$ for all $x \in]0, +\infty[$.
– Deduce that the function $]0, +\infty[\ni x \mapsto F(x) := \int_1^x \frac{f(t)}{t} dt$ is stricly convex on $]0, +\infty[$.

4. Let I be an open interval of $\mathbb{R}$, let $f : I \to \mathbb{R}$ be three times continuously differentiable on I, with $f'' > 0$ on I. Consider the function φ defined by

$$I \times I \ni (x, y) \mapsto \varphi(x, y) := f(x) - f(y) - f'(y)(x - y) .$$

We use the notation $\theta := f''$.
– Prove the equivalence of the following statements:
  (i) φ is convex (jointly in x and y);
  (ii) $\theta(y) + \theta'(y)(y - x) \geqslant \dfrac{\theta^2(y)}{\theta(x)}$ for all $(x, y) \in I \times I$;
  (iii) $\dfrac{1}{\theta}$ is concave on I.
– Check that φ is convex in the following cases:
  · $I = \mathbb{R}$ and $f(x) = ax^2 + bx + c$ ($a > 0$);
  · $I =]0, +\infty[$ and $f(x) = x \log x - x$;
  · $I =]0, 1[$ and $f(x) = x \log x + (1 - x) \log(1 - x)$.

5. Prove the strict convexity of the function $x \mapsto f(x) := \log\left(\frac{1}{1 - \|x\|^2}\right)$ on the set $\{x \in \mathbb{R}^n : \|x\| < 1\}$.

6. Denote by $\mathcal{O} := \{x = (\xi_1, \ldots, \xi_n) : \xi_i > 0 \text{ for } i = 1, \ldots, n\}$ the positive orthant of $\mathbb{R}^n$. A function $f : \mathcal{O} \to \mathbb{R}$ is said to be of $\mathcal{G}$-type when it has the form

$$f(\xi_1, \ldots, \xi_n) = \sum_{k=1}^p c_k (\xi_1)^{a_{1k}} \cdots (\xi_n)^{a_{nk}} ;$$

here the c_k's are nonnegative coefficients and the a_{ik}'s are real numbers (possibly negative). Show that, if f is of $\mathcal{G}$-type, then

$$\mathbb{R}^n \ni (\eta_1, \ldots, \eta_n) = y \mapsto g(y) := f\left(e^{\eta_1}, \ldots, e^{\eta_n}\right)$$

is convex.

7. Let Ω be a convex open set in $\mathbb{R}^n$. The function $f : \Omega \to]0, +\infty[$ is said to be *logarithmically convex* when $\log f$ is convex on Ω.
   Assuming that f is twice differentiable on Ω, show that the following statements are equivalent:
  (i) f is logarithmically convex;
  (ii) $f(x) \nabla^2 f(x) - \nabla f(x) [\nabla f(x)]^\top \succcurlyeq 0$ for all $x \in \Omega$;
  (iii) the function $x \mapsto e^{\langle a, x \rangle} f(x)$ is convex for all $a \in \mathbb{R}^n$.

8*. Suppose that the convex functions f and g are such that $f + g$ is affine. Show that f and g are actually affine.

9. For a nonempty closed convex set $C \subset \mathbb{R}^n \times \mathbb{R}$, show that

$$[C = \text{epi } f \text{ with } f \in \overline{\text{Conv}} \, \mathbb{R}^n] \quad \Longleftrightarrow \quad [(0,1) \in C_\infty \text{ and } (0,-1) \notin C_\infty].$$

10. Show that the function log is concave on $]0, +\infty[$ and prove the inequality between arithmetic and geometric means:

$$\prod_{i=1}^{m} (x_i)^{1/m} \leqslant \frac{1}{m} \sum_{i=1}^{m} x_i \quad \text{for } x_1, \ldots, x_m > 0.$$

11. For $f \in \text{Conv } \mathbb{R}^n$, prove the equality $\inf_{x \in \mathbb{R}^n} f(x) = \inf_{x \in \text{ri dom } f} f(x)$.

12. Here f and c are two finite-valued convex functions and $C := \{x : c(x) < 0\}$ is assumed nonempty. Show that C is convex. What are its boundary and its closure?

Show that the function $x \mapsto f(x) - \log(-c(x))$ is convex on C; what is its domain? Same questions with the function $x \mapsto f(x) + [\max\{0, c(x)\}]^2$.

13. Let $\mathbb{R}^2 \ni (\xi, \eta) \mapsto f(\xi, \eta) := \xi^2(1 + e^\eta)$. Compute $\nabla^2 f$. Is f convex, strictly convex, strongly convex on the set $\{(\xi, \eta) \in \mathbb{R}^2 : \eta > 0\}$?

14*. Let C be a nonempty closed convex set, with nonempty complement C^c. Show that $f := -d_{C^c} + i_C$ is closed convex.

Show that f remains convex if $f(x)$ is replaced by $d_C(x)$ for $x \notin C$; i.e. f becomes $-d_{C^c} + d_C$.

15*. For $\varphi :]0, +\infty[\to \mathbb{R}$ convex, define the function $S_\varphi :]0, +\infty[^n \times]0, +\infty[^n$ by

$$(p^1, \ldots, p^n; q^1, \ldots, q^n) = (p; q) \mapsto S_\varphi(p, q) := \sum_{i=1}^{n} q^i \varphi\left(\frac{p^i}{q^i}\right).$$

Check that S_φ is convex. Setting $P := \sum_{i=1}^n p^i$ and $Q := \sum_{i=1}^n q^i$, deduce the inequality

$$S_\varphi(p, q) \geqslant Q\varphi\left(\frac{P}{Q}\right).$$

Defining $\psi :]0, +\infty[\to \mathbb{R}$ by $\psi(x) := x\varphi\left(\frac{1}{x}\right)$,

– How do S_φ and S_ψ compare?

– Compute ψ and S_φ when $\varphi(t)$ has the following expressions: $t \log t$, $\left(1 - \sqrt{t}\right)^2$, t^α with $\alpha > 1$, $(t - 1)^2$, $|t - 1|$.

16*. Let $f :]-1, +1[\to]-\infty, 0[$ be convex. Prove the convexity of the function

$$]-1, +1[\times \mathbb{R} \ni (\xi, \eta) \mapsto g(\xi, \eta) := -\frac{\eta^2}{f(\xi)}.$$

Let $f :]a, b[\to]-\infty, 0[$ be convex. Prove the convexity of the function

$$]0, +\infty[\times]a, b[\ni (\xi, \eta) \mapsto h(\xi, \eta) := \xi \log\left(-\frac{\xi}{f(\eta)}\right).$$

17. Show that the following function lies in $\overline{\mathrm{Conv}}\,\mathrm{S}(\mathbb{R}^n)$:

$$\mathrm{S}(\mathbb{R}^n) \ni M \mapsto f(M) := \begin{cases} \dfrac{-1}{\mathrm{tr}\,(M^{-1})} & \text{if } M \succ 0, \\ +\infty & \text{otherwise}. \end{cases}$$

*18**.* Let $A \in \mathrm{S}_n(\mathbb{R})$ be positive definite. Prove the convexity of the function

$$\mathbb{R}^n \ni x \mapsto f(x) := \langle Ax, x \rangle \langle A^{-1}x, x \rangle.$$

19.* Let $f := \min\{f_1, \ldots, f_m\}$, where $f_1, \ldots, f_m$ are convex finite-valued. Find two convex functions g and h such that $f = g - h$.

20.* Let $f : [0, +\infty[\to [0, +\infty[$ be convex with $f(0) = 0$. Its mean-value on $[0, x]$ is defined by

$$F(0) := 0, \quad F(x) := \frac{1}{x} \int_0^x f(t)\mathrm{d}t \ \text{ for } x > 0.$$

Show that F is convex.

21.* Let $C \neq \emptyset$ be closed convex in $\mathbb{R}^n$. Show that $(\mathrm{d}_C)'_\infty = \mathrm{d}_{(C_\infty)}$.

22. Let $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$. Show that $f'_\infty(d) = \liminf\left\{tf\left(\frac{d'}{t}\right) : d' \to d, t \downarrow 0\right\}$ for all $d \in \mathbb{R}^n$.

23. Show that $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$ is 0-oercive if and only if there is $\varepsilon > 0$ such that $f'_\infty(d) \geqslant \varepsilon$ for all d of norm 1.

24. Consider the polyhedral function $x \mapsto f(x) := \max_{i=1,\ldots,k} (\langle s_i, x \rangle + r_i)$. Show that $f'_\infty(d) = \max_{i=1,\ldots,k} \langle s_i, d \rangle$.

25.* Let $f : \mathbb{R} \to \mathbb{R}$ be continuous. For $h > 0$ define

$$\mathbb{R} \ni x \mapsto f_h(x) := \frac{1}{2h} \int_{x-h}^{x+h} f(t)\mathrm{d}t.$$

Show that f is convex if and only if $f_h(x) \geqslant f(x)$ for all $x \in \mathbb{R}$ and all $h > 0$.
    Give a geometric interpretation of this result.

*26**.* Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and differentiable over $\mathbb{R}^n$. Prove that the following three properties are equivalent (L being a positive constant):
   (i) $\|\nabla f(x) - \nabla f(x')\| \leqslant L\|x - x'\|$ for all x, x' in $\mathbb{R}^n$;
   (ii) $f(x') \geqslant f(x) + \langle \nabla f(x), x' - x \rangle + \frac{1}{2L}\|\nabla f(x) - \nabla f(x')\|^2$ for all x, x' in $\mathbb{R}^n$;
   (iii) $\langle \nabla f(x) - \nabla f(x'), x - x' \rangle \geqslant \frac{1}{L}\|\nabla f(x) - \nabla f(x')\|^2$ for all x, x' in $\mathbb{R}^n$.

C. Sublinearity and Support Functions

Introduction In classical real analysis, the simplest functions are *linear*. In convex analysis, the next simplest convex functions (apart from the affine functions, widely used in §B.1.2), are so-called *sublinear*. We give three motivations for their study.

(i) A suitable generalization of linearity. A linear function ℓ from $\mathbb{R}^n$ to $\mathbb{R}$, or a linear form on $\mathbb{R}^n$, is primarily defined as a function satisfying for all $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$ and $(t_1, t_2) \in \mathbb{R} \times \mathbb{R}$:

$$\ell(t_1 x_1 + t_2 x_2) = t_1 \ell(x_1) + t_2 \ell(x_2). \tag{0.1}$$

A corresponding definition for a sublinear function σ from $\mathbb{R}^n$ into $\mathbb{R}$ is: for all $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$ and $(t_1, t_2) \in \mathbb{R}^+ \times \mathbb{R}^+$,

$$\sigma(t_1 x_1 + t_2 x_2) \leqslant t_1 \sigma(x_1) + t_2 \sigma(x_2). \tag{0.2}$$

A first observation is that requiring an inequality in (0.2), rather than an equality, allows infinite values for σ without destroying the essence of the concept of sublinearity. Of course, (0.2) is less stringent than (0.1), but more stringent than the definition of a convex function: the inequality must hold in (0.2) even if $t_1 + t_2 \neq 1$. This confirms that sublinear functions, which generalize linear functions, are particular instances of convex functions.

Remark 0.1 Note that (0.1) and (0.2) can be made more similar by restricting t_1 and t_2 in (0.1) to be positive – this leaves unchanged the definition of a linear function.

The prefix "sub" comes from the inequality-sign "$\leqslant$" in (0.2). It also suggests that sublinearity is less demanding than linearity, but this is a big piece of luck. In fact, draw the graph of a convex and of a concave function and ask a non-mathematician: "which is convex?". He will probably give the wrong answer. Yet, if convex functions were defined the other way round, (0.2) should have the "$\geqslant$" sign. The associated concept would be superlinearity, an unfortunate wording which suggests *more* or *better* than linear. □

In a word, sublinear functions are reasonable candidates for "simplest non-trivial convex functions". Whether they are interesting candidates will be seen in (ii) and (iii). Here, let us just mention that their epigraphs are convex cones, the next simplest convex epigraphs after half-spaces.

(ii) Tangential approximation of convex functions. To say that a function $f : \mathbb{R}^n \to \mathbb{R}$ is differentiable at x is to say that there is a linear function ℓ_x which approximates $f(x + h) - f(x)$ to first order, i.e.

$$f(x+h) - f(x) = \ell_x(h) + o(\|h\|) \,.$$

This fixes the rate of change of f when x is moved along a line d: with $\varepsilon(t) \to 0$ if $t \to 0$,

$$\frac{f(x+td) - f(x)}{t} = \ell_x(d) + \varepsilon(t) \quad \text{for all } t \neq 0 \,.$$

Geometrically, the graph of f has a tangent hyperplane at $(x, f(x)) \in \mathbb{R}^n \times \mathbb{R}$; and this hyperplane is the graph of the affine function $h \mapsto f(x) + \ell_x(h)$.

When f is merely convex, its graph may have no tangent hyperplane at a given $(x, f(x))$. Nevertheless, under reasonable assumptions, $f(x+h) - f(x)$ can still be approximated to first order by a function which is sublinear: there exists a sublinear function $h \mapsto \sigma_x(h)$ such that

$$f(x+h) - f(x) = \sigma_x(h) + o(\|h\|) \,.$$

This will be seen in Chap. D.

Geometrically, $\operatorname{gr} \sigma_x$ is no longer a hyperplane but rather a cone, which is therefore tangent to $\operatorname{gr} f$ (the word "tangent" should be understood here in its intuitive meaning of a tangent surface, as opposed to tangent cones of Chap. A; neither $\operatorname{gr} \sigma_x$ nor $\operatorname{gr} f$ are convex). Thus, one can say that differentiable functions are "tangentially linear", while convex functions are "tangentially sublinear". See Fig. 0.1, which displays the graphs of a differentiable and of a convex function. The graph of ℓ_x is the thick line L, while the graph of σ_x is made up of the two thick half-lines S_1 and S_2.
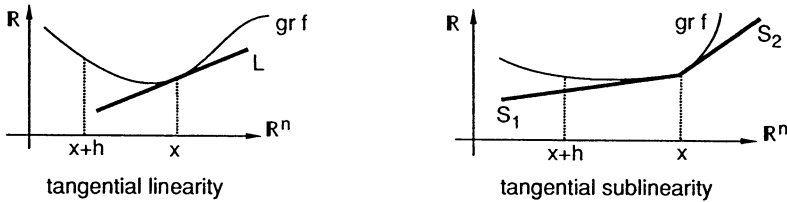
Fig. 0.1. Two concepts of tangency

(iii) Nice correspondence with nonempty closed convex sets. In the Euclidean space $(\mathbb{R}^n, \langle \cdot, \cdot \rangle)$, a linear form ℓ can be represented by a vector: there is a unique $s \in \mathbb{R}^n$ such that

$$\ell(x) = \langle s, x \rangle \quad \text{for all } x \in \mathbb{R}^n \,. \tag{0.3}$$

The definition (0.3) of a linear function is more geometric than (0.1), and just as accurate. A large part of the present chapter will be devoted to generalizing the above representation theorem to sublinear functions.

First observe that, given a nonempty set $S \subset \mathbb{R}^n$, the function $\sigma_S : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ defined by

$$\sigma_S(x) := \sup\{\langle s, x \rangle \ : \ s \in S\} \tag{0.4}$$

is sublinear. It is called the *support function* of S, already encountered in Sects A.4.1 and B.1.3(a). When S is bounded, its support function is finite everywhere; otherwise, σ_S can take on the value $+\infty$ but it remains lower semi-continuous. Furthermore, it is easy to check that σ_S is also the support function of the closure of S, and even of the closed convex hull of S. It is therefore logical to consider support functions of nonempty closed convex sets only.

Now, a key result is that the mapping $S \mapsto \sigma_S$ is then bijective: a lower semi-continuous (i.e. closed) sublinear function is the support function of a *uniquely determined* nonempty closed convex set. Thus, (0.4) establishes the announced representation, just as (0.3) does in the linear case. Note that the linear case is covered: it corresponds to S being a singleton $\{s\}$ in (0.4).

This correspondence between nonempty closed convex sets of $\mathbb{R}^n$ and closed sublinear functions allows fruitful and enlightening geometric interpretations when studying these functions. Vice versa, it provides powerful analytical tools for the study of these sets. In particular, when closed convex sets are combined (intersected, added, etc.) to form new convex sets, we will show how their support functions are correspondingly combined: the mapping (0.4) is an *isomorphism*, with respect to a number of structures.

1 Sublinear Functions

1.1 Definitions and First Properties

Definition 1.1.1 A function $\sigma : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ is said to be *sublinear* if it is convex and positively homogeneous (of degree 1): $\sigma \in \operatorname{Conv} \mathbb{R}^n$ and

$$\sigma(tx) = t\sigma(x) \quad \text{for all } x \in \mathbb{R}^n \text{ and } t > 0. \tag{1.1.1}$$

$\square$

Remark 1.1.2 Inequality in (1.1.1) would be enough to define positive homogeneity: a function σ is positively homogeneous if and only if it satisfies

$$\sigma(tx) \leqslant t\sigma(x) \quad \text{for all } x \in \mathbb{R}^n \text{ and } t > 0. \tag{1.1.2}$$

In fact, (1.1.2) implies ($tx \in \mathbb{R}^n$ and $t^{-1} > 0$!)

$$\sigma(x) = \sigma(t^{-1}tx) \leqslant t^{-1}\sigma(tx)$$

which, together with (1.1.2), shows that σ is positively homogeneous. $\square$

We deduce from (1.1.1) that $\sigma(0) = t\sigma(0)$ for all $t > 0$. This leaves only two possible values for $\sigma(0)$: 0 and $+\infty$. However, most of the sublinear functions to be encountered in the sequel do satisfy $\sigma(0) = 0$. According to our Definition B.1.1.3 of convex functions, σ should be finite somewhere; otherwise $\operatorname{dom}\sigma$ would be empty. Now, if $\sigma(x) < +\infty$, (1.1.1) shows that $\sigma(tx) < +\infty$ for all $t > 0$. In other words, $\operatorname{dom}\sigma$ is a cone, convex because σ is itself convex. Note that, being

convex, σ is continuous relatively to ri dom σ, but discontinuities may occur on the boundary-rays of dom σ, including at 0.

The following result is a geometric characterization of sublinear functions.

Proposition 1.1.3 *A function* $\sigma : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ *is sublinear if and only if its epigraph* epi σ *is a nonempty convex cone in* $\mathbb{R}^n \times \mathbb{R}$.

Proof. We know that σ is a convex function if and only if epi σ is a nonempty convex set in $\mathbb{R}^n \times \mathbb{R}$ (Proposition B.1.1.6). Therefore, we just have to prove the equivalence between positive homogeneity and epi σ being a cone.

Let σ be positively homogeneous. For $(x, r) \in$ epi σ, the relation $\sigma(x) \leqslant r$ gives

$$\sigma(tx) = t\sigma(x) \leqslant tr \quad \text{for all } t > 0,$$

so epi σ is a cone. Conversely, if epi σ is a cone in $\mathbb{R}^n \times \mathbb{R}$, the property $(x, \sigma(x)) \in$ epi σ implies $(tx, t\sigma(x)) \in$ epi σ, i.e.

$$\sigma(tx) \leqslant t\sigma(x) \quad \text{for all } t > 0.$$

From Remark 1.1.2, this is just positive homogeneity. $\square$

Another important concept in analysis is *subadditivity*: a function σ is subadditive when it satisfies

$$\sigma(x_1 + x_2) \leqslant \sigma(x_1) + \sigma(x_2) \quad \text{for all } (x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n \tag{1.1.3}$$

– watch the difference with (0.2). Here again, the inequality is understood in $\mathbb{R} \cup \{+\infty\}$. Together with positive homogeneity, the above axiom gives another characterization (analytical, rather than geometrical) of sublinear functions.

Proposition 1.1.4 *A function* $\sigma : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$, *not identically equal to* $+\infty$, *is sublinear if and only if one of the following two properties holds:*

$$\sigma(t_1 x_1 + t_2 x_2) \leqslant t_1 \sigma(x_1) + t_2 \sigma(x_2) \quad \text{for all } x_1, x_2 \in \mathbb{R}^n \text{ and } t_1, t_2 > 0, \tag{1.1.4}$$

or
$$\sigma \text{ is positively homogeneous and subadditive}. \tag{1.1.5}$$

Proof. [*sublinearity* $\Rightarrow$ *(1.1.4)*] For $x_1, x_2 \in \mathbb{R}^n$ and $t_1, t_2 > 0$, set $t := t_1 + t_2 > 0$; we have

$$\begin{aligned}
\sigma(t_1 x_1 + t_2 x_2) &= \sigma\left(t\left[\tfrac{t_1}{t} x_1 + \tfrac{t_2}{t} x_2\right]\right) \\
&= t\sigma\left(\tfrac{t_1}{t} x_1 + \tfrac{t_2}{t} x_2\right) && \text{[positive homogeneity]} \\
&\leqslant t\left[\tfrac{t_1}{t}\sigma(x_1) + \tfrac{t_2}{t}\sigma(x_2)\right], && \text{[convexity]}
\end{aligned}$$

and (1.1.4) is proved.

[*(1.1.4)* $\Rightarrow$ *(1.1.5)*] A function satisfying (1.1.4) is obviously subadditive (take $t_1 = t_2 = 1$) and satisfies (take $x_1 = x_2 = x$, $t_1 = t_2 = 1/2\,t$)

$$\sigma(tx) \leqslant t\sigma(x) \quad \text{for all } x \in \mathbb{R}^n \text{ and } t > 0\,,$$

which is just positive homogeneity because of Remark 1.1.2.

[(1.1.5) $\Rightarrow$ sublinearity] Take $t_1, t_2 > 0$ with $t_1 + t_2 = 1$ and apply successively subadditivity and positive homogeneity:

$$\sigma(t_1 x_1 + t_2 x_2) \leqslant \sigma(t_1 x_1) + \sigma(t_2 x_2) = t_1 \sigma(x_1) + t_2 \sigma(x_2)\,,$$

hence σ is convex. $\qquad\qquad\square$

Corollary 1.1.5 *If σ is sublinear, then*

$$\sigma(x) + \sigma(-x) \geqslant 0 \quad \text{for all } x \in \mathbb{R}^n\,. \qquad\qquad (1.1.6)$$

Proof. Take $x_2 = -x_1$ in (1.1.3) and remember that $\sigma(0) \geqslant 0$. $\qquad\square$

It is worth mentioning that, to become sublinear, a positively homogeneous function just needs to be subadditive as well (rather than convex, as suggested by Definition 1.1.1); then, of course, it becomes convex at the same time. Figure 1.1.1 summarizes the connections between the classes of functions given so far. Note for completeness that a convex and subadditive function need not be sublinear: think of $f(x) \equiv 1$.
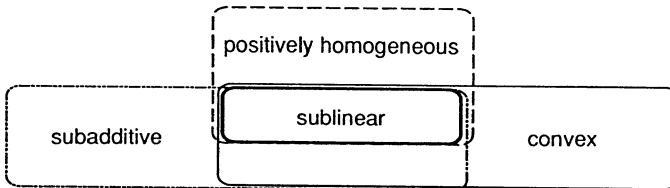
Fig. 1.1.1. Various classes of functions

Similarly, one can ask when a sublinear function becomes linear. For a linear function, (1.1.6) holds as an equality, and the next result implies that this is exactly what makes the difference.

Proposition 1.1.6 *Let σ be sublinear and suppose that there exist $x_1, \ldots, x_m$ in* dom σ *such that*

$$\sigma(x_j) + \sigma(-x_j) = 0 \quad \text{for } j = 1, \ldots, m\,. \qquad\qquad (1.1.7)$$

Then σ is linear on the subspace spanned by $x_1, \ldots, x_m$.

Proof. With $x_1, \ldots, x_m$ as stated, each $-x_j$ is in dom σ. Let $x := \sum_{j=1}^m t_j x_j$ be an arbitrary linear combination of $x_1, \ldots, x_m$; we must prove that $\sigma(x) = \sum_{j=1}^m t_j \sigma(x_j)$. Set

$$J_1 := \{j : t_j > 0\}, \quad J_2 := \{j : t_j < 0\}$$

and obtain (as usual, $\sum_\emptyset = 0$):

$$
\begin{aligned}
\sigma(x) &= \sigma\big(\textstyle\sum_{J_1} t_j x_j + \sum_{J_2} (-t_j)(-x_j)\big) \\
&\leqslant \textstyle\sum_{J_1} t_j \sigma(x_j) + \sum_{J_2} (-t_j)\sigma(-x_j) && \text{[from (1.1.4)]} \\
&= \textstyle\sum_{J_1} t_j \sigma(x_j) + \sum_{J_2} t_j \sigma(x_j) = \sum_{j=1}^m t_j \sigma(x_j) && \text{[from (1.1.7)]} \\
&= -\textstyle\sum_{J_1} t_j \sigma(-x_j) - \sum_{J_2}(-t_j)\sigma(x_j) && \text{[from (1.1.7)]} \\
&\leqslant -\sigma\big(-\textstyle\sum_{j=1}^m t_j x_j\big) && \text{[from (1.1.4)]} \\
&= -\sigma(-x) \leqslant \sigma(x). && \text{[from (1.1.6)]}
\end{aligned}
$$

In summary, we have proved $\sigma(x) \leqslant \sum_{j=1}^m t_j \sigma(x_j) \leqslant -\sigma(-x) \leqslant \sigma(x)$. $\qquad\square$

Thanks to this result, we are entitled to define

$$U := \{x \in \mathbb{R}^n \ : \ \sigma(x) + \sigma(-x) = 0\} \tag{1.1.8}$$

which is a subspace of $\mathbb{R}^n$: the subspace in which σ is *linear*, sometimes called the *lineality space* of σ. Note that U nonempty corresponds to $\sigma(0) = 0$ (even if U reduces to $\{0\}$).

What is interesting in this concept is its geometric interpretation. If V is another subspace such that $U \cap V = \{0\}$, there holds by definition $\sigma(x) + \sigma(-x) > 0$ for all $0 \neq x \in V$. This means that, if $0 \neq d \in V$, σ is "V-shaped" along d: for $t > 0$, $\sigma(td) = \alpha t$ and $\sigma(-td) = \beta t$, for some α and β in $\mathbb{R} \cup \{+\infty\}$ such that $\alpha + \beta > 0$; whereas $\alpha + \beta$ would be 0 if d were in U. See Fig. 1.1.2 for an illustration. For d of norm 1, the number $\alpha + \beta$ above could be called the "lack of linearity" of σ along d: when restricted to the line d, the graph of σ makes an angle; when finite, the number $\alpha + \beta$ measures how acute this angle is.
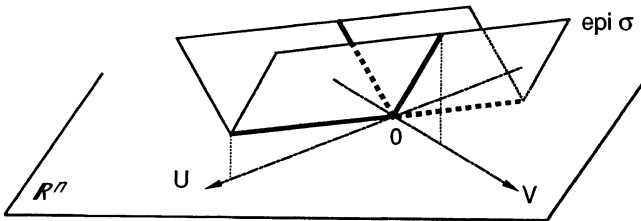
Fig. 1.1.2. Subspace of linearity of a sublinear function

Figure 1.1.2 suggests that $\operatorname{gr} \sigma$ is a hyperplane not only in U, but also in the translations of U: the restriction of σ to $\{y\} + U$ is affine, for any fixed y. This comes from the next result.

Proposition 1.1.7 *Let σ be sublinear. If $x \in U$, i.e. if*

$$\sigma(x) + \sigma(-x) = 0, \tag{1.1.9}$$

then there holds

$$\sigma(x + y) = \sigma(x) + \sigma(y) \quad \text{for all } y \in \mathbb{R}^n . \tag{1.1.10}$$

Proof. In view of subadditivity, we just have to prove "$\geqslant$" in (1.1.10). Start from the identity $y = x + y - x$; apply successively subadditivity and (1.1.9) to obtain

$$\sigma(y) \leqslant \sigma(x + y) + \sigma(-x) = \sigma(x + y) - \sigma(x).$$ □

1.2 Some Examples

We start with some simple situations. If K is a nonempty convex cone, its indicator function

$$i_K(x) := \begin{cases} 0 & \text{if } x \in K, \\ +\infty & \text{if not} \end{cases}$$

is clearly sublinear. In $\mathbb{R}^n \times \mathbb{R}$, the epigraph epi i_K is made up of all the copies of K, shifted upwards. Likewise, its distance function

$$d_K(x) := \inf \{\|y - x\| : y \in K\}$$

is also sublinear: nothing in the picture is essentially changed when both x and y are multiplied by $t > 0$. Another example is the function from $\mathbb{R}^2$ to $\mathbb{R} \cup \{+\infty\}$

$$\sigma(x) = \sigma(\xi, \eta) := \begin{cases} -2\sqrt{\xi\eta} & \text{if } \xi, \eta \geqslant 0 \\ +\infty & \text{if not.} \end{cases}$$

Its positive homogeneity is clear, its convexity is not particularly difficult to check (see Example B.4.3.3), it is therefore sublinear. A good exercise is to try to visualize its epigraph.

Example 1.2.1 Let $f \in \text{Conv } \mathbb{R}^n$; its perspective $\tilde{f}$ of §B.2.2, which is convex, is clearly positively homogeneous (from $\mathbb{R}^{n+1}$ to $\mathbb{R} \cup \{+\infty\}$); it is an important instance of sublinear function. For example, in $\mathbb{R}^2$

$$\tilde{f}(u, \xi) := \begin{cases} \frac{1}{2}\xi^2/u & \text{if } u > 0, \\ +\infty & \text{if not} \end{cases} \tag{1.2.1}$$

is the perspective of $\xi \mapsto f(\xi) = 1/2\,\xi^2$.

Note that $\tilde{f}(0, 0) = +\infty$. The closure of $\tilde{f}$ can be computed with the help of Example B.3.2.3: clearly enough, the asymptotic function of f is $i_{\{0\}}$. Hence $(\text{cl } \tilde{f})(0, 0) = 0$, while $\tilde{f}$ coincides with its closure everywhere else. □

Example 1.2.2 (Norms) We recall that a norm $\|\cdot\|$ on $\mathbb{R}^n$ is a function from $\mathbb{R}^n$ to $[0, +\infty[$ satisfying the following properties:

 (i) $\|x\| = 0$ if and only if $x = 0$;
 (ii) $\|tx\| = |t|\,\|x\|$ for all $x \in \mathbb{R}^n$ and $t \in \mathbb{R}$;
(iii) $\|x_1 + x_2\| \leqslant \|x_1\| + \|x_2\|$ for all $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$.

Clearly, $\|\cdot\|$ is a positive (except at 0) and finite sublinear function which, moreover, is symmetric i.e. $\|-x\| = \|x\|$ for all x. It is linear on no line: the subspace U of (1.1.8) is reduced to $\{0\}$.

Conversely, if σ is a sublinear function from $\mathbb{R}^n$ into $[0, +\infty[$ which is linear on no line, i.e. such that

$$\sigma(x) + \sigma(-x) > 0 \quad \text{for all } x \neq 0,$$

then $\|x\| := \max \{\sigma(x), \sigma(-x)\}$ is a norm on $\mathbb{R}^n$.   □

Example 1.2.3 (Quadratic Semi-Norms) Take a symmetric positive semi-definite operator Q from $\mathbb{R}^n$ to $\mathbb{R}^n$ and define

$$f(x) := \sqrt{\langle Qx, x \rangle} \quad \text{for all } x \in \mathbb{R}^n .$$

Convexity of f (i.e. its subadditivity, i.e. the Cauchy-Schwarz inequality) is rather tedious to prove directly. Consider, however, the convex set

$$E_Q := \{x \in \mathbb{R}^n : \langle Qx, x \rangle \leqslant 1\}.$$

Then f can be obtained as follows:

$$\begin{aligned}
f(x) &= \inf \{\lambda > 0 : \langle Qx, x \rangle \leqslant \lambda^2\} \\
&= \inf \{\lambda > 0 : \langle Q\tfrac{x}{\lambda}, \tfrac{x}{\lambda} \rangle \leqslant 1\} \\
&= \inf \{\lambda > 0 : \tfrac{x}{\lambda} \in E_Q\}
\end{aligned}$$

and we will see below that this establishes convexity – hence sublinearity – of f.

Observe in passing that E_Q is the sublevel-set at level 1 of both f and $f^2 = \langle Q\cdot, \cdot \rangle$. Decompose the space as $\mathbb{R}^n = \operatorname{Ker} Q \oplus \operatorname{Im} Q$: the intersection of E_Q with $\operatorname{Im} Q$ is an elliptic set centered at the origin, say $\widetilde{E}_Q$. The entire E_Q is the cylinder $\widetilde{E}_Q + \operatorname{Ker} Q$, whose asymptotic cone is just the subspace $\operatorname{Ker} Q$. If and only if $\operatorname{Ker} Q = \{0\}$, i.e. Q is positive definite, is E_Q compact; it is an elliptic body. On the other hand, f is finite, nonnegative, symmetric because E_Q has center 0; and f is zero on the asymptotic cone $\operatorname{Ker} Q$ of E_Q. Theorem 1.2.5 below establishes the convexity of f, which is therefore a semi-norm, actually a norm if Q is positive definite.   □

The mapping $E_Q \mapsto f$, introduced in Example 1.2.3, is important in the context of sublinear functions; let us put it in perspective.

Definition 1.2.4 (Gauge) Let C be a closed convex set containing the origin. The function γ_C defined by

$$\gamma_C(x) := \inf \{\lambda > 0 : x \in \lambda C\} \qquad (1.2.2)$$

is called the *gauge* of C. As usual, we set $\gamma_C(x) := +\infty$ if $x \in \lambda C$ for no $\lambda > 0$.   □

Geometrically, γ_C can be obtained as follows: shift C ($\subset \mathbb{R}^n$) in the hyperplane $\mathbb{R}^n \times \{1\}$ of the graph-space $\mathbb{R}^n \times \mathbb{R}$ (by contrast to a perspective-function, the
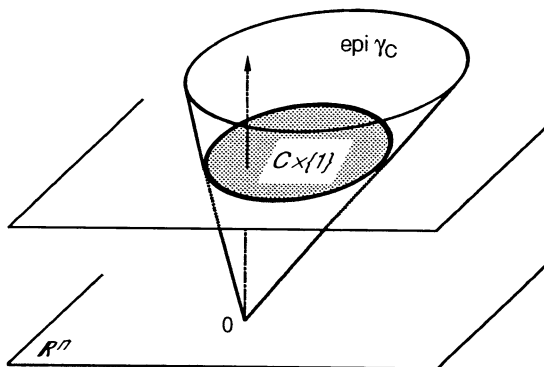
Fig. 1.2.1. The epigraph of a gauge

present shift is vertical, along the axis of function-values). Then the epigraph of γ_C is the cone generated by this shifted copy of C; see Fig. 1.2.1.

The next result summarizes the main properties of a gauge. Each statement should be read with Fig. 1.2.1 in mind, even though the picture is slightly misleading, due to closure problems.

Theorem 1.2.5 *Let C be a closed convex set containing the origin. Then*

(i) *its gauge γ_C is a nonnegative closed sublinear function;*

(ii) *γ_C is finite everywhere if and only if 0 lies in the interior of C;*

(iii) *C_∞ being the asymptotic cone of C,*

$$\{x \in \mathbb{R}^n \ : \ \gamma_C(x) \leqslant r\} = rC \quad \text{for all } r > 0,$$
$$\{x \in \mathbb{R}^n \ : \ \gamma_C(x) = 0\} = C_\infty .$$

Proof. [*(i) and (iii)*] Nonnegativity and positive homogeneity are obvious from the definition of γ_C; also, $\gamma_C(0) = 0$ because $0 \in C$. We prove convexity via a geometric interpretation of (1.2.2). Let

$$K_C := \text{cone}\,(C \times \{1\}) = \{(\lambda c, \lambda) \in \mathbb{R}^n \times \mathbb{R} \ : \ c \in C, \ \lambda \geqslant 0\}$$

be the convex conical hull of $C \times \{1\} \subset \mathbb{R}^n \times \mathbb{R}$. It is convex (beware that K_C need not be closed) and γ_C is clearly given by

$$\gamma_C(x) = \inf\,\{\lambda \ : \ (x, \lambda) \in K_C\} .$$

Thus, γ_C is the lower-bound function of §B.1.3(g), constructed on the convex set K_C; this establishes the convexity of γ_C, hence its sublinearity.

Now we prove

$$\{x \in \mathbb{R}^n \ : \ \gamma_C(x) \leqslant 1\} = C . \tag{1.2.3}$$

This will imply the first part in (iii), thanks to positive homogeneity. Then the second part will follow because of (A.2.2.2): $C_\infty = \cap\{rC : r > 0\}$ and closedness of γ_C will also result from (iii) via Proposition B.1.2.2.

So, to prove (1.2.3), observe first that $x \in C$ implies from (1.2.2) that certainly $\gamma_C(x) \leqslant 1$. Conversely, let x be such that $\gamma_C(x) \leqslant 1$; we must prove that $x \in C$. For this we prove that $x_k := (1 - 1/k)x \in C$ for $k = 1, 2, \ldots$ (and then, the desired property will come from the closedness of C). By positive homogeneity, $\gamma_C(x_k) = (1 - 1/k)\gamma_C(x) < 1$, so there is $\lambda_k \in \,]0, 1[$ such that $x_k \in \lambda_k C$, or equivalently $x_k/\lambda_k \in C$. Because C is convex and contains the origin, $\lambda_k(x_k/\lambda_k) + (1 - \lambda_k)0 = x_k$ is in C, which is what we want.

[(ii)] Assume $0 \in \operatorname{int} C$. There is $\varepsilon > 0$ such that for all $x \neq 0$, $x_\varepsilon := \varepsilon x/\|x\| \in C$; hence $\gamma_C(x_\varepsilon) \leqslant 1$ because of (1.2.3). We deduce by positive homogeneity

$$\gamma_C(x) = \frac{\|x\|}{\varepsilon}\gamma_C(x_\varepsilon) \leqslant \frac{\|x\|}{\varepsilon} \,;$$

this inequality actually holds for all $x \in \mathbb{R}^n$ ($\gamma_C(0) = 0$) and γ_C is a finite function.

Conversely, suppose γ_C is finite everywhere. By continuity (Theorem B.3.1.2), γ_C has an upper bound $L > 0$ on the unit ball:

$$\|x\| \leqslant 1 \quad \Longrightarrow \quad \gamma_C(x) \leqslant L \quad \Longrightarrow \quad x \in LC \,,$$

where the last implication comes from (iii). In other words, $B(0, 1/L) \subset C$.   □

Since γ_C is the lower-bound function of the cone K_C ($= K_C + \{0\} \times \mathbb{R}^+$) of Fig. 1.2.1, we know from (B.1.3.6) that $K_C \subset \operatorname{epi}\gamma_C \subset \operatorname{cl} K_C$; but γ_C has a closed epigraph, therefore

$$\operatorname{epi}\gamma_C = \operatorname{cl} K_C = \overline{\operatorname{cone}}(C \times \{1\}) \,. \tag{1.2.4}$$

Since $C_\infty = \{0\}$ if and only if C is compact (Proposition A.2.2.3), we obtain another consequence of (iii):

Corollary 1.2.6 *C is compact if and only if $\gamma_C(x) > 0$ for all $x \neq 0$.*   □

Example 1.2.7 The quadratic semi-norms of Example 1.2.3 can be generalized: let $f \in \overline{\operatorname{Conv}}\,\mathbb{R}^n$ have nonnegative values and be positively homogeneous of degree 2, i.e.

$$0 \leqslant f(tx) = t^2 f(x) \quad \text{for all } x \in \mathbb{R}^n \text{ and all } t > 0 \,.$$

Then, $\sqrt{f}$ is convex; in fact

$$\begin{aligned} \sqrt{f}(x) &= \inf\{\lambda > 0 \,:\, \sqrt{f(x)} \leqslant \lambda\} \\ &= \inf\{\lambda > 0 \,:\, f(x) \leqslant \lambda^2\} \\ &= \inf\{\lambda > 0 \,:\, \tfrac{x}{\lambda} \in S_1(f)\} \,, \end{aligned}$$

which reveals the sublevel-set

$$S_1(f) = \{x \in \mathbb{R}^n \,:\, f(x) \leqslant 1\} =: C \,.$$

In other words, $\sqrt{f}$ is the gauge of a closed convex set C containing the origin.   □

Gauges are examples of sublinear functions which are closed. This is not the case of all sublinear functions: see the function $\tilde{f}$ of (1.2.1); another example in $\mathbb{R}^2$ is

$$h(\xi, \eta) := \begin{cases} 0 & \text{if } \eta > 0, \\ |\xi| & \text{if } \eta = 0, \\ +\infty & \text{if } \eta < 0. \end{cases}$$

By taking the closure, or lower semi-continuous hull, of a sublinear function σ, we get a new function defined by

$$\operatorname{cl} \sigma(x) := \liminf_{x' \to x} \sigma(x') \tag{1.2.5}$$

which is (i) closed by construction, (ii) convex (Proposition B.1.2.6) and (iii) positively homogeneous, as is immediately seen from (1.2.5). For example, to close the above h, one must set $h(\xi, 0) = 0$ for all ξ. We retain from this observation that, when we close a sublinear function, we obtain a new function which is closed, of course, but which inherits sublinearity. The subclass of sublinear functions that are also closed is extremely important; in fact most of our study will be restricted to these.

Note that, for a closed sublinear function σ,

$$\sigma(0) \leqslant \lim_{t \downarrow 0} \sigma(tx) = 0 \quad \text{for all } x \in \operatorname{dom} \sigma,$$

so certainly $\sigma(0) = 0$; otherwise, $\operatorname{dom} \sigma$ would be empty, a situation that we reject from our definitions. Another observation is that a closed sublinear function σ coincides with its asymptotic function:

$$\sigma'_\infty = \sigma \quad \text{if } \sigma \text{ is closed and sublinear}$$

(take $x_0 = 0$ in the definition of Proposition B.3.2.1). In particular, if σ is finite everywhere, then Proposition B.3.2.6 tells us that it is Lipschitzian, and its best Lipschitz constant is

$$\sup \{\sigma(d) \ : \ \|d\| = 1\}. \tag{1.2.6}$$

1.3 The Convex Cone of All Closed Sublinear Functions

Similarly to convex functions, sublinear functions, closed or not, can be combined to give new sublinear functions.

Proposition 1.3.1 (i) *If σ_1 and σ_2 are [closed] sublinear and t_1, t_2 are positive numbers, then $\sigma := t_1 \sigma_1 + t_2 \sigma_2$ is [closed] sublinear, if not identically $+\infty$.*

(ii) *If $\{\sigma_j\}_{j \in J}$ is a family of [closed] sublinear functions, then $\sigma := \sup_{j \in J} \sigma_j$ is [closed] sublinear, if not identically $+\infty$.*

Proof. Concerning convexity and closedness, everything is known from §B.2. Note in passing that a closed sublinear function is zero (hence finite) at zero. As for positive homogeneity, it is straightforward. $\square$

Proposition 1.3.2 *Let $\{\sigma_j\}_{j\in J}$ be a family of sublinear functions all minorized by some linear function. Then*

(i) $\sigma := \mathrm{co}\,(\inf_{j\in J}\sigma_j)$ *is sublinear.*

(ii) *If $J = \{1,\ldots,m\}$ is a finite set, we obtain the infimal convolution*

$$\mathrm{co}\min\{\sigma_1,\ldots,\sigma_m\} = \sigma_1 \,\mathbin{\rotatebox[origin=c]{180}{$\triangledown$}}\cdots\mathbin{\rotatebox[origin=c]{180}{$\triangledown$}}\, \sigma_m\,.$$

Proof. [*(i)*] Once again, the only thing to prove for (i) is positive homogeneity. Actually, it suffices to multiply x and each x_j by $t > 0$ in a formula giving $\mathrm{co}\,(\inf_j \sigma_j)(x)$, say (B.2.5.3).

[*(ii)*] By definition, computing $\mathrm{co}\,(\min_j \sigma_j)(x)$ amounts to solving the minimization problem in the m couples of variables $(x_j,\alpha_j) \in \mathrm{dom}\,\sigma_j \times \mathbb{R}$

$$\left|\begin{array}{l} \inf \sum_{j=1}^m \alpha_j\sigma_j(x_j) \qquad \alpha_j \geqslant 0 \\ \sum_{j=1}^m \alpha_j = 1\,,\ \sum_{j=1}^m \alpha_j x_j = x\,. \end{array}\right. \tag{1.3.1}$$

In view of positive homogeneity, the variables α_j play no role by themselves: the relevant variables are actually the products $\alpha_j x_j$ and (1.3.1) can be written – denoting $\alpha_j x_j$ again by x_j:

$$\mathrm{co}\left(\min_j \sigma_j\right)(x) = \inf\left\{\sum_{j=1}^m \sigma_j(x_j)\ :\ \sum_{j=1}^m x_j = x\right\}.$$

We recognize the infimal convolution of the σ_j's.   □

From Proposition 1.3.1(i), the collection of all closed sublinear functions has an *algebraic* structure: it is a convex cone contained in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$. It contains another convex cone, namely the collection of finite sublinear functions.

A *topological* structure can be defined on the latter cone. In linear analysis, one defines the Euclidean distance between two linear forms $\ell_1 = \langle s_1,\cdot\rangle$ and $\ell_2 = \langle s_2,\cdot\rangle$:

$$\|\ell_1 - \ell_2\| := \|s_1 - s_2\| = \max_{\|x\|\leqslant 1} |\ell_1(x) - \ell_2(x)|\,.$$

A distance can also be defined on the convex cone of *everywhere finite* sublinear functions (the extended-valued case is somewhat more delicate, just as with unbounded sets; see some explanations in §0.5.2), which of course contains the vector space of linear forms.

Theorem 1.3.3 *For σ_1 and σ_2 in the set Φ of sublinear functions that are finite everywhere, define*

$$\Delta(\sigma_1,\sigma_2) := \max_{\|x\|\leqslant 1} |\sigma_1(x) - \sigma_2(x)|\,. \tag{1.3.2}$$

Then Δ is a distance on Φ.

Proof. Clearly $\Delta(\sigma_1, \sigma_2) < +\infty$ and $\Delta(\sigma_1, \sigma_2) = \Delta(\sigma_2, \sigma_1)$. Now positive homogeneity of σ_1 and σ_2 gives for all $x \neq 0$

$$\begin{aligned}
|\sigma_1(x) - \sigma_2(x)| &= \|x\| \left| \sigma_1\left(\tfrac{x}{\|x\|}\right) - \sigma_2\left(\tfrac{x}{\|x\|}\right) \right| \\
&\leqslant \|x\| \max_{\|u\|=1} |\sigma_1(u) - \sigma_2(u)| \\
&\leqslant \|x\| \Delta(\sigma_1, \sigma_2) .
\end{aligned}$$

In addition, $\sigma_1(0) = \sigma_2(0) = 0$, so

$$|\sigma_1(x) - \sigma_2(x)| \leqslant \|x\| \Delta(\sigma_1, \sigma_2) \quad \text{for all } x \in \mathbb{R}^n$$

and $\Delta(\sigma_1, \sigma_2) = 0$ if and only if $\sigma_1 = \sigma_2$.

As for the triangle inequality, we have for arbitrary $\sigma_1, \sigma_2, \sigma_3$ in Φ

$$|\sigma_1(x) - \sigma_3(x)| \leqslant |\sigma_1(x) - \sigma_2(x)| + |\sigma_2(x) - \sigma_3(x)| \quad \text{for all } x \in \mathbb{R}^n ,$$

so there holds

$$\begin{aligned}
\Delta(\sigma_1, \sigma_3) &\leqslant \max_{\|x\|\leqslant 1} \left[|\sigma_1(x) - \sigma_2(x)| + |\sigma_2(x) - \sigma_3(x)| \right] \\
&\leqslant \max_{\|x\|\leqslant 1} |\sigma_1(x) - \sigma_2(x)| + \max_{\|x\|\leqslant 1} |\sigma_2(x) - \sigma_3(x)| ,
\end{aligned}$$

which is the required inequality. $\qquad\square$

The index-set in (1.3.2) can be replaced by the unit sphere $\|x\| = 1$, just as in (1.2.6); and the distance between an arbitrary $\sigma \in \Phi$ and the zero-function (which is in Φ) is just the value (1.2.6). The function $\Delta(\cdot, 0)$ acts like a *norm* on the convex cone Φ.

Example 1.3.4 Consider $\|\cdot\|_1$ and $\|\cdot\|_\infty$, the ℓ_1- and ℓ_∞-norms on $\mathbb{R}^n$. They are finite sublinear (Example 1.2.2) and there holds

$$\Delta(\|\cdot\|_1, \|\cdot\|_\infty) = \frac{n-1}{\sqrt{n}} .$$

To accept this formula, consider that, for symmetry reasons, the maximum in the definition (1.3.2) of Δ is achieved at $x = (1/\sqrt{n}, \dots, 1/\sqrt{n})$. $\qquad\square$

The convergence associated with this new distance function turns out to be the natural one:

Theorem 1.3.5 *Let (σ_k) be a sequence of finite sublinear functions and let σ be a finite function. Then the following are equivalent when $k \to +\infty$:*

(i) *(σ_k) converges pointwise to σ;*
(ii) *(σ_k) converges to σ uniformly on each compact set of $\mathbb{R}^n$;*
(iii) *$\Delta(\sigma_k, \sigma) \to 0$.*

Proof. First, the (finite) function σ is of course sublinear whenever it is the pointwise limit of sublinear functions. The equivalence between (i) and (ii) comes from the general Theorem B.3.1.4 on the convergence of convex functions.

Now, (ii) clearly implies (iii). Conversely $\Delta(\sigma_k, \sigma) \to 0$ is the uniform convergence on the unit ball, hence on any ball of radius $L > 0$ (the maximand in (1.3.2) is positively homogeneous), hence on any compact set. $\qquad\square$

2 The Support Function of a Nonempty Set

2.1 Definitions, Interpretations

Definition 2.1.1 (Support Function) Let S be a nonempty set in $\mathbb{R}^n$. The function $\sigma_S : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ defined by

$$\mathbb{R}^n \ni x \mapsto \sigma_S(x) := \sup \{\langle s, x \rangle \: : \: s \in S\} \qquad (2.1.1)$$

is called the *support function* of S. $\square$

For a given S, the support function therefore depends on the scalar product: changing $\langle \cdot, \cdot \rangle$ changes σ_S. In (2.1.1), the space where s runs and the space where σ_S acts are dual to each other.

The supremum in (2.1.1) may be finite or infinite, achieved on S or not. In this context, S can be interpreted as an index set: $\sigma_S(\cdot)$ is the supremum of the collection of linear forms $\langle s, \cdot \rangle$ over S. We obtain immediately:

Proposition 2.1.2 *A support function is closed and sublinear.*

Proof. This results from Proposition 1.3.1(ii) (a linear form is closed and convex!). Observe in particular that a support function is null (hence $< +\infty$) at the origin. $\square$

The domain of σ_S is a convex cone, closed or not. Actually, $x \in \operatorname{dom} \sigma_S$ means that, for some $r := \sigma_S(x)$:

$$S \subset \{s \in \mathbb{R}^n \: : \: \langle s, x \rangle \leqslant r\} \qquad (2.1.2)$$

i.e. S is contained in a closed half-space "opposite" to x.

Proposition 2.1.3 *The support function of S is finite everywhere if and only if S is bounded.*

Proof. Let S be bounded, say $S \subset B(0, L)$ for some $L > 0$. Then

$$\langle s, x \rangle \leqslant \|s\| \, \|x\| \leqslant L \|x\| \quad \text{for all } s \in S \,,$$

which implies $\sigma_S(x) \leqslant L \|x\|$ for all $x \in \mathbb{R}^n$.

Conversely, finiteness of the convex σ_S implies its continuity on the whole space (Theorem B.3.1.2), hence its local boundedness: for some L,

$$\langle s, x \rangle \leqslant \sigma_S(x) \leqslant L \quad \text{for all } (s, x) \in S \times B(0, 1) \,.$$

If $s \neq 0$, we can take $x = s/\|s\|$ in the above relation, which implies $\|s\| \leqslant L$. $\square$

Observing that

$$-\sigma_S(-x) = -\sup_{s \in S} [-\langle s, x \rangle] = \inf_{s \in S} \langle s, x \rangle \,,$$

the number $\sigma_S(x) + \sigma_S(-x)$ of (1.1.6) is particularly interesting here:

Definition 2.1.4 (Breadth of a Set) The *breadth* of a nonempty set S along $x \neq 0$ is

$$\sigma_S(x) + \sigma_S(-x) = \sup_{s \in S} \langle s, x \rangle - \inf_{s \in S} \langle s, x \rangle,$$

a number in $[0, +\infty]$. It is 0 if and only if S lies entirely in some affine hyperplane orthogonal to x; such a hyperplane is expressed as

$$\{y \in \mathbb{R}^n \ : \ \langle y, x \rangle = \sigma_S(x)\},$$

which in particular contains S. The intersection of all these hyperplanes is just the affine hull of S.   □

If x has norm 1 and is interpreted as a *direction*, the breadth of S measures how "thick" S is along x: it is the distance between the two hyperplanes orthogonal to x and "squeezing" S. This observation calls for a more general comment: a sublinear function $x \mapsto \sigma(x)$ being positively homogeneous, the norm of its argument x has little importance. This argument should always be thought of as an *oriented direction*, i.e. a normalized vector of $\mathbb{R}^n$. Accordingly, we will generally use from now on the notation $\sigma(d)$, more suggestive for a support function than $\sigma(x)$.

Here, we give two geometric constructions which help interpreting a support function.

Interpretation 2.1.5 (Construction in $\mathbb{R}^n$) Given $S \subset \mathbb{R}^n$ and $d \neq 0$, consider for each $r \in \mathbb{R}$ the closed half-space alluded to in (2.1.2):

$$H_{d,r}^- := \{z \in \mathbb{R}^n \ : \ \langle z, d \rangle \leqslant r\}. \tag{2.1.3}$$

If (2.1.2) holds, we can find r large enough so that $S \subset H_{d,r}^-$. The value $\sigma_S(d)$ is the smallest of those r: decreasing r as much as possible while keeping S in $H_{d,r}^-$ means "leaning" onto S the affine hyperplane $H_{d,r} := \{z \in \mathbb{R}^n \ : \ \langle z, d \rangle = r\}$. See Fig. 2.1.1 for an illustration.
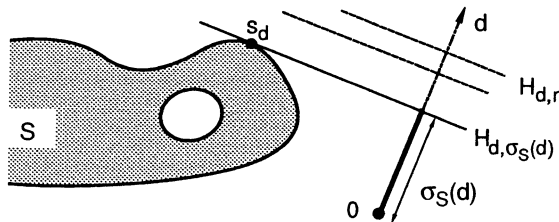
Fig. 2.1.1. Supporting hyperplanes and support functions

If (2.1.2) does not hold, however, this operation is impossible: S is "unbounded in the oriented direction" d and $\sigma_S(d) = +\infty$. Take for example $S := \mathbb{R}^+ \times \{0\}$ in $\mathbb{R}^2$. For $d = (1, 1)$ say (and assuming that $\langle \cdot, \cdot \rangle$ is the usual dot-product), no closed half-space of the form (2.1.3) can contain S, even if r is increased to $+\infty$.

If S is compact, the supremum of the (continuous) function $\langle \cdot, d \rangle$ is achieved on S, no matter how d is chosen. This means that, somewhere on the hyperplane $H_{d,\sigma_S(d)}$ there is some s_d which is also in S, actually a boundary point of S; more accurately, s_d lies in the face of S exposed by d (assuming S convex). $\square$

Figure 2.1.1 suggests (and Proposition 2.2.1 below confirms) that the support functions of S and of $\overline{\text{co}}\, S$ coincide. Note also that the distance from the origin 0 to the "optimal" hyperplane $H_{d,\sigma_S(d)}$ is $|\sigma_S(d/\|d\|)|$. This is easily confirmed: project the origin onto $H_{d,\sigma_S(d)}$ to obtain the vector t^*d such that $\langle d, t^*d \rangle = \sigma_S(d)$. Then the distance from 0 to $H_{d,\sigma_S(d)}$ is $\|t^*d\|$.

Interpretation 2.1.6 (Construction in $\mathbb{R}^{n+1}$) In the graph-space $\mathbb{R}^n \times \mathbb{R}$, we shift S down to $\mathbb{R}^n \times \{-1\}$ and consider the convex conical hull K_S of this shifted copy of S. Then the polar cone $(K_S)^\circ$ of K_S is nothing else than the epigraph of σ_S. Indeed

$$K_S = \mathbb{R}^+ \,\text{co}\,(S \times \{-1\}) = \text{co}\,[\mathbb{R}^+ (S \times \{-1\})],$$

so that

$$\begin{aligned}
(K_S)^\circ &= \{(d,r) : t\langle s, d\rangle - tr \leqslant 0 \text{ for all } s \in S \text{ and } t > 0\} \\
&= \{(d,r) : \langle s, d\rangle \leqslant r \text{ for all } s \in S\} \\
&= \{(d,r) : \sup_{s \in S}\langle s, d\rangle \leqslant r\} = \text{epi}\,\sigma_S\,.
\end{aligned}$$

This is illustrated on Fig. 2.1.2. On this picture, $0 \in S$; this implies $\sigma_S(d) \geqslant 0$ for all d, which is obvious just from the definition (2.1.1) of σ_S. $\square$
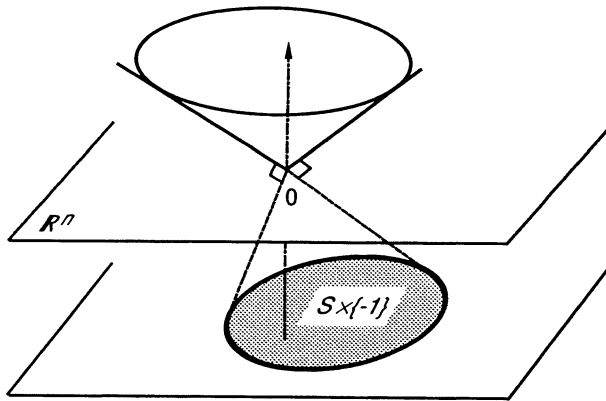
Fig. 2.1.2. The epigraph of a support function

2.2 Basic Properties

First, we list some properties of support functions that are directly derived from their definition.

Proposition 2.2.1 *For $S \subset \mathbb{R}^n$ nonempty, there holds $\sigma_S = \sigma_{\mathrm{cl}\,S} = \sigma_{\mathrm{co}\,S}$; whence*

$$\sigma_S = \sigma_{\overline{\mathrm{co}}\,S} \, . \tag{2.2.1}$$

Proof. The continuity [resp. linearity, hence convexity] of the function $\langle s, \cdot \rangle$, which is maximized over S, implies that $\sigma_S = \sigma_{\mathrm{cl}\,S}$ [resp. $\sigma_S = \sigma_{\mathrm{co}\,S}$]. Knowing that $\overline{\mathrm{co}}\,S = \mathrm{cl}\,\mathrm{co}\,S$ (Proposition A.1.4.2), (2.2.1) follows immediately. $\qquad\square$

This result is of utmost importance: it says that the concept of support function *does not distinguish* a set S from its closed convex hull. Thus, when dealing with support functions, it makes no difference if we restrict ourselves to the case of closed convex sets.

As a result of (2.1.1) and (2.2.1), we can write

$$s \in \overline{\mathrm{co}}\,S \quad \Longrightarrow \quad [\langle s, d \rangle \leqslant \sigma_S(d) \quad \text{for all } d \in \mathbb{R}^n] \, .$$

Now, what about the converse? Can it be that the above (infinite) set of inequalities still holds if s is not in $\overline{\mathrm{co}}\,S$? The answer is no:

Theorem 2.2.2 *For the nonempty $S \subset \mathbb{R}^n$ and its support function σ_S, there holds*

$$s \in \overline{\mathrm{co}}\,S \quad \Longleftrightarrow \quad [\langle s, d \rangle \leqslant \sigma_S(d) \quad \text{for all } d \in X] \, , \tag{2.2.2}$$

where the set X can be indifferently taken as: the whole of $\mathbb{R}^n$, the unit ball $B(0,1)$ or its boundary the unit sphere $\widetilde{B}$, or $\mathrm{dom}\,\sigma_S$.

Proof. First, the equivalence between all the choices for X is clear enough; in particular due to positive homogeneity. Because "$\Rightarrow$" is Proposition 2.2.1, we have to prove "$\Leftarrow$" only, with $X = \mathbb{R}^n$ say.

So suppose that $s \notin \overline{\mathrm{co}}\,S$. Then $\{s\}$ and $\overline{\mathrm{co}}\,S$ can be strictly separated (Theorem A.4.1.1): there exists $d_0 \in \mathbb{R}^n$ such that

$$\langle s, d_0 \rangle > \sup \{ \langle s', d_0 \rangle \ : \ s' \in \overline{\mathrm{co}}\,S \} = \sigma_S(d_0) \, ,$$

where the last equality is (2.2.1). Our result is proved by contradiction. $\qquad\square$

As a result, a closed convex set is completely determined by its support function: between the classes of closed convex sets and of support functions, there is a correspondence which is bijective, as illustrated on Fig. 2.2.1.
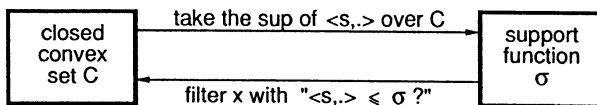
Fig. 2.2.1. Correspondence between closed convex sets and support functions

Thus, whether a given point s belongs to a given closed convex set S can be checked with the help of (2.2.2), which holds as an equivalence. Actually, more can

be said: the support function *filters* the interior, the relative interior and the affine hull of a closed convex set.

This property is best understood with Fig. 1.1.2 in mind. Let V be the subspace parallel to aff S, and $U := V^{\perp}$. Indeed, U is just given in (1.1.8) with $\sigma = \sigma_S$: U can be viewed either as the subspace where the sublinear function σ_S is linear, or where the supported set S is flat; by contrast, σ_S is V-shaped in V, while S is thick along V. When drawn in the geometric space of convex sets, Fig. 1.1.2 becomes Fig. 2.2.2, which is very helpful to follow the next proof.

Theorem 2.2.3 *Let S be a nonempty closed convex set in $\mathbb{R}^n$. Then*

(i) $s \in$ aff S *if and only if*

$$\langle s, d \rangle = \sigma_S(d) \quad \text{for all } d \text{ with } \sigma_S(d) + \sigma_S(-d) = 0; \qquad (2.2.3)$$

(ii) $s \in$ ri S *if and only if*

$$\langle s, d \rangle < \sigma_S(d) \quad \text{for all } d \text{ with } \sigma_S(d) + \sigma_S(-d) > 0; \qquad (2.2.4)$$

(iii) *in particular, $s \in$ int S if and only if*

$$\langle s, d \rangle < \sigma_S(d) \quad \text{for all } d \neq 0. \qquad (2.2.5)$$
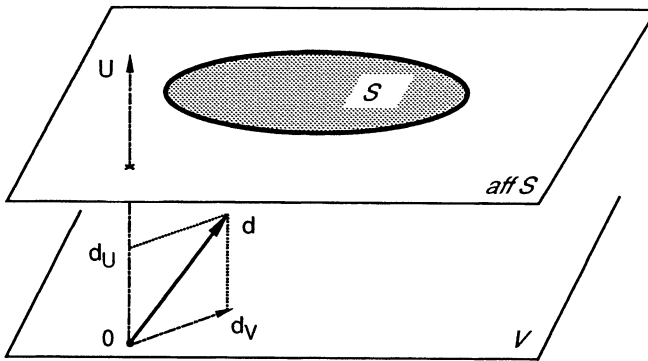
Fig. 2.2.2. Affine hulls and orthogonal spaces

Proof. [*(i)*] Let first $s \in S$. We have already seen in Definition 2.1.4 that

$$-\sigma_S(-d) \leqslant \langle s, d \rangle \leqslant \sigma_S(d) \quad \text{for all } d \in \mathbb{R}^n .$$

If the breadth of S along d is zero, we obtain a pair of equalities: for such d, there holds

$$\langle s, d \rangle = \sigma_S(d) ,$$

an equality which extends by affine combination to any element $s \in \text{aff } S$.

Conversely, let s satisfy (2.2.3). A first case is when the only d described in (2.2.3) is $d = 0$; as a consequence of our observations in Definition 2.1.4, there is no affine hyperplane containing S, i.e. aff $S = \mathbb{R}^n$ and there is nothing to prove. Otherwise, there does exist a hyperplane H containing S; it is defined by

$$H := \{ p \in \mathbb{R}^n \ : \ \langle p, d_H \rangle = \sigma_S(d_H) \} , \tag{2.2.6}$$

for some $d_H \neq 0$. We proceed to prove $\langle s, \cdot \rangle \leqslant \sigma_H$.

In fact, the breadth of S along d_H is certainly 0, hence $\langle s, d_H \rangle = \sigma_S(d_H)$ because of (2.2.3), while (2.2.6) shows that $\sigma_S(d_H) = \sigma_H(d_H)$. On the other hand, it is obvious that $\sigma_H(d) = +\infty$ if d is not collinear to d_H. In summary, we have proved $\langle s, d \rangle \leqslant \sigma_H(d)$ for all d, i.e. $s \in H$. We conclude that our s is in any affine manifold containing S: $s \in \text{aff } S$.

[(iii)] In view of positive homogeneity, we can normalize d in (2.2.5). For $s \in \text{int } S$, there exists $\varepsilon > 0$ such that $s + \varepsilon d \in S$ for all d in the unit sphere $\widetilde{B}$. Then, from the very definition (2.1.1),

$$\sigma_S(d) \geqslant \langle s + \varepsilon d, d \rangle = \langle s, d \rangle + \varepsilon \quad \text{for all } d \in \widetilde{B} .$$

Conversely, let $s \in \mathbb{R}^n$ be such that

$$\sigma_S(d) - \langle s, d \rangle > 0 \quad \text{for all } d \in \widetilde{B}$$

which implies, because σ_S is closed and the unit sphere is compact:

$$0 < \varepsilon := \inf \{ \sigma_S(d) - \langle s, d \rangle \ : \ d \in \widetilde{B} \} \leqslant +\infty .$$

Thus

$$\langle s, d \rangle + \varepsilon \leqslant \sigma_S(d) \quad \text{for all } d \in \widetilde{B} .$$

Now take u with $\|u\| < \varepsilon$. From the Cauchy-Schwarz inequality, we have for all $d \in \widetilde{B}$

$$\langle s + u, d \rangle = \langle s, d \rangle + \langle u, d \rangle \leqslant \langle s, d \rangle + \varepsilon \leqslant \sigma_S(d)$$

and this implies $s + u \in S$ because of Theorem 2.2.2: $s \in \text{int } S$ and (iii) is proved.

[(ii)] Look at Fig. 2.2.2 again: decompose $\mathbb{R}^n = V \oplus U$, where V is the subspace parallel to aff S and $U = V^\perp$. In the decomposition $d = d_V + d_U$, $\langle \cdot, d_U \rangle$ is constant over S, so S has 0-breadth along d_U and

$$\sigma_S(d) = \sup_{s \in S} \langle s, d_V + d_U \rangle = \langle s, d_U \rangle + \sup_{s \in S} \langle s, d_V \rangle$$

for any $s \in S$. With these notations, a direction described as in (2.2.4) is a d such that

$$\sigma_S(d) + \sigma_S(-d) = \sigma_S(d_V) + \sigma_S(-d_V) > 0 .$$

Then, (ii) is just (iii) written in the subspace V. $\qquad \qquad \square$

We already know that the domain of σ_S is a convex cone, which consists of all oriented directions "along which S is bounded" (remember Interpretation 2.1.5). This can be made more explicit.

Proposition 2.2.4 *Let S be a nonempty closed convex set in $\mathbb{R}^n$. Then* $\operatorname{cl} \operatorname{dom} \sigma_S$ *and the asymptotic cone S_∞ of S are mutually polar cones.*

Proof. Recall from §A.3.2 that, if K_1 and K_2 are two closed convex cones, then $K_1 \subset K_2$ if and only if $(K_1)^\circ \supset (K_2)^\circ$.

Let $p \in S_\infty$. Fix s_0 arbitrary in S and use the fact that $S_\infty = \cap_{t>0} t(S - s_0)$ (§A.2.2): for all $t > 0$, we can find $s_t \in S$ such that $p = t(s_t - s_0)$. Now, for $q \in \operatorname{dom} \sigma_S$, there holds

$$\langle p, q \rangle = t \langle s_t - s_0, q \rangle \leqslant t \left[\sigma_S(q) - \langle s_0, q \rangle \right] < +\infty$$

and letting $t \downarrow 0$ shows that $\langle p, q \rangle \leqslant 0$. In other words, $\operatorname{dom} \sigma_S \subset (S_\infty)^\circ$; then $\operatorname{cl} \operatorname{dom} \sigma_S \subset (S_\infty)^\circ$ since the latter is closed.

Conversely, let $q \in (\operatorname{dom} \sigma_S)^\circ$, which is a cone, hence $tq \in (\operatorname{dom} \sigma_S)^\circ$ for any $t > 0$. Thus, given $s_0 \in S$, we have for arbitrary $p \in \operatorname{dom} \sigma_S$

$$\langle s_0 + tq, p \rangle = \langle s_0, p \rangle + t \langle q, p \rangle \leqslant \langle s_0, p \rangle \leqslant \sigma_S(p) \,,$$

so $s_0 + tq \in S$ by virtue of Theorem 2.2.2. In other words: $q \in (S - s_0)/t$ for all $t > 0$ and $q \in S_\infty$. □

2.3 Examples

Let us start with elementary situations. The simplest example of a support function is that of a singleton $\{s\}$. Then $\sigma_{\{s\}}$ is merely $\langle s, \cdot \rangle$, we have a first illustration of the introduction (iii) to this chapter: the concept of a linear form $\langle s, \cdot \rangle$ can be generalized to s not being a singleton, which amounts to generalizing linearity to closed sublinearity (more details will be given in §3). The case when S is the unit ball $B(0, 1)$ is also rather simple:

$$\sigma_{B(0,1)}(d) \geqslant \left\langle \frac{d}{\|d\|}, d \right\rangle = \|d\| \quad (\text{if } d \neq 0)$$

and, for $s \in B(0, 1)$, the Cauchy-Schwarz inequality implies $\langle s, d \rangle \leqslant \|d\|$. Altogether,

$$\sigma_{B(0,1)}(d) = \|d\| \,. \tag{2.3.1}$$

Our next example is the simplest possible illustration of Proposition 2.2.4, namely when S_∞ is S itself:

Example 2.3.1 (Cones, Half-Spaces, Subspaces) Let K be a closed convex cone of $\mathbb{R}^n$. Then

$$\sigma_K(d) = \begin{cases} 0 & \text{if } \langle s, d \rangle \leqslant 0 \text{ for all } s \in K, \\ +\infty & \text{otherwise} \,. \end{cases}$$

In other words, σ_K is the indicator function of the polar cone K°. Note the symmetry: since $K^{\circ\circ} = K$, the support function of K° is the indicator of K. In summary:

$$i_K = \sigma_{K^\circ} \quad \text{and} \quad \sigma_K = i_{K^\circ}.$$

Two particular cases are of interest. One is when K is a half-space:

$$K := \{s \in \mathbb{R}^n \ : \ \langle s, v \rangle \leqslant 0\};$$

then it is clear enough that

$$\sigma_K(d) = \begin{cases} 0 & \text{if } d = tv \text{ with } t \geqslant 0, \\ +\infty & \text{otherwise}. \end{cases} \tag{2.3.2}$$

Needless to say, the support function of the half-line $\mathbb{R}^+ v$ (the polar of K) is in turn the indicator of K.

The other interesting case is that of a subspace. Let $A : \mathbb{R}^n \to \mathbb{R}^m$ be a linear operator and H be defined by

$$H := \operatorname{Ker} A = \{s \in \mathbb{R}^n \ : \ As = 0\}.$$

Then the support function of H is the indicator of the orthogonal subspace $H^\perp$:

$$\sigma_H(d) = i_{H^\perp}(d) = \begin{cases} 0 & \text{if } \langle s, d \rangle = 0 \text{ for all } s \in H, \\ +\infty & \text{otherwise}. \end{cases}$$

The subspace $H^\perp$ can be defined with the help of the adjoint of A:

$$H^\perp = (\operatorname{Ker} A)^\perp = \operatorname{Im} A^* = \{A^* \lambda \ : \ \lambda \in \mathbb{R}^m\}.$$

If A or H are defined in terms of linear constraints

$$H := \{s \in \mathbb{R}^n \ : \ \langle s, a_j \rangle = 0 \text{ for } j = 1, \ldots, m\},$$

then $H^\perp = \{\sum_{j=1}^m \lambda_j a_j : \lambda \in \mathbb{R}^m\}$.

All these calculations are useful when dealing with closed convex polyhedra, expressed as intersections of half-spaces and subspaces.

Figure 2.3.1 illustrates a modification in which our cone K is modified to $K' := K \cap B(0, 1)$. The calculus rules of §3.3 will prove what is suggested by the picture: the support function of K' is the distance function to K° (check the similarity of the appropriate triangles, and note that $\sigma_{K'}(d) = 0$ when $d \in K^\circ$). $\qquad \square$

Example 2.3.2 The asymptotic cone of the set

$$S := \{s = (\rho, \tau) \in \mathbb{R}^2 \ : \ \rho > 0, \ \tau \geqslant 1/\rho\} \tag{2.3.3}$$

is $S_\infty = \{(\rho, \tau) \in \mathbb{R}^2 : \rho \geqslant 0, \ \tau \geqslant 0\}$ and, from Proposition 2.2.4, the closure of $\operatorname{dom} \sigma_S$ is $\{(\xi, \eta) : \xi \leqslant 0, \ \eta \leqslant 0\}$.
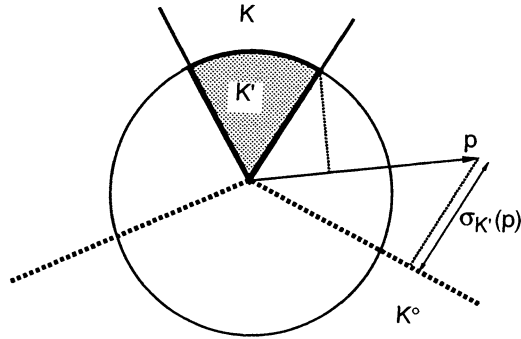
Fig. 2.3.1. Support function of a truncated cone

The exact status of the boundary of $\operatorname{dom}\sigma_S$ (i.e. when $\xi\eta = 0$) is not specified by Proposition 2.2.4: is σ_S finite there? The computation of σ_S can be done directly from the definitions (2.1.1) and (2.3.3). However the following geometric argument yields simpler calculations (see Fig. 2.3.2): for given $d = (\xi, \eta) \neq (0,0)$, consider the hyperplane $H_{d,\sigma_S(d)} = \{(\alpha, \beta) : \xi\alpha + \eta\beta = \sigma_S(d)\}$. It has to be tangent to the boundary of S, defined by the equation $\alpha\beta = 1$. So, the discriminant $\sigma_S^2(d) - 4\xi\eta$ of the equation in α

$$\xi\alpha + \eta\frac{1}{\alpha} = \sigma_S(d)$$

must be 0. We obtain directly $\sigma_S(\xi, \eta) = -2\sqrt{\xi\eta}$ for $\xi < 0, \eta < 0$ (the sign is "$-$" because $0 \notin S$; remember Theorem 2.2.2). Finally, Proposition 2.1.2 tells us that the closed function $(\xi, \eta) \mapsto \sigma_S(\xi, \eta)$ has to be 0 when $\xi\eta = 0$. All this is confirmed by Fig. 2.3.2. $\square$
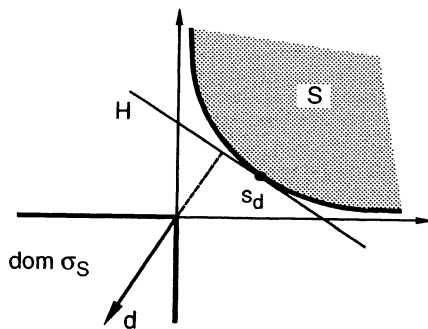
Fig. 2.3.2. A support function

Remark 2.3.3 Two features concerning the boundary of $\operatorname{dom}\sigma_S$ are worth mentioning on the above example: the supremum in (2.1.1) is not attained when $d \in \operatorname{bd}\operatorname{dom}\sigma_S$ (the point s_d of Fig. 2.1.1 is sent to infinity when d approaches $\operatorname{bd}\operatorname{dom}\sigma_S$), and $\operatorname{dom}\sigma_S$ is closed.

These are not the only possible cases: Example 2.3.1 shows that the supremum in (2.1.1) can well be attained for all $d \in \operatorname{dom} \sigma_S$; and in the example

$$S := \left\{ (\rho, \tau) \ : \ \tau \geqslant \tfrac{1}{2}\rho^2 \right\},$$

$\operatorname{dom} \sigma_S$ is not closed. The difference is that, now, S has no asymptote "at finite distance".

□

Example 2.3.4 Just as in Example 1.2.3, let Q be a symmetric positive definite operator from $\mathbb{R}^n$ to $\mathbb{R}^n$. Its sublevel-set

$$E_Q := \left\{ s \in \mathbb{R}^n \ : \ \langle Qs, s \rangle \leqslant 1 \right\}$$

is elliptic, with support function

$$d \mapsto \sigma_{E_Q}(d) := \max \left\{ \langle s, d \rangle \ : \ \langle Qs, s \rangle \leqslant 1 \right\}. \tag{2.3.4}$$

Calling $Q^{1/2}$ the square root of Q, the change of variable $p = Q^{1/2}s$ in (2.3.4) gives

$$\sigma_{E_Q}(d) = \max \left\{ \langle p, Q^{-1/2}d \rangle \ : \ \|p\|^2 \leqslant 1 \right\}$$

whose unique solution for $d \neq 0$ (again Cauchy-Schwarz!) is $p = \dfrac{Q^{-1/2}d}{\|Q^{-1/2}d\|}$ and finally

$$\sigma_{E_Q}(d) = \left\| Q^{-1/2}d \right\| = \sqrt{\langle d, Q^{-1}d \rangle}. \tag{2.3.5}$$

Observe in this example the "duality" between the gauge $x \mapsto \sqrt{\langle Qx, x \rangle}$ of E_Q and its support function (2.3.5).

When Q is merely symmetric positive semi-definite, E_Q becomes an elliptic cylinder, whose asymptotic cone is $\operatorname{Ker} Q$ (remember Example 1.2.3). Then Proposition 2.2.4 tells us that

$$\operatorname{cl} \operatorname{dom} \sigma_{E_Q} = (\operatorname{Ker} Q)^\circ = (\operatorname{Ker} Q)^\perp = \operatorname{Im} Q.$$

When $d \in \operatorname{Im} Q$, $\sigma_{E_Q}(d)$ is finite indeed and (2.3.5) does hold, $Q^{-1}d$ denoting now any element p such that $Qp = d$. We leave this as an exercise.   □

3 The Isomorphism Between Closed Convex Sets and Closed Sublinear Functions

3.1 The Fundamental Correspondence

We have seen in Proposition 2.1.2 that a support function is closed and sublinear. What about the converse? Are there closed sublinear functions which support no set in $\mathbb{R}^n$? The answer is no: any closed sublinear function *can be viewed* as a support function. The key lies in the representation of a closed convex function f via affine functions minorizing it: when the starting f is also positively homogeneous, the underlying affine functions can be assumed linear.

Theorem 3.1.1 *Let σ be a closed sublinear function; then there is a linear function minorizing σ. In fact, σ is the supremum of the linear functions minorizing it. In other words, σ is the support function of the nonempty closed convex set*

$$S_\sigma := \{s \in \mathbb{R}^n \; : \; \langle s, d \rangle \leqslant \sigma(d) \text{ for all } d \in \mathbb{R}^n \}. \tag{3.1.1}$$

Proof. Being convex, σ is minorized by some affine function (Proposition B.1.2.1): for some $(s, r) \in \mathbb{R}^n \times \mathbb{R}$,

$$\langle s, d \rangle - r \leqslant \sigma(d) \quad \text{for all } d \in \mathbb{R}^n . \tag{3.1.2}$$

Because $\sigma(0) = 0$, the above r is nonnegative. Also, by positive homogeneity,

$$\langle s, d \rangle - \tfrac{1}{t} r \leqslant \sigma(d) \quad \text{for all } d \in \mathbb{R}^n \text{ and all } t > 0 .$$

Letting $t \to +\infty$, we see that σ is actually minorized by a linear function:

$$\langle s, d \rangle \leqslant \sigma(d) \quad \text{for all } d \in \mathbb{R}^n . \tag{3.1.3}$$

Now observe that the minorization (3.1.3) is sharper than (3.1.2): when expressing the closed convex σ as the supremum of all the affine functions minorizing it (Proposition B.1.2.8), we can restrict ourselves to linear functions. In other words

$$\sigma(d) = \sup \{ \langle s, d \rangle \; : \; \text{the linear } \langle s, \cdot \rangle \text{ minorizes } \sigma \} ;$$

in the above index-set, we just recognize S_σ. $\qquad\square$

One of the important points in this result is the *nonemptiness* of S_σ in (3.1.1); we have here the analytical form of Hahn-Banach theorem: there *exists* a linear function minorizing the closed sublinear function σ; compare this with the geometric form given in Theorem A.4.1.1.

Another way of expressing Theorem 3.1.1 is that the closed convex set epi σ is the intersection of the closed half-spaces containing it; but since epi σ is actually a cone, these half-spaces can be assumed to have *linear* hyperplanes as boundaries (remember Corollary A.4.2.7). A connection between S_σ and the cone polar to epi σ is thus introduced; Chap. D will exploit this remark.

The main consequence of this important theorem is an assessment of closed sublinear functions. Section 2.2 has established a bijection from closed convex sets onto support functions. Thanks to Theorem 3.1.1, this bijection is actually onto *closed sublinear functions*, which is of course much more satisfactory: the latter class of functions is defined in abstracto, while the former class was ad hoc, as far as this bijection was concerned.

Thus, the wording "support function" in Fig. 2.2.1 can everywhere be replaced by "closed sublinear". This replacement can be done in Theorem 2.2.2 as well:

Corollary 3.1.2 *For a nonempty closed convex set S and a closed sublinear function σ, the following are equivalent:*

(i) σ is the support function of S,

(ii) $S = \{s : \langle s, d \rangle \leqslant \sigma(d)$ for all $d \in X\}$, where the set X can be indifferently taken as: the whole of $\mathbb{R}^n$, the unit ball $B(0, 1)$ or its boundary, or $\mathrm{dom}\,\sigma$.

Proof. The case $X = \mathbb{R}^n$ is just Theorem 3.1.1. The other cases are then clear. $\quad\Box$

Remember the outer construction of §A.4.2(b): a closed convex set S is geometrically characterized as an intersection of half-spaces, which in turn can be characterized in terms of the support function of S. Each $(d, r) \in \mathbb{R}^n \times \mathbb{R}$ defines (for $d \neq 0$) the half-space $H^-_{d,r}$ via (2.1.3). This half-space contains S if and only if $r \geqslant \sigma(d)$, and Corollary 3.1.2 expresses that

$$S = \cap\{s : \langle s, d \rangle \leqslant r \text{ for all } d \in \mathbb{R}^n \text{ and } r \geqslant \sigma(d)\},$$

in which the couple (d, r) plays the role of an index, running in the index-set epi $\sigma \subset \mathbb{R}^n \times \mathbb{R}$ (compare with the discussion after Definition 2.1.1). Of course, this index-set can be reduced down to $\mathbb{R}^n$: the above formula can be simplified to

$$S = \cap\{s : \langle s, d \rangle \leqslant \sigma(d) \text{ for all } d \in X\}$$

where X can be taken as in Corollary 3.1.2.

Recall from §A.2.4 that an exposed face of a convex set S is defined as the set of points of S which maximize some (nonzero) linear form. This concept appears as particularly welcome in the context of support functions:

Definition 3.1.3 (Direction Exposing a Face) Let C be a nonempty closed convex set, with support function σ. For given $d \neq 0$, the set

$$F_C(d) := \{x \in C : \langle x, d \rangle = \sigma(d)\}$$

is called the exposed face of C associated with d, or the *face exposed* by d. $\quad\Box$

For a unified notation, the entire C can be considered as the face exposed by 0. On the other hand, a given d may expose no face at all (when C is unbounded).

Symmetrically to Definition 3.1.3, one can ask what are those $d \in \mathbb{R}^n$ such that $\langle \cdot, d \rangle$ is maximized at a given $x \in C$. We obtain nothing other than the normal cone $N_C(x)$ to C at x, as is obvious from its Definition A.5.2.3. The following result is simply a restatement of Proposition A.5.3.3.

Proposition 3.1.4 *For x in a nonempty closed convex set C, it holds*

$$x \in F_C(d) \quad \Longleftrightarrow \quad d \in N_C(x)\,. \qquad\qquad \Box$$

When d describes the set of normalized directions, the corresponding exposed faces exactly describe the boundary of C:

Proposition 3.1.5 *For a nonempty closed convex set C, it holds*

$$\mathrm{bd}\,C = \cup\{F_C(d) : d \in X\}$$

where X can be indifferently taken as: $\mathbb{R}^n \backslash \{0\}$, the unit sphere $\widetilde{B}$, or $\mathrm{dom}\,\sigma_C \backslash \{0\}$.

Proof. Observe from Definition 3.1.3 that the face exposed by $d \neq 0$ does not depend on $\|d\|$. This establishes the equivalence between the first two choices for X. As for the third choice, it is due to the fact that $F_C(d) = \emptyset$ if $d \notin \operatorname{dom} \sigma_C$.

Now, if x is interior to C and $d \neq 0$, then $x + \varepsilon d \in C$ and x cannot be a maximizer of $\langle \cdot, d \rangle$: x is not in the face exposed by d. Conversely, take x on the boundary of C. Then $N_C(x)$ contains a nonzero vector d; by Proposition 3.1.4, $x \in F_C(d)$. □

3.2 Example: Norms and Their Duals, Polarity

Let $\| \cdot \|$ be an arbitrary norm on $\mathbb{R}^n$. It is a positive (except at 0) closed sublinear function and its sublevel-set

$$B := \{x \in \mathbb{R}^n : \|x\| \leqslant 1\} \tag{3.2.1}$$

is particularly interesting. It is the unit ball associated with the norm, a symmetric, convex, compact set containing the origin as an interior point; $\| \cdot \|$ is the gauge of B (§1.2). On the other hand, why not take the set whose support function is $\| \cdot \|$? In view of Corollary 3.1.2, it is defined by

$$\{s \in \mathbb{R}^n : \langle s, x \rangle \leqslant \|x\| \text{ for all } x \in \mathbb{R}^n\} =: B^*. \tag{3.2.2}$$

It is an easy exercise to check that B^* is also symmetric, convex, compact; and it contains the origin as an interior point (Theorem 2.2.3(iii)).

Now, we have two closed convex sets B and B^*. We can generate two more closed sublinear functions: take the support function σ_B of B and the gauge γ_{B^*} of B^*. It turns out that we then obtain the same function, which actually is a norm, denoted by $\| \cdot \|^*$: the so-called *dual norm* of $\| \cdot \|$. The game finishes there: the two sets that $\| \cdot \|^*$ supports and is the gauge of, respectively, are B and B^*.

Proposition 3.2.1 *Let B and B^* be defined by (3.2.1) and (3.2.2), where $\| \cdot \|$ is a norm on $\mathbb{R}^n$. The support function of B and the gauge of B^* are the same function $\| \cdot \|^*$ defined by*

$$\|s\|^* := \max\{\langle s, x \rangle : \|x\| \leqslant 1\}. \tag{3.2.3}$$

Furthermore, $\| \cdot \|^*$ is a norm on $\mathbb{R}^n$. The support function of its unit ball B^* and the gauge of its supported set B are the same function $\| \cdot \|$: there holds

$$\|x\| = \max\{\langle s, x \rangle : \|s\|^* \leqslant 1\}. \tag{3.2.4}$$

Proof. It is a particular case of the results 3.2.4 and 3.2.5 below. □

Note the following symmetric relation ("Cauchy-Schwarz")

$$\langle s, x \rangle \leqslant \|s\|^* \, \|x\| \quad \text{for all } (s, x) \in \mathbb{R}^n \times \mathbb{R}^n, \tag{3.2.5}$$

which comes directly from (3.2.3), using positive homogeneity. It expresses the duality correspondence between the two Banach spaces $(\mathbb{R}^n, \| \cdot \|)$ and $(\mathbb{R}^n, \| \cdot \|^*)$.

Furthermore, equality holds in (3.2.5) when $s \neq 0$ and $x \neq 0$ form an associated pair via Proposition 3.1.4:

$$\frac{s}{\|s\|_*} \in F_{B^*}(x) \quad \text{or equivalently} \quad \frac{x}{\|x\|} \in F_B(s).$$

Thus, a norm automatically defines another norm (its dual); and the operation is symmetric: the dual of the dual norm is the norm itself.
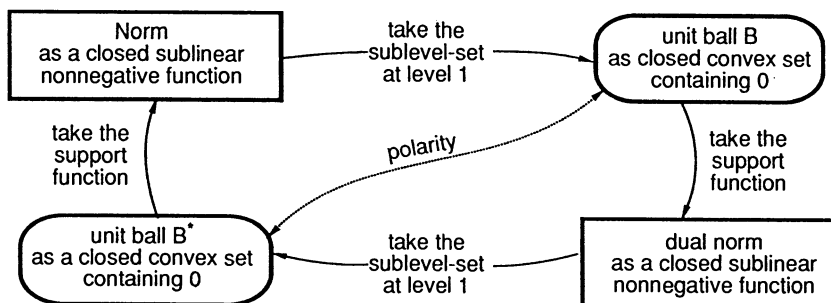
Fig. 3.2.1. Dual norms and polar sets

Remark 3.2.2 The operation (3.2.3) – (3.2.4) establishes a "duality" correspondence within a subclass of closed sublinear functions: those that are symmetric, finite everywhere, and positive (except at 0) – in short, norms.

This analytic operation has its counterpart in the geometric world: starting from a closed convex set which is symmetric, bounded and contains the origin as an interior point – in short, a "unit ball" – such as B, one constructs via gauges and support functions another closed convex set B^* which has the same properties. This correspondence is called *polarity*, demonstrated by Fig. 3.2.1: the polar (set) of B is

$$B^* := \{s \, : \, \langle s, x \rangle \leqslant 1 \text{ for all } x \in B\}. \tag{3.2.6}$$

As can be seen with a separation argument, the polar of B^* is symmetrically (proceed as for Theorem A.4.2.6 and remember that $0 \in B$)

$$(B^*)^* := \{x \, : \, \langle s, x \rangle \leqslant 1 \text{ for all } s \in B^*\} = B. \tag{3.2.7}$$
□

We leave it as an exercise to draw the unit balls of the ℓ_1- and ℓ_∞-norms on $\mathbb{R}^n$:

$$\|x\|_1 := \sum_{i=1}^{n} |x^i| \quad \text{and} \quad \|x\|_\infty := \max\{|x^1|, ..., |x^n|\}$$

(proceed as in Interpretation 2.1.5: a picture in $\mathbb{R}^n$ will do). Observe on the picture thus obtained that they are in polarity correspondence if the scalar product is the usual dot-product $\langle x, y \rangle = x^\top y$.

Another situation is illustrated by the "hexagonal norm" of Fig. 3.2.2. Observe how elongation in one direction corresponds to contraction for the polar. Also: a facet of one of the sets is exposed by a vertex in the polar.
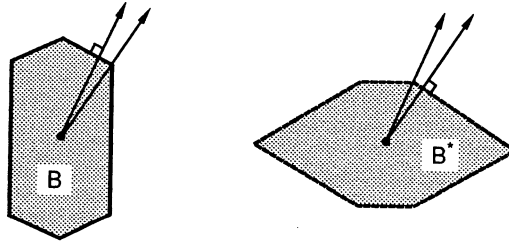
Fig. 3.2.2. Hexagonal unit-balls

Example 3.2.3 Other important norms are the quadratic norms, defined by

$$\|x\|_Q := \sqrt{\langle Qx, x \rangle}$$

where Q is a symmetric positive definite linear operator. They are important because they derive from a scalar product on $\mathbb{R}^n$, namely:

$$\langle x, y \rangle_Q := \langle Qx, y \rangle .$$

We refer to Example 2.3.4, more precisely formula (2.3.5), to compute the corresponding dual norm

$$(\|s\|_Q)^* = \sqrt{\langle s, Q^{-1}s \rangle} = \|s\|_{Q^{-1}} .$$

When $Q = I_n$, we get back the Euclidean norm $\langle \cdot, \cdot \rangle^{1/2}$. A comparison of (2.3.1) and (3.2.3) shows that it is self-dual: $\| \cdot \|^* = \| \cdot \|$. Among all the possible norms on $\mathbb{R}^n$, it is the only one having this property (once the scalar product is chosen!).   □

Actually, polarity neither relies upon symmetry, nor boundedness, nor on having 0 as an interior point. To take gauges and support functions resulting in (3.2.6), (3.2.7), the only important property is after all that 0 be in the closed convex set under consideration (B or B^*). In other words, the polarity relations (3.2.6), (3.2.7) establish an involution between sets that are merely closed convex, and contain the origin. More precisely, we have the following result:

Proposition 3.2.4 *Let C be a closed convex set containing the origin. Its gauge γ_C is the support function of a closed convex set containing the origin, namely*

$$C^\circ := \{s \in \mathbb{R}^n \ : \ \langle s, d \rangle \leqslant 1 \text{ for all } d \in C\}, \tag{3.2.8}$$

which defines the polar *(set) of C.*

Proof. We know that γ_C (which, by Theorem 1.2.5(i), is closed, sublinear and non-negative) is the support function of some closed convex set containing the origin, say D; from (3.1.1),

$$D = \{s \in \mathbb{R}^n \ : \ \langle s, d \rangle \leqslant r \text{ for all } (d, r) \in \text{epi } \gamma_C\}.$$

As seen in (1.2.4), epi γ_C is the closed convex conical hull of $C \times \{1\}$; we can use positive homogeneity to write

$$D = \{s \in \mathbb{R}^n \ : \ \langle s, d \rangle \leqslant 1 \text{ for all } d \text{ such that } \gamma_C(d) \leqslant 1\}.$$

In view of Theorem 1.2.5(iii), the above index-set is just C; in other words, $D = C^\circ$. $\square$
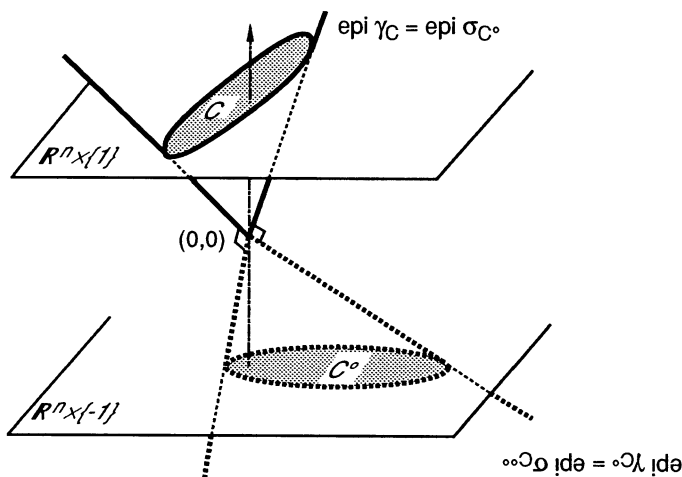
Fig. 3.2.3. Gauges and supports

Geometrically, the above proof is illustrated by Fig. 3.2.3, in which dual elements are drawn in dashed lines: $D = C^\circ$ is obtained by cutting the polar cone $(\text{epi } \gamma_C)^\circ$ at the level -1. Turn the picture upside down: cutting the polar cone $(\text{epi } \gamma_{C^\circ})^\circ$ at the level which has now become -1, we obtain $(C^\circ)^\circ$. But the polarity between closed convex cones is involutive: the picture shows that $(\text{epi } \gamma_{C^\circ})^\circ$ is our original cone epi γ_C. In other words, $C^{\circ\circ} = C$, Proposition 3.2.4 has its dual version:

Corollary 3.2.5 *Let C be a closed convex set containing the origin. Its support function σ_C is the gauge of C°.* $\square$

Remark 3.2.6 The elementary operation making up polarity is a one-to-one mapping between nonzero vectors and affine hyperplanes not containing the origin, via the equation inspired from (3.2.8):

$$s \mapsto H(s) := H_{s,1} = \{y \in \mathbb{R}^n : \langle s, y \rangle = 1\}. \tag{3.2.9}$$

Direct calculations show for example that the polar of the half-space

$$H^- := \{y = (\xi, \eta) \in \mathbb{R}^2 : \xi \leqslant 2\}$$

is the segment

$$(H^-)^\circ = \{(\rho, 0) : 0 \leqslant \rho \leqslant 1/2\}.$$

This simple example suggests the following comment: if σ is a given nonnegative closed sublinear function, it is the gauge of a set G which can be immediately constructed: along $0 \neq s \in \mathbb{R}^n$, plot the point $g(s) = s/\sigma(s) \in [0, +\infty]s$. Then G is the union of the segments $[0, g(s)]$, with s describing the unit sphere. If, along the same s, we plot the point $\sigma(s)s$, we likewise get a description of the set S supported by σ, but in a much less direct way: G is now *enveloped* by the affine hyperplane orthogonal to s and containing the point $\sigma(s)s$; now, *differentiation* is involved.
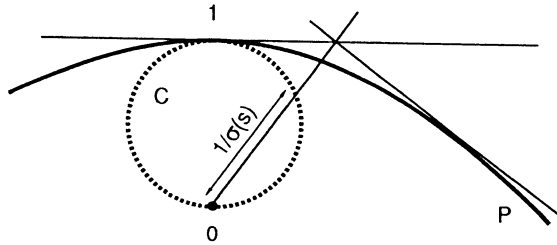
Fig. 3.2.4. Description of mutually polar sets

An expert in geometry will for example see on Fig. 3.2.4 that the polar of the circle

$$C = \{(\rho, \tau) : \rho^2 + (\tau - 1/2)^2 \leqslant 1/4\}$$

has a parabolic boundary. We leave it as an exercise to compute the gauge of C, and to realize that it is the support function of

$$P = \{(\xi, \eta) : \xi^2 \leqslant 1 - \eta\}.$$

Constructing a set from its gauge thus appears to be substantially easier than it is from its support function. Furthermore, to make a support function, we need a scalar product, while a gauge just needs an origin in $\mathbb{R}^n$. These advantages, however, are balanced by the rich calculus which can be developed with support functions, and which will be the subject of §3.3.   □

It is clear from (3.2.8) that $\langle s, d \rangle \leqslant 1$ for all $(d, s) \in C \times C^\circ$; this implies in particular that no nonzero $s \in C^\circ$ can be in the asymptotic cone of C. Furthermore, the property $\langle s, d \rangle = 1$ means that d exposes in C° a face $F_{C^\circ}(d)$ containing s; and s exposes likewise in $C^{\circ\circ} = C$ a face $F_C(s)$ containing d. Because the boundary of a closed convex set is described by its exposed faces (Proposition 3.1.5), the following result is then natural; compare it with Fig. 3.2.2.

Proposition 3.2.7 *Let C be a nonempty compact convex set having 0 in its interior, so that C° enjoys the same properties. Then, for all d and s in $\mathbb{R}^n$, the following statements are equivalent (the notation (3.2.9) is used)*

(i) *$H(s)$ is a supporting hyperplane to C at d;*

(ii) *$H(d)$ is a supporting hyperplane to C° at s;*

(iii) *$d \in \operatorname{bd} C$, $s \in \operatorname{bd} C^\circ$ and $\langle s, d \rangle = 1$;*

(iv) *$d \in C$, $s \in C^\circ$ and $\langle s, d \rangle = 1$.*

Proof. Left as an exercise; the assumptions are present to make sure that every nonzero vector in $\mathbb{R}^n$ does expose a face in each set. $\square$

Finally, suppose that C in (3.2.8) is a (closed convex) cone. By positive homogeneity, the number "1" can be replaced by any positive number, and even by "0" (remember the proof of Theorem 3.1.1). We recognize the definition of polarity between closed convex cones. Remembering Example 2.3.1, we see that, for a closed convex cone K, $\sigma_{K^\circ} = \gamma_K$, hence $\gamma_K = i_K$, which could be checked directly from Definition 1.2.4.

3.3 Calculus with Support Functions

From §1.3, the set of sublinear functions has a structure allowing calculus. Likewise, a calculus exists with subsets of $\mathbb{R}^n$. Then a natural question is: to what extent are these structures in correspondence via the supporting operation? In other words, to what extent is the supporting operation an isomorphism? The answer turns out to be very rich indeed.

We start with the order relation

Theorem 3.3.1 *Let S_1 and S_2 be nonempty closed convex sets; call σ_1 and σ_2 their support functions. Then*

$$S_1 \subset S_2 \quad \Longleftrightarrow \quad \sigma_1(d) \leqslant \sigma_2(d) \text{ for all } d \in \mathbb{R}^n .$$

Proof. Apply the equivalence stated in Corollary 3.1.2:

$$
\begin{aligned}
S_1 \subset S_2 &\Longleftrightarrow s \in S_2 \text{ for all } s \in S_1 \\
&\Longleftrightarrow \sigma_2(d) \geqslant \langle s, d \rangle \text{ for all } s \in S_1 \text{ and all } d \in \mathbb{R}^n \\
&\Longleftrightarrow \sigma_2(d) \geqslant \sup_{s \in S_1} \langle s, d \rangle \text{ for all } d \in \mathbb{R}^n .
\end{aligned}
$$
$\square$

In a way, the above result generalizes Theorem 2.2.2. The next statement goes with Propositions 1.3.1 and 1.3.2.

Theorem 3.3.2

(i) *Let σ_1 and σ_2 be the support functions of the nonempty closed convex sets S_1 and S_2. If t_1 and t_2 are positive, then*

$$t_1 \sigma_1 + t_2 \sigma_2 \text{ is the support function of } \operatorname{cl}(t_1 S_1 + t_2 S_2) .$$

(ii) *Let* $\{\sigma_j\}_{j\in J}$ *be the support functions of the family of nonempty closed convex sets* $\{S_j\}_{j\in J}$. *Then*

$$\sup_{j\in J}\sigma_j \text{ is the support function of } \overline{\mathrm{co}}\,\{\cup S_j \;:\; j \in J\}\,.$$

(iii) *Let* $\{\sigma_j\}_{j\in J}$ *be the support functions of the family of closed convex sets* $\{S_j\}_{j\in J}$. *If*

$$S := \bigcap_{j\in J} S_j \neq \emptyset\,,$$

then

$$\sigma_S = \overline{\mathrm{co}}\,\{\inf \sigma_j \;:\; j \in J\}\,.$$

Proof. [*(i)*] Call S the closed convex set $\mathrm{cl}(t_1 S_1 + t_2 S_2)$. By definition, its support function is

$$\sigma_S(d) = \sup\,\{\langle t_1 s_1 + t_2 s_2, d\rangle \;:\; s_1 \in S_1,\; s_2 \in S_2\}\,.$$

In the above expression, s_1 and s_2 run independently in their index sets S_1 and S_2, t_1 and t_2 are positive, so

$$\sigma_S(d) = t_1 \sup_{s\in S_1}\langle s, d\rangle + t_2 \sup_{s\in S_2}\langle s, d\rangle\,.$$

[*(ii)*] The support function of $S := \cup_{j\in J} S_j$ is

$$\sup_{s\in\cup S_j}\langle s, d\rangle = \sup_{j\in J}\,[\sup_{s_j\in S_j}\langle s_j, d\rangle] = \sup_{j\in J}\sigma_j(d)\,.$$

This implies (ii) since $\sigma_S = \sigma_{\overline{\mathrm{co}}\,S}$.

[*(iii)*] The set $S := \cap S_j$ being nonempty, it has a support function σ_S. Now, from Corollary 3.1.2,

$$\begin{aligned}
s \in S &\Longleftrightarrow s \in S_j \text{ for all } j \in J\\
&\Longleftrightarrow \langle s, \cdot\rangle \leqslant \sigma_j \text{ for all } j \in J\\
&\Longleftrightarrow \langle s, \cdot\rangle \leqslant \inf_{j\in J}\sigma_j \quad \Longleftrightarrow \quad \langle s, \cdot\rangle \leqslant \overline{\mathrm{co}}(\inf_{j\in J}\sigma_j)
\end{aligned}$$

where the last equivalence comes directly from the Definition B.2.5.3 of a closed convex hull. Again Corollary 3.1.2 tells us that the closed sublinear function $\overline{\mathrm{co}}(\inf \sigma_j)$ is just the support function of S.   □

Observe in (i) that, if S_2 is bounded, then $t_1 S_1 + t_2 S_2$ is automatically closed.

As for (iii), we have seen in Proposition 1.3.2(ii) that, if $J = \{1, \ldots, m\}$ is a finite set, then the "co" operation can be replaced by the infimal convolution: there holds

$$\sigma_{S_1\cap\ldots\cap S_m} = \mathrm{cl}\,(\sigma_1 \,\diamond\, \ldots \,\diamond\, \sigma_m)\,. \tag{3.3.1}$$

This last formula is a simplification of (iii), but the closure operation should not be forgotten, and it is something really complicated; these issues will be addressed more thoroughly in §E.2.3.

Returning to the end of Example 2.3.1, let K be a closed convex cone and take $K' := K \cap B(0,1)$. In view of the above observation, the support function of K' is given by an inf-convolution:

$$\sigma_{K'}(d) = \mathrm{cl}\left\{\inf_y[\sigma_K(y) + \sigma_B(d-y)]\right\}.$$

Since $\sigma_K = i_{K^\circ}$, the infimum forces y to be in K°, in which case σ_K vanishes; knowing that $\sigma_{B(0,1)} = \|\cdot\|$, the infimum is

$$\inf\left\{\|d-y\| \ : \ y \in K^\circ\right\}.$$

Here, we are in a favourable case: this infimum is actually a minimum – achieved at the projection $p_{K^\circ}(d)$ – and the result is a finite convex function, hence continuous; the closure operation is useless and can be omitted. In a word,

$$\sigma_{K \cap B(0,1)} = d_{K^\circ}. \tag{3.3.2}$$

Positive homogeneity can also be exploited in Theorem 3.3.2(i) to write

$$\sigma_{tS}(d) = \sigma_S(td) \quad \text{for all } d \in \mathbb{R}^n \text{ and } t > 0,$$

a formula which also holds for negative t (just write the definition). More generally:

Proposition 3.3.3 *Let $A : \mathbb{R}^n \to \mathbb{R}^m$ be a linear operator, with adjoint A^* (for some scalar product $\langle\!\langle \cdot, \cdot \rangle\!\rangle$ in $\mathbb{R}^m$). For $S \subset \mathbb{R}^n$ nonempty, we have*

$$\sigma_{\mathrm{cl}\,A(S)}(y) = \sigma_S(A^*y) \quad \text{for all } y \in \mathbb{R}^m.$$

Proof. Just write the definitions

$$\sigma_{A(S)}(y) = \sup_{s \in S} \langle\!\langle As, y \rangle\!\rangle = \sup_{s \in S} \langle s, A^*y \rangle$$

and use Proposition 2.2.1 to obtain the result. □

Taking an image-function (see §B.2.4) is another operation involving a linear operator. Its status is slightly more delicate.

Proposition 3.3.4 *Let $A : \mathbb{R}^m \to \mathbb{R}^n$ be a linear operator, with adjoint A^* (for some scalar product $\langle\!\langle \cdot, \cdot \rangle\!\rangle$ in $\mathbb{R}^m$). Let σ be the support function of a nonempty closed convex set $S \subset \mathbb{R}^m$. If σ is minorized on the inverse image*

$$\overset{-1}{A}(d) = \{p \in \mathbb{R}^m \ : \ Ap = d\} \tag{3.3.3}$$

of each $d \in \mathbb{R}^n$, then the support function of the set $\left(\overset{-1}{A^*}\right)(S)$ is the closure of the image-function $A\sigma$.

Proof. Our assumption is tailored to guarantee $A\sigma \in \mathrm{Conv}\,\mathbb{R}^n$ (Theorem B.2.4.2). The positive homogeneity of $A\sigma$ is clear: for $d \in \mathbb{R}^n$ and $t > 0$,

$$(A\sigma)(td) = \inf_{Ap=td} \sigma(p) = \inf_{A(p/t)=d} t\sigma(p/t) = t \inf_{Aq=d} \sigma(q) = t(A\sigma)(d).$$

Thus, the closed sublinear function $\mathrm{cl}\,(A\sigma)$ supports some set S'; by definition, $s \in S'$ if and only if

$$\langle s, d \rangle \leqslant \inf \{\sigma(p) \,:\, Ap = d\} \quad \text{for all } d \in \mathbb{R}^n \;;$$

but this just means

$$\langle s, Ap \rangle \leqslant \sigma(p) \quad \text{for all } p \in \mathbb{R}^m \;,$$

i.e. $A^* s \in S$, because $\langle s, Ap \rangle = \langle\!\langle A^* s, p \rangle\!\rangle$.   □

The inverse image $\big(\overset{-1}{A^*}\big)(S)$ of the closed set S under the continuous mapping A^* is closed. By contrast, $A\sigma$ need not be a closed function. As a particular case, suppose that S is bounded (σ_S is finite everywhere) and that A is surjective; then $A\sigma$ is finite everywhere as well, which means that $\big(\overset{-1}{A^*}\big)(S)$ is compact.

Remark 3.3.5 The assumption made in Proposition 3.3.4 means exactly that the function $A\sigma$ is nowhere $-\infty$; in other words, its closure $\mathrm{cl}\,(A\sigma)$ is the support function of a *nonempty* set: $\big(\overset{-1}{A^*}\big)(S) \neq \emptyset$. This last property can be rewritten as

$$S \cap \mathrm{Im}\,A^* \neq \emptyset \quad \text{or} \quad 0 \in S - \mathrm{Im}\,A^* = S + (\mathrm{Ker}\,A)^\perp \;. \tag{3.3.4}$$
    □

It has already been mentioned that taking an image-function is an important operation, from which several other operations can be constructed. We give two examples inspired from those at the end of §B.2.4:

– Let S_1 and S_2 be two nonempty closed convex sets of $\mathbb{R}^n$, with support functions σ_1 and σ_2 respectively. With $\mathbb{R}^m = \mathbb{R}^n \times \mathbb{R}^n$, take $A(x, y) := x + y$ and $\sigma(d_1, d_2) := \sigma_1(d_1) + \sigma_2(d_2)$; observe that σ is the support function of $S := S_1 \times S_2$, associated with the scalar product

$$\langle\!\langle (s_1, s_2), (d_1, d_2) \rangle\!\rangle := \langle s_1, d_1 \rangle + \langle s_2, d_2 \rangle \;.$$

Then we obtain $A\sigma = \sigma_1 \,\dot{+}\, \sigma_2$. On the other hand, the adjoint of A is clearly given by

$$A^* x = (x, x) \in \mathbb{R}^n \times \mathbb{R}^n \quad \text{for all } x \in \mathbb{R}^n \;,$$

so that the inverse image of S under A^* is nothing but $S_1 \cap S_2$: we recover (3.3.1).
– Let σ be the support function of some nonempty closed convex set $S \subset \mathbb{R}^n \times \mathbb{R}^p$ and let $A(x, y) := x$, so that the image of σ under A is defined by

$$\mathbb{R}^n \ni x \mapsto (A\sigma)(x) = \inf \{\sigma(x, y) \,:\, y \in \mathbb{R}^p\} \;.$$

Now A^* is

$$\mathbb{R}^n \ni x \mapsto A^* x = (x, 0) \in \mathbb{R}^n \times \mathbb{R}^p$$

and $\mathrm{cl}\,(A\sigma)$ is the support function of the slice $\{x \in \mathbb{R}^n \,:\, (x, 0) \in S\}$. This last set must not be confused with the projection of S onto $\mathbb{R}^n$, whose support function is $x \mapsto \sigma_S(x, 0)$ (Proposition 3.3.3).

Having studied the isomorphism with respect to order and algebraic structures, we pass to topologies. Theorem 1.3.3 has defined a distance Δ on the set of finite sublinear functions. Likewise, the *Hausdorff distance* Δ_H can be defined for nonempty closed sets (see §0.5). When restricted to nonempty compact convex sets, Δ_H plays the role of the distance introduced in Theorem 1.3.3:

Theorem 3.3.6 *Let S and S' be two nonempty compact convex sets of $\mathbb{R}^n$. Then*

$$\Delta(\sigma_S, \sigma_{S'}) := \max_{\|d\| \leqslant 1} |\sigma_S(d) - \sigma_{S'}(d)| = \Delta_H(S, S') . \qquad (3.3.5)$$

Proof. As mentioned in §0.5.1, for all $r \geqslant 0$, the property

$$\max \{ d_S(d) : d \in S' \} \leqslant r \qquad (3.3.6)$$

simply means $S' \subset S + B(0, r)$.

Now, the support function of $B(0, 1)$ is $\| \cdot \|$ – see (2.3.1). Calculus rules on support functions therefore tell us that (3.3.6) is also equivalent to

$$\sigma_{S'}(d) \leqslant \sigma_S(d) + r\|d\| \quad \text{for all } d \in \mathbb{R}^n ,$$

which in turn can be written

$$\max_{\|d\| \leqslant 1} \left[\sigma_{S'}(d) - \sigma_S(d) \right] \leqslant r .$$

In summary, we have proved

$$\max_{d \in S'} d_S(d) = \max_{\|d\| \leqslant 1} \left[\sigma_{S'}(d) - \sigma_S(d) \right]$$

and symmetrically

$$\max_{d \in S} d_{S'}(d) = \max_{\|d\| \leqslant 1} \left[\sigma_S(d) - \sigma_{S'}(d) \right] ;$$

the result follows. $\square$
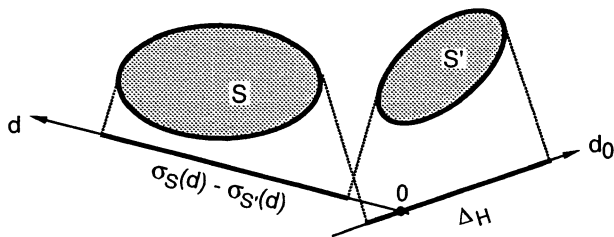
Fig. 3.3.1. Hausdorff distances

Naturally, the max in (3.3.5) is attained at some d_0: for S and S' convex compact, there exists d_0 of norm 1 such that

$$\Delta_H (S, S') = \Delta (\sigma_S, \sigma_{S'}) = |\sigma_S (d_0) - \sigma_{S'} (d_0)| \ .$$

Figure 3.3.1 illustrates a typical situation. When $S' = \{0\}$, we obtain the number

$$\Delta_H (\{0\}, S) = \max_{s \in S} \|s\| = \max_{\|d\|=1} \sigma_S (d) ,$$

already seen in (1.2.6); it is simply the distance from 0 to the most remote hyperplane $H_{d,\sigma_S(d)}$ touching S (see again the end of Interpretation 2.1.5).

Using (3.3.5), it becomes rather easy to compute the distance in Example 1.3.4, which becomes the Hausdorff distance (in fact an excess) between the corresponding unit balls.

When speaking of limits of nonempty convex compact sets to a nonempty convex compact set, the following result is a further illustration of our isomorphism.

Proposition 3.3.7 *A convex-compact-valued and locally bounded multifunction F : $\mathbb{R}^n \longrightarrow \mathbb{R}^n$ is outer [resp. inner] semi-continuous at $x_0 \in$ int dom F if and only if its support function $x \mapsto \sigma_{F(x)}(d)$ is upper [resp. lower] semi-continuous at x_0 for all d of norm 1.*

Proof. Calculus with support functions tells us that our definition (0.5.2) of outer semi-continuity is equivalent to

$$\forall \varepsilon > 0, \exists \delta > 0 : y \in B(x_0, \delta) \implies \sigma_{F(y)}(d) \leqslant \sigma_{F(x_0)}(d) + \varepsilon \|d\| \text{ for all } d \in \mathbb{R}^n$$

and division by $\|d\|$ shows that this is exactly upper semi-continuity of the support function for $\|d\| = 1$. Same proof for inner/lower semi-continuity.   □

Thus, a convex-compact-valued, locally bounded mapping F is both outer and inner semi-continuous at x_0 if and only if its support function $\sigma_{F(\cdot)}(d)$ is continuous at x_0 for all d. In view of Theorem 1.3.5, $\sigma_{F(\cdot)}(d)$ is then continuous at x_0 uniformly for $d \in B(0,1)$; and Theorem 3.3.6 tells us that this property in turn means:

$$\Delta_H (F(x), F(x_0)) \to 0 \quad \text{when} \quad x \to x_0 \ .$$

The following interpretation in terms of sequences is useful.

Corollary 3.3.8 *Let (S_k) be a sequence of nonempty convex compact sets and S a nonempty convex compact set. When $k \to +\infty$, the following are equivalent*

  (i) *$S_k \to S$ in the Hausdorff sense, i.e. $\Delta_H(S_k, S) \to 0$;*
 (ii) *$\sigma_{S_k} \to \sigma_S$ pointwise;*
(iii) *$\sigma_{S_k} \to \sigma_S$ uniformly on each compact set of $\mathbb{R}^n$.*   □

Let us sum up this Section 3.3: when combining/comparing closed convex sets, one knows what happens to their support functions (apply the results 3.3.1 – 3.3.3). Conversely, when closed sublinear functions are combined/compared, one knows what happens to the sets they support. The various rules involved are summarized in Table 3.3.1. Each S_j is a nonempty closed convex set, with support function σ_j.

This table deserves some comments.

- Generally speaking, it helps to remember that when a set increases, its support function increases (first line); hence the "crossing" of closed convex hulls in the last two lines.

- The rule of the last line comes directly from the definition (2.1.1) of a support function, if each S_j is thought of as a singleton.

- Most of these rules are still applicable without closed convexity of each S_j (remembering that $\sigma_S = \sigma_{\overline{co}\,S}$). For example, the equivalence in the first line requires closed convexity of S_2 only. We mention one trap, however: when intersecting sets, each set must be closed and convex. A counter-example in one dimension is $S_1 := \{0,1\}$, $S_2 := \{0,2\}$; the support functions do not see the difference between $S_1 \cap S_2 = \{0\}$ and $\mathrm{co}\,S_1 \cap \mathrm{co}\,S_2 = [0,1]$.

Table 3.3.1. Calculus rules for support functions

Closed convex sets	Closed sublinear functions
$S_1 \subset S_2$	$\sigma_1 \leqslant \sigma_2$
$\Delta_H(S_1, S_2)$ (S_i bounded)	$\Delta(\sigma_1, \sigma_2)$ (σ_i finite)
Hausdorff convergence (on bounded sets)	uniform/compact or pointwise convergence (on finite functions)
tS ($t > 0$)	$t\sigma$
$\mathrm{cl}\,(S_1 + S_2)$	$\sigma_1 + \sigma_2$
$\mathrm{cl}\,A(S)$ (A linear)	$\sigma \circ A^*$
$(\overset{-1}{A^*})(S)$ (A linear)	$\mathrm{cl}\,(A\sigma)$
$\cap_{j \in J} S_j$ (nonempty)	$\overline{co}\,(\inf_{j \in J} \sigma_j)$ (minorized)
$\overline{co}(\cup_{j \in J} S_j)$	$\sup_{j \in J} \sigma_j$

Example 3.3.9 (Maximal Eigenvalues) Recall from §B.1.3(e) that, if the eigenvalues of a symmetric matrix A are denoted by $\lambda_1(A) \geqslant \cdots \geqslant \lambda_n(A)$, the function

$$S_n(\mathbb{R}) \ni A \mapsto f_m(A) := \sum_{j=1}^{m} \lambda_j(A)$$

is convex – and finite everywhere. Its positive homogeneity is obvious, therefore it is the support function of a certain convex compact set C_m of symmetric matrices. Let us compute the set C_1 when the scalar product in $S_n(\mathbb{R})$ is the standard dot-product of $\mathbb{R}^{n \times n}$:

$$\langle\!\langle A, B \rangle\!\rangle := \mathrm{tr}\,AB = \sum_{i,j=1}^{n} A_{ij} B_{ij} \,.$$

Indeed, we know that

$$\lambda_1(A) = \sup_{x^\top x = 1} x^\top A x = \sup_{x^\top x = 1} \langle\!\langle xx^\top, A \rangle\!\rangle \,.$$

Hence C_1 is the closed convex hull of the set of matrices $\{xx^\top : x^\top x = 1\}$, which is clearly compact. Actually, its Hausdorff distance to $\{0\}$ is

$$\Delta_H(\{0\}, C_1) = \max_{x^\top x = 1} \sqrt{\langle\!\langle xx^\top, xx^\top \rangle\!\rangle} = 1 \,.$$

Incidentally, $\lambda_1(\cdot)$ is therefore nonexpansive in $S_n(\mathbb{R})$.

We leave it as an exercise to prove the following nicer representation of C_1:

$$C_1 = \operatorname{co} \{xx^\top : x^\top x = 1\} = \{M \in S_n(\mathbb{R}) : \lambda_n(M) \geqslant 0, \operatorname{tr} M = 1\} \,,$$

generalizing to $S_n(\mathbb{R})$ the expression of the unit-simplex in $\mathbb{R}^n$: C_1 could be called the unit *spectraplex*. $\qquad\square$

3.4 Example: Support Functions of Closed Convex Polyhedra

Polyhedral sets are encountered all the time, and thus deserve special study. They are often defined by finitely many affine constraints, i.e. obtained as intersections of closed half-spaces; in view of Table 3.3.1, this explains that the infimal convolution encountered in Proposition 1.3.2 is fairly important.

Example 3.4.1 (Compact Convex Polyhedra) First of all, the support function of a polyhedron defined as

$$P := \operatorname{co} \{p_1, \ldots, p_m\} \tag{3.4.1}$$

is trivially

$$d \mapsto \sigma_P(d) = \max \{\langle p_i, d \rangle : i = 1, \ldots, m\} \,.$$

There is no need to invoke Theorem 3.3.2 for this: a linear function $\langle \cdot, d \rangle$ attains its maximum on an extreme point of P (Proposition A.2.4.6), even if this extreme point is not the entire face exposed by d. $\qquad\square$

Example 3.4.2 (Closed Convex Polyhedral Cones) Taking again Example 2.3.1, suppose that the cone K is given as a finite intersection of half-spaces:

$$K = \cap\{K_j : j = 1, \ldots, m\} \,, \tag{3.4.2}$$

where

$$K_j := H^-_{a_j,0} := \{s \in \mathbb{R}^n : \langle a_j, s \rangle \leqslant 0\} \tag{3.4.3}$$

(the a_j's are assumed nonzero). We use Proposition 1.3.2:

$$\sigma_K(d) = \operatorname{cl inf} \Big\{ \sum_{j=1}^m \sigma_{K_j}(d_j) : \sum_{j=1}^m d_j = d \Big\} \,.$$

Only those d_j in K_j° – namely nonnegative multiples of a_j, see (2.3.2) – count to yield the infimum; their corresponding support vanishes and we obtain

$$\sigma_K(d) = \begin{cases} 0 & \text{if } d = \sum_{j=1}^m t_j a_j, \ t_j \geqslant 0 \text{ for } j = 1, \ldots, m, \\ +\infty & \text{otherwise} \,. \end{cases}$$

Here, we are lucky: the closure operation is useless because the righthand side is already a closed convex function. Note that we recognize Farkas' Lemma A.4.3.3: $K° = \text{dom } \sigma_K$ is the conical hull of the a_j's, which is closed thanks to the fact that there are *finitely many* generators.   □

Example 3.4.3 (Extreme Points and Directions) Suppose our polyhedron is defined in the spirit of 3.4.1, but unbounded:

$$S := \text{co}\{p_1, \ldots, p_m\} + \text{cone}\{a_1, \ldots, a_\ell\}.$$

Then it suffices to observe that $S = P + K°$, with P of (3.4.1) and K of (3.4.2), (3.4.3). Using Table 3.3.1 and knowing that $K°° = K$ – hence $\sigma_{K°} = i_K$:

$$\sigma_S(d) = \begin{cases} \max_{i=1,\ldots,m} \langle p_i, d \rangle & \text{if } \langle a_j, d \rangle \leqslant 0 \text{ for } j = 1, \ldots, \ell, \\ +\infty & \text{otherwise.} \end{cases}$$ □

The representations of Examples 3.4.1 and 3.4.3 are not encountered so frequently. Our next examples, dealing with half-spaces, represent the vast majority of situations.

Example 3.4.4 (Inequality Constraints) Perturb Example 3.4.2 to express the support function of $S := \cap H_{a_j, b_j}^-$, with

$$H_{a,b}^- := \{s \in \mathbb{R}^n : \langle s, a \rangle \leqslant b\} \quad (a \neq 0).$$

Here, we deal with translations of the K_j's: $H_{a_j, b_j}^- = \frac{b_j}{\|a_j\|^2} a_j + K_j$ so, with the help of Table 3.3.1:

$$\sigma_{H_{a_j, b_j}^-}(d) = \begin{cases} tb_j & \text{if } d = ta_j, \, t \geqslant 0, \\ +\infty & \text{otherwise.} \end{cases}$$

Provided that $S \neq \emptyset$, our support function σ_S is therefore the closure of the function

$$d \mapsto \begin{cases} \inf\left\{\sum_{j=1}^m t_j b_j : \sum_{j=1}^m t_j a_j = d, \, t_j \geqslant 0\right\} & \text{if } d \in \text{cone}\{a_1, \ldots, a_m\}, \\ +\infty & \text{otherwise.} \end{cases}$$ □

Now we have a sudden complication: the domain of σ is still the closed convex cone $K°$, but the status of the closure operation is no longer quite clear. Also, it is not even clear whether the above infimum is attained. Actually, all this results from Farkas' Lemma of §A.4.3; before giving the details, let us adopt different notation.

Example 3.4.5 (Closed Convex Polyhedra in Standard Form) Even though Example 3.4.4 uses the Definition A.4.2.5 of a closed convex polyhedron, the following "standard" description is often used. Let A be a linear operator from $\mathbb{R}^n$ to $\mathbb{R}^m$, $b \in \text{Im } A \subset \mathbb{R}^m$, $K \subset \mathbb{R}^n$ a closed convex polyhedral cone (K is usually characterized as in Example 3.4.2). Then S is given by

$$S := \{s \in \mathbb{R}^n \; : \; As = b, s \in K\} = (\{s_0\} + H) \cap K \,, \qquad (3.4.4)$$

where s_0 is some point in $\mathbb{R}^n$ satisfying $As_0 = b$, and $H := \operatorname{Ker} A$.

With the expression of σ_H given in Example 2.3.1, we see that the support function of $\{s_0\} + H$ is finite only on $\operatorname{Im} A^*$, where it is equal to

$$\sigma_{\{s_0\}}(d) + \sigma_H(d) = \langle s_0, d \rangle = \langle\!\langle b, z \rangle\!\rangle \quad \text{for } d = A^* z, \; z \in \mathbb{R}^m$$

(here, $\langle\!\langle \cdot, \cdot \rangle\!\rangle$ denotes the scalar product in $\mathbb{R}^m$). Thus, σ_S is the closure of the infimal convolution

$$\left(\sigma_{\{s_0\}} + \sigma_H\right) \mathbin{\mathpalette\@divot\relax} \sigma_K = \left(\sigma_{\{s_0\}} + \sigma_H\right) \mathbin{\mathpalette\@divot\relax} i_{K^\circ} \,, \qquad (3.4.5)$$

which can be made explicit as the function

$$d \mapsto \inf\left\{\langle\!\langle b, z \rangle\!\rangle \; : \; (z, y) \in \mathbb{R}^m \times K^\circ, \; A^* z + y = d\right\}.$$

Of course, this formula clearly reveals

$$\operatorname{dom} \sigma_S = \operatorname{dom} \sigma_H + \operatorname{dom} i_{K^\circ} = \operatorname{Im} A^* + K^\circ \,.$$

In the pure *standard form*, $\mathbb{R}^n$ and $\mathbb{R}^m$ are both equipped with the standard dot-product – A being a matrix with m rows and n columns – and K is the nonnegative orthant; K° is therefore the nonpositive orthant. Our "standard" S of (3.4.4) is now

$$\{s \in \mathbb{R}^n \; : \; As = b, \; s \geqslant 0\} \,, \qquad (3.4.6)$$

assumed nonempty. Then (3.4.5) becomes

$$\inf\{b^\top z \; : \; A^\top z \geqslant d\} \,, \qquad (3.4.7)$$

a function of d which is by no means simpler than in 3.4.4 – only the notation is different. In summary, the support function

$$\sigma_S(d) = \sup\{s^\top d \; : \; As = b, \; s \geqslant 0\} \qquad (3.4.8)$$

of the set (3.4.6) is the closure of (3.4.7), considered as a function of $d \in \mathbb{R}^n$.

Now, invoke Farkas' Lemma: write the equivalent statements (i)'' and (ii)'' from the end of §A.4, with (x, ρ, α, r) changed to $(-z, -d, s, -\sigma)$:

$$\{z \in \mathbb{R}^n \; : \; A^\top z \geqslant d\} \subset \{z \in \mathbb{R}^n \; : \; b^\top z \geqslant \sigma\} \qquad (3.4.9)$$

is equivalent to

$$\exists s \geqslant 0 \quad \text{such that} \quad As = b, \; s^\top d \geqslant \sigma \,. \qquad (3.4.10)$$

In other words: the largest σ for which (3.4.9) holds – i.e. the value (3.4.7) – is also the largest σ for which (3.4.10) holds – i.e. $\sigma_S(d)$. The closure operation can be omitted and we do have

$$\sigma_S(d) = \inf\{b^\top z \; : \; A^\top z \geqslant d\} \quad \text{for all } d \in \mathbb{R}^n \,.$$

Another interesting consequence can be noted. Take d such that $\sigma_S(d) < +\infty$: if we put $\sigma = \sigma_S(d)$ in (3.4.9), we obtain a true statement, i.e. (3.4.10) is also true. This means that the supremum in (3.4.8) is attained when it is finite.   □

It is worth noting that Example 3.4.5 describes general polyhedral functions, up to notational changes. As such, it discloses results of general interest, namely:

– A linear function which is bounded from below on a closed convex polyhedron attains its minimum on this polyhedron.
– The infimum of a linear function under affine constraints is a closed sublinear function of the righthand side; said otherwise, an image of a polyhedral function is closed: in Example 3.4.5, the polyhedral function in question is

$$\mathbb{R}^m \times \mathbb{R}^n \ni (y, z) \mapsto b^\top z + i_K(y),$$

and (3.4.7) gives its image under the linear mapping $[\, A^\top \,|\, 0 \,]$.

Exercises

1*. Let C_1 and C_2 be two nonempty closed convex sets in $\mathbb{R}^n$ and let S be bounded. Show that $C_1 + S = C_2 + S$ implies $C_1 = C_2$.

2**. Let P be a compact convex polyhedron on $\mathbb{R}^n$ with nonempty interior. Show that P has at least $n + 1$ facets.

3*. Let $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ be positively homogeneous, not identically $+\infty$ and minorized by a linear function. Show that $\overline{\mathrm{co}}\, f(x)$ is the supremum of $\sigma(x)$ over all closed sublinear functions σ minorizing f.

4*. Let f and g be two gauges. Show that $f \,\dot{\vee}\, g$ is still a gauge and that

$$\{x \,:\, (f \,\dot{\vee}\, g)(x) < 1\} = \mathrm{co}\left(\{x \,:\, f(x) < 1\} \cup \{x \,:\, g(x) < 1\}\right).$$

5**. Let C be a compact convex set with $0 \in \mathrm{int}\, C$, so that its polar set C° enjoys the same properties. Show that the following relations are equivalent:

  (i) the hyperplane $H_{s,1}$ supports C at $x \in C$;
  (ii) the hyperplane $H_{x,1}$ supports C° at $s \in C^\circ$;
 (iii) $x \in \mathrm{bd}\, C$, $s \in \mathrm{bd}\, C^\circ$, $\langle s, x \rangle = 1$.

6. Draw a picture to compute the support function σ_P of the parabolic set $P := \{(\xi, \eta) \in \mathbb{R}^2 \,:\, \eta \geqslant \frac{1}{2}\eta^2\}$. What is $\mathrm{dom}\, \sigma_P$? Show that σ_P is not upper semi-continuous on its domain.

7. Let $\|\cdot\|$ be a norm on $\mathbb{R}^n$ and denote by $S := \{x \in \mathbb{R}^n \,:\, \|x\| = 1\}$ the associated unit sphere. What is the convex hull of S?

8. Let H be a hyperplane in $\mathbb{R}^n$ and suppose the set S is contained in one of the corresponding half-spaces: $S \subset H_-$. Show that $\overline{\mathrm{co}}\,(S \cap H) = (\overline{\mathrm{co}}\, S) \cap H$. Compare with Exercise A.21.

9. Let $x \in C$, where C is a closed convex set in $\mathbb{R}^n$. Show that $x \in \mathrm{ri}\, C$ if and only if the normal cone $\mathrm{N}_C(x)$ is a subspace.

10. Let $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ and $A \subset \mathbb{R}^n$ be given (none of which is assumed convex). Assume f is minorized by some affine function and set $f_A := f + i_A$.

– Show that $\overline{co}\,\mathrm{epi}\,f_A \subset \{(x, r) \in \mathbb{R}^n \times \mathbb{R} : x \in \overline{co}\,A, \ r \geqslant \overline{co}\,f(x)\}$.
– Construct an example for which equality need not hold, even if A is closed convex.
– Show that equality does hold if f is affine.

11*. Recall that $M \succcurlyeq 0$ means: $M \in S_n(\mathbb{R})$ is positive semidefinite. Show that the *unit spectraplex* $\{M \in S_n(\mathbb{R}) : M \succcurlyeq 0, \mathrm{tr}\, M = 1\}$ is a convex compact set, whose support function is $M \mapsto \lambda_{\max}(M)$.

12*. In the Euclidean space of $n \times n$ real matrices (equipped with the scalar product $\langle M, P \rangle = \mathrm{tr}\,(M^\top P) = \sum_{ij} M_{ij} P_{ij}$ and associated norme $\|\cdot\|$), denote by S the set of orthogonal matrices ($M^\top M = M M^\top = I_n$).
– Show that S is contained in the sphere $\widetilde{B}(0, \sqrt{n}) = \{M : \|M\| = \sqrt{n}\}$; deduce that S is compact.
– Show that the support function of S is $M \mapsto \sigma_S(M) = \mathrm{tr}\,(M^\top M)^{1/2}$ [hint: use the polar decomposition of M].
– Express the values of the support functions $\sigma_S(M)$ and $\sigma_{B(0,\sqrt{n})}(M)$ in terms of the eigenvalues of $M^\top M$.

13. Let $K := \{M \in S_n(\mathbb{R}) : M \succcurlyeq 0\}$ be the cone of symmetric positive semidefinite matrices. What are its interior $\mathrm{int}\,K$ and boundary $\mathrm{bd}\,K$?
    Show that the distance of a symmetric positive semidefinite matrix to the boundary of K is its smallest eigenvalue: $d_{\mathrm{bd}\,K}(M) = \lambda_{\min}(M)$ for $M \in K$.

14*. Let $(S_k)_k$ be a nested family of compact sets in $\mathbb{R}^n$ ($S_{k+1} \subset S_k$ for all k). Show that $\mathrm{co}\,(\cap_k S_k) = \cap_k \mathrm{co}\,S_k$. Exhibit an example showing that the result becomes wrong if the S_k's are unbounded.

15*. Let $F : \mathbb{R} \longrightarrow \mathbb{R}^n$ be a multifunction such that $F(t)$ is closed convex for each t. Show that the graph of F is a convex set in $\mathbb{R} \times \mathbb{R}^n$ if and only if the function $t \mapsto \sigma_{F(t)}(d)$ is concave for all $d \in \mathbb{R}^n$.
    Let C be closed convex in $\mathbb{R}^n$ and $\sigma' : \mathbb{R}^n \to \mathbb{R}$ be nonnegative and positively homogeneous: $+\infty > \sigma'(\lambda d) = \lambda \sigma'(d) \geqslant 0$ for all $(\lambda, d) \in \mathbb{R}_+ \times \mathbb{R}^n$. Show that, for any $t \geqslant 0$, the set $F(t) := \{x \in \mathbb{R}^n : \langle \cdot, x \rangle \leqslant \sigma_C(\cdot) + t\sigma'(\cdot)\}$ is closed convex and contains C. What is its support function? Show that this defines a multifunction $F : \mathbb{R}_+ \longrightarrow \mathbb{R}^n$ whose graph is convex.

D. Subdifferentials of Finite Convex Functions

Introduction We have mentioned in our preamble to Chap. C that sublinearity permits the approximation of convex functions to first order around a given point. In fact, we will show here that, if $f : \mathbb{R}^n \to \mathbb{R}$ is convex and $x \in \mathbb{R}^n$ is fixed, then the function

$$d \mapsto f'(x, d) := \lim_{t \downarrow 0} \frac{f(x + td) - f(x)}{t}$$

exists and is finite *sublinear*. Furthermore, f' *approximates* f around x in the sense that

$$f(x + h) = f(x) + f'(x, h) + o(\|h\|) \, . \tag{0.1}$$

In view of the correspondence between finite sublinear functions and compact convex sets (which formed a large part of Chap. C), $f'(x, \cdot)$ can be expressed for all $d \in \mathbb{R}^n$ as

$$f'(x, d) = \sigma_S(d) = \max \left\{ \langle s, d \rangle \, : \, s \in S \right\}$$

for some nonempty compact convex set S. This S is called the *subdifferential* of f at x and is traditionally denoted by $\partial f(x)$. When f is differentiable at x, with gradient $\nabla f(x)$, (0.1) shows that $f'(x, \cdot)$ becomes linear and S contains only the element $\nabla f(x)$. Thus, the concept of subdifferential generalizes that of gradient, just as sublinearity generalizes linearity.

The "subdifferentiation" thus introduced is supposed to generalize the ordinary differentiation; one should therefore not be surprised to find counterparts of most of the results encountered in differential calculus: first-order Taylor expansions, mean-value theorems, calculus rules, etc. The importance of calculus rules increases in the framework of convex analysis: some operations on convex functions destroy differentiability (and thereby find no place in differential calculus) but preserve convexity. An important example is the max-operation; indeed, we will give a detailed account of the calculus rules for the subdifferential of max-functions.

This chapter deals with *finite-valued* convex functions exclusively: it is essential for practitioners to have a good command of subdifferential calculus, and this framework is good enough. Furthermore, its generalization to the extended-valued case then becomes easier to assimilate. Unless otherwise specified, therefore:

$$\boxed{f : \ \mathbb{R}^n \to \mathbb{R} \text{ is convex} \, .}$$

This implies the continuity and local Lipschitz continuity of f. We note from (0.1), however, that the concept of subdifferential is essentially local; for an extended-valued f, most results in this chapter remain true at a point $x \in \text{int dom} f$ (assumed nonempty). It can be considered as an exercise to check those results in this generalized setting.

1 The Subdifferential: Definitions and Interpretations

1.1 First Definition: Directional Derivatives

Let x and d be fixed in $\mathbb{R}^n$ and consider the difference quotient of f at x in the direction d:

$$q(t) := \frac{f(x + td) - f(x)}{t} \quad \text{for } t > 0. \tag{1.1.1}$$

The function $t \mapsto q(t)$ is:
– increasing; this is Theorem 0.6.1, the criterion of increasing slopes;
– bounded near 0; this comes from the local Lipschitz property of f (§B.3.1).
For t decreasing to 0, $q(t)$ has therefore a limit and the following definition makes sense.

Definition 1.1.1 (Directional Derivative) The *directional derivative* of f at x in the direction d is

$$f'(x, d) := \lim \{q(t) : t \downarrow 0\} = \inf \{q(t) : t > 0\}. \tag{1.1.2}$$
□

If φ denotes the one-dimensional function $t \mapsto \varphi(t) := f(x + td)$, then

$$f'(x, d) = \mathrm{D}_+\varphi(0) \tag{1.1.3}$$

is nothing other than the right-derivative of φ at 0 (Theorem 0.6.3). Changing d to $-d$ in (1.1.1), one obtains

$$f'(x, -d) = \lim_{t \downarrow 0} \frac{f(x - td) - f(x)}{t} = \lim_{\tau \uparrow 0} \frac{f(x + \tau d) - f(x)}{-\tau}$$

which is *not* the left-derivative of φ at 0 but rather its negative counterpart:

$$f'(x, -d) = -\mathrm{D}_-\varphi(0). \tag{1.1.4}$$

Proposition 1.1.2 *For fixed x, the function $f'(x, \cdot)$ is finite sublinear.*

Proof. Let d_1, d_2 in $\mathbb{R}^n$, and positive α_1, α_2 with $\alpha_1 + \alpha_2 = 1$. From the convexity of f:
$$f(x + t(\alpha_1 d_1 + \alpha_2 d_2)) - f(x) =$$
$$f(\alpha_1(x + td_1) + \alpha_2(x + td_2)) - \alpha_1 f(x) - \alpha_2 f(x) \leqslant$$
$$\leqslant \alpha_1[f(x + td_1) - f(x)] + \alpha_2[f(x + td_2) - f(x)]$$

for all t. Dividing by $t > 0$ and letting $t \downarrow 0$, we obtain

$$f'(x, \alpha_1 d_1 + \alpha_2 d_2) \leqslant \alpha_1 f'(x, d_1) + \alpha_2 f'(x, d_2)$$

which establishes the convexity of f' with respect to d. Its positive homogeneity is clear: for $\lambda > 0$

$$f'(x, \lambda d) = \lim_{t \downarrow 0} \lambda \frac{f(x + \lambda t d) - f(x)}{\lambda t} = \lambda \lim_{\tau \downarrow 0} \frac{f(x + \tau d) - f(x)}{\tau} = \lambda f'(x, d) .$$

Finally suppose $\|d\| = 1$. As a finite convex function, f is Lipschitz continuous around x (Theorem B.3.1.2); in particular there exist $\varepsilon > 0$ and $L > 0$ such that

$$|f(x + td) - f(x)| \leqslant Lt \quad \text{for } 0 \leqslant t \leqslant \varepsilon .$$

Hence, $|f'(x, d)| \leqslant L$ and we conclude with positive homogeneity:

$$|f'(x, d)| \leqslant L\|d\| \quad \text{for all } d \in \mathbb{R}^n . \tag{1.1.5}$$

$\square$

Remark 1.1.3 From the end of the above proof, a local Lipschitz constant L of f around x is transferred to $f'(x, \cdot)$ via (1.1.5). In view of (C.1.2.6), this same L is a *global* Lipschitz constant for $f'(x, \cdot)$. This is even true of $f'(y, \cdot)$ for y close to x: with δ and L such that f has the Lipschitz constant L on $B(x, \delta)$,

$$\|y - x\| < \delta \quad \Longrightarrow \quad |f'(y, d_1) - f'(y, d_2)| \leqslant L\|d_1 - d_2\| \quad \text{for all } d_1, d_2 \in \mathbb{R}^n . \quad \square$$

A consequence of Proposition 1.1.2 is that $f'(x, \cdot)$ is a support function, so the following suggests itself:

Definition 1.1.4 (Subdifferential I) The subdifferential $\partial f(x)$ of f at x is the nonempty compact convex set of $\mathbb{R}^n$ whose support function is $f'(x, \cdot)$, i.e.

$$\partial f(x) := \{s \in \mathbb{R}^n : \langle s, d \rangle \leqslant f'(x, d) \text{ for all } d \in \mathbb{R}^n\} . \tag{1.1.6}$$

A vector $s \in \partial f(x)$ is called a *subgradient* of f at x. $\qquad\qquad \square$

A first observation is therefore that the concept of subdifferential is *attached* to a scalar product, just because the concept of support is so; changing the scalar product changes ∂f.

All the properties of the correspondence between compact convex sets and finite sublinear functions can be reformulated for $\partial f(x)$ and $f'(x, \cdot)$. For example, the breadth of $\partial f(x)$ (cf. Definition C.2.1.4) along a normalized direction d is

$$f'(x, d) + f'(x, -d) = D_+ \varphi(0) - D_- \varphi(0) \geqslant 0$$

and represents the "lack of differentiability" of the function φ alluded to in (1.1.3), (1.1.4); remember Proposition B.4.2.1.

Remark 1.1.5 In particular, for d in the subspace U of linearity of $f'(x, \cdot)$ – the lineality space of $f'(x, \cdot)$, see (C.1.1.8) –, the corresponding φ is differentiable at 0. The restriction of $f'(x, \cdot)$ to U is linear (Proposition C.1.1.6) and equals $\langle s, \cdot \rangle$, no matter how s is chosen in $\partial f(x)$. In words, U is the set of h for which $h \mapsto f(x + h)$ behaves as a function differentiable at $h = 0$. See Fig. 1.1.1: $\partial f(x)$ is entirely contained in a hyperplane parallel to U; said otherwise, U is the set of directions along which $\partial f(x)$ has 0-breadth.

We also recall from Definition C.2.1.4 that

$$-f'(x, -d) \leqslant \langle s, d \rangle \leqslant f'(x, d) \quad \text{for all } (s, d) \in \partial f(x) \times \mathbb{R}^n . \qquad \Box$$
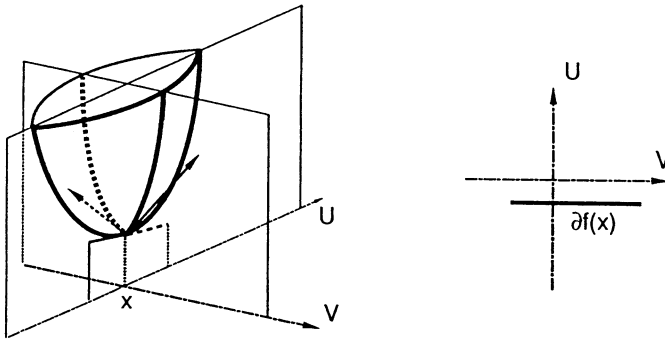
Fig. 1.1.1. Linearity-space of the directional derivative

It results directly from Chap. C that Definition 1.1.4 can also be looked at from the other side: (1.1.6) is equivalent to $f'(x, d) = \sup \{\langle s, d \rangle : s \in \partial f(x)\}$. Remembering that $\partial f(x)$ is compact, this supremum is attained at some s – which depends on d! In other words: for any $d \in \mathbb{R}^n$, there is some $s_d \in \partial f(x)$ such that

$$f(x + td) = f(x) + t\langle s_d, d \rangle + t\varepsilon_d(t) \quad \text{for } t \geqslant 0 . \qquad (1.1.7)$$

Here $\varepsilon_d(t) \to 0$ for $t \downarrow 0$, and we will see later that ε_d can actually be made independent of the normalized d; as for s_d, it is a subgradient giving the largest $\langle s, d \rangle$. Thus, from its very construction, the subdifferential contains all the necessary information for a first-order description of f.

As a finite convex function, $d \mapsto f'(x, d)$ has itself directional derivatives and subdifferentials. These objects at $d = 0$ are of particular interest; the case $d \neq 0$ will be considered later.

Proposition 1.1.6 *The finite sublinear function* $d \mapsto \sigma(d) := f'(x, d)$ *satisfies*

$$\sigma'(0, \delta) = f'(x, \delta) \qquad \text{for all } \delta \in \mathbb{R}^n ; \qquad (1.1.8)$$

$$\sigma(\delta) = \sigma(0) + \sigma'(0, \delta) = \sigma'(0, \delta) \text{ for all } \delta \in \mathbb{R}^n ; \qquad (1.1.9)$$

$$\partial\sigma(0) = \partial f(x) . \qquad (1.1.10)$$

Proof. Because σ is positively homogeneous and $\sigma(0) = 0$,

$$\frac{\sigma(t\delta) - \sigma(0)}{t} = \sigma(\delta) = f'(x,\delta) \quad \text{for all } t > 0\,.$$

This implies immediately (1.1.8) and (1.1.9). Then (1.1.10) follows from uniqueness of the supported set. □

One should not be astonished by (1.1.8): tangency is a self-reproducing operation. Since the graph of $f'(x,\cdot)$ is made up of the (half-)lines tangent to gr f at $(x, f(x))$, the same set must be obtained when taking the (half-)lines tangent to gr $f'(x,\cdot)$. As for (1.1.9), it simply expresses that, when developing a sublinear function to first order at 0, there is no error of linearization: (1.1.7) holds with $\varepsilon_d \equiv 0$ in that case.

1.2 Second Definition: Minorization by Affine Functions

The previous Definition 1.1.4 of the subdifferential involved two steps: first, calculating the directional derivative, and then determining the set that it supports. It is however possible to give a direct definition, with no reference to differentiation.

Definition 1.2.1 (Subdifferential II) The subdifferential of f at x is the set of vectors $s \in \mathbb{R}^n$ satisfying

$$f(y) \geqslant f(x) + \langle s, y - x \rangle \quad \text{for all } y \in \mathbb{R}^n \,. \tag{1.2.1}$$

□

Of course, we have to prove that our new definition coincides with 1.1.4. This will be done in Theorem 1.2.2 below. First, we make a few remarks illustrating the difference between Definitions 1.1.4 and 1.2.1.

– The present definition is *unilateral*: an inequality is required in (1.2.1), expressing the fact that the affine function $y \mapsto f(x) + \langle s, y - x \rangle$ *minorizes* f and *coincides* with f for $y = x$.
– It is a *global* definition, in the sense that (1.2.1) involves all y in $\mathbb{R}^n$.
– These two observations do suggest that 1.2.1 deviates from the concept of differentiation, namely:

(i) no remainder term shows up in (1.2.1), and
(ii) every y counts, not only those close to x.

Actually, the proof below will show that nothing changes if:

(i') an extra $o(\|y - x\|)$ is added in (1.2.1), or
(ii') (1.2.1) is required to hold for y close to x only.

Of course, these two properties (i') and (ii') rely on convexity of f; more precisely on monotonicity of the difference quotient.

– All subgradients are described by (1.2.1) at the same time. By contrast, $f'(x,d) = \langle s_d, d \rangle$ plots, for $d \neq 0$, only the boundary of $\partial f(x)$, one exposed face at a time. The whole subdifferential is then obtained by convexification – remember Proposition C.3.1.5.

Theorem 1.2.2 *The definitions 1.1.4 and 1.2.1 are equivalent.*

Proof. Let s satisfy (1.1.6), i.e.

$$\langle s, d \rangle \leqslant f'(x,d) \quad \text{for all } d \in \mathbb{R}^n . \tag{1.2.2}$$

The second equality in (1.1.2) makes it clear that (1.2.2) is equivalent to

$$\langle s, d \rangle \leqslant \frac{f(x+td) - f(x)}{t} \quad \text{for all } d \in \mathbb{R}^n \text{ and } t > 0 . \tag{1.2.3}$$

When d describes $\mathbb{R}^n$ and t describes $\mathbb{R}_*^+$, $y := x + td$ describes $\mathbb{R}^n$ and we realize that (1.2.3) is just (1.2.1). □

The above proof is deeper than it looks: because the difference quotient is increasing, the inequality of (1.2.3) holds whenever it holds for all $(d,t) \in B(0,1) \times]0, \varepsilon]$. Alternatively, this means that nothing is changed in (1.2.1) if y is restricted to a neighborhood of x.

It is interesting to note that, in terms of first-order approximation of f, (1.2.1) brings some additional information to (1.1.7): it says that the remainder term $\varepsilon_d(t)$ is nonnegative for all $t \geqslant 0$. On the other hand, (1.1.7) says that, for some specific s (depending on y), (1.2.1) holds almost as an equality for y close to x.

Now, the path "directional derivative $\to$ subdifferential" adopted in §1.1 can be reproduced backwards: the set defined in (1.2.1) is

– nonempty (Proposition B.1.2.1),
– closed and convex (immediate from the definitions),
– bounded, due to a simple Lipschitz argument: for given $0 \neq s \in \partial f(x)$, take in (1.2.1) $y = x + \delta s/\|s\|$ ($\delta > 0$ arbitrary) to obtain

$$f(x) + L\delta \geqslant f(y) \geqslant f(x) + \delta\|s\| ,$$

where the first inequality comes from the Lipschitz property B.3.1.2, written on the compact set $B(x, \delta)$.

As a result, this set of (1.2.1) has a finite-valued support function. Theorem 1.2.2 simply tells us that this support function is precisely the directional derivative $f'(x, \cdot)$ of (1.1.2).

Remark 1.2.3 A (finite) sublinear function σ has a subdifferential, just as any other convex function. Its subdifferential at 0 is defined by

$$\partial\sigma(0) = \{s : \langle s, d \rangle \leqslant \sigma(d) \text{ for all } d \in \mathbb{R}^n \} ,$$

in which we recognize Theorem C.2.2.2. This permits a more compact way than (C.2.2.2) to construct a set from its support function: a (finite) sublinear function *is the support* of its subdifferential at 0:

$$\sigma(d) = \sup\{\langle s, d\rangle \ : \ s \in \partial\sigma(0)\}.$$

In Fig. C.2.2.1, for example, the wording "filter with $\langle s, \cdot\rangle \leqslant \sigma?$" can be replaced by the more elegant "take the subdifferential of σ at 0". ☐

1.3 Geometric Constructions and Interpretations

Definition 1.2.1 means that the elements of $\partial f(x)$ are the slopes of the hyperplanes supporting the epigraph of f at $(x, f(x)) \in \mathbb{R}^n \times \mathbb{R}$. In terms of tangent and normal cones, this is expressed by the following result, which could serve as a third definition of the subdifferential and directional derivative.

Proposition 1.3.1

(i) *A vector $s \in \mathbb{R}^n$ is a subgradient of f at x if and only if $(s, -1) \in \mathbb{R}^n \times \mathbb{R}$ is normal to* epi f *at $(x, f(x))$. In other words:*

$$N_{\text{epi} f}(x, f(x)) = \{(\lambda s, -\lambda) \ : \ s \in \partial f(x), \lambda \geqslant 0\}.$$

(ii) *The tangent cone to the set* epi f *at $(x, f(x))$ is the epigraph of the directional-derivative function $d \mapsto f'(x, d)$:*

$$T_{\text{epi} f}(x, f(x)) = \{(d, r) \ : \ r \geqslant f'(x, d)\}.$$

Proof. [(*i*)] Apply Definition A.5.2.3 to see that $(s, -1) \in N_{\text{epi} f}(x, f(x))$ means

$$\langle s, y - x\rangle + (-1)[r - f(x)] \leqslant 0 \quad \text{for all } y \in \mathbb{R}^n \text{ and } r \geqslant f(y)$$

and the equivalence with (1.2.1) is clear. The formula follows since the set of normals forms a cone containing the origin.

[(*ii*)] The tangent cone to epi f is the polar of the above normal cone, i.e. the set of $(d, r) \in \mathbb{R}^n \times \mathbb{R}$ such that

$$\langle \lambda s, d\rangle + (-\lambda)r \leqslant 0 \quad \text{for all } s \in \partial f(x) \text{ and } \lambda \geqslant 0.$$

Barring the trivial case $\lambda = 0$, we divide by $\lambda > 0$ to obtain

$$r \geqslant \max\{\langle s, d\rangle \ : \ s \in \partial f(x)\} = f'(x, d).$$ ☐

Figure 1.3.1 illustrates this result. The right part of the picture represents the normal cone N and tangent cone T to epi f at $(x, f(x))$. The intersection of N with the space $\mathbb{R}^n$ at level -1 is just $\partial f(x) \times \{-1\}$. On the left part of the picture, the origin is translated to $(x, f(x))$ and the translated T is tangent to epi f. Note that the boundary of $T + (x, f(x))$ is also a (nonconvex) cone, "tangent", in the intuitive
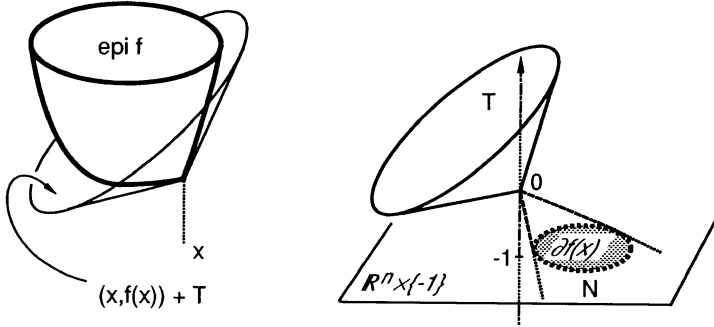
Fig. 1.3.1. Tangents and normals to the epigraph

sense of the term, to the graph of f; it is the graph of $f'(x, \cdot)$, translated at $(x, f(x))$. Compare Figs. A.5.1.1 and 1.3.1.

Proposition 1.3.1 and its associated Fig. 1.3.1 refer to Interpretation C.2.1.6, with a supported set drawn in $\mathbb{R}^n \times \mathbb{R}$. One can also use Interpretation C.2.1.5, in which the supported set was drawn in $\mathbb{R}^n$. In this framework, the sublevel-set passing through x

$$Sf(x) := S_{f(x)}(f) = \{y \in \mathbb{R}^n \ : \ f(y) \leqslant f(x)\} \tag{1.3.1}$$

is particularly interesting, and is closely related to $\partial f(x)$:

Lemma 1.3.2 *For the convex function* $f: \mathbb{R}^n \to \mathbb{R}$ *and the sublevel-set* (1.3.1)*, we have*

$$T_{Sf(x)}(x) \subset \{d : f'(x, d) \leqslant 0\} . \tag{1.3.2}$$

Proof. Take arbitrary $y \in Sf(x)$, $t > 0$, and set $d := t(y - x)$. Then, using the second equality in (1.1.2),

$$0 \geqslant t[f(y) - f(x)] = \frac{f(x + d/t) - f(x)}{1/t} \geqslant f'(x, d) .$$

So we have proved

$$\mathbb{R}^+ [Sf(x) - x] \subset \{d \ : \ f'(x, d) \leqslant 0\} \tag{1.3.3}$$

(note: the case $d = 0$ is covered since $0 \in Sf(x) - x$).

Because $f'(x, \cdot)$ is a closed function, the righthand set in (1.3.3) is closed. Knowing that $T_{Sf(x)}(x)$ is the closure of the lefthand side in (1.3.3) (Proposition A.5.2.1), we deduce the result by taking the closure of both sides in (1.3.3). □

The converse inclusion in (1.3.2) need not hold: for a counter-example, take $f(x) = 1/2 \|x\|^2$. The sublevel-set $Sf(0)$ is then $\{0\}$ and $f'(0, d) = 0$ for all d. In this case, (1.3.2) reads $\{0\} \subset \mathbb{R}^n$! To prove the converse inclusion, an additional assumption is definitely needed – for example the one considered in the following technical result.

Proposition 1.3.3 *Let $g : \mathbb{R}^n \to \mathbb{R}$ be convex and suppose that $g(x_0) < 0$ for some $x_0 \in \mathbb{R}^n$. Then*

$$\operatorname{cl} \{z \: : \: g(z) < 0\} = \{z \: : \: g(z) \leqslant 0\}, \tag{1.3.4}$$

$$\{z \: : \: g(z) < 0\} = \operatorname{int} \{z \: : \: g(z) \leqslant 0\}. \tag{1.3.5}$$

It follows

$$\operatorname{bd} \{z \: : \: g(z) \leqslant 0\} = \{z \: : \: g(z) = 0\}. \tag{1.3.6}$$

Proof. Because g is (lower semi-) continuous, the inclusion "$\subset$" automatically holds in (1.3.4). Conversely, let $\bar{z}$ be arbitrary with $g(\bar{z}) \leqslant 0$ and, for $k > 0$, set

$$z_k := \tfrac{1}{k}x_0 + (1 - \tfrac{1}{k})\bar{z}.$$

By convexity of g, $g(z_k) < 0$, so (1.3.4) is established by letting $k \to +\infty$.

Now, take the interior of both sides in (1.3.4). The "int cl" on the left is actually an "int" (Proposition A.2.1.8), and this "int"-operation is useless because g is (upper semi-) continuous: (1.3.5) is established. □

The existence of x_0 in this result is indeed a qualification assumption, often called *Slater assumption*. When this x_0 exists, taking closures, interiors and boundaries of sublevel-sets amounts to imposing "$\leqslant$", "$<$" and "$=$" in their definitions. Needless to say, convexity is essential for such an equivalence: with $n = 1$, think of $g(z) := \min \{0, |z| - 1\}$.

We are now in a position to characterize the tangential elements to a sublevel-set.

Theorem 1.3.4 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and suppose $0 \notin \partial f(x)$. Then, $Sf(x)$ being the sublevel-set (1.3.1),*

$$\mathrm{T}_{Sf(x)}(x) = \{d \in \mathbb{R}^n \: : \: f'(x,d) \leqslant 0\} \tag{1.3.7}$$

$$\operatorname{int} [\mathrm{T}_{Sf(x)}(x)] = \{d \in \mathbb{R}^n \: : \: f'(x,d) < 0\} \neq \emptyset. \tag{1.3.8}$$

Proof. From the very definition (1.1.6), our assumption means that $f'(x,d) < 0$ for some d, and (1.1.2) then implies that $f(x + td) < f(x)$ for $t > 0$ small enough: our d is of the form $(x + td - x)/t$ with $x + td \in Sf(x)$ and we have proved

$$\{d \: : \: f'(x,d) < 0\} \subset \mathbb{R}^+ [Sf(x) - x] \subset \mathrm{T}_{Sf(x)}(x). \tag{1.3.9}$$

Now, we can apply (1.3.4) with $g = f'(x, \cdot)$:

$$\operatorname{cl} \{d \: : \: f'(x,d) < 0\} = \{d \: : \: f'(x,d) \leqslant 0\},$$

so (1.3.7) is proved by closing the sets in (1.3.9) and using (1.3.2). Finally, take the interior of both sides in (1.3.7) and apply (1.3.5) with $g = f'(x, \cdot)$ to prove (1.3.8). □

The above result can be formulated in terms of normal cones.

Theorem 1.3.5 *Let* $f : \mathbb{R}^n \to \mathbb{R}$ *be convex and suppose* $0 \notin \partial f(x)$. *Then a direction* d *is normal to* $Sf(x)$ *at* x *if and only if there is some* $t \geqslant 0$ *and some* $s \in \partial f(x)$ *such that* $d = ts$:
$$N_{Sf(x)}(x) = \mathbb{R}^+ \partial f(x).$$

Proof. Write (1.3.7) as
$$T_{Sf(x)}(x) = \{d \in \mathbb{R}^n : \langle s, d \rangle \leqslant 0 \text{ for all } s \in \partial f(x)\}$$
$$= \{d \in \mathbb{R}^n : \langle \lambda s, d \rangle \leqslant 0 \text{ for all } \lambda \geqslant 0 \text{ and } s \in \partial f(x)\} = [\mathbb{R}^+ \partial f(x)]^\circ.$$

The result follows by taking the polar cone of both sides, and observing that the assumption implies closedness of $\mathbb{R}^+ \partial f(x)$ (Proposition A.1.4.7):
$$N_{Sf(x)}(x) = \mathrm{cl}\,[\mathbb{R}^+ \partial f(x)] = \mathbb{R}^+ \partial f(x). \qquad \square$$

Remark 1.3.6 The assumption $0 \notin \partial f(x)$, required by the above two results, can be formulated in a number of equivalent ways:

– In view of Definition 1.1.4, it means $f'(x, d_0) < 0$ for some d_0.
– Using the other definition (1.2.1), there is some x_0 such that $f(x_0) < f(x)$.
– The latter implies that the same assumption holds everywhere on the level-set $f(\cdot) = f(x)$.

As a result, the existence of one point x with $0 \notin \partial f(x)$ allows the computation of the tangent and normal cone to the corresponding sublevel-set $Sf(x)$ at all its points: on its boundary – which, thanks to (1.3.6), is the level-set $f(\cdot) = f(x)$ – as well as on its interior (trivial case). $\qquad \square$

Figure 1.3.2 illustrates these results. It is similar to Fig. 1.3.1, except that it is drawn in $\mathbb{R}^n$. Its left part represents the horizontal cut of Fig. 1.3.1 at level $f(x)$, i.e.
$$T_{\mathrm{epi}\,f}(x, f(x)) \cap \{(d, r) \in \mathbb{R}^n \times \mathbb{R} : r = 0\}$$

which is $T_{Sf(x)}(x) \times \{0\}$; considered as a set in $\mathbb{R}^n$, it is the tangent cone to the sublevel-set $Sf(x)$. In the right part of the picture, we have also drawn the subdifferential, neglecting its vertical component drawn in Fig. 1.3.1. The cone $N_{Sf(x)}(x)$ generated by this subdifferential appears to be the projection (onto the same horizontal space) of the normal cone $N_{\mathrm{epi}\,f}(x, f(x))$.
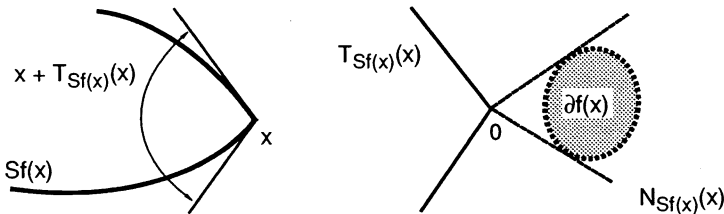
Fig. 1.3.2. Tangent and normal cones to a sublevel-set

All this confirms how a subdifferential generalizes a gradient: if $\partial f(x)$ is the singleton $\{\nabla f(x)\}$, the level-set $f(\cdot) = f(x)$ has a tangent hyperplane at $(x, f(x))$;

the cone $T_{Sf(x)}(x)$ is the half-space opposite to $\nabla f(x)$; the cone $N_{Sf(x)}(x)$ is the half-line $\mathbb{R}^+ \nabla f(x)$. When $\partial f(x)$ becomes "fatter", $Sf(x)$ becomes "narrower" around x.

Drawing the normal cone to the sublevel-set, i.e. considering only nonnegative multiples of the subgradients, does describe the sublevel-set locally around x; but it also destroys some information, namely the magnitudes of these gradients. This information contains the magnitudes of the directional derivatives, and Fig. 1.3.3 shows how to recover it, even in $\mathbb{R}^n$. It is similar to Fig. 1.3.2 and should be compared to Fig. C.2.1.1: the supporting hyperplane $H := \{s \in \mathbb{R}^n : \langle d, s - x \rangle = f'(x, d)\}$ is orthogonal to d; its (algebraic) distance to x is $f'(x, d)$, if d is normalized. The half-line $x + \mathbb{R}^+ d$ would cut the sublevel-set $S_{f(x)-1}(f)$ at the (algebraic) distance $f'(x, d)$ if $t \mapsto f(x + td)$ were an affine function of $t \geqslant 0$.
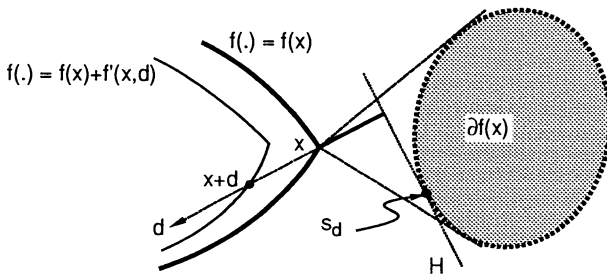
Fig. 1.3.3. Rate of decrease along a direction

2 Local Properties of the Subdifferential

In this section, we study some properties of $\partial f(x)$, considered as a generalization of the concept of gradient, at a given fixed x.

2.1 First-Order Developments

As already mentioned, a finite convex function enjoys a "directional first-order approximation" (1.1.7), and an important result is that the convergence in (1.1.7) is *uniform* in d on any bounded set: ε_d can be taken independent of the normalized d.

Lemma 2.1.1 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and $x \in \mathbb{R}^n$. For any $\varepsilon > 0$, there exists $\delta > 0$ such that $\|h\| \leqslant \delta$ implies*

$$|f(x + h) - f(x) - f'(x, h)| \leqslant \varepsilon \|h\|. \tag{2.1.1}$$

Proof. Suppose for contradiction that there is $\varepsilon > 0$ and a sequence (h_k) with $\|h_k\| =: t_k \leqslant 1/k$ such that

$$|f(x + h_k) - f(x) - f'(x, h_k)| > \varepsilon t_k \quad \text{for } k = 1, 2, \dots$$

Extracting a subsequence if necessary, assume that $h_k/t_k \to d$ for some d of norm 1. Then take a local Lipschitz constant L of f (see Remark 1.1.3) and expand:

$$\begin{aligned}
\varepsilon t_k &< |f(x + h_k) - f(x) - f'(x, h_k)| \\
&\leqslant |f(x + h_k) - f(x + t_k d)| + \\
&\quad + |f(x + t_k d) - f(x) - f'(x, t_k d)| + |f'(x, t_k d) - f'(x, h_k)| \\
&\leqslant 2L \|h_k - t_k d\| + |f(x + t_k d) - f(x) - t_k f'(x, d)|.
\end{aligned}$$

Divide by $t_k > 0$ and pass to the limit to obtain the contradiction $\varepsilon \leqslant 0$.   □

Another way of writing (2.1.1) is the first-order expansion

$$f(x + h) = f(x) + f'(x, h) + o(\|h\|), \tag{2.1.2}$$

or also

$$\lim_{t \downarrow 0,\, d' \to d} \frac{f(x + td') - f(x)}{t} = f'(x, d).$$

Remark 2.1.2 Convexity plays a little role for Lemma 2.1.1. Apart from the existence of a directional derivative, the proof uses only Lipschitz properties of f and $f'(x, \cdot)$.

This remark is of general interest. When defining a concept of "derivative" $D : \mathbb{R}^n \to \mathbb{R}$ attached to some function $f : \mathbb{R}^n \to \mathbb{R}$ at $x \in \mathbb{R}^n$, one considers the property

$$\frac{f(x + h) - f(x) - D(h)}{\|h\|} \to 0, \tag{2.1.3}$$

with h tending to 0. In classical calculus, a linear D is postulated, while here we rather have a sublinear $D(\cdot) = f'(x, \cdot)$. Actually, even without specifying the class of functions that D belongs to, several sorts of derivatives can be defined, depending on the type of convergence allowed for h:

  (i) the point of view of Gâteaux: (2.1.3) holds for $h = td$, d fixed in $\mathbb{R}^n$, $t \to 0$ in $\mathbb{R}$;
  (ii) the point of view of Fréchet: (2.1.3) holds for $h \to 0$;
  (i') the directional point of view: as in (i), but with $t > 0$;
  (ii') the directional point of view of Dini: $h = td'$, with $t \downarrow 0$ and $d' \to d$ in $\mathbb{R}^n$.

It should be clear from the proof of Lemma 2.1.1 that, once the approximating function D is specified, (say linear or sublinear), these four types of convergence are equivalent when f is Lipschitzian around x.   □

Compare (2.1.2) with the radial development (1.1.7), which can be written with any subgradient s_d maximizing $\langle s, d \rangle$. Such an s_d is an arbitrary element in the face of $\partial f(x)$ exposed by d (remember Definition C.3.1.3). Equivalently, d lies in the normal cone to $\partial f(x)$ at s_d; or also (Proposition A.5.3.3), s_d is the projection of $s_d + d$ onto $\partial f(x)$. Thus, the following is just a restatement of Lemma 2.1.1.

Corollary 2.1.3 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. At any x,*

$$f(x + h) = f(x) + \langle s, h \rangle + o(\|h\|)$$

whenever one of the following equivalent properties holds:

$$s \in \mathrm{F}_{\partial f(x)}(h) \quad \Longleftrightarrow \quad h \in \mathrm{N}_{\partial f(x)}(s) \quad \Longleftrightarrow \quad s = \mathrm{p}_{\partial f(x)}(s + h).\quad \square$$

In words: as long as the increment h varies in a portion of $\mathbb{R}^n$ that is in some fixed normal cone to $\partial f(x)$, f looks differentiable; any subgradient in the corresponding exposed face can be considered as a "local gradient", active only on that cone; see Fig. 2.1.1. When h moves to another normal cone, it is another "local gradient" that prevails.
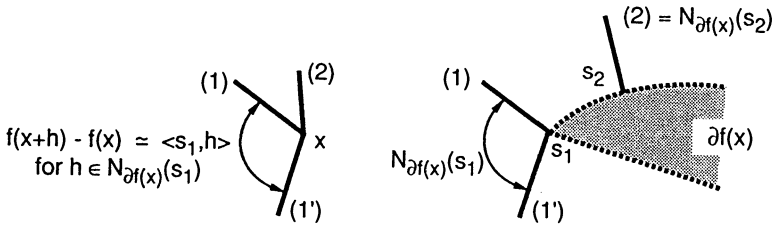
Fig. 2.1.1. Apparent differentiability in normal cones

Because $\partial f(x)$ is compact, any nonzero $h \in \mathbb{R}^n$ exposes a nonempty face. When h describes $\mathbb{R}^n \setminus \{0\}$, the corresponding exposed faces cover the boundary of $\partial f(x)$: this is Proposition C.3.1.5. An important special case is of course that of $\partial f(x)$ with only one exposed face, i.e. only one element. This means that there is some *fixed* $s \in \mathbb{R}^n$ such that

$$\lim_{t \downarrow 0} \frac{f(x + td) - f(x)}{t} = \langle s, d \rangle \quad \text{for all } d \in \mathbb{R}^n ,$$

which expresses precisely the directional differentiability of f at x. Change d and t to $-d$ and $-t$ to see that f is indeed *Gâteaux differentiable* at x. From Corollary 2.1.3 (see again Remark 2.1.2), this is further equivalent to

$$f(x + h) - f(x) = \langle s, h \rangle + o(\|h\|) \quad \text{for all } h \in \mathbb{R}^n ,$$

i.e. f is *Fréchet differentiable* at x. We will simply say that our function is differentiable at x, a non-ambiguous terminology.

We summarize our observations:

Corollary 2.1.4 *If the convex f is (Gâteaux) differentiable at x, its only subgradient at x is its gradient $\nabla f(x)$. Conversely, if $\partial f(x)$ contains only one element s, then f is (Fréchet) differentiable at x, with $\nabla f(x) = s$.*   □

Note the following consequence of Proposition C.1.1.6: if $\{d_1, \ldots, d_k\}$ is a set of vectors generating the whole space and if $f'(x, d_i) = -f'(x, -d_i)$ for $i = 1, \ldots, k$, then f is differentiable at x. In particular (take $\{d_i\}$ as the canonical basis of $\mathbb{R}^n$), the *existence alone* of the partial derivatives

$$\frac{\partial f}{\partial \xi^i}(x) = f'(x, e_i) = -f'(x, -e_i) \quad \text{for } i = 1, \ldots, n$$

guarantees the differentiability of the convex f at $x = (\xi^1, \dots, \xi^n)$. See again Proposition B.4.2.1.

For the general case where $\partial f(x)$ is not a singleton, we mention here another way of defining exposed faces: the function $f'(x, \cdot)$ being convex, it has subdifferentials in its own right (Proposition 1.1.6 studied the subdifferential at 0 only). These subdifferentials are precisely the exposed faces of $\partial f(x)$.

Proposition 2.1.5 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. For all x and d in $\mathbb{R}^n$, we have*

$$F_{\partial f(x)}(d) = \partial[f'(x, \cdot)](d).$$

Proof. If $s \in \partial f(x)$ then $f'(x, d') \geqslant \langle s, d' \rangle$ for all $d' \in \mathbb{R}^n$, simply because $f'(x, \cdot)$ is the support function of $\partial f(x)$. If, in addition, $\langle s, d \rangle = f'(x, d)$, we get

$$f'(x, d') \geqslant f'(x, d) + \langle s, d' - d \rangle \quad \text{for all } d' \in \mathbb{R}^n \qquad (2.1.4)$$

which proves the inclusion $F_{\partial f(x)}(d) \subset \partial[f'(x, \cdot)](d)$.

Conversely, let s satisfy (2.1.4). Set $d'' := d' - d$ and deduce from subadditivity

$$f'(x, d) + f'(x, d'') \geqslant f'(x, d') \geqslant f'(x, d) + \langle s, d'' \rangle \quad \text{for all } d'' \in \mathbb{R}^n$$

which implies $f'(x, \cdot) \geqslant \langle s, \cdot \rangle$, hence $s \in \partial f(x)$. Also, putting $d' = 0$ in (2.1.4) shows that $\langle s, d \rangle \geqslant f'(x, d)$. Altogether, we have $s \in F_{\partial f(x)}(d)$. $\qquad \square$
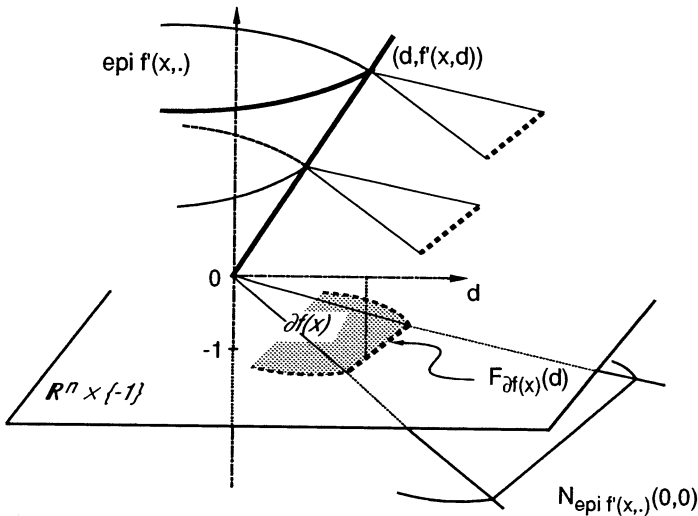
Fig. 2.1.2. Faces of subdifferentials

This result is illustrated in Fig. 2.1.2. Observe in particular that the subdifferential of $f'(x, \cdot)$ at the point td does not depend on $t > 0$; but when t reaches 0, this subdifferential explodes to the entire $\partial f(x)$.

Definition 2.1.6 A point x at which $\partial f(x)$ has more than one element – i.e. at which f is not differentiable – is called a *kink* (or corner-point) of f.   □

We know that f is differentiable almost everywhere (Theorem B.4.2.3). The set of kinks is therefore of zero-measure (remember Remark A.2.1.7). In most examples in practice, this set is the union of a finite number of algebraic surfaces in $\mathbb{R}^n$.

Example 2.1.7 Let $q_1, \ldots, q_m$ be m convex quadratic functions and take f as the max of the q_j's. Given $x \in \mathbb{R}^n$, let $J(x) := \{j \leqslant m \ : \ q_j(x) = f(x)\}$ denote the set of active q_j's at x.

It is clear that f is differentiable at each x such that $J(x)$ reduces to a singleton $\{j(x)\}$: by continuity, $J(y)$ is still this singleton $\{j(x)\}$ for all y close enough to x, so f and its gradient coincide around this x with the smooth $q_{j(x)}$ and its gradient. Thus, our f has all its kinks in the union of the $1/2\,m(m-1)$ surfaces

$$\Sigma_{ij} := \{x \in \mathbb{R}^n \ : \ q_i(x) = q_j(x)\} \quad \text{for } i \neq j\,.$$

Figure 2.1.3 gives an idea of what this case could look like, in $\mathbb{R}^2$. The dashed lines represent portions of Σ_{ij} at which $q_i = q_j < f$.   □
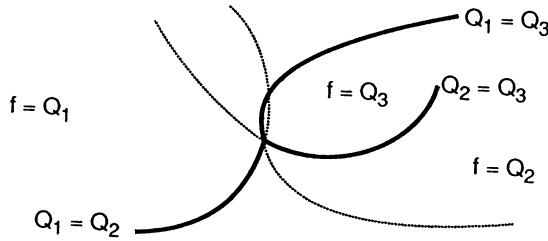
Fig. 2.1.3. A maximum of convex quadratic functions

2.2 Minimality Conditions

We start with a fundamental result, coming directly from the definitions of the sub-differential.

Theorem 2.2.1 *For $f : \mathbb{R}^n \to \mathbb{R}$ convex, the following three properties are equivalent:*

(i) *f is minimized at x over $\mathbb{R}^n$, i.e.: $f(y) \geqslant f(x)$ for all $y \in \mathbb{R}^n$;*
(ii) *$0 \in \partial f(x)$;*
(iii) *$f'(x, d) \geqslant 0$ for all $d \in \mathbb{R}^n$.*

Proof. The equivalence (i) ⇔ (ii) [resp. (ii) ⇔ (iii)] is obvious from (1.2.1) [resp. (1.1.6)].   □

Naturally, x can be called "stationary", or critical, if $0 \in \partial f(x)$. Observe that the equivalence (i) $\Leftrightarrow$ (iii) says: f is minimal at x if and only if its tangential approximation $f'(x, \cdot)$ is minimal at 0; a statement which makes sense, and which calls for two remarks.

– When x is a local minimum of f (i.e. $f(x) \leqslant f(y)$ for all y in a neighborhood of x), (iii) holds; thus, convexity implies that a local minimum is automatically global (i.e. $f(x) \leqslant f(y)$ for all y).

– In the smooth case, the corresponding statement would be "the tangential approximation $\langle \nabla f(x), \cdot \rangle$ of $f(\cdot) - f(x)$, which is linear, is identically 0". Here, the tangential approximation $f'(x, \cdot)$ need not be 0; but there does exist a minorizing linear function, not necessarily tangential, which is identically 0.

Remark 2.2.2 The property "$0 \in \partial f(x)$" is a generalization of the usual stationarity condition "$\nabla f(x) = 0$" of the smooth case. Even though the gradient exists at almost every x, one should not think that a convex function has almost certainly a 0-gradient at a minimum point. As a matter of fact, the "probability" that a given x is a kink is 0; but this probability may not stay 0 if some more information is known about x, for example that it is a stationary point. As a rule, the minimum points of a convex function are indeed kinks.   □

2.3 Mean-Value Theorems

Given two distinct points x and y, and knowing the subdifferential of f on the whole line-segment $]x, y[$, can we evaluate $f(y) - f(x)$? Or also, is it possible to express f as the integral of its subdifferential? This is the aim of mean-value theorems.

Of course, the problem reduces to that of one-dimensional convex functions, since

$$f(y) - f(x) = \varphi(1) - \varphi(0)$$

where

$$\varphi(t) := f(ty + (1 - t)x) \quad \text{for all } t \in [0, 1] \tag{2.3.1}$$

is the trace of f on the line-segment $[x, y]$. The key question, however, is to express the subdifferential of φ at t in terms of the subdifferential of f at $ty + (1 - t)x$ in the surrounding space $\mathbb{R}^n$. The next lemma anticipates the calculus rules to be given in §4. Here and below, we use the following notation:

$$x_t := ty + (1 - t)x$$

where x and y are considered as fixed in $\mathbb{R}^n$.

Lemma 2.3.1 *The subdifferential of φ defined by* (2.3.1) *is*

$$\partial \varphi(t) = \{ \langle s, y - x \rangle \ : \ s \in \partial f(x_t) \}$$

or, more symbolically:

$$\partial \varphi(t) = \langle \partial f(x_t), y - x \rangle .$$

Proof. In terms of right- and left-derivatives (see Theorem 0.6.3), we have

$$D_+\varphi(t) = \lim_{\tau\downarrow 0} \frac{f(x_t + \tau(y - x)) - f(x_t)}{\tau} = f'(x_t, y - x),$$

$$D_-\varphi(t) = \lim_{\tau\uparrow 0} \frac{f(x_t + \tau(y - x)) - f(x_t)}{\tau} = -f'(x_t, -(y - x));$$

so, knowing that

$$f'(x_t, y - x) = \max_{s\in\partial f(x_t)} \langle s, y - x\rangle,$$

$$-f'(x_t, -(y - x)) = \min_{s\in\partial f(x_t)} \langle s, y - x\rangle,$$

we obtain $\partial\varphi(t) := [D_-\varphi(t), D_+\varphi(t)] = \{\langle s, y - x\rangle : s \in \partial f(x)\}.$ □

Remark 2.3.2 The derivative of φ exists except possibly on a countable set in $\mathbb{R}$. One should not think that, with this pretext, f is differentiable except possibly at countably many points of $]x, y[$. For example, with $f(\xi, \eta) := |\xi|$, $x := (0, 0)$, $y := (0, 1)$, f is differentiable nowhere on $[x, y]$. What Lemma 2.3.1 ensures, however, is that for almost all t, $\partial f(x_t)$ has a zero breadth in the direction $y - x$: $f'(x_t, y - x) + f'(x_t, x - y) = 0$. □

The function φ above is the key to proving mean-value theorems.

Theorem 2.3.3 *Let* $f : \mathbb{R}^n \to \mathbb{R}$ *be convex. Given two points* $x \neq y$ *in* $\mathbb{R}^n$, *there exist* $t \in]0, 1[$ *and* $s \in \partial f(x_t)$ *such that*

$$f(y) - f(x) = \langle s, y - x\rangle. \tag{2.3.2}$$

In other words,

$$f(y) - f(x) \in \bigcup_{t\in]0,1[} \{\langle \partial f(x_t), y - x\rangle\}. \qquad\qquad □$$

Proof. Start from the function φ of (2.3.1) and, as usual in this context, consider the auxiliary function

$$\psi(t) := \varphi(t) - \varphi(0) - t[\varphi(1) - \varphi(0)],$$

which is clearly convex. Computing directional derivatives gives easily $\partial\psi(t) = \partial\varphi(t) - [\varphi(1) - \varphi(0)]$. Now ψ is continuous on $[0, 1]$, it has been constructed so that $\psi(0) = \psi(1) = 0$, so it is minimal at some $t \in]0, 1[$: at this t, $0 \in \partial\psi(t)$ (Theorem 2.2.1). In view of Lemma 2.3.1, this means that there is $s \in \partial f(x_t)$ such that

$$\langle s, y - x\rangle = \varphi(1) - \varphi(0) = f(y) - f(x). \qquad\qquad □$$

The mean-value theorem can also be given in integral form.

Theorem 2.3.4 *Let* $f : \mathbb{R}^n \to \mathbb{R}$ *be convex. For* $x, y \in \mathbb{R}^n$,

$$f(y) - f(x) = \int_0^1 \langle \partial f(x_t), y - x\rangle \, dt. \tag{2.3.3}$$
□

The meaning of (2.3.3) is as follows: if $\{s_t : t \in [0,1]\}$ is any selection of subgradients of f on the line-segment $[x, y]$, i.e. $s_t \in \partial f(x_t)$ for all $t \in [0, 1]$, then $\int_0^1 \langle s_t, y - x \rangle \mathrm{d}t$ is independent of the selection and its value is $f(y) - f(x)$.

Example 2.3.5 Mean-value theorems can be applied to nondifferentiable functions in the same way as they are in (ordinary) differential calculus. As an example, let f, g be two convex functions such that $f(x) = g(x)$ and $f(y) \leqslant g(y)$ for all y in a neighborhood of x. It follows from the definitions that $\partial f(x) \subset \partial g(x)$.

Conversely, suppose that f, g are such that $\partial f(x) \subset \partial g(x)$ for all $x \in \mathbb{R}^n$; can we compare f and g? Same question if $\partial f(x) \cap \partial g(x) \neq \emptyset$ for all x. The answer lies in (2.3.3): the difference $f - g$ is a constant function (take the same selection in the integral (2.3.3)!).

An amusing particular case is when f and g are (finite) sublinear functions: then $f \leqslant g$ if and only if $\partial f(0) \subset \partial g(0)$ (f and g are the support functions of $\partial f(0)$ and $\partial g(0)$ respectively! See Remark 1.2.3 again). $\qquad\square$

3 First Examples

Example 3.1 (Support Functions) Let C be a nonempty convex compact set, with support function σ_C. The first-order differential elements of σ_C at the origin are obtained immediately from the definitions:
$$\partial \sigma_C(0) = C \quad \text{and} \quad (\sigma_C)'(0, \cdot) = \sigma_C .$$
Read this with Proposition 1.1.6 in mind: any convex compact set C can be considered as the subdifferential of some finite convex function f at some point x. The simplest instance is $f = \sigma_C$, $x = 0$.

On the other hand, the first-order differential elements of a support function σ_C at $x \neq 0$ are given in Proposition 2.1.5:
$$\partial \sigma_C(x) = \mathrm{F}_C(x) \quad \text{and} \quad (\sigma_C)'(x, \cdot) = \sigma_{\mathrm{F}_C(x)} . \tag{3.1}$$

The expression of $(\sigma_C)'(x, d)$ above is a bit tricky: it is the optimal value of the following optimization problem (s is the variable, x and d are fixed, the objective function is linear, there is one linear constraint in addition to those describing C):
$$\left| \begin{array}{l} \max \langle d, s \rangle , \quad s \in C , \\ \langle s, x \rangle = \sigma_C(x) . \end{array} \right.$$

As a particular case, take a norm $\| \cdot \|$. As seen already in §C.3.2, it is the gauge of its unit ball B, and it is the support function of the unit ball B^* associated with the dual norm $\| \cdot \|^*$. Hence
$$\partial \| \cdot \|(0) = B^* = \left\{ s \in \mathbb{R}^n : \max_{\|d\| \leqslant 1} \langle s, d \rangle \leqslant 1 \right\} .$$
More generally, for x not necessarily zero, (3.1) can be rewritten as
$$\partial \| \cdot \|(x) = \left\{ s \in B^* : \langle s, x \rangle = \max_{u \in B^*} \langle u, x \rangle = \|x\| \right\} . \tag{3.2}$$

All the points s in (3.2) have dual norm 1; they form the face of B^* exposed by x. $\qquad\square$

Example 3.2 (Gauges) Suppose now that, rather than being compact, the C of Example 3.1 is closed and contains the origin as an interior point. Then, another example of a convex finite function is its gauge γ_C (Theorem C.1.2.5). Taking into account the correspondence between supports and gauges – see Proposition C.3.2.4 and Corollary C.3.2.5 – (3.1) can be copied if we replace C by its polar set

$$C^\circ := \{x \,:\, \langle s, x \rangle \leqslant 1 \text{ for all } s \in C\}.$$

So we obtain

$$\partial \gamma_C(0) = C^\circ, \quad (\gamma_C)'(0, \cdot) = \gamma_C,$$

and of course, Proposition 2.1.5 applied to C° gives at $x \neq 0$

$$\partial \gamma_C(x) = \mathrm{F}_{C^\circ}(x) \quad \text{and} \quad (\gamma_C)'(x, \cdot) = \sigma_{\mathrm{F}_{C^\circ}(x)}.$$

Gauges and support functions of elliptic sets deserve a more detailed study. Given a symmetric positive semi-definite operator Q, define

$$\mathbb{R}^n \ni x \mapsto f(x) := \sqrt{\langle Qx, x \rangle}$$

which is just the gauge of the sublevel-set $\{x : f(x) \leqslant 1\}$. From elementary calculus,

$$\partial f(x) = \{\nabla f(x)\} = \left\{\frac{Qx}{f(x)}\right\} \quad \text{for } x \notin \mathrm{Ker}\, Q$$

while, for $x \in \mathrm{Ker}\, Q$, $s \in \partial f(x)$ if and only if, for all $y \in \mathbb{R}^n$

$$\langle s, y - x \rangle \leqslant \sqrt{\langle Qy, y \rangle} = \sqrt{\langle Q(y-x), y-x \rangle} = \left\| Q^{1/2}(y-x) \right\|.$$

From the Cauchy-Schwarz inequality (remember Example C.2.3.4), we see that $\partial f(x)$ is the image by $Q^{1/2}$ of the unit ball $B(0, 1)$. $\square$

Example 3.3 (Distance Functions) Let again C be closed and convex. Another finite convex function is the distance to C: $\mathrm{d}_C(x) := \min\{\|y - x\| : y \in C\}$, in which the min is attained at the projection $\mathrm{p}_C(x)$ of x onto C. The subdifferential of d_C is

$$\partial \mathrm{d}_C(x) = \begin{cases} \mathrm{N}_C(x) \cap B(0,1) & \text{if } x \in C, \\ \left\{\frac{x - \mathrm{p}_C(x)}{\|x - \mathrm{p}_C(x)\|}\right\} & \text{if } x \notin C, \end{cases} \tag{3.3}$$

a formula illustrated by Fig. C.2.3.1 when C is a closed convex cone. To prove it, we distinguish two cases. When $x \notin C$, hence $\mathrm{d}_C(x) > 0$, we can apply standard differential calculus:

$$\nabla \mathrm{d}_C(x) = \nabla \sqrt{\mathrm{d}_C^2}(x) = \frac{\nabla \mathrm{d}_C^2(x)}{2 \mathrm{d}_C(x)},$$

so we have to prove $\nabla \mathrm{d}_C^2(x) = 2[x - \mathrm{p}_C(x)]$. For this, consider $\Delta := \mathrm{d}_C^2(x + h) - \mathrm{d}_C^2(x)$. Because $\mathrm{d}_C^2(x) \leqslant \|x - \mathrm{p}_C(x + h)\|^2$, we have

$$\Delta \geqslant \|x + h - \mathrm{p}_C(x + h)\|^2 - \|x - \mathrm{p}_C(x + h)\|^2 = \|h\|^2 + 2\langle h, x - \mathrm{p}_C(x + h)\rangle.$$

Inverting the role of x and $x + h$, we obtain likewise

$$\Delta \leqslant \|x + h - \mathrm{p}_C(x)\|^2 - \|x - \mathrm{p}_C(x)\|^2 = \|h\|^2 + 2\langle h, x - \mathrm{p}_C(x)\rangle.$$

Remembering from (A.3.1.6) that p_C is nonexpansive: $\Delta = 2\langle x - \mathrm{p}_C(x), h \rangle + o(\|h\|)$, and this is what we wanted.

Now, for $x \in C$, let $s \in \partial \mathrm{d}_C(x)$, i.e. $\mathrm{d}_C(x') \geqslant \langle s, x' - x \rangle$ for all $x' \in \mathbb{R}^n$. This implies in particular $\langle s, x' - x \rangle \leqslant 0$ for all $x' \in C$, hence $s \in N_C(x)$; and taking $x' = x + s$, we obtain

$$\|s\|^2 \leqslant \mathrm{d}_C(x + s) \leqslant \|x + s - x\| = \|s\|.$$

Conversely, let $s \in N_C(x) \cap B(0, 1)$ and, for all $x' \in \mathbb{R}^n$, write

$$\langle s, x' - x \rangle = \langle s, x' - \mathrm{p}_C(x') \rangle + \langle s, \mathrm{p}_C(x') - x \rangle.$$

The last scalar product is nonpositive because $s \in N_C(x)$ and, with the Cauchy-Schwarz inequality, the property $\|s\| \leqslant 1$ gives

$$\langle s, x' - \mathrm{p}_C(x') \rangle \leqslant \|x' - \mathrm{p}_C(x')\| = \mathrm{d}_C(x').$$

Altogether, $s \in \partial \mathrm{d}_C(x)$. Note that the set of kinks of d_C is exactly the boundary of C.

Consider the closed convex cone $K := N_C(x)$, whose polar cone is $K^\circ = T_C(x)$. As seen in Chap. C, more particularly in (C.3.3.2), the support function of $K' = K \cap B(0, 1)$ is the distance to the tangent cone $T_C(x)$. From (3.3), we see that

$$\mathrm{d}'_C(x, \cdot) = \mathrm{d}_{T_C(x)} \quad \text{for all } x \in C. \qquad \qquad \square$$

Example 3.4 (Piecewise Affine Functions) Consider the function

$$\mathbb{R}^n \ni x \mapsto f(x) := \max \{ f_j(x) : j = 1, \ldots, m \} \qquad (3.4)$$

where each f_j is affine: $f_j(x) := r_j + \langle s_j, x \rangle$ for $j = 1, \ldots, m$.

To compute the first-order differential elements of f at a given x, it is convenient to translate the origin at this x. Thus, we rewrite (3.4) as

$$f(y) = f(x) + \max \{ -e_j + \langle s_j, y - x \rangle : j = 1, \ldots, m \} \qquad (3.5)$$

where we have set, for $j = 1, \ldots, m$

$$e_j := f(x) - r_j - \langle s_j, x \rangle = f(x) - f_j(x) \geqslant 0 \qquad (3.6)$$

(look at the left part of Fig. 3.1 to visualize e_j).
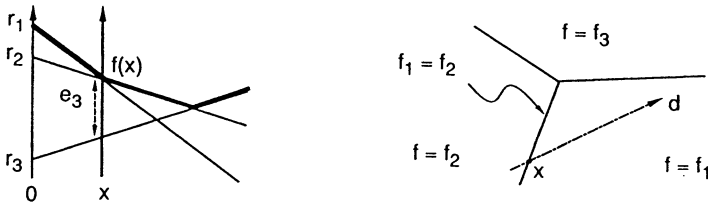
Fig. 3.1. Piecewise affine functions: translation of the origin and directional derivative

Now consider $f(x + td)$, as illustrated on the right part of Fig. 3.1, representing the space $\mathbb{R}^n$ around x. For $t > 0$ small enough, those j such that $e_j > 0$ do not count. Accordingly, set

$$J(x) := \{ j : e_j = 0 \} = \{ j : f_j(x) = f(x) \}.$$

Rewriting (3.5) again as

$$f(x + td) = f(x) + t \max \{ \langle s_j, d \rangle : j \in J(x) \} \quad \text{for small } t > 0,$$

it becomes obvious that

$$f'(x, d) = \max \left\{ \langle s_j, d \rangle \; : \; j \in J(x) \right\}.$$

From the calculus rule C.3.3.2(ii) and Definition 1.1.4, this means exactly that

$$\partial f(x) = \mathrm{co} \left\{ s_j \; : \; j \in J(x) \right\}. \tag{3.7}$$

This result will be confirmed in §4.4. We have demonstrated it here nevertheless in intuitive terms, because it illustrates the effect of the (important) max operation on the subdifferential.   □

4 Calculus Rules with Subdifferentials

Calculus with subdifferentials of convex functions is important for the theory, just as in ordinary differential calculus. Its role is illustrated in Fig. 4.0.1: if f is constructed from some other convex functions f_j, the problem is to compute ∂f in terms of the ∂f_j's.
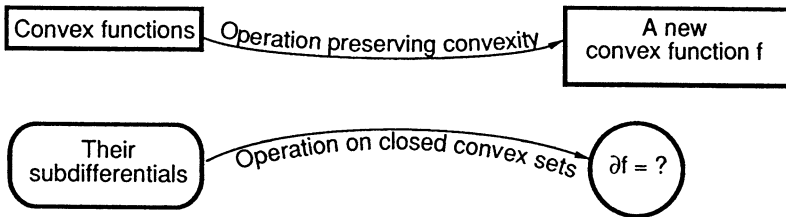
Fig. 4.0.1. Subdifferential calculus

To develop our calculus rules, the two definitions 1.1.4 and 1.2.1 will be used. Calculus with support functions (§C.3.3) will therefore be an essential tool.

4.1 Positive Combinations of Functions

Theorem 4.1.1 *Let f_1, f_2 be two convex functions from $\mathbb{R}^n$ to $\mathbb{R}$ and t_1, t_2 be positive. Then*

$$\partial(t_1 f_1 + t_2 f_2)(x) = t_1 \partial f_1(x) + t_2 \partial f_2(x) \quad \text{for all } x \in \mathbb{R}^n. \tag{4.1.1}$$

Proof. Apply Theorem C.3.3.2(i): $t_1 \partial f_1(x) + t_2 \partial f_2(x)$ is a compact convex set whose support function is

$$t_1 f_1'(x, \cdot) + t_2 f_2'(x, \cdot). \tag{4.1.2}$$

On the other hand, the support function of $\partial(t_1 f_1 + t_2 f_2)(x)$ is by definition the directional derivative $(t_1 f_1 + t_2 f_2)'(x, \cdot)$ which, from elementary calculus, is just (4.1.2). Therefore the two (compact convex) sets in (4.1.1) coincide, since they have the same support function.   □

Remark 4.1.2 Needless to say, the sign of t_1 and t_2 in (4.1.1) is important to obtain a resulting function which is convex. There is a deeper reason, though: take $f_1(x) = f_2(x) = \|x\|$, $t_1 = -t_2 = 1$. We obtain $f_1 - f_2 \equiv 0$, yet $t_1 \partial f_1(0) + t_2 \partial f_2(0) = B(0,2)$, a gross overestimate of $\{0\}$! $\qquad\square$

To illustrate this calculus rule, consider $f : \mathbb{R}^p \times \mathbb{R}^q \to \mathbb{R}$ defined by

$$f(x_1, x_2) = f_1(x_1) + f_2(x_2),$$

with f_1 and f_2 convex on $\mathbb{R}^p$ and $\mathbb{R}^q$ respectively. First, call

$$\mathbb{R}^p \times \mathbb{R}^q \ni (x_1, x_2) \mapsto \tilde{f}_1(x_1, x_2) = f_1(x_1)$$

the extension of f_1; its subdifferential is obviously $\partial \tilde{f}_1(x_1, x_2) = \partial f_1(x_1) \times \{0\}$. Then Theorem 4.1.1 gives, after the same extension is made with f_2,

$$\partial f(x_1, x_2) = \partial f_1(x_1) \times \{0\} + \{0\} \times \partial f_2(x_2) = \partial f_1(x_1) \times \partial f_2(x_2). \qquad (4.1.3)$$

Remark 4.1.3 Given a convex function $f : \mathbb{R}^n \to \mathbb{R}$, an interesting trick is to view its epigraph as a sublevel-set of a certain convex function, namely:

$$\mathbb{R}^n \times \mathbb{R} \ni (x, r) \mapsto g(x, r) := f(x) - r.$$

Clearly enough, epi f is the sublevel-set $S_0(g)$. The directional derivatives of g are easy to compute: $g'(x, f(x); d, \rho) = f'(x, d) - \rho$ for all $(d, \rho) \in \mathbb{R}^n \times \mathbb{R}$ and (4.1.3) gives for all $x \in \mathbb{R}^n$

$$\partial g(x, f(x)) = \partial f(x) \times \{-1\} \not\ni 0.$$

We can therefore apply Theorems 1.3.4 and 1.3.5 to g, which gives back the formulae of Proposition 1.3.1:

$$\begin{aligned} T_{\text{epi } f}(x, f(x)) &= \{(d, \rho) : f'(x, d) \leqslant \rho\}, \\ \text{int } T_{\text{epi } f}(x, f(x)) &= \{(d, \rho) : f'(x, d) < \rho\} \neq \emptyset, \\ N_{\text{epi } f}(x, f(x)) &= \mathbb{R}^+[\partial f(x) \times \{-1\}]. \end{aligned} \qquad\square$$

4.2 Pre-Composition with an Affine Mapping

Theorem 4.2.1 *Let* $A : \mathbb{R}^n \to \mathbb{R}^m$ *be an affine mapping* $(Ax = A_0 x + b$, *with* A_0 *linear and* $b \in \mathbb{R}^m$) *and let* g *be a finite convex function on* $\mathbb{R}^m$. *Then*

$$\partial(g \circ A)(x) = A_0^* \partial g(Ax) \quad \text{for all } x \in \mathbb{R}^n. \qquad (4.2.1)$$

Proof. Form the difference quotient giving rise to $(g \circ A)'(x, d)$ and use the relation $A(x + td) = Ax + t A_0 d$ to obtain

$$(g \circ A)'(x, d) = g'(Ax, A_0 d) \quad \text{for all } d \in \mathbb{R}^n.$$

From Proposition C.3.3.3, the righthand side in the above equality is the support function of the convex compact set $A_0^* \partial g(Ax)$. $\qquad\square$

This result is illustrated by Lemma 2.3.1: with fixed $x, y \in \mathbb{R}^n$, consider the affine mapping $A : \mathbb{R} \to \mathbb{R}^n$

$$At := x + t(y - x).$$

Then $A_0 t = t(y - x)$, and the adjoint of A_0 is defined by $A_0^*(s) = \langle y - x, s \rangle$ for all $s \in \mathbb{R}^n$. Twisting the notation, replace (n, m, x, g) in Theorem 4.2.1 by $(1, n, t, f)$: this gives the subdifferential $\partial \varphi$ of Lemma 2.3.1.

As another illustration, let us come back to the example of §4.1. Needless to say, the validity of (4.1.3) relies crucially on the "decomposed" form of f. Indeed, take a convex function $f : \mathbb{R}^p \times \mathbb{R}^q \to \mathbb{R}$ and the affine mapping

$$\mathbb{R}^p \ni x_1 \mapsto Ax_1 = (x_1, x_2) \in \mathbb{R}^p \times \mathbb{R}^q.$$

Its linear part is $x_1 \mapsto A_0 x_1 = (x_1, 0)$ and $A_0^*(s_1, s_2) = s_1$. Then consider the partial function

$$f \circ A : \mathbb{R}^p \ni x_1 \mapsto f_{x_2}^{(1)}(x_1) = f(x_1, x_2).$$

According to Theorem 4.2.1,

$$\partial f_{x_2}^{(1)}(x_1) = \{ s_1 \in \mathbb{R}^p : \exists s_2 \in \mathbb{R}^q \text{ such that } (s_1, s_2) \in \partial f(x_1, x_2) \}$$

is the projection of $\partial f(x_1, x_2)$ onto $\mathbb{R}^p$. Naturally, we can construct likewise the projection of $\partial f(x_1, x_2)$ onto $\mathbb{R}^q$, which yields the inclusion

$$\partial f(x_1, x_2) \subset \partial f_{x_2}^{(1)}(x_1) \times \partial f_{x_1}^{(2)}(x_2). \tag{4.2.2}$$

Remark 4.2.2 Beware that equality in (4.2.2) need not hold, except in special cases; for example in the decomposable case (4.1.3), or also when one of the projections is a singleton, i.e. when the partial function $f_{x_2}^{(1)}$, say, is differentiable. For a counter-example, take $p = q = 1$ and

$$f(x_1, x_2) = |x_1 - x_2| + \tfrac{1}{2}(x_1 + 1)^2 + \tfrac{1}{2}(x_2 + 1)^2.$$

This function has a unique minimum at $(-1, -1)$ but, at $(0, 0)$, we have $\partial f_0^{(i)}(0) = [0, 2]$ for $i = 1, 2$; hence the righthand side of (4.2.2) contains $(0, 0)$. Yet, f is certainly not minimal there, $\partial f(0, 0)$ is actually the line-segment $\{2(\alpha, 1 - \alpha) : \alpha \in [0, 1]\}$.   □

4.3 Post-Composition with an Increasing Convex Function of Several Variables

As seen in §B.5.2.1(d), post-composition with an increasing one-dimensional convex function preserves convexity. A relevant object is the subdifferential of the result; we somewhat generalize the problem, by considering a vector-valued version of this operation.

Let $f_1, \ldots, f_m$ be m convex functions from $\mathbb{R}^n$ to $\mathbb{R}$; they define a mapping F by

$$\mathbb{R}^n \ni x \mapsto F(x) := (f_1(x), \ldots, f_m(x)) \in \mathbb{R}^m.$$

Equip $\mathbb{R}^m$ with the dot-product and let $g : \mathbb{R}^m \to \mathbb{R}$ be convex and increasing componentwise, i.e

$$y^i \geq z^i \text{ for } i = 1, \ldots, m \quad \Longrightarrow \quad g(y) \geq g(z).$$

Establishing the convexity of the function

$$\mathbb{R}^n \ni x \mapsto (g \circ F)(x) := g(f_1(x), \ldots, f_m(x))$$

is an elementary exercise. Another easy observation is that, if $(\rho^1, \ldots, \rho^m) \in \partial g(y)$, then each ρ^i is nonnegative: indeed, $\{e_1, \ldots, e_m\}$ being the canonical basis in $\mathbb{R}^m$,

$$g(y) \geqslant g(y - e_j) \geqslant g(y) + \sum_{i=1}^{m} \rho^i(-e_j)^i = g(y) - \rho^j .$$

Theorem 4.3.1 *Let f, F and g be defined as above. For all $x \in \mathbb{R}^n$,*

$$\partial(g \circ F)(x) = \left\{ \sum_{i=1}^{m} \rho^i s_i : (\rho^1, \ldots, \rho^m) \in \partial g(F(x)), \atop s_i \in \partial f_i(x) \text{ for } i = 1, \ldots, m \right\} . \tag{4.3.1}$$

Proof. [*Preamble*] Our aim is to show the formula via support functions, hence we need to establish the convexity and compactness of the righthand side in (4.3.1) – call it S. Boundedness and closedness are easy, coming from the fact that a subdifferential (be it ∂g or ∂f_i) is bounded and closed. As for convexity, pick two points in S and form their convex combination

$$s = \alpha \sum_{i=1}^{m} \rho^i s_i + (1 - \alpha) \sum_{i=1}^{m} \rho'^i s_i' = \sum_{i=1}^{m} [\alpha \rho^i s_i + (1 - \alpha)\rho'^i s_i'] ,$$

where $\alpha \in {]}0, 1{[}$. Remember that each ρ^i and ρ'^i is nonnegative and the above sum can be restricted to those terms such that $\rho''^i := \alpha \rho^i + (1 - \alpha)\rho'^i > 0$. Then we write each such term as

$$\rho''^i \left[\frac{\alpha \rho^i}{\rho''^i} s_i + \frac{(1 - \alpha)\rho'^i}{\rho''^i} s_i' \right] .$$

It suffices to observe that $\rho''^i \in \partial g(F(x))$, so the bracketed expression is in $\partial f_i(x)$; thus $s \in S$.

[*Step 1*] Now let us compute the support function σ_S of S. For $d \in \mathbb{R}^n$, we denote by $F'(x, d) \in \mathbb{R}^m$ the vector whose components are $f_i'(x, d)$ and we proceed to prove

$$\sigma_S(d) = g'\big(F(x), F'(x, d)\big) . \tag{4.3.2}$$

For any $s = \sum_{i=1}^{m} \rho^i s_i \in S$, we write $\langle s, d \rangle$ as

$$\sum_{i=1}^{m} \rho^i \langle s_i, d \rangle \leqslant \sum_{i=1}^{m} \rho^i f_i'(x, d) \leqslant g'\big(F(x), F'(x, d)\big) ; \tag{4.3.3}$$

the first inequality uses $\rho^i \geqslant 0$ and the definition of $f_i'(x, \cdot) = \sigma_{\partial f_i}(x)$; the second uses the definition $g'(F(x), \cdot) = \sigma_{\partial g(F(x))}$.

   On the other hand, the compactness of $\partial g(F(x))$ implies the existence of an m-tuple $(\bar{\rho}^i) \in \partial g(F(x))$ such that

$$g'\big(F(x), F'(x,d)\big) = \sum_{i=1}^{m} \bar{\rho}^i f_i'(x,d) \,,$$

and the compactness of each $\partial f_i(x)$ yields likewise an $\bar{s}_i \in \partial f_i(x)$ such that

$$f_i'(x,d) = \langle \bar{s}_i, d \rangle \quad \text{for } i = 1, \ldots, m \,.$$

Altogether, we have exhibited an $\bar{s} = \sum_{i=1}^{m} \bar{\rho}^i \bar{s}_i \in S$ such that equality holds in (4.3.3), so (4.3.2) is established.

[*Step 2*] It remains to prove that the support function (4.3.2) is really the directional derivative $(g \circ F)'(x,d)$. For $t > 0$, expand $F(x + td)$, use the fact that g is locally Lipschitzian, and then expand $g(F(x + td))$:

$$g(F(x+td)) = g\big(F(x) + tF'(x,d) + o(t)\big) = g\big(F(x) + tF'(x,d)\big) + o(t)$$
$$= g(F(x)) + tg'\big(F(x), F'(x,d)\big) + o(t) \,.$$

From there, it follows

$$(g \circ F)'(x,d) := \lim_{t \downarrow 0} \frac{g(F(x+td)) - g(F(x))}{t} = g'\big(F(x), F'(x,d)\big) \,. \qquad \square$$

Let us give some illustrations:

– When g is differentiable at $F(x)$, (4.3.1) has a classical flavour:

$$\partial(g \circ F)(x) = \sum_{i=1}^{m} \frac{\partial g}{\partial y^i}(F(x)) \partial f_i(x) \,.$$

In particular, with $g(y^1, \ldots, y^m) = \tfrac{1}{2} \sum_{i=1}^{m} (y^{i+})^2$ (r^+ denoting $\max\{0, r\}$), we obtain

$$\partial \left[\tfrac{1}{2} \sum_{i=1}^{m} (f_i^+)^2 \right] = \sum_{i=1}^{m} f_i^+ \partial f_i \,.$$

– Take $g(y^1, \ldots, y^m) = \sum_{i=1}^{m} (y^i)^+$ and use the following notation:

$$I_0(x) := \{i : f_i(x) = 0\}, \quad I_+(x) := \{i : f_i(x) > 0\} \,.$$

Then

$$\partial \big(\textstyle\sum_{i=1}^{m} f_i^+ \big)(x) = \sum_{i \in I_+(x)} \partial f_i(x) + \sum_{i \in I_0(x)} [0,1] \partial f_i(x) \,.$$

– Finally, we consider once more our important max operation (generalizing Example 3.4, and to be generalized in §4.4):

Corollary 4.3.2 *Let $f_1, \ldots, f_m$ be m convex functions from $\mathbb{R}^n$ to $\mathbb{R}$ and define*

$$f := \max \{f_1, \ldots, f_m\} \,.$$

Denoting by $I(x) := \{i : f_i(x) = f(x)\}$ the active index-set, we have

$$\partial f(x) = \mathrm{co} \{ \cup \partial f_i(x) \ : \ i \in I(x) \} \,. \tag{4.3.4}$$

Proof. Take $g(y) = \max\{y^1, \ldots, y^m\}$, whose subdifferential was computed in (3.7): $\{e_i\}$ denoting the canonical basis of $\mathbb{R}^m$,

$$\partial g(y) = \text{co}\{e_i \; : \; i \text{ such that } y^i = g(y)\}.$$

Then, using the notation of Theorem 4.3.1, we write $\partial g(F(x))$ as

$$\left\{(\rho^1, \ldots, \rho^m) \; : \; \rho^i = 0 \text{ for } i \notin I(x), \; \rho^i \geqslant 0 \text{ for } i \in I(x), \; \sum_{i=1}^{m} \rho^i = 1\right\},$$

and (4.3.1) gives

$$\partial f(x) = \left\{\sum_{i \in I(x)} \rho^i \partial f_i(x) \; : \; \rho^i \geqslant 0 \text{ for } i \in I(x), \; \sum_{i \in I(x)} \rho^i = 1\right\}.$$

Remembering Example A.1.3.5, it suffices to recognize in the above expression the convex hull announced in (4.3.4)   □

4.4 Supremum of Convex Functions

We come now to an extremely important calculus rule, generalizing Corollary 4.3.2. It has no equivalent in classical differential calculus, and is used very often. In this subsection, we study the following situation: J is an arbitrary index-set, $\{f_j\}_{j \in J}$ is a collection of convex functions from $\mathbb{R}^n$ to $\mathbb{R}$, and we assume that

$$f(x) := \sup\{f_j(x) \; : \; j \in J\} < +\infty \text{ for all } x \in \mathbb{R}^n. \tag{4.4.1}$$

We already know that f is convex (Proposition B.2.1.2) and we are interested in computing its subdifferential. At a given x, call

$$J(x) := \{j \in J : f_j(x) = f(x)\} \tag{4.4.2}$$

the active index-set (possibly empty).

Let us start with an elementary result.

Lemma 4.4.1 *With the notation (4.4.1), (4.4.2),*

$$\partial f(x) \supset \overline{\text{co}}\{\cup \partial f_j(x) \; : \; j \in J(x)\}. \tag{4.4.3}$$

Proof. Take $j \in J(x)$ and $s \in \partial f_j(x)$; from the definition (1.2.1) of the subdifferential,

$$f(y) \geqslant f_j(y) \geqslant f_j(x) + \langle s, y - x \rangle \quad \text{for all } y \in \mathbb{R}^n,$$

so $\partial f(x)$ contains $\partial f_j(x)$. Being closed and convex, it also contains the closed convex hull appearing in (4.4.3).   □

Conversely, when is it true that the subdifferentials $\partial f_j(x)$ at the active indices j "fill up" the whole of $\partial f(x)$? This question is much more delicate, and requires some additional assumption, for example as follows:

Theorem 4.4.2 *With the notation* (4.4.1), (4.4.2), *assume that J is a compact set (in some metric space), on which the functions $j \mapsto f_j(x)$ are upper semi-continuous for each $x \in \mathbb{R}^n$. Then*

$$\partial f(x) = \operatorname{co} \{\cup \partial f_j(x) \; : \; j \in J(x)\}. \tag{4.4.4}$$

Proof. [*Step 0*] Our assumptions make $J(x)$ nonempty and compact. Denote by S the curly bracketed set in (4.4.4); because of (4.4.3), S is bounded, let us check that it is closed. Take a sequence $(s_k) \subset S$ converging to s; to each s_k, we associate some $j_k \in J(x)$ such that $s_k \in \partial f_{j_k}(x)$, i.e.

$$f_{j_k}(y) \geqslant f_{j_k}(x) + \langle s_k, y - x \rangle \quad \text{for all } y \in \mathbb{R}^n .$$

Let $k \to \infty$; extract a subsequence so that $j_k \to j \in J(x)$; we have $f_{j_k}(x) \equiv f(x) = f_j(x)$; and by upper semi-continuity of the function $f_{(\cdot)}(y)$, we obtain

$$f_j(y) \geqslant \limsup f_{j_k}(y) \geqslant f_j(x) + \langle s, y - x \rangle \quad \text{for all } y \in \mathbb{R}^n ,$$

which shows $s \in \partial f_j(x) \subset S$. Altogether, S is compact and its convex hull is also compact (Theorem A.1.4.3).

In view of Lemma 4.4.1, it suffices to prove the "$\subset$"-inclusion in (4.4.4); for this, we will establish the corresponding inequality between support functions which, in view of the calculus rule C.3.3.2(ii), is: for all $d \in \mathbb{R}^n$,

$$f'(x, d) \leqslant \sigma_S(d) = \sup \{f'_j(x, d) \; : \; j \in J(x)\}. \tag{4.4.5}$$

[*Step 1*] Let $\varepsilon > 0$; from the definition (1.1.2) of $f'(x, d)$,

$$\frac{f(x + td) - f(x)}{t} > f'(x, d) - \varepsilon \quad \text{for all } t > 0. \tag{4.4.6}$$

For $t > 0$, set

$$J_t := \left\{ j \in J \; : \; \frac{f_j(x + td) - f(x)}{t} \geqslant f'(x, d) - \varepsilon \right\}.$$

The definition of $f(x + td)$ shows with (4.4.6) that J_t is nonempty. Because J is compact and $f_{(\cdot)}(x + td)$ is upper semi-continuous, J_t is visibly compact. Observe that J_t is a superlevel-set of the function

$$0 < t \mapsto \frac{f_j(x + td) - f_j(x)}{t} + \frac{f_j(x) - f(x)}{t},$$

which is nondecreasing: the first fraction is the slope of a convex function, and the second fraction has a nonpositive numerator. Thus, $J_{t_1} \subset J_{t_2}$ for $0 < t_1 \leqslant t_2$.

[*Step 2*] By compactness, we deduce the existence of some $j^* \in \cap_{t>0} J_t$ (for each $\tau \in \;]0, t]$, pick some $j_\tau \in J_\tau \subset J_t$; take a cluster point for $\tau \downarrow 0$: it is in J_t). We therefore have

$$f_{j^*}(x + td) - f(x) \geqslant t[f'(x, d) - \varepsilon] \quad \text{for all } t > 0,$$

hence $j^* \in J(x)$ (continuity of the convex function f_{j^*} for $t \downarrow 0$). In this inequality, we can replace $f(x)$ by $f_{j^*}(x)$, divide by t and let $t \downarrow 0$ to obtain

$$\sigma_S(d) \geqslant f'_{j^*}(x, d) \geqslant f'(x, d) - \varepsilon.$$

Since $d \in \mathbb{R}^n$ and $\varepsilon > 0$ were arbitrary, (4.4.5) is established. $\qquad\square$

Some comments on the additional assumption are worth mentioning. First, the result concerns $\partial f(x)$, for which it is sufficient to know f only around x. It therefore applies if we have some neighborhood V of x, in which f is representable as

$$f(y) = \sup \{f_j(y) : j \in J(V)\} \quad \text{for all } y \in V,$$

where $J(V)$ is a compact set on which $j \mapsto f_j(y)$ is upper semi-continuous whenever $y \in V$. Secondly, this assumption deals with j only but this is somewhat misleading. The convexity of each f_j actually implies that f is *jointly* upper semi-continuous on $J \times \mathbb{R}^n$.

Finally, the set J is usually a subset of some $\mathbb{R}^p$ and our assumption then implies three properties: J is closed, bounded, and the $f_{(\cdot)}$ are upper semi-continuous. Let us examine what happens when *one* of these properties does not hold.

If J is not closed, we may first have $J(x) = \emptyset$, in which case the formula is of no help. This case does not cause much trouble, though: nothing is changed if J is replaced by its closure, setting

$$\hat{f}_j(x) := \limsup_{j' \to j} f_{j'}(x)$$

for $j \in (\text{cl } J)\backslash J$. A simple example is

$$\mathbb{R} \ni x \mapsto f_j(x) := x - j \quad \text{with} \quad j \in J =]0, 1]. \tag{4.4.7}$$

Closing J places us in a situation in which applying Theorem 4.4.2 is trivial. The other two properties (upper semi-continuity and boundedness) are more fundamental:

Example 4.4.3 [*Upper Semi-Continuity*] Complete (4.4.7) by appending 0 to J and set $f_0(x) \equiv 0$; then $f(x) = x^+$ and

$$J(x) = \begin{cases} \{0\} & \text{if } x \leqslant 0, \\ \emptyset & \text{if } x > 0. \end{cases}$$

Here, $j \mapsto f_j(x)$ is upper semi-continuous at $x = 0$ only; $J(0)$ yields $\partial f_0(0) = \{0\} \subset \partial f(0) = [0, 1]$ and nothing more.

Observe that, introducing the upper semi-continuous hull $\hat{f}_{(\cdot)}(x)$ of $f_{(\cdot)}(x)$ does not change f (and hence ∂f), but it changes the data, since the family now contains the additional function $\hat{f}_0(x) = x^+$. With the functions $\hat{f}_j$ instead of f_j, formula (4.4.4) works.

[*Boundedness*] Take essentially the same functions as in the previous example, but with other notation:

$$f_0(x) \equiv 0, \quad f_j(x) = x - \tfrac{1}{j} \quad \text{for } j = 1, 2, \ldots$$

Now $J = \mathbb{N}$ is closed, and upper semi-continuity of $f_{(\cdot)}(x)$ is automatic; $f(x)$ and $J(x)$ are as before and the same discrepancy occurs at $x = 0$. $\qquad\square$

A special case of Theorem 4.4.2 is when each f_j is differentiable (see Corollary 2.1.4 to remember what it means exactly).

Corollary 4.4.4 *The notation and assumptions are those of Theorem* 4.4.2. *Assume also that each f_j is differentiable; then*

$$\partial f(x) = \mathrm{co}\left\{\nabla f_j(x) \,:\, j \in J(x)\right\}.\qquad\qquad\square$$

A geometric proof of this result was given in Example 3.4, in the simpler situation of finitely many affine functions f_j. Thus, in the framework of Corollary 4.4.4, and whenever there are only finitely many active indices at x, $\partial f(x)$ is a compact convex polyhedron, generated by the active gradients at x.

The case of $J(x)$ being a singleton deserves further comments. We rewrite Corollary 4.4.4 in this case, using a different notation.

Corollary 4.4.5 *For some compact set $Y \subset \mathbb{R}^p$, let $g : \mathbb{R}^n \times Y \to \mathbb{R}$ be a function satisfying the following properties:*

– *for each $x \in \mathbb{R}^n$, $g(x, \cdot)$ is upper semi-continuous;*
– *for each $y \in Y$, $g(\cdot, y)$ is convex and differentiable;*
– *the function $f := \sup_{y \in Y} g(\cdot, y)$ is finite-valued on $\mathbb{R}^n$;*
– *at some $x \in \mathbb{R}^n$, $g(x, \cdot)$ is maximized at a unique $y(x) \in Y$.*

Then f is differentiable at this x, and its gradient is

$$\nabla f(x) = \nabla_x g(x, y(x)) \qquad\qquad (4.4.8)$$

(where $\nabla_x g(x, y)$ denotes the gradient of the function $g(\cdot, y)$ at x). $\qquad\square$

Computing f at a given x amounts to solving a certain maximization problem, to obtain a solution y^*, say (which depends on x!). Then a practical rule is: to obtain the gradient of f, simply differentiate g with respect to x, the variable y being set to this value y^*. If, by any chance, no other $y \in Y$ maximizes g at this x, one does get $\nabla f(x)$. If not, at least a subgradient is obtained (Lemma 4.4.1).

Remark 4.4.6 When Y is a finite set, Corollary 4.4.5 can be easily accepted (see Corollary 4.3.2): when x varies, y^* stays locally the same, just because each $g(\cdot, y)$ is continuous.

When Y is infinite, however, a really baffling phenomenon occurs: although f is a fairly complicated function, its gradient exists (!) and is given by the very simple formula (4.4.8) (!!) This result is perhaps easier to accept if one looks at the following naive calculation.

Suppose that we are in a very favourable situation: Y is some space in which differentiation is possible; $g(\cdot, \cdot)$ is a smooth function; the problem $\max g(x + h, \cdot)$ has a unique solution for h close to 0; and finally, this unique solution $y(\cdot)$ is a smooth function. Then write formally

$$\nabla f(x) = \nabla_x g(x, y(x)) + \nabla_y g(x, y(x)) \nabla y(x)$$

and here comes the trickery: because $y(x)$ is a maximal point, $g(x, \cdot)$ is stationary at $y(x)$; the second term is therefore zero. $\qquad\square$

4.5 Image of a Function Under a Linear Mapping

Let $g : \mathbb{R}^m \to \mathbb{R}$ be a convex function and $A : \mathbb{R}^m \to \mathbb{R}^n$ a linear operator. We recall from §B.2.4 that the associated function

$$\mathbb{R}^n \ni x \mapsto (Ag)(x) := \inf \{g(y) \ : \ Ay = x\} \tag{4.5.1}$$

is convex, provided that, for all x, g is bounded from below on $\overset{-1}{A}(x)$; besides, Ag is finite everywhere if A is surjective. Analogously to (4.4.2), we denote by

$$Y(x) := \{y \in \mathbb{R}^m \ : \ Ay = x, \ g(y) = (Ag)(x)\} \tag{4.5.2}$$

the set of minimizers in (4.5.1).

Theorem 4.5.1 *With the notation* (4.5.1), (4.5.2), *assume A is surjective. Let x be such that $Y(x)$ is nonempty. Then, for arbitrary $y \in Y(x)$,*

$$\partial(Ag)(x) = \{s \in \mathbb{R}^n : A^*s \in \partial g(y)\} = \left(\overset{-1}{A^*}\right)[\partial g(y)] \tag{4.5.3}$$

(and this set is thus independent of the particular optimal y).

Proof. By definition, $s \in \partial(Ag)(x)$ if and only if $(Ag)(x') \geqslant (Ag)(x) + \langle s, x' - x \rangle$ for all $x' \in \mathbb{R}^n$, which can be rewritten

$$(Ag)(x') \geqslant g(y) + \langle s, x' - Ay \rangle \quad \text{for all } x' \in \mathbb{R}^n$$

where y is arbitrary in $Y(x)$. Furthermore, because A is surjective and by definition of Ag, this last relation is equivalent to

$$g(y') \geqslant g(y) + \langle s, Ay' - Ay \rangle = g(y) + \langle A^*s, y' - y \rangle \quad \text{for all } y' \in \mathbb{R}^m$$

which means that $A^*s \in \partial g(y)$. □

The surjectivity of A implies first that $(Ag)(x) < +\infty$ for all x, but it has a more interesting consequence:

Corollary 4.5.2 *Make the assumptions of Theorem 4.5.1. If g is differentiable at some $y \in Y(x)$, then Ag is differentiable at x.*

Proof. Surjectivity of A is equivalent to injectivity of A^*: in (4.5.3), we have an equation in s: $A^*s = \nabla g(y)$, whose solution is unique, and is therefore $\nabla(Ag)(x)$. □

A first example of image-function is when A is a restriction in a product space: g being a convex function on $\mathbb{R}^n \times \mathbb{R}^m$, consider the marginal function, obtained by partial minimization of g:

$$\mathbb{R}^n \ni x \mapsto f(x) := \inf \{g(x, y) \ : \ y \in \mathbb{R}^m\}. \tag{4.5.4}$$

This f is put under the form Ag, if we choose $A : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ defined by $A(x, y) := x$.

Corollary 4.5.3 *Suppose that the subdifferential of g in* (4.5.4) *is associated with a scalar product* $\langle\!\langle \cdot, \cdot \rangle\!\rangle$ *preserving the structure of a product space:*

$$\langle\!\langle (x, y), (x', y') \rangle\!\rangle = \langle x, x' \rangle_n + \langle y, y' \rangle_m \quad \text{for } x, x' \in \mathbb{R}^n \text{ and } y, y' \in \mathbb{R}^m .$$

At a given $x \in \mathbb{R}^n$, *take an arbitrary* y *solving* (4.5.4). *Then*

$$\partial f(x) = \{s \in \mathbb{R}^n : (s, 0) \in \partial_{(x,y)} g(x, y)\} .$$

Proof. With our notation, $A^* s = (s, 0)$ for all $s \in \mathbb{R}^n$. It suffices to apply Theorem 4.5.1 (the symbol $\partial_{(x,y)} g$ is used as a reminder that we are dealing with the subdifferential of g with respect to the variable $(\cdot, \cdot) \in \mathbb{R}^n \times \mathbb{R}^m$).   □

If g is differentiable on $\mathbb{R}^n \times \mathbb{R}^m$ *and* is minimized "at finite distance" in (4.5.4), then the resulting f is differentiable (see Remark 4.5.2). In fact,

$$\nabla_{(x,y)} g(x, y) = (\nabla_x g(x, y), \nabla_y g(x, y)) \in \mathbb{R}^n \times \mathbb{R}^m$$

and the second component $\nabla_y g(x, y)$ is 0 just because y minimizes g at the given x. We do obtain $\nabla f(x) = \nabla_x g(x, y)$ with y solving (4.5.4).

A geometric explanation of this differentiability property appears on Fig. 4.5.1: the shadow of a smooth convex epigraph is normally a smooth convex epigraph.
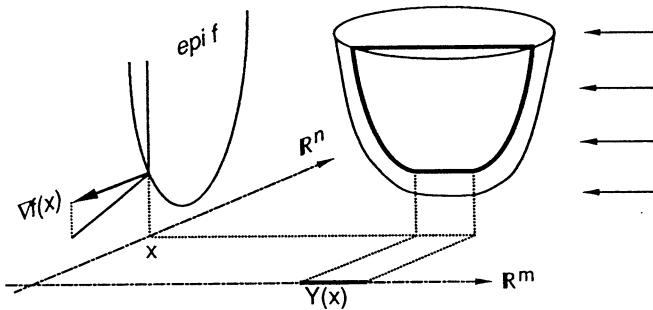
Fig. 4.5.1. The gradient of a marginal function

Remark 4.5.4 The following counter-example emphasizes the necessity of a minimizer y to apply Theorem 4.5.1: the function

$$\mathbb{R}^2 \ni (x, y) \mapsto g(x, y) := \sqrt{x^2 + e^{2y}}$$

is convex (check it), perfectly smooth (C^∞), but "minimal at infinity" (for all x). The resulting marginal function $f(x) = |x|$ is not a smooth function.   □

Another important instance of an image-function was seen in §B.2.3: the infimal convolution of two functions, defined by

$$(f_1 \,\dot{+}\, f_2)(x) := \inf \{f_1(y_1) + f_2(y_2) : y_1, y_2 \in \mathbb{R}^n, \ y_1 + y_2 = x\} . \quad (4.5.5)$$

Recall from the end of §B.2.4 that this operation can be put in the form (4.5.1), by considering

$$\mathbb{R}^n \times \mathbb{R}^n \ni (y_1, y_2) \mapsto g(y_1, y_2) := f_1(y_1) + f_2(y_2) \in \mathbb{R},$$
$$\mathbb{R}^n \times \mathbb{R}^n \ni (y_1, y_2) \mapsto A(y_1, y_2) := y_1 + y_2 \in \mathbb{R}^n.$$

Corollary 4.5.5 *Let f_1 and $f_2 : \mathbb{R}^n \to \mathbb{R}$ be two convex functions minorized by a common affine function. For given x, let (y_1, y_2) be such that the inf-convolution is exact at $x = y_1 + y_2$, i.e.: $(f_1 \mathbin{\dot{\vee}} f_2)(x) = f_1(y_1) + f_2(y_2)$. Then*

$$\partial(f_1 \mathbin{\dot{\vee}} f_2)(x) = \partial f_1(y_1) \cap \partial f_2(y_2). \tag{4.5.6}$$

Proof. First observe that $A^* s = (s, s)$. Also, apply Definition 1.2.1 to see that $(s_1, s_2) \in \partial g(y_1, y_2)$ if and only if $s_1 \in \partial f_1(y_1)$ and $s_2 \in \partial f_2(y_2)$. Then (4.5.6) is just the copy of (4.5.3) in the present context.   □

Once again, we obtain a regularity result (among others): $\nabla(f_1 \mathbin{\dot{\vee}} f_2)(x)$ exists whenever there is an optimal (y_1, y_2) in (4.5.5) for which *either* f_1 *or* f_2 is differentiable. For an illustration, see again Example B.2.3.9, more precisely (B.2.3.6).

Remark 4.5.6 In conclusion, let us give a warning: the max-operation (§4.4) does not destroy differentiability if *uniqueness* of the argmax holds. By contrast, the differentiability of a min-function (§4.5) has nothing to do with uniqueness of the argmin.
    This may seem paradoxical, since maximization and minimization are just the same operation, as far as *differentiability* of the result f is concerned. Observe, however, that the differentiability obtained in §4.5 relies heavily on the *joint* convexity of the underlying function g of Theorem 4.5.1 and Corollary 4.5.3. This last property has little relevance in §4.4.   □

5 Further Examples

With the help of the calculus rules developed in §4, we can study more sophisticated examples than those of §3. They will reveal, in particular, the important role of §4.4: max-functions appear all the time.

5.1 Largest Eigenvalue of a Symmetric Matrix

We adopt the notation of §B.1.3(e): in the space $S_n(\mathbb{R})$ of symmetric $n \times n$ matrices equipped with $\langle\!\langle \cdot, \cdot \rangle\!\rangle$, the function

$$S_n(\mathbb{R}) \ni M \mapsto \lambda_1(M) \in \mathbb{R}$$

is convex and can be represented as

$$\lambda_1(M) = \max \left\{ u^\top M u \ : \ u \in \mathbb{R}^n, \ u^\top u = 1 \right\}. \tag{5.1.1}$$

Furthermore, the set of optimal u in (5.1.1) is the set of normalized eigenvectors associated with the resulting λ_1.

Thus, Corollary 4.4.4 will give the subdifferential of λ_1 if the gradient of the function $M \mapsto u^\top M u$ can be computed. This is easy: by direct calculation, we have

$$u^\top M u = \langle\!\langle uu^\top, M \rangle\!\rangle \tag{5.1.2}$$

so the linear function $u \mapsto u^\top M u$ supports the singleton $\{uu^\top\}$, a rank-one matrix of $S_n(\mathbb{R})$ (its kernel is the subspace orthogonal to u). The subdifferential of λ_1 at M is therefore the convex hull of all these matrices:

$$\partial\lambda_1(M) = \operatorname{co}\left\{uu^\top \ : \ u^\top u = 1, \ Mu = \lambda_1(M)u\right\}. \tag{5.1.3}$$

Naturally, this is the face of $\partial\lambda_1(0)$ exposed by M, where $\partial\lambda_1(0)$ was given in Example C.3.3.9. It is the singleton $\{\nabla\lambda_1(M) = uu^\top\}$ – where u is a normalized corresponding eigenvector – if and only if the maximal eigenvalue λ_1 of M is simple.

The directional derivatives of λ_1 can of course be computed: the support function of (5.1.3) is, using (5.1.2) to reduce superfluous notation,

$$P \mapsto \lambda_1'(M, P) = \max\left\{u^\top P u \ : \ u \text{ normalized eigenvector for } \lambda_1(M)\right\}.$$

Remark 5.1.1 It is tempting to think of the problem as follows.
– There are a finite number of eigenvalues;
– each one is a root of the characteristic polynomial of M, whose coefficients are smooth functions of M,
– therefore, each eigenvalue is a smooth function of M.
– As a result, λ_1 is a max of finitely many smooth functions.

If this reasoning were correct, Corollary 4.4.4 would tell us that $\partial\lambda_1(M)$ is a compact convex polyhedron; this is certainly not the case of the set (5.1.3)! The flaw is that the roots of a polynomial cannot be enumerated. When they are all distinct, each can be followed by continuity; but when two roots coincide, this continuity argument vanishes.   □

As an example, let us study the cone of symmetric negative semi-definite matrices, i.e. the sublevel-set $K^- := \{M \ : \ \lambda_1(M) \leqslant 0\}$. Its boundary is the set of matrices $M \in K^-$ that are singular ($-I_n$ has $\lambda_1(-I_n) < 0$, so Proposition 1.3.3 applies). For such M, Proposition 1.3.4 characterizes the tangent and normal cones to K^- at M:

$$\begin{aligned}
\mathrm{T}_{K^-}(M) &= \{P \ : \ u^\top P u \leqslant 0 \quad \text{for all } u \in \operatorname{Ker} M\}, \\
\mathrm{N}_{K^-}(M) &= \operatorname{co}\{uu^\top \ : \ u \in \operatorname{Ker} M\}.
\end{aligned}$$

If $M \neq 0$, Example A.5.2.6(a) gives a more handy expression for the normal cone:

$$\mathrm{N}_{K^-}(M) = \{P \in S_n(\mathbb{R}) \ : \ P \succcurlyeq 0, \ \langle\!\langle M, P \rangle\!\rangle = 0\}.$$

In problems involving largest eigenvalues, the variable matrix M is often imposed a certain pattern. For example, one considers matrices with fixed off-diagonal elements, only their diagonal being free. In that case, a symmetric matrix M_0 is

given and the function to be studied is $\lambda_1(M_0 + D)$, where D is an arbitrary diagonal $n \times n$ matrix. Identifying the set of such diagonals with $\mathbb{R}^n$, we thus obtain the function

$$f(x) = f(\xi^1, \dots, \xi^n) := \lambda_1\left(M_0 + \operatorname{diag}\left(\xi^1, \dots, \xi^n\right)\right).$$

This f is λ_1 pre-composed with an affine mapping whose linear part is the operator $A_0 : \mathbb{R}^n \to \mathrm{S}_n(\mathbb{R})$ defined by

$$\mathbb{R}^n \ni x = (\xi^1, \dots, \xi^n) \mapsto A_0(x) := \operatorname{diag}\left(\xi^1, \dots, \xi^n\right) \in \mathrm{S}_n(\mathbb{R}).$$

We have

$$\langle\!\langle A_0(x), M \rangle\!\rangle = \sum_{i=1}^n \xi^i M_{ii} \quad \text{for all } x \in \mathbb{R}^n \text{ and } M \in \mathrm{S}_n(\mathbb{R}).$$

Knowing that $\mathbb{R}^n$ is equipped with the usual dot-product, the adjoint of A_0 is therefore defined by

$$x^\top A_0^* M = \sum_{i=1}^n \xi^i M_{ii} \quad \text{for all } x \in \mathbb{R}^n \text{ and } M \in \mathrm{S}_n(\mathbb{R}).$$

Thus, $A_0^* : \mathrm{S}_n(\mathbb{R}) \to \mathbb{R}^n$ appears as the operator that takes an $n \times n$ matrix and makes an n-vector with its diagonal elements. Because the $(i,j)^{\text{th}}$ element of the matrix $uu^\top$ is $u^i u^j$, (5.1.3) gives with the calculus rule (4.2.1)

$$\partial f(x) = \operatorname{co}\left\{ \left((u^1)^2, \dots, (u^n)^2\right) \ : \ u \text{ normalized eigenvector at } f(x)\right\}.$$

5.2 Nested Optimization

Convex functions that are themselves the result of some optimization problem are encountered fairly often. Let us mention, among others, optimization problems issuing from game theory, all kinds of decomposition schemes, semi-infinite programming, optimal control problems in which the state equation is replaced by an inclusion, etc. We consider two examples below: partially linear least-squares problems, and Lagrangian relaxation.

(a) Partially Linear Least-Squares Problems In our first example, there are three vector spaces: $\mathbb{R}^n$, $\mathbb{R}^m$ and $\mathbb{R}^p$, each equipped with its dot-product. A matrix $A(x)$: $\mathbb{R}^m \to \mathbb{R}^p$ and a vector $b(x) \in \mathbb{R}^p$ are given, both depending on the parameter $x \in \mathbb{R}^n$. Then one considers the function

$$\mathbb{R}^n \times \mathbb{R}^m \ni (x, y) \mapsto g(x, y) := \tfrac{1}{2}\|A(x)y - b(x)\|^2. \qquad (5.2.1)$$

The problem is to minimize g, assumed for simplicity to be convex jointly in x and y.

It makes good sense to minimize g hierarchically: first with respect to y (x being fixed), and then minimize the result with respect to x. In other words, the function

$$\mathbb{R}^n \ni x \mapsto f(x) := \min\{g(x,y) : y \in \mathbb{R}^m\}$$

is important for the sake of minimizing g.

Since y is given by a linear least-squares system

$$A^\top(x)[A(x)y - b(x)] = 0 \tag{5.2.2}$$

which has always a solution (not necessarily unique), we are right in the framework of Corollary 4.5.3. Without any assumption on the rank of $A(x)$,

$$\nabla f(x) = [A'(y) - b'](Ay - b). \tag{5.2.3}$$

Here y is any solution of (5.2.2); b' is the matrix whose k^{th} row is the derivative of b with respect to the k^{th} component ξ^k of x; $A'(y)$ is the matrix whose k^{th} row is $y^\top(A'_k)^\top$; A'_k is the derivative of A with respect to ξ^k. Thus, $\nabla f(x)$ is available as soon as (5.2.2) is solved.

(b) Lagrangian Relaxation Our second example is of utmost practical importance for many applications. Consider an optimization problem briefly described as follows. Given a set U and $n+1$ functions $c_0, c_1, \ldots, c_n$ from U to $\mathbb{R}$, consider the problem

$$\left| \begin{array}{l} \sup c_0(u), \quad u \in U, \\ c_i(u) = 0, \text{ for } i = 1, \ldots, n. \end{array} \right. \tag{5.2.4}$$

Associated with this problem is the *Lagrange function*, which depends on $u \in U$ and $x = (\xi^1, \ldots, \xi^n) \in \mathbb{R}^n$, and is defined by

$$g(x,u) := c_0(u) + \sum_{i=1}^n \xi^i c_i(u).$$

Then another important object is the so-called *dual function*

$$f(x) := \sup\{g(x,u) : u \in U\}, \tag{5.2.5}$$

and a relevant problem towards solving (5.2.4) is to minimize f. Needless to say, this dual function is convex, being the supremum of the affine functions $g(\cdot, u)$. Also, $\nabla_x g(x,u) = c(u)$, if $c(u)$ is the vector with components $c_i(u)$, $i = 1, \ldots, n$. In the good cases, when the hypotheses of Theorem 4.4.2 bring no trouble, the subdifferential is given by Corollary 4.4.4:

$$\partial f(x) = \text{co}\{c(u) : u \in U(x)\} \tag{5.2.6}$$

where $U(x) := \{x \in U : g(x,u) = f(x)\}$ is the optimal set in (5.2.5).

According to (5.2.6), the subgradients of f are obtained from the constraint-values at those u solving (5.2.5); at least, the inclusion "$\supset$" always holds in (5.2.6). An $\bar{x}$ minimizing f is characterized by the following condition: for some positive integer $p \leqslant n+1$, there exist $u_1, \ldots, u_p$ in $U(\bar{x})$ and a set of convex multipliers $\alpha = (\alpha_1, \ldots, \alpha_p) \in \Delta_p$ such that

$$g(\bar{x}, u_k) = f(\bar{x}) \quad \text{and} \quad \sum_{k=1}^{p} \alpha_k c(u_k) = 0 \in \mathbb{R}^n \,.$$

In particular, if $g(\bar{x}, \cdot)$ happens to have a unique maximum $\bar{u}$, then $p = 1$, which means that $c(\bar{u}) = 0$.

At a non-optimal x, Theorem 1.3.4 tells us that

$$\text{int}\, T_{Sf(x)}(x) = \{d \in \mathbb{R}^n \ : \ d^\top c(u) < 0 \text{ for all } u \in U(x)\}$$

is the set of d's satisfying $f'(x, d) < 0$; these are called "descent directions".

5.3 Best Approximation of a Continuous Function on a Compact Interval

Let T be a compact interval of $\mathbb{R}$ and φ_0 a real-valued continuous function defined on T. Furthermore, n functions $\varphi_1, \ldots, \varphi_n$ are given in the space $C(T)$ of real-valued continuous functions on T; usually, they are linearly independent. We are interested in finding a linear combination of the φ_i's which approximates φ_0 best, in the sense of the max-norm. In other words, we want to minimize over $\mathbb{R}^n$ the *error-function*

$$f(x) := \max\{|g(x, t)| \ : \ t \in T\} \tag{5.3.1}$$

where g denotes the function (affine in x)

$$g(x, t) := \sum_{i=1}^{n} \xi^i \varphi_i(t) - \varphi_0(t) = [\varphi(t)]^\top x - \varphi_0(t)\,. \tag{5.3.2}$$

Minimizing f is one of the simplest instances of *semi-infinite programming*: optimization problems with finitely many variables but infinitely many constraints.

The error-function is convex and, once more, enters the framework of Corollary 4.4.4 (observe in particular that $|g| = \max\{g, -g\}$). We neglect the case of an x with $f(x) = 0$, which is not particularly interesting: $g(x, \cdot) \equiv 0$, x is optimal anyway. Denoting by H the (usually n-dimensional) subspace of $C(T)$ generated by the φ_i's, we therefore assume $\varphi_0 \notin H$.

Fix an x with $f(x) > 0$, and call $T(x) \subset T$ the set of t yielding the max in (5.3.1); $T(x)$ is nonempty from our assumptions. At each such t, we can define $\varepsilon(t) \in \{-1, +1\}$ by

$$\varepsilon(t)g(x, t) = f(x) \quad \text{for all } t \in T(x)\,.$$

Then, $\partial f(x)$ is the convex combination of all the n-vectors $\varepsilon(t)\varphi(t)$, where t describes $T(x)$. Deriving an optimality condition is then easy with Corollary 4.4.4 and Theorem 4.1.1:

Theorem 5.3.1 *With the notations* (5.3.1), (5.3.2), *suppose* $\varphi_0 \notin H$. *A necessary and sufficient condition for* $\bar{x} = \left(\bar{\xi}^1, \ldots, \bar{\xi}^n\right) \in \mathbb{R}^n$ *to minimize* f *of* (5.3.1) *is that, for some positive integer* $p \leqslant n + 1$, *there exist* p *points* $t_1, \ldots, t_p$ *in* T, p *integers* $\varepsilon_1, \ldots, \varepsilon_p$ *in* $\{-1, +1\}$ *and* p *positive numbers* $\alpha_1, \ldots, \alpha_p$ *such that*

$$\sum_{i=1}^{n} \bar{\xi}^i \varphi_i(t_k) - \varphi_0(t_k) = \varepsilon_k f(\bar{x}) \quad \text{for } k = 1, \ldots, p,$$

$$\sum_{k=1}^{p} \alpha_k \varepsilon_k \varphi_i(t_k) = 0 \quad \text{for } i = 1, \ldots, n$$

$$\left(\text{or equivalently:} \quad \sum_{k=1}^{p} \alpha_k \varepsilon_k \psi(t_k) = 0 \quad \text{for all } \psi \in H \right). \qquad \square$$

Indeed, this example is formally identical to Lagrangian relaxation; the possible differences are usually in the assumptions on T, which plays the role of U in §5.2(b).

6 The Subdifferential as a Multifunction

Section 2 was mainly concerned with properties of the "static" set $\partial f(x)$. Here, we study the properties of this set varying with x, and also with f.

6.1 Monotonicity Properties of the Subdifferential

We have seen in §B.4.1 that the gradient mapping of a differentiable convex function is *monotone*, a concept generalizing to several dimensions that of a nondecreasing function. Now, this monotonicity has its formulation even in the absence of differentiability.

Proposition 6.1.1 *The subdifferential mapping is monotone in the sense that, for all x_1 and x_2 in $\mathbb{R}^n$,*

$$\langle s_2 - s_1, x_2 - x_1 \rangle \geqslant 0 \quad \text{for all } \ s_i \in \partial f(x_i), \ i = 1, 2. \qquad (6.1.1)$$

Proof. The subgradient inequalities

$$f(x_2) \geqslant f(x_1) + \langle s_1, x_2 - x_1 \rangle \quad \text{for all } s_1 \in \partial f(x_1)$$
$$f(x_1) \geqslant f(x_2) + \langle s_2, x_1 - x_2 \rangle \quad \text{for all } s_2 \in \partial f(x_2)$$

give the result simply by addition. $\qquad \square$

A convex function can be "more or less non-affine", according to how much its graph deviates from a hyperplane. We recall, for example, that f is strongly convex with modulus $c > 0$ on a convex set C when, for all x_1, x_2 in C, and all $\alpha \in \,]0, 1[$, it holds

$$f(\alpha x_2 + (1 - \alpha)x_1) \leqslant \alpha f(x_2) + (1 - \alpha)f(x_1) - \frac{c}{2}\alpha(1 - \alpha)\|x_2 - x_1\|^2. \quad (6.1.2)$$

It turns out that this "degree of non-affinity" is also measured by how much (6.1.1) deviates from equality: the next result is to be compared to Theorems B.4.1.1(ii)-(iii) and B.4.1.4, in a slightly different setting: f is now assumed convex but not differentiable.

Theorem 6.1.2 *A necessary and sufficient for a convex function $f : \mathbb{R}^n \to \mathbb{R}$ to be strongly convex with modulus $c > 0$ on a convex set C is: for all x_1, x_2 in C,*

$$f(x_2) \geqslant f(x_1) + \langle s, x_2 - x_1 \rangle + \frac{c}{2} \|x_2 - x_1\|^2 \quad \text{for all } s \in \partial f(x_1), \quad (6.1.3)$$

or equivalently

$$\langle s_2 - s_1, x_2 - x_1 \rangle \geqslant c\|x_2 - x_1\|^2 \quad \text{for all } s_i \in \partial f(x_i), \ i = 1, 2. \quad (6.1.4)$$

Proof. For x_1, x_2 given in C and $\alpha \in \ ]0, 1[$, we will use the notation

$$x^\alpha := \alpha x_2 + (1 - \alpha)x_1 = x_1 + \alpha(x_2 - x_1)$$

and we will prove $(6.1.3) \Rightarrow (6.1.2) \Rightarrow (6.1.4) \Rightarrow (6.1.3)$.

[(6.1.3) $\Rightarrow$ (6.1.2)] Write (6.1.3) with x_1 replaced by $x^\alpha \in C$: for $s \in \partial f(x^\alpha)$,

$$f(x_2) \geqslant f(x^\alpha) + \langle s, x_2 - x^\alpha \rangle + \frac{c}{2} \|x_2 - x^\alpha\|^2 \,,$$

or equivalently

$$f(x_2) \geqslant f(x^\alpha) + (1 - \alpha)\langle s, x_2 - x_1 \rangle + \frac{c}{2}(1 - \alpha)^2 \|x_2 - x_1\|^2 \,.$$

Likewise,

$$f(x_1) \geqslant f(x^\alpha) + \alpha\langle s, x_1 - x_2 \rangle + \frac{c}{2}\alpha^2 \|x_1 - x_2\|^2 \,.$$

Multiply these last two inequalities by α and $(1 - \alpha)$ respectively, and add to obtain

$$\alpha f(x_2) + (1 - \alpha)f(x_1) \geqslant f(x^\alpha) + \frac{c}{2}\|x_2 - x_1\|^2 \left[\alpha(1 - \alpha)^2 + (1 - \alpha)\alpha^2\right].$$

Then realize after simplification that this is just (6.1.2).

[(6.1.2) $\Rightarrow$ (6.1.4)] Write (6.1.2) as

$$\frac{f(x^\alpha) - f(x_1)}{\alpha} + \frac{c}{2}(1 - \alpha)\|x_2 - x_1\|^2 \leqslant f(x_2) - f(x_1)$$

and let $\alpha \downarrow 0$ to obtain $f'(x_1, x_2 - x_1) + \frac{c}{2}\|x_2 - x_1\|^2 \leqslant f(x_2) - f(x_1)$, which implies (6.1.3). Then, copying (6.1.3) with x_1 and x_2 interchanged and adding yields (6.1.4) directly.

[(6.1.4) $\Rightarrow$ (6.1.3)] Apply Theorem 2.3.4 to the one-dimensional convex function $\mathbb{R} \ni \alpha \mapsto \varphi(\alpha) := f(x^\alpha)$:

$$f(x_2) - f(x_1) = \varphi(1) - \varphi(0) = \int_0^1 \langle s^\alpha, x_2 - x_1 \rangle \, \mathrm{d}\alpha \quad (6.1.5)$$

where $s^\alpha \in \partial f(x^\alpha)$ for $\alpha \in [0, 1]$. Then take s_1 arbitrary in $\partial f(x_1)$ and apply (6.1.4): $\langle s^\alpha - s_1, x^\alpha - x_1 \rangle \geqslant c\|x^\alpha - x_1\|^2$ i.e., using the value of x^α,

$$\alpha\langle s^\alpha, x_2 - x_1 \rangle \geqslant \alpha\langle s_1, x_2 - x_1 \rangle + c\alpha^2 \|x_2 - x_1\|^2 \,.$$

The result follows by using this inequality to minorize the integral in (6.1.5). $\qquad \square$

Monotonicity properties of ∂f characterize strictly convex functions in just the same way as they do for strongly convex functions.

Proposition 6.1.3 *A necessary and sufficient condition for a convex function f : $\mathbb{R}^n \to \mathbb{R}$ to be strictly convex on a convex set C is: for all $x_1, x_2 \in C$ with $x_2 \neq x_1$,*

$$f(x_2) > f(x_1) + \langle s, x_2 - x_1 \rangle \quad \text{for all } s \in \partial f(x_1)$$

or equivalently

$$\langle s_2 - s_1, x_2 - x_1 \rangle > 0 \quad \text{for all } s_i \in \partial f(x_i), \ i = 1, 2.$$

Proof. Copy the proof of Theorem 6.1.2 with $c = 0$ and the relevant "$\geqslant$"-signs replaced by strict inequalities. The only delicate point is in the [(6.1.2) $\Rightarrow$ (6.1.4)]-stage: use monotonicity of the difference quotient.   □

6.2 Continuity Properties of the Subdifferential

Our aim here is to study continuity properties of the subdifferential: to what extent can we say that $\partial f(x)$ "varies continuously" with x, or with f? We are therefore dealing with continuity properties of multifunctions, and we refer to §0.5 for the basic terminology.

We already know that $\partial f(x)$ is compact convex for each x, and the next two results concern "global" properties.

Proposition 6.2.1 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. The graph of its subdifferential mapping is closed in $\mathbb{R}^n \times \mathbb{R}^n$.*

Proof. Let (x_k, s_k) be a sequence in $\operatorname{gr} \partial f$ converging to $(x, s) \in \mathbb{R}^n \times \mathbb{R}^n$. We must prove that $(x, s) \in \operatorname{gr} \partial f$, which is easy. We have for all k

$$f(y) \geqslant f(x_k) + \langle s_k, y - x_k \rangle \quad \text{for all } y \in \mathbb{R}^n \, ;$$

pass to the limit on k, using continuity of f and of the scalar product.   □

Proposition 6.2.2 *The mapping ∂f is locally bounded, i.e. the image $\partial f(B)$ of a bounded set $B \subset \mathbb{R}^n$ is a bounded set in $\mathbb{R}^n$.*

Proof. For arbitrary x in B and $s \neq 0$ in $\partial f(x)$, the subgradient inequality implies in particular $f(x + s/\|s\|) \geqslant f(x) + \|s\|$. On the other hand, f is Lipschitz-continuous on the bounded set $B + B(0, 1)$ (Theorem B.3.1.2). Hence $\|s\| \leqslant L$ for some L.   □

Remark 6.2.3 Combining these two results, we obtain a bit more than compact-valuedness of ∂f, namely: the image by ∂f of a *compact* set is compact. In fact, for (x_k) in a compact set, with a subsequence $(x_{k'})$, say, converging to x, take $s_k \in \partial f(x_k)$ and extract the subsequence $(s_{k'})$. From Proposition 6.2.2, a subsequence

of $(s_{k'})$ converges to some $s \in \mathbb{R}^n$; from Proposition 6.2.1, $s \in \partial f(x)$. As another consequence, we obtain for example: the image by ∂f of a *compact connected* set is compact connected.

On the other hand, the image by ∂f of a *convex* set is certainly *not* convex (except for $n = 1$, where convexity and connectedness coincide): take for f the ℓ_1-norm on $\mathbb{R}^2$; the image by ∂f of the unit simplex Δ_2 is the union of two segments which are not collinear.

Concerning the graph of ∂f, the same type of results hold: if $K \subset \mathbb{R}^n$ is compact connected, the set

$$\{(x, s) \in \mathbb{R}^n \times \mathbb{R}^n \; : \; x \in K, \; s \in \partial f(x)\}$$

is compact connected in $\mathbb{R}^n \times \mathbb{R}^n$. Also, it is a "skinny" set because $\partial f(x)$ is a singleton almost everywhere (Theorem B.4.2.3).   □

Thanks to local boundedness, our mapping ∂f takes its values in a compact set when the argument x itself varies in a compact set; the "nice" form (0.5.2) of outer and inner semi-continuity can then be used.

Theorem 6.2.4 *The subdifferential mapping of a convex function* $f \; : \; \mathbb{R}^n \to \mathbb{R}$ *is outer semi-continuous at any* $x \in \mathbb{R}^n$, *i.e.*

$$\forall \varepsilon > 0, \exists \delta > 0 : \quad y \in B(x, \delta) \implies \partial f(y) \subset \partial f(x) + B(0, \varepsilon). \qquad (6.2.1)$$

Proof. Assume for contradiction that, at some x, there are $\varepsilon > 0$ and a sequence $(x_k, s_k)_k$ with

$$\begin{aligned} x_k \to x \quad \text{for } k \to \infty \qquad &\text{and} \\ s_k \in \partial f(x_k), \; s_k \notin \partial f(x) + B(0, \varepsilon) \quad &\text{for } k = 1, 2, \ldots \end{aligned} \qquad (6.2.2)$$

A subsequence of the bounded (s_k) (Proposition 6.2.2) converges to $s \in \partial f(x)$ (Proposition 6.2.1). This contradicts (6.2.2), which implies $s \notin \partial f(x) + B(0, 1/2\varepsilon)$.   □

In terms of directional derivatives, we recover a natural result, if we remember that $f'(\cdot, d)$ is an infimum of continuous functions $[f(\cdot + td) - f(\cdot)]/t$ over $t > 0$:

Corollary 6.2.5 *For* $f \; : \; \mathbb{R}^n \to \mathbb{R}$ *convex, the function* $f'(\cdot, d)$ *is upper semi-continuous: at all* $x \in \mathbb{R}^n$,

$$f'(x, d) = \limsup_{y \to x} f'(y, d) \quad \text{for all } d \in \mathbb{R}^n.$$

Proof. Use Theorem 6.2.4, in conjunction with Proposition C.3.3.7.   □

Remark 6.2.6 If f is differentiable at x, then Theorem 6.2.4 reads as follows: all the subgradients at y tend to $\nabla f(x)$ when y tends to x. The inner semi-continuity then follows: ∂f is actually continuous at x. In particular, if f is differentiable on an open set Ω, then it is continuously differentiable on Ω.

In the general case, however, inner semi-continuity is hopeless: for $n = 1$ and $f(x) := |x|$, ∂f is not inner semi-continuous at 0: $\partial f(0) = [-1, +1]$ is much larger than, say, $\partial f(x) = \{1\}$ when $x > 0$.   □

All the previous results concerned the behaviour of $\partial f(x)$ as varying with x. This behaviour is essentially the same when f varies as well.

Theorem 6.2.7 *Let (f_k) be a sequence of (finite) convex functions converging pointwise to $f : \mathbb{R}^n \to \mathbb{R}$ and let (x_k) converge to $x \in \mathbb{R}^n$. For any $\varepsilon > 0$,*

$$\partial f_k(x_k) \subset \partial f(x) + B(0, \varepsilon) \quad \text{for } k \text{ large enough}.$$

Proof. Let $\varepsilon > 0$ be given. Recall (Theorem B.3.1.4) that the pointwise convergence of (f_k) to f implies its uniform convergence on every compact set of $\mathbb{R}^n$.

First, we establish boundedness: for $s_k \neq 0$ arbitrary in $\partial f_k(x_k)$, we have

$$f_k(x_k + s_k/\|s_k\|) \geqslant f_k(x_k) + \|s_k\|.$$

The uniform convergence of (f_k) to f on $B(x, 2)$ implies for k large enough

$$\|s_k\| \leqslant f(x_k + s_k/\|s_k\|) - f(x_k) + \varepsilon,$$

and the Lipschitz property of f on $B(x, 2)$ ensures that (s_k) is bounded.

Now suppose for contradiction that, for some infinite subsequence, there is some $s_k \in \partial f_k(x_k)$ which is not in $\partial f(x) + B(0, \varepsilon)$. Any cluster point of this (s_k) – and there is at least one – is out of $\partial f(x) + B(0, 1/2 \, \varepsilon)$. Yet, with y arbitrary in $\mathbb{R}^n$, write

$$f_k(y) \geqslant f_k(x_k) + \langle s_k, y - x_k \rangle$$

and pass to the limit (on a further subsequence such that $s_k \to s$): pointwise [resp. uniform] convergence of (f_k) to f at y [resp. around x], and continuity of the scalar product give $f(y) \geqslant f(x) + \langle s, y - x \rangle$. Because y was arbitrary, we obtain the contradiction $s \in \partial f(x)$. □

The differentiable case is worth mentioning:

Corollary 6.2.8 *Let (f_k) be a sequence of (finite) differentiable convex functions converging pointwise to the differentiable $f : \mathbb{R}^n \to \mathbb{R}$. Then ∇f_k converges to ∇f uniformly on every compact set of $\mathbb{R}^n$.*

Proof. Take S compact; suppose for contradiction that there exist $\varepsilon > 0$ and a sequence $(x_k) \subset S$ such that

$$\|\nabla f_k(x_k) - \nabla f(x_k)\| > \varepsilon \quad \text{for } k = 1, 2, \ldots$$

Extracting a subsequence if necessary, we may suppose $x_k \to x \in S$; Theorem 6.2.7 assures that the sequences $(\nabla f_k(x_k))$ and $(\nabla f(x_k))$ both converge to $\nabla f(x)$, implying $0 \geqslant \varepsilon$. □

6.3 Subdifferentials and Limits of Subgradients

One of the main results of the previous section was the outer semi-continuity of the subdifferential: (6.2.1) just means that this latter set contains all the possible limits of subgradients calculated at all neighboring points.

It turns out that some sort of converse can be proved. Indeed consider sequences (y_k) such that f is differentiable at each y_k. Recall from Theorem B.4.2.3 that f is differentiable except possibly on a set of measure zero; call it Δ^c, i.e.

$$y \in \Delta \quad \Longleftrightarrow \quad \partial f(y) = \{\nabla f(y)\}.$$

Thus, even if our given x is not in Δ, we can construct a sequence $(y_k) \subset \Delta$ with $y_k \to x$. The corresponding sequence $(\nabla f(y_k))$ is bounded (by a Lipschitz constant of f around x), so we can extract some cluster point; according to §6.2, any such cluster point lies in $\partial f(x)$. Then we ask the question: how much of $\partial f(x)$ do we cover with all the possible subgradients obtained with this limiting process?

Example 3.4 can be used to illustrate the construction above: with the x of the right part of Fig. 3.1, we let (y_k) be any sequence tending to x and keeping away from the kinky line where $f_1 = f_2$. For example, with the d of the picture, we can take $y_k := x + (-1)^k d/k$, in which case the corresponding sequence $(\nabla f(y_k))$ has two cluster points s_1 and s_2 – and our set of limits is complete: no other sequence of gradients can produce another limit.

Observe in this example that $\partial f(x)$ is the convex hull of the cluster points s_1 and s_2 thus obtained. It turns out that this is always the case. Setting

$$\gamma f(x) := \{s \; : \; \exists (y_k) \subset \Delta \text{ with } y_k \to x, \; \nabla f(y_k) \to s\}, \qquad (6.3.1)$$

it is rather clear that $\gamma f(x) \subset \text{co}\, \gamma f(x) \subset \partial f(x)$. The next result, given without proof, establishes the converse inclusion.

Theorem 6.3.1 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. With the notation* (6.3.1), $\partial f(x) = \text{co}\, \gamma f(x)$ *for all $x \in \mathbb{R}^n$.*   □

In summary, the subdifferential can be reconstructed as the convex hull of all possible limits of gradients at points y_k tending to x. This gives a fourth possible definition of the subdifferential, in addition to 1.1.4, 1.2.1 and 1.3.1.

Remark 6.3.2 As a consequence of the above characterization, we indicate a practical trick to compute subgradients: pretend that the function is differentiable and proceed as usual in differential calculus. This "rule of thumb" can be quite useful in some situations.

For an illustration, consider the example of §5.1. Once the largest eigenvalue λ is computed, together with an associated eigenvector u, just pretend that λ has multiplicity 1 and differentiate formally the equation $Mu = \lambda u$. A differential dM induces the differentials $d\lambda$ and du satisfying $M du + dM\, u = \lambda du + d\lambda\, u$.

We need to eliminate du; for this, premultiply by $u^\top$ and use symmetry of M, i.e. $u^\top M = \lambda u^\top$:

$$\lambda u^\top du + u^\top dM\, u = \lambda u^\top du + d\lambda\, u^\top u.$$

Observing that $u^\top dM\, u = uu^\top dM$, $d\lambda$ is obtained as a linear form of dM:

$$d\lambda = \frac{uu^\top}{u^\top u}\, dM = S\, dM.$$

Moral: if λ is differentiable at M, its gradient is the matrix S; if not, we find the expression of a subgradient (depending on the particular u).

This trick is only heuristic, however, and Remark 4.1.2 warns us against it: for $n = 1$, take $f = f_1 - f_2$ with $f_1(x) = f_2(x) = |x|$; if we "differentiate" naively f_1 and f_2 at 0 and subtract the "derivatives" thus obtained, there are 2 chances out of 3 of ending up with a wrong result.   □

Example 6.3.3 Let d_C be the distance-function to a nonempty closed convex set C. From Example 3.3, we know several things: the kinks of d_C form the set $\Delta^c = \operatorname{bd} C$; for $x \in \operatorname{int} C$, $\nabla d_C(x) = 0$; and

$$\nabla d_C(x) = \frac{x - \operatorname{p}_C(x)}{\|x - \operatorname{p}_C(x)\|} \quad \text{for } x \notin C. \tag{6.3.2}$$

Now take $x_0 \in \operatorname{bd} C$; we give some indications to construct $\partial d_C(x_0)$ via Theorem 6.3.1 (draw a picture).

– First, $\gamma d_C(x_0)$ contains all the limits of vectors of the form (6.3.2), for $x \to x_0$, $x \notin C$. These can be seen to make up the intersection of the normal cone $\operatorname{N}_C(x_0)$ with the unit sphere (the multifunction $x \longmapsto \partial d_C(x)$ is outer semi-continuous).

– If $\operatorname{int} C \neq \emptyset$, append $\{0\}$ to this set; the description of $\gamma d_C(x_0)$ is now complete.

– As seen in Example 3.3, the convex hull of the result must be the truncated cone $\operatorname{N}_C(x_0) \cap B(0, 1)$. This is rather clear in the second case, when $0 \in \gamma d_C(x_0)$; but it is also true even if $\operatorname{int} C = \emptyset$: in fact $\operatorname{N}_C(x_0)$ contains the subspace orthogonal to $\operatorname{aff} C$, which in this case contains two opposite vectors of norm 1. □

To conclude, let us give another result specifying the place in $\partial f(x)$ that can be reached by sequences of (sub)gradients as above.

Proposition 6.3.4 *Let x and $d \neq 0$ be given in $\mathbb{R}^n$. For any sequence $(t_k, s_k, d_k) \subset \mathbb{R}_*^+ \times \mathbb{R}^n \times \mathbb{R}^n$ satisfying*

$$t_k \downarrow 0, \quad s_k \in \partial f(x + t_k d_k), \quad d_k \to d, \quad \text{for } k = 1, 2, \ldots$$

and any cluster point s of (s_k), there holds

$$s \in \partial f(x) \quad \text{and} \quad \langle s, d \rangle = f'(x, d).$$

Proof. The first property comes from the results in §6.2. For the second, use the monotonicity of ∂f:

$$0 \leqslant \langle s_k - s', x + t_k d_k - x \rangle = t_k \langle s_k - s', d_k \rangle \quad \text{for all } s' \in \partial f(x).$$

Divide by $t_k > 0$ and pass to the limit to get $f'(x, d) \leqslant \langle s, d \rangle$. The converse inequality being trivial, the proof is complete. □

In other words, tending to x "tangentially" to some direction d produces a subgradient which is not arbitrary, but which lies in the face $\operatorname{F}_{\partial f(x)}(d)$ exposed by the direction d. When this direction describes the unit sphere, each exposed face of $\partial f(x)$ is visited (Proposition C.3.1.5). From this observation, we realize that the whole boundary of $\partial f(x)$ is thus visited.

Exercises

1. For $f : \mathbb{R}^n \to \mathbb{R}$ convex and $s \in \partial f(x)$, show that $f(x + ts) \geqslant f(x) + t\|s\|^2$ for all $t \geqslant 0$. Deduce that, for such an f, the image by ∂f of a compact set is compact [hint: use bounds of f].

2*. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. For given x, establish the following relations:

$$\min\{c \geqslant 0 \ : \ f(x+d) \geqslant f(x) - c\|d\| \text{ for all } d \in \mathbb{R}^n\} = \min_{s \in \partial f(x)} \|s\| \,;$$

$$\frac{1}{\varepsilon} \sup\{f(x+d) - f(x) \ : \ \|d\| \leqslant \varepsilon\} \geqslant \max_{s \in \partial f(x)} \|s\|, \quad \text{for all } \varepsilon > 0 \,;$$

$$\partial f(x) = \bigcap_{\varepsilon > 0} \overline{\mathrm{co}}\,\{\partial f(x+d) \ : \ \|d\| < \varepsilon\}\,.$$

3*. Let C_1 and C_2 be convex and set $C := C_1 + C_2$. For $x_i \in C_i$, $i = 1, 2$, take $x := x_1 + x_2 \in C$.
– Show that $\mathrm{N}_C(x) = \mathrm{N}_{C_1}(x_1) \cap \mathrm{N}_{C_2}(x_2)$.
– Deduce an expression of $\partial \mathrm{d}_C(x)$ using $\partial \mathrm{d}_{C_i}(x_i)$, $i = 1, 2$.

4*. Let $\varphi : \mathbb{R} \to \mathbb{R}$ be nondecreasing and set $f(x) := \int_0^x \varphi(t)\mathrm{d}t$. Show that $\partial f(x) = [\varphi(x^-), \varphi(x^+)]$ (the segments whose endpoints are the left- and right-limits of φ at x).

5*. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. For given $x \in \mathbb{R}^n$, define

$$\ell_f(x) := \min\left\{0, \sup_{x' \neq x} \frac{f(x') - f(x)}{\|x' - x\|}\right\}.$$

Show that $\ell_f(x) = \min\{\|s\| : s \in \partial f(x)\}$.

6**. Let $f : \mathbb{R}^2 \to \mathbb{R}$ be convex and twice continuously differentiable, with $\det \nabla^2 f(x) = 1$ for all $x \in \mathbb{R}^2$. Show that ∇f is actually constant over $\mathbb{R}^2$ (and hence f is a strongly convex quadratic function); this is a theorem due to K. Jörgens (1954).

7*. Let $f, g : \mathbb{R}^n \to \mathbb{R}$ be two convex functions such that $\partial f(x) \cap \partial g(x) \neq \emptyset$ for all $x \in \mathbb{R}^n$. Show that $f - g$ is actually constant.

8. Let $f, g : \mathbb{R}^n \to \mathbb{R}$ be two convex functions such that $f + g$ is differentiable. Show that f and g are actually differentiable.

9*. Let $f :]a, b[\to]0, +\infty[$ be concave. Show that $g := 1/f$ is convex and

$$\partial(-f)(x) = [f(x)]^2 \partial g(x) \quad \text{for all } x \in]a, b[\,.$$

10**. Consider the affine function

$$\mathbb{R}^n \ni x \mapsto f_t(x) := tx + \sqrt{1 - t^2}\,,$$

where t is a parameter in $T := [-1, +1]$.
    What is the function $\sup_{t \in T} f_t$? Illustrate on this example the formula giving the subdifferential of a supremum.

11**. Consider the convex functions defined on $\mathbb{R}^2$ by

$$x = (\xi, \eta) \mapsto \begin{cases} f_1(x) := \xi^2, \\ f_2(x) := \alpha\xi^2 + \eta, \end{cases}$$

where $\alpha \in]0, 1[$ is a fixed parameter. Set $f(x) := \max\{f_1(x), f_2(x)\}$.
– Show directly that 0 minimizes f; compute $\partial f(0)$ and realize that $0 \in \partial f(0)$.
– Compute directly the directional derivative $f'(0, \cdot)$; realize that it coincides with $\sigma_{\partial f(0)}(\cdot)$.
– Show that the "second-order difference quotient"

$$\mathbb{R}^2 \ni 0 \neq h \mapsto q_2(h) := \frac{f(h) - f(0) - f'(0, h)}{\frac{1}{2}\|h\|^2}$$

has a limit when $h = td$ with d fixed in $\mathbb{R}^2$ and $t \downarrow 0$. Compute this limit as a function of d and α. Study more particularly the case $d = (\pm 1, 0)$.
– For given ξ, compute $\eta(\xi)$ solving $f_1(\xi, \eta) = f_2(\xi, \eta)$. Show that $q_2(\xi, \eta(\xi))$ has a limit when $\xi \to 0$.
– Compute the function

$$\mathbb{R} \ni \xi \mapsto \ell(\xi) := \min\{f(\xi, \eta) : \eta \in \mathbb{R}\}.$$

Show that $\ell(\xi) = f(\xi, \eta)$ for any $\eta \leqslant \eta(\xi)$. Give the first and second derivatives of ℓ at 0.

12*. Let f and g be two convex increasing functions from $\mathbb{R}$ to $\mathbb{R}^+$.
– Show that fg is convex increasing;
– check that $\partial(fg)(x) = f(x)\partial g(x) + g(x)\partial f(x)$ for all $x \in \mathbb{R}$.

13. Let C be convex compact and $d \neq 0$. Show that $\partial\sigma_C(d)$ is the face of C exposed by d.

14. Let $f : \mathbb{R}^n \to \mathbb{R}$ be such that $\partial f(x) \neq \emptyset$ for all $x \in \mathbb{R}^n$. Show that f is convex.

15*. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and take $r < f(x)$. Denote by $(\bar{x}, \bar{r})$ the projection of (x, r) onto epi f. Show that $f(\bar{x}) = \bar{r}$ and

$$\frac{x - \bar{x}}{f(\bar{x}) - r} \in \partial f(\bar{x})$$

(draw a picture).

16*. This exercise illustrates §5.2(b). For given positive numbers V, $v_1, \ldots, v_m$, $c_1, \ldots, c_m$, consider the convex, finite-valued function $f : \mathbb{R} \to \mathbb{R}$ defined by

$$f(x) := \max\left\{ Vx + \sum_{i=1}^{m}(c_i - xv_i)u_i : u_i \in \{0, 1\} \text{ for } i = 1, \ldots, m \right\}. \quad (*_x)$$

(i) Compute the subdifferential of f, depending on the position of x with respect to the ratios c_i/v_i.

(ii) Show that f and ∂f coincide with the values and subdifferentials of the function defined by

$$\max \left\{ Vx + \sum_{i=1}^{m} (c_i - xv_i)u_i \ : \ u_i \in [0,1] \text{ for } i = 1, \dots, m \right\}. \qquad (**_x)$$

(iii) Find the necessary and sufficient condition on V so that f is bounded from below and attains its minimum.

(iv) Express the optimality condition for an $\bar{x}$ minimizing f, in terms of the u's optimal in $(*_{\bar{x}})$ [resp. $(**_{\bar{x}})$].

*17**. This is also an illustration of §5.2(b), in a more complicated situation. All spaces $\mathbb{R}^k$ considered are equipped with the standard dot product. For $x \in \mathbb{R}^n$ we form

– the symmetric $m \times m$ matrix $Q(x) := Q_0 + \sum_{i=1}^{n} x_i Q_i$,
– the m-vector $b(x) := b_0 + \sum_{i=1}^{n} x_i b_i$,
– the number $c(x) := c_0 + \sum_{i=1}^{n} x_i c_i$,

where each $Q_i \in S_m(\mathbb{R})$, $b_i \in \mathbb{R}^m$ and $c_i \in \mathbb{R}$ is given. Then we define the function $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$ by

$$x \mapsto f(x) := \sup \left\{ c(x) + 2b(x)^\top u - u^\top Q(x)u \ : \ u \in \mathbb{R}^m \right\}.$$

What is its domain? Show that $\mathrm{int}\,\mathrm{dom}\, f = \Omega := \{x \in \mathbb{R}^n : Q(x) \succ 0\}$. Compute $\partial f(x)$ for $x \in \Omega$; deduce that f is continuously differentiable on Ω.

Now let H be a subspace of $\mathbb{R}^m$ and define the function

$$x \mapsto f_H(x) := \sup \left\{ c(x) + 2b(x)^\top u - u^\top Q(x)u \ : \ u \in H \right\}.$$

What is the domain of f_H and its interior Ω_H? Show that $\inf f_H \leqslant \inf f$.

Suppose H is given explicitly as $H = \{x : Ax = 0\}$, with $A : \mathbb{R}^m \to \mathbb{R}^p$ linear. Define the function $\varphi \in \overline{\mathrm{Conv}}\, \mathbb{R}^{n \times p}$, whose value at (x, y) is

$$\varphi(x,y) := \sup \left\{ c(x) + 2b(x)^\top u - u^\top Q(x)u + 2y^\top A u \ : \ u \in \mathbb{R}^m \right\}.$$

What is $\mathrm{dom}\,\varphi$? Show that φ is differentiable on $\Omega \times \mathbb{R}^p$ and deduce

$$f_H(x) = \inf_{y \in \mathbb{R}^p} \varphi(x,y)\,, \quad \text{for all } x \in \Omega\,,$$

hence

$$\inf_{(x,y) \in \mathbb{R}^n \times \mathbb{R}^p} \varphi(x,y) = \inf_{x \in \Omega} f_H(x) \geqslant \inf_{x \in \mathbb{R}^n} f_H(x)\,.$$

18. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex piecewise affine. Show that, given $x \in \mathbb{R}^n$, there is $\varepsilon > 0$ such that $\partial f(x') \subset \partial f(x)$ whenever $x' \in B(x, \varepsilon)$.

E. Conjugacy in Convex Analysis

Introduction In classical real analysis, the gradient of a differentiable function f : $\mathbb{R}^n \to \mathbb{R}$ plays a key role – to say the least. Considering this gradient as a mapping $x \mapsto s(x) = \nabla f(x)$ from (some subset X of) $\mathbb{R}^n$ to (some subset S of) $\mathbb{R}^n$, an interesting object is then its inverse: to a given $s \in S$, associate the $x \in X$ such that $s = \nabla f(x)$. This question may be meaningless: not all mappings are invertible! but could for example be considered locally, taking for $X \times S$ a neighborhood of some $(x_0, s_0 = \nabla f(x_0))$, with $\nabla^2 f$ continuous and invertible at x_0 (use the local inverse theorem).

Let us skip for the moment all such technicalities. Geometrically, we want to find x such that the hyperplane in $\mathbb{R}^n \times \mathbb{R}$ defined by the given $(s, -1)$, and passing through $(x, f(x))$, is tangent to gr f at x; the problem is meaningful when this x exists and is unique. Its construction is rather involved but analytically, an amazing fact is that the new mapping $x(\cdot) = (\nabla f)^{-1}$ thus defined is itself a *gradient* mapping: say $x(\cdot) = \nabla h$, with $h : S \subset \mathbb{R}^n \to \mathbb{R}$. Even more surprising: this function h has a simple expression, namely

$$S \ni s \mapsto h(s) = \langle s, x(s) \rangle - f(x(s)) \,. \tag{0.1}$$

To explain this, do a formal a posteriori calculation in (0.1): a differential $\mathrm{d}s$ induces the differentials $\mathrm{d}x$ and $\mathrm{d}h$, which are linked by the relation

$$\mathrm{d}h = \langle s, \mathrm{d}x \rangle + \langle \mathrm{d}s, x \rangle - \langle \nabla f(x), \mathrm{d}x \rangle = \langle s, \mathrm{d}x \rangle + \langle \mathrm{d}s, x \rangle - \langle s, \mathrm{d}x \rangle = \langle \mathrm{d}s, x \rangle$$

and this *defines* x as the gradient of h with respect to s.

Remark 0.1 When ∇f is invertible, the function

$$s \mapsto h(s) = \langle s, (\nabla f)^{-1}(s) \rangle - f\big((\nabla f)^{-1}(s)\big)$$

is called the *Legendre transform* (relative to f). It should never be forgotten that, from the above motivation itself, it is not really the function h which is primarily interesting, but rather its gradient ∇h. □

The gradients of f and of h are inverse to each other by definition, so they establish a reciprocal correspondence:

$$s = \nabla f(x) \quad \Longleftrightarrow \quad x = \nabla h(s) \,. \tag{0.2}$$

In particular, applying the Legendre transform to h, we have to get back f. This symmetry appears in the expression of h itself: (0.1) tells us that, for x and s related by (0.2), $f(x) + h(s) = \langle s, x \rangle$.

Once again, the above statements are rather formal, insofar as they implicitly assume that $(\nabla f)^{-1}$ is well-defined. Convex analysis, however, provides a nice framework to give this last operation a precise meaning.

First of all, observe that the mapping $x \mapsto \nabla f(x)$ is now replaced by a *set-valued* mapping $x \longmapsto \partial f(x)$ – see Chap. D. To invert it is to find x such that $\partial f(x)$ contains a given s; and we can accept a nonunique such x: a set-valued $(\partial f)^{-1}$ will be obtained, but the price has been already paid anyway.

Second, the construction of $x(s)$ is now much simpler: $s \in \partial f(x)$ means that $0 \in \partial f(x) - \{s\}$ and, thanks to convexity, the last property means that x minimizes $f - \langle s, \cdot \rangle$ over $\mathbb{R}^n$. In other words, to find $x(s)$, we have to solve

$$\inf \left\{ f(x) - \langle s, x \rangle \ : \ x \in \mathbb{R}^n \right\}. \tag{0.3}$$

The Legendre transform – in the classical sense of the term – is well-defined when this problem has a unique solution. At any rate, note that the above *infimal value* is always well-defined (possibly $-\infty$); it is a (nonambiguous) function of s, and this is the basis for a nice generalization.

Let us sum up: if f is convex, (0.3) is a possible way of defining the Legendre transform when it exists unambiguously. It is easy to see that the latter holds when f satisfies three properties:

– differentiability – so that there is something to invert;
– strict convexity – to have uniqueness in (0.3);
– $\nabla f(\mathbb{R}^n) = \mathbb{R}^n$ – so that (0.3) does have a solution for all $s \in \mathbb{R}^n$; this latter property essentially means that, when $\|x\| \to \infty$, $f(x)$ increases faster than any linear function: f is *1-coercive*.

In all other cases, there is no well-defined Legendre transform; but then, the transformation implied by (0.3) can be taken as a new definition, generalizing the initial inversion of ∇f. We can even extend this definition to *nonconvex* f, namely to any function such that (0.3) is meaningful! Finally, an important observation is that the infimal value in (0.3) is a concave, and even closed, function of s; this is Proposition B.2.1.2: the infimand is affine in s, f has little importance, $\mathbb{R}^n$ is nothing more than an index set.

The concept of conjugacy in convex analysis results from all the observations above, and has a number of uses:

– It is often useful to simplify algebraic calculations;
– it plays an important role in deriving duality schemes for convex minimization problems;
– it is also a basic operation for formulating variational principles in optimization (convex or not);
– and this has applications in other areas of applied mathematics, such as probability, statistics, nonlinear elasticity, economics, etc.

1 The Convex Conjugate of a Function

1.1 Definition and First Examples

As suggested by (0.3), conjugating a function f essentially amounts to minimizing a perturbation of it. There are two degenerate situations we want to avoid:

– the result is $+\infty$ for some s; observe that this is the case if and only if the result is $+\infty$ for all s;

– the result is $-\infty$ for all s.

Towards this end, we assume throughout that $f : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ (not necessarily convex) satisfies

$$\boxed{f \not\equiv +\infty, \text{ and there is an affine function minorizing } f \text{ on } \mathbb{R}^n .} \qquad (1.1.1)$$

Note in particular that this implies $f(x) > -\infty$ for all x. As before, we use the notation $\operatorname{dom} f := \{x : f(x) < +\infty\} \neq \emptyset$. We know from Proposition B.1.2.1 that (1.1.1) holds for example if $f \in \operatorname{Conv} \mathbb{R}^n$.

Definition 1.1.1 The *conjugate* of a function f satisfying (1.1.1) is the function f^* defined by

$$\mathbb{R}^n \ni s \mapsto f^*(s) := \sup \left\{ \langle s, x \rangle - f(x) \ : \ x \in \operatorname{dom} f \right\}. \qquad (1.1.2)$$

For simplicity, we may also let x run over the whole space instead of $\operatorname{dom} f$.

   The mapping $f \mapsto f^*$ will often be called the *conjugacy* operation, or the *Legendre-Fenchel* transform. □

   A very first observation is that a conjugate function is associated with a scalar product on $\mathbb{R}^n$: f^* changes if $\langle \cdot, \cdot \rangle$ changes. Of course, note also with relation to (0.3) that $f^*(s) = - \inf \{f(x) - \langle s, x \rangle : x \in \operatorname{dom} f\}$.

   As an immediate consequence of (1.1.2), we have for all $(x, s) \in \operatorname{dom} f \times \mathbb{R}^n$

$$f^*(s) + f(x) \geqslant \langle s, x \rangle . \qquad (1.1.3)$$

Furthermore, this inequality is obviously true if $x \notin \operatorname{dom} f$: it does hold for all $(x, s) \in \mathbb{R}^n \times \mathbb{R}^n$ and is called *Fenchel's inequality*. Another observation is that $f^*(s) > -\infty$ for all $s \in \mathbb{R}^n$; also, if $x \mapsto \langle s_0, x \rangle + r_0$ is an affine function smaller than f, we have

$$-f^*(s_0) = \inf_x [f(x) - \langle s_0, x \rangle] \geqslant r_0, \quad \text{i.e.} \quad f^*(s_0) \leqslant -r_0 < +\infty .$$

In a word, f^* satisfies (1.1.1).

   Thus, $\operatorname{dom} f^*$ is the set of slopes of all the possible affine functions minorizing f over $\mathbb{R}^n$. Likewise, any $s \in \operatorname{dom} f$ is the slope of an affine function smaller than f^*.

Theorem 1.1.2 *For f satisfying* (1.1.1), *the conjugate f^* is a closed convex function: $f^* \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$.*

Proof. See Example B.2.1.3. □

Example 1.1.3 (Convex Quadratic Functions) Let Q be a symmetric positive definite linear operator on $\mathbb{R}^n$, $b \in \mathbb{R}^n$ and consider

$$f(x) := \tfrac{1}{2}\langle x, Qx \rangle + \langle b, x \rangle \quad \text{for all } x \in \mathbb{R}^n . \tag{1.1.4}$$

A straightforward calculation gives the optimal $x = Q^{-1}(s - b)$ in the defining problem (1.1.2), and the resulting f^* is

$$f^*(s) = \tfrac{1}{2}\langle s - b, Q^{-1}(s - b) \rangle \quad \text{for all } s \in \mathbb{R}^n .$$

In particular, the function $1/2 \,\| \cdot \|^2$ is its own conjugate.

Needless to say, the Legendre transformation is present here, in a particularly simple setting: ∇f is the affine mapping $x \mapsto Qx + b$, its inverse is the affine mapping $s \mapsto Q^{-1}(s - b)$, the gradient of f^*. What we have done is a parametrization of $\mathbb{R}^n$ via the change of variable $s = \nabla f(x)$; with respect to the new variable, f is given by

$$f(x(s)) = \tfrac{1}{2}\langle Q^{-1}(s - b), s - b \rangle + \langle b, Q^{-1}(s - b) \rangle .$$

Adding $f^*(s)$ on both sides gives $f(x(s)) + f^*(s) = \langle Q^{-1}(s - b), s \rangle = \langle x(s), s \rangle$, which illustrates (0.1).

Taking $b = 0$ and applying Fenchel's inequality (1.1.3) gives

$$\langle x, Qx \rangle + \langle s, Q^{-1}s \rangle \geqslant 2\langle s, x \rangle \quad \text{for all } (x, s) ,$$

a generalization of the well-known inequality (obtained for $Q = c\mathrm{I}, c > 0$):

$$c\|x\|^2 + \tfrac{1}{c}\|s\|^2 \geqslant 2\langle s, x \rangle \quad \text{for all } c > 0 \text{ and } x, s \in \mathbb{R}^n . \quad □$$

When Q in the above example is merely positive semi-definite, a meaning can still be given to $(\nabla f)^{-1}$ provided that two problems are taken care of:
– first, s must be restricted to $\nabla f(\mathbb{R}^n)$, which is the affine manifold $b + \mathrm{Im}\, Q$;
– second, $\nabla f(x + y) = \nabla f(x)$ for all $y \in \mathrm{Ker}\, Q$, so $(\nabla f)^{-1}(s)$ can be defined only up to the subspace $\mathrm{Ker}\, Q$.
Using the definition of the conjugate function, we then obtain the following:

Example 1.1.4 (Convex Degenerate Quadratic Functions) Take again the quadratic function (1.1.4), but with Q symmetric positive semi-definite. The supremum in (1.1.2) is finite only if $s - b \in (\mathrm{Ker}\, Q)^\perp$, i.e. $s - b \in \mathrm{Im}\, Q$; it is attained at an x such that $s - \nabla f(x) = 0$ (optimality condition D.2.2.1). In a word, we obtain

$$f^*(s) = \begin{cases} +\infty & \text{if } s \notin b + \mathrm{Im}\, Q , \\ \tfrac{1}{2}\langle x, s - b \rangle & \text{otherwise} , \end{cases}$$

where x is any element satisfying $Qx + b = s$. This formulation can be condensed to

$$f^*(Qx + b) = \tfrac{1}{2}\langle x, Qx \rangle \quad \text{for all } x \in \mathbb{R}^n ,$$

which is one more illustration of (0.1): add $f(x)$ to both sides and obtain

$$f(x) + f^*(Qx + b) = \langle x, Qx + b \rangle = \langle x, \nabla f(x) \rangle .$$

It is also interesting to express f^* in terms of a *pseudo-inverse*: for example, Q^- denoting the Moore-Penrose pseudo-inverse of Q (see §0.3.4), f^* can be written

$$f^*(s) = \begin{cases} \tfrac{1}{2}\langle s - b, Q^-(s - b) \rangle & \text{if } s \in b + \operatorname{Im} Q , \\ +\infty & \text{otherwise} . \end{cases} \qquad \square$$

In the above example, take $b = 0$ and let $Q = \mathrm{p}_H$ be the orthogonal projection onto a subspace H of $\mathbb{R}^n$: $f(x) = 1/2\,\langle x, \mathrm{p}_H x \rangle = 1/2\,\|\mathrm{p}_H x\|^2$. Then $\operatorname{Im} Q = H$ and $Q^- = Q$, so

$$\left(\tfrac{1}{2}\|\mathrm{p}_H \cdot \|^2 \right)^*(s) = \begin{cases} \tfrac{1}{2}\|s\|^2 & \text{if } s \in H , \\ +\infty & \text{otherwise} . \end{cases}$$

Another interesting example is when Q is a rank-one operator, i.e. $Q = uu^\top$, with $0 \neq u \in \mathbb{R}^n$ (we assume the usual dot-product for $\langle \cdot, \cdot \rangle$). Then $\operatorname{Im} Q = \mathbb{R}u$ and, for $x \in \operatorname{Im} Q$, $Qx = \|u\|^2 x$. Therefore,

$$\left[\tfrac{1}{2}\left(u^\top \cdot \right)^2 \right]^*(s) = \begin{cases} \frac{1}{2\|u\|^2} s & \text{if } s \text{ and } u \text{ are collinear} , \\ +\infty & \text{otherwise} . \end{cases}$$

More on conjugacy involving quadratic functions will be seen in Example 2.1.4 and in §3.2.

Example 1.1.5 Let i_C be the indicator function of a nonempty set $C \subset \mathbb{R}^n$. Then

$$(\mathrm{i}_C)^*(s) = \sup_{x \in \operatorname{dom} \mathrm{i}_C} \left[\langle s, x \rangle - \mathrm{i}_C(x) \right] = \sup_{x \in C} \langle s, x \rangle$$

is just the support function of C. If C is a closed convex cone, we conclude from Example C.2.3.1 that $(\mathrm{i}_C)^*$ is the indicator of its polar C°. If C is a subspace, $(\mathrm{i}_C)^*$ is the indicator of its orthogonal $C^\perp$.

If $C = \mathbb{R}^n$, $C^\circ = \{0\}$; $\mathrm{i}_C \equiv 0$ and the conjugate of i_C is 0 at 0, $+\infty$ elsewhere (this is the indicator of $\{0\}$, or the support of $\mathbb{R}^n$).

A similar example is the nonconvex function $x \mapsto f(x) = \|x\|^{1/2}$, where $\| \cdot \|$ is some norm. Then $f^*(0) = 0$; but if $s \neq 0$, take x of the form $x = ts$ to realize that

$$f^*(s) \geqslant \sup_{t \geqslant 0} \left[t\|s\|^2 - \sqrt{t\,\|s\|} \right] = +\infty .$$

In other words, f^* is still the indicator function of $\{0\}$. The conjugacy operation has ignored the difference between f and the zero function, simply because f increases at infinity more slowly than any linear function. Note at this point that the conjugate of f^* (the "biconjugate" of f) is the zero-function. $\qquad \square$

1.2 Interpretations

The computation of f^* can be illustrated in the graph-space $\mathbb{R}^n \times \mathbb{R}$. For given $s \in \mathbb{R}^n$, consider the family of affine functions $x \mapsto \langle s, x \rangle - r$, parametrized by $r \in \mathbb{R}$. They correspond to affine hyperplanes $H_{s,r}$ orthogonal to $(s, -1) \in \mathbb{R}^{n+1}$; see Fig. 1.2.1. From (1.1.1), $H_{s,r}$ is below gr f whenever (s, r) is properly chosen, namely $s \in \text{dom } f^*$ and r is large enough. To construct $f^*(s)$, we lift $H_{s,r}$ as much as possible subject to supporting gr f. Then, admitting that there is contact at some $(x, f(x))$, we write $\langle s, x \rangle - r = f(x)$, or rather, $r = \langle s, x \rangle - f(x)$, to see that $r = f^*(s)$. This means that the best $H_{s,r}$ intersects the vertical axis $\{0\} \times \mathbb{R}$ at the altitude $-f^*(s)$.
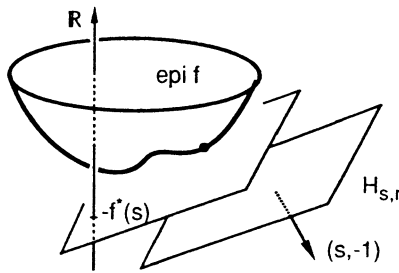
Fig. 1.2.1. Computation of f^* in the graph-space

Naturally, the horizontal hyperplanes $H_{0,r}$ correspond to minimizing f:

$$-f^*(0) = \inf \left\{ f(x) \; : \; x \in \mathbb{R}^n \right\}.$$

The picture illustrates another definition of f^*: being normal to $H_{s,r}$, the vector $(s, -1)$ is also normal to gr f (more precisely to epi f) at the contact when it exists.

Proposition 1.2.1 *There holds for all $x \in \mathbb{R}^n$*

$$f^*(s) = \sigma_{\text{epi } f}(s, -1) = \sup \left\{ \langle s, x \rangle - r \; : \; (x, r) \in \text{epi } f \right\}. \tag{1.2.1}$$

It follows that the support function of epi f *has the expression*

$$\sigma_{\text{epi } f}(s, -u) = \begin{cases} u f^* \left(\frac{1}{u} s \right) & \text{if } u > 0, \\ \sigma_{\text{epi } f}(s, 0) = \sigma_{\text{dom } f}(s) & \text{if } u = 0, \\ +\infty & \text{if } u < 0. \end{cases} \tag{1.2.2}$$

Proof. In (1.2.1), the right-most term can be written

$$\sup_x \; \sup_{r \geqslant f(x)} \; [\langle s, x \rangle - r] = \sup_x \left[\langle s, x \rangle - f(x) \right]$$

and the first equality is established. As for (1.2.2), the case $u < 0$ is trivial; when $u > 0$, use the positive homogeneity of support functions to get

$$\sigma_{\text{epi } f}(s, -u) = u\sigma_{\text{epi } f}\left(\tfrac{1}{u}s, -1\right) = uf^*\left(\tfrac{1}{u}s\right) ;$$

finally, for $u = 0$, we have by definition

$$\sigma_{\text{epi } f}(s, 0) = \sup\left\{\langle s, x\rangle \ : \ (x, r) \in \text{epi } f \text{ for some } r \in \mathbb{R}\right\},$$

and we recognize $\sigma_{\text{dom } f}(s)$.   □

Assume $f \in \overline{\text{Conv}}\,\mathbb{R}^n$. This result, illustrated in Fig. 1.2.1, confirms that the contact-set between the optimal hyperplane and epi f is the face of epi f exposed by the given $(s, -1)$. From (1.2.2), we see that $\sigma_{\text{epi } f}$ and the perspective-function of f^* (§B.2.2) coincide for $u \neq 0$ – up to the change of variable $u \mapsto -u$. As a closed function, $\sigma_{\text{epi } f}$ therefore coincides (still up to the change of sign in u) with the closure of the perspective of f^*. As for $u = 0$, we obtain a relation which will be used on several occasions:

Proposition 1.2.2 *For $f \in \overline{\text{Conv}}\,\mathbb{R}^n$,*

$$\sigma_{\text{dom } f}(s) = \sigma_{\text{epi } f}(s, 0) = (f^*)'_\infty(s) \quad \text{for all } s \in \mathbb{R}^n . \tag{1.2.3}$$

Proof. Use direct calculations; or see Proposition B.2.2.2 and the calculations in Example B.3.2.3.   □

The additional variable u introduced in Proposition 1.2.1 gives a second geometric construction: Fig. 1.2.1 plays the role of Fig. C.2.1.1, knowing that $\mathbb{R}^n$ is replaced here by $\mathbb{R}^{n+1}$. Here again, note in passing that the closed convex hull of epi f gives the same f^*. Now, playing the same game as we did for Interpretation C.2.1.6, f^* can be looked at from the point of view of projective geometry, $\mathbb{R}^n$ being identified with $\mathbb{R}^n \times \{-1\}$.

Consider the set epi $f \times \{-1\} \subset \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}$, i.e. the copy of epi f, translated down vertically by one unit. It generates a cone $K_f \subset \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}$:

$$K_f := \{t(x, r, -1) \ : \ t > 0, \ (x, r) \in \text{epi } f\} . \tag{1.2.4}$$

Now take the polar cone $(K_f)^\circ \subset \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}$. We know from Interpretation C.2.1.6 that it is the epigraph (in $\mathbb{R}^{n+1} \times \mathbb{R}$!) of the support function of epi f. In view of Proposition 1.2.1, its intersection with $\mathbb{R}^n \times \{-1\} \times \mathbb{R}$ is therefore epi f^*. A short calculation confirms this:

$$\begin{aligned}
(K_f)^\circ &= \{(s, \alpha, \beta) \in \mathbb{R}^n \times \mathbb{R} \times \mathbb{R} \ : \ t\langle s, x\rangle + t\alpha r - t\beta \leqslant 0 \\
&\qquad\qquad\qquad\qquad \text{for all } (x, r) \in \text{epi } f \text{ and } t > 0\} \\
&= \{(s, \alpha, \beta) \ : \ \langle s, x\rangle + \alpha r \leqslant \beta \text{ for all } (x, r) \in \text{epi } f\} .
\end{aligned}$$

Imposing $\alpha = -1$, we just obtain the translated epigraph of the function described in (1.2.1). Remembering Example 1.1.5, we see how the polarity correspondence is the basis for conjugacy.

Figure 1.2.2 illustrates the construction, with $n = 1$ and epi f drawn horizontally; this epi f plays the role of S in Fig. C.2.1.2. Note once more the symmetry:
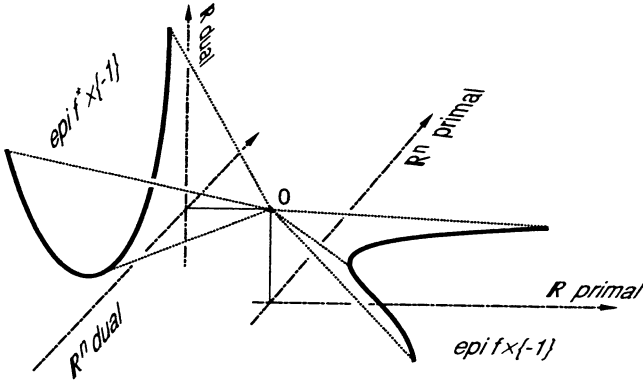
Fig. 1.2.2. A projective view of conjugacy

K_f [resp. $(K_f)^\circ$] is defined in such a way that its intersection with the hyperplane $\mathbb{R}^n \times \mathbb{R} \times \{-1\}$ [resp. $\mathbb{R}^n \times \{-1\} \times \mathbb{R}$] is just epi f [resp. epi f^*].

Finally, we mention a simple economic interpretation. Suppose $\mathbb{R}^n$ is a set of goods and its dual $(\mathbb{R}^n)^*$ a set of prices: to produce the goods x costs $f(x)$, and to sell x brings an income $\langle s, x \rangle$. The net benefit associated with x is then $\langle s, x \rangle - f(x)$, whose supremal value $f^*(s)$ is the best possible profit, resulting from the given set of unit prices s. Incidentally, this last interpretation opens the way to *nonlinear* conjugacy, in which the selling price would be a nonlinear (but concave) function of x.

Remark 1.2.3 This last interpretation confirms the warning already made on several occasions: $\mathbb{R}^n$ should not be confused with its dual; the arguments of f and of f^* are not comparable, an expression like $x + s$ is meaningless (until an isomorphism is established between the Euclidean space $\mathbb{R}^n$ and its dual).

On the other hand, f and f^*-values are comparable, indeed: for example, they can be added to each other – which is just done by Fenchel's inequality (1.1.3)! This is explicitly due to the particular value "-1" in (1.2.1), which goes together with the "-1" of (1.2.4). □

1.3 First Properties

Some properties of the conjugacy operation $f \mapsto f^*$ come immediately from its definition.

(a) Elementary Calculus Rules Direct arguments prove easily a first result:

Proposition 1.3.1 *The functions f, f_j appearing below are assumed to satisfy* (1.1.1).

(i) *The conjugate of the function $g(x) := f(x) + r$ is $g^*(s) = f^*(s) - r$.*
(ii) *With $t > 0$, the conjugate of the function $g(x) := t f(x)$ is $g^*(s) = t f^*(s/t)$.*
(iii) *With $t \neq 0$, the conjugate of the function $g(x) := f(tx)$ is $g^*(s) = f^*(s/t)$.*
(iv) *More generally: if A is an invertible linear operator, $(f \circ A)^* = f^* \circ (A^{-1})^*$.*

(v) *The conjugate of the function* $g(x) := f(x - x_0)$ *is* $g^*(s) = f^*(s) + \langle s, x_0 \rangle$.

(vi) *The conjugate of the function* $g(x) := f(x) + \langle s_0, x \rangle$ *is* $g^*(s) = f^*(s - s_0)$.

(vii) *If* $f_1 \leqslant f_2$, *then* $f_1^* \geqslant f_2^*$.

(viii) *"Convexity" of the conjugation: if* $\operatorname{dom} f_1 \cap \operatorname{dom} f_2 \neq \emptyset$ *and* $\alpha \in {]0, 1[}$,

$$[\alpha f_1 + (1 - \alpha) f_2]^* \leqslant \alpha f_1^* + (1 - \alpha) f_2^* \, ;$$

(ix) *The Legendre-Fenchel transform preserves decomposition: with*

$$\mathbb{R}^n := \mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_m} \ni x \mapsto f(x) := \sum_{j=1}^m f_j(x_j)$$

and assuming that $\mathbb{R}^n$ *has the scalar product of a product-space, there holds*

$$f^*(s_1, \ldots, s_m) = \sum_{j=1}^m f_j^*(s_j) \, . \qquad \square$$

Among these results, (iv) deserves comment: it gives the effect of a change of variables on the conjugate function; this is of interest for example when the scalar product is put in the form $\langle x, y \rangle = (Ax)^\top Ay$, with A invertible. Using Example 1.1.3, an illustration of (vii) is: the only function f satisfying $f = f^*$ is $1/2 \, \| \cdot \|^2$ (start from Fenchel's inequality); and also, for symmetric positive definite Q and P: $Q \preccurlyeq P$ implies $P^{-1} \preccurlyeq Q^{-1}$; and an illustration of (viii) is: $[\alpha Q + (1 - \alpha) P]^{-1} \preccurlyeq \alpha Q^{-1} + (1 - \alpha) P^{-1}$.

Our next result expresses how the conjugate is transformed, when the starting function is restricted to a subspace.

Proposition 1.3.2 *Let* f *satisfy* (1.1.1), *let* H *be a subspace of* $\mathbb{R}^n$, *and call* p_H *the operator of orthogonal projection onto* H. *Suppose that there is a point in* H *where* f *is finite. Then* $f + \operatorname{i}_H$ *satisfies* (1.1.1) *and its conjugate is*

$$(f + \operatorname{i}_H)^* = (f \circ \operatorname{p}_H)^* \circ \operatorname{p}_H \, . \qquad (1.3.1)$$

Proof. When y describes $\mathbb{R}^n$, $\operatorname{p}_H y$ describes H so we can write, knowing that p_H is symmetric:

$$
\begin{aligned}
(f + \operatorname{i}_H)^*(s) &:= \sup \{ \langle s, x \rangle - f(x) \,:\, x \in H \} \\
&= \sup \{ \langle s, \operatorname{p}_H y \rangle - f(\operatorname{p}_H y) \,:\, y \in \mathbb{R}^n \} \\
&= \sup \{ \langle \operatorname{p}_H s, y \rangle - f(\operatorname{p}_H y) \,:\, y \in \mathbb{R}^n \} \, . \qquad \square
\end{aligned}
$$

When conjugating $f + \operatorname{i}_H$, we disregard f outside H, i.e. we consider f as a function defined on the space H; this is reflected in the "$f \circ \operatorname{p}_H$"-part of (1.3.1). We then obtain a "partial" conjugate, say $\hat{f}^* \in \overline{\operatorname{Conv}} \, H$; the last "$\circ \operatorname{p}_H$" of (1.3.1) says that, to recover the whole conjugate $(f + \operatorname{i}_H)^*$, which is a function of $\overline{\operatorname{Conv}} \, \mathbb{R}^n$, we just have to translate horizontally the graph of $\hat{f}^*$ along $H^\perp$.

Remark 1.3.3 Thus, if a subspace H and a function g (standing for $f + \operatorname{i}_H$) are such that $\operatorname{dom} g \subset H$, then $g^*(\cdot + s) = g^*$ for all $s \in H^\perp$. It is interesting to note that this property has a converse: if, for some $x_0 \in \operatorname{dom} g$, we have $g(x_0 + y) = g(x_0)$ for all $y \in H$, then $\operatorname{dom} g^* \subset H^\perp$. The proof is immediate: take x_0 as above and, for $s \notin H^\perp$, take $y_0 \in H$ such that $\langle s, y_0 \rangle = \alpha \neq 0$; there holds $\langle s, \lambda y_0 \rangle - g(x_0 + \lambda y_0) = \lambda \alpha - g(x_0)$ for all $\lambda \in \mathbb{R}$; hence

$$
\begin{aligned}
g^*(s) &\geqslant \sup_\lambda [\langle s, x_0 + \lambda y_0 \rangle - g(x_0 + \lambda y_0)] \\
&= \langle s, x_0 \rangle + \sup_\lambda [\lambda \alpha - g(x_0)] = +\infty \, . \qquad \square
\end{aligned}
$$

The above formulae can be considered from a somewhat opposite point of view: suppose that $\mathbb{R}^n$, the space on which our function f is defined, is embedded in some larger space $\mathbb{R}^{n+p}$. Various corresponding extensions of f to the whole of $\mathbb{R}^{n+p}$ are then possible. One is to set $f := +\infty$ outside $\mathbb{R}^n$ (which is often relevant when minimizing f), the other is the horizontal translation $f(x + y) = f(x)$ for all $y \in \mathbb{R}^p$.

These two possibilities are, in a sense, dual to each other (see Proposition 1.3.4 below). Such a duality, however, holds only if the extended scalar product preserves the structure of $\mathbb{R}^n \times \mathbb{R}^p$ as a product-space: so is the case of the decomposition $H + H^\perp$ appearing in Proposition 1.3.2, because $\langle s, x \rangle = \langle \mathrm{p}_H s, \mathrm{p}_H x \rangle + \langle \mathrm{p}_{H^\perp} s, \mathrm{p}_{H^\perp} x \rangle$.

Considering affine manifolds instead of subspaces gives a useful result:

Proposition 1.3.4 *For f satisfying (1.1.1), let a subspace V contain the subspace parallel to* aff dom f *and set $U := V^\perp$. For any $z \in$ aff dom f and any $s \in \mathbb{R}^n$ decomposed as $s = s_U + s_V$, there holds*

$$f^*(s) = \langle s_U, z \rangle + f^*(s_V).$$

Proof. In (1.1.2), the variable x can range through $z + V \supset$ aff dom f:

$$\begin{aligned}
f^*(s) &= \sup_{v \in V}[\langle s_U + s_V, z + v \rangle - f(z + v)] \\
&= \langle s_U, z \rangle + \sup_{v \in V}[\langle s_V, z + v \rangle - f(z + v)] \\
&= \langle s_U, z \rangle + f^*(s_V).
\end{aligned}$$ $\quad\square$

(b) The Biconjugate of a Function What happens if we take the conjugate of f^* again? Remember that f^* satisfies automatically (1.1.1) if f does. We can therefore compute the *biconjugate* function of f: for all $x \in \mathbb{R}^n$,

$$f^{**}(x) := (f^*)^*(x) = \sup\{\langle s, x \rangle - f^*(s) : s \in \mathbb{R}^n\}. \qquad (1.3.2)$$

This operation appears as fundamental. The function f^{**} thus defined turns out to be the *closed convex hull* $\overline{\mathrm{co}}\, f$ of the function f, already seen in Definition B.2.5.3:

Theorem 1.3.5 *For f satisfying (1.1.1), the function f^{**} of (1.3.2) is the pointwise supremum of all the affine functions on $\mathbb{R}^n$ majorized by f. In other words*

$$\mathrm{epi}\, f^{**} = \overline{\mathrm{co}}\,(\mathrm{epi}\, f). \qquad (1.3.3)$$

Proof. Call $\Sigma \subset \mathbb{R}^n \times \mathbb{R}$ the set of pairs (s, r) defining affine functions $\langle s, \cdot \rangle - r$ majorized by f:

$$\begin{aligned}
(s, r) \in \Sigma &\iff f(x) \geqslant \langle s, x \rangle - r \quad \text{for all } x \in \mathbb{R}^n \\
&\iff r \geqslant \sup\{\langle s, x \rangle - f(x) : x \in \mathbb{R}^n\} \\
&\iff r \geqslant f^*(s) \quad \text{(and } s \in \mathrm{dom}\, f^*!).
\end{aligned}$$

Then we obtain, for $x \in \mathbb{R}^n$,

$$\begin{aligned}
\sup_{(s,r) \in \Sigma}[\langle s, x \rangle - r] &= \sup\{\langle s, x \rangle - r : s \in \mathrm{dom}\, f^*, -r \leqslant -f^*(s)\} \\
&= \sup\{\langle s, x \rangle - f^*(s) : s \in \mathrm{dom}\, f^*\} = f^{**}(x).
\end{aligned}$$

Geometrically, the epigraphs of the affine functions associated with $(s, r) \in \Sigma$ are the (non-vertical) closed half-spaces containing epi f. From §B.2.5, the epigraph of their supremum is the closed convex hull of epi f, and this proves (1.3.3). $\quad\square$

Note: the biconjugate of an $f \in \text{Conv } \mathbb{R}^n$ is not exactly f itself but its *closure*: $f^{**} = \text{cl } f$. Thanks to Theorem 1.3.5, the general notation

$$\overline{\text{co}} \, f := f^{**} \qquad (1.3.4)$$

can be – and will be – used for a function simply satisfying (1.1.1); it reminds one more directly that f^{**} is the closed convex hull of f:

$$f^{**}(x) = \sup_{r,s} \left\{ \langle s, x \rangle - r \ : \ \langle s, y \rangle - r \leqslant f(y) \text{ for all } y \in \mathbb{R}^n \right\}. \qquad (1.3.5)$$

Corollary 1.3.6 *If g is a function satisfying $\overline{\text{co}} \, f \leqslant g \leqslant f$, then $g^* = f^*$. The function f is equal to its biconjugate f^{**} if and only if $f \in \overline{\text{Conv}} \, \mathbb{R}^n$.*

Proof. Immediate.   □

Thus, the conjugacy operation defines an *involution* on the set of closed convex functions. When applied to strictly convex quadratic functions, it corresponds to the inversion of symmetric positive definite operators (Example 1.1.3 with $b = 0$). When applied to indicators of closed convex cones, it corresponds to polarity (Example 1.1.5). For general $f \in \overline{\text{Conv}} \, \mathbb{R}^n$, it is the correspondence illustrated in Fig. 1.3.1 (a correspondence also based on polarity, see Fig. 1.2.2). Of course, this involution property implies a lot of symmetry, already alluded to, for pairs of conjugate functions.

$$(\, . \,)^*$$

$$f \in \overline{\text{Conv}} \, \mathbf{R}^n \qquad \xrightarrow{\qquad\qquad\qquad} \qquad f^* \in \overline{\text{Conv}} \, \mathbf{R}^n$$
$$\xleftarrow{\qquad\qquad\qquad}$$

$$(\, . \,)^*$$
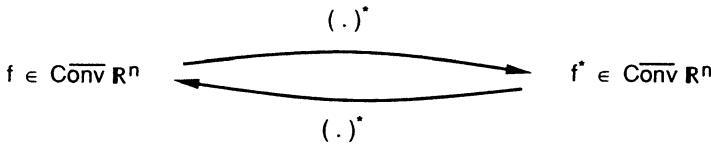
Fig. 1.3.1. The *-involution

(c) Conjugacy and Coercivity A basic question in (1.1.2) is whether the supremum is going to be $+\infty$. This depends only on the behaviour of f at infinity, so we extend to non-convex situations the concepts seen in Definition B.3.2.5:

Definition 1.3.7 A function f satisfying (1.1.1) is said to be *coercive* [resp. *1-coercive*] when

$$\lim_{\|x\| \to +\infty} f(x) = +\infty \qquad \left[\text{resp.} \quad \lim_{\|x\| \to +\infty} \frac{f(x)}{\|x\|} = +\infty \right]. \qquad \square$$

Proposition 1.3.8 *If f satisfying (1.1.1) is 1-coercive, then $f^*(s) < +\infty$ for all $s \in \mathbb{R}^n$.*

Proof. Let s be given. The 1-coercivity of f implies the existence of a number R such that $f(x) \geqslant \|s\| \, \|x\|$ (hence $\langle s, x \rangle - f(x) \leqslant 0$) whenever $\|x\| \geqslant R$. As a result, we have in (1.1.2)

$$\sup\left\{\langle s,x\rangle - \|s\|\,\|x\| \ : \ \|x\| \geqslant R\right\} \leqslant 0\,.$$

On the other hand, (1.1.1) implies an upper bound

$$\sup\left\{\langle s,x\rangle - f(x) \ : \ \|x\| \leqslant R\right\} \leqslant M\,.$$

Altogether, $f^*(s) \leqslant \max\{0, M\}$. □

For a converse to this property, we have the following:

Proposition 1.3.9 *For f satisfying* (1.1.1)*, the following holds:*

(i) *If $x_0 \in \operatorname{int} \operatorname{dom} f$ then $f^* - \langle x_0, \cdot\rangle$ is 0-coercive;*

(ii) *in particular, if f is finite over $\mathbb{R}^n$, then f^* is 1-coercive.*

Proof. We know from (1.2.3) that $\sigma_{\operatorname{dom} f} = (f^*)'_\infty$ so, using Theorem C.2.2.3(iii), $x_0 \in \operatorname{int} \operatorname{dom} f \subset \operatorname{int}(\operatorname{co} \operatorname{dom} f)$ implies $(f^*)'_\infty(s) - \langle x_0, s\rangle > 0$ for all $s \neq 0$. By virtue of Proposition B.3.2.4, this means exactly that $f^* - \langle x_0, \cdot\rangle$ has compact sublevel-sets; (i) is proved.

Then, as demonstrated in Definition B.3.2.5, 0-coercivity of $f^* - \langle x_0, \cdot\rangle$ for all x_0 means 1-coercivity of f^*. □

Piecing together, we see that the 1-coercivity of a function f implies that f^* is finite everywhere, and this in turn implies 1-coercivity of $\overline{\operatorname{co}} f$.

Remark 1.3.10 If we assume in particular $f \in \overline{\operatorname{Conv}} \mathbb{R}^n$, (i) and (ii) become equivalences:

$$x_0 \in \operatorname{int} \operatorname{dom} f \iff f^* - \langle x_0, \cdot\rangle \text{ is 0-coercive}\,,$$
$$\operatorname{dom} f = \mathbb{R}^n \iff f^* \text{ is 1-coercive}\,.$$ □

1.4 Subdifferentials of Extended-Valued Functions

For a function f satisfying (1.1.1), consider the following set:

$$\partial f(x) := \{s \in \mathbb{R}^n \ : \ f(y) \geqslant f(x) + \langle s, y - x\rangle \text{ for all } y \in \mathbb{R}^n\}\,. \qquad (1.4.1)$$

When f happens to be convex and finite-valued, this is just the subdifferential of f at x, defined in Definition D.1.2.1; but (1.4.1) can be used in a much more general framework. We therefore keep the terminology *subdifferential* for the set of (1.4.1), and *subgradients* for its elements. Note here that $\partial f(x)$ is empty if $x \notin \operatorname{dom} f$: take $y \in \operatorname{dom} f$ in (1.4.1).

Theorem 1.4.1 *For f satisfying* (1.1.1) *and ∂f defined by* (1.4.1)*, $s \in \partial f(x)$ if and only if*

$$f^*(s) + f(x) - \langle s, x\rangle = 0 \ (\text{or} \leqslant 0)\,. \qquad (1.4.2)$$

Proof. To say that s lies in the set (1.4.1) is to say that

$$\langle s, y\rangle - f(y) \leqslant \langle s, x\rangle - f(x) \quad \text{for all } y \in \operatorname{dom} f\,,$$

i.e. $f^*(s) \leqslant \langle s, x\rangle - f(x)$; but this is indeed an equality, in view of Fenchel's inequality (1.1.3). □

As before, $\partial f(x)$ is closed and convex: it is a sublevel-set of the closed convex function $f^* - \langle \cdot, x \rangle$, namely at the level $-f(x)$. A subgradient of f at x is the slope of an affine function minorizing f and coinciding with f at x; $\partial f(x)$ can therefore be empty: for example if epi f has a vertical tangent hyperplane at $(x, f(x))$; or also if f is not convex-like near x.

Theorem 1.4.2 *Let* $f \in \mathrm{Conv}\, \mathbb{R}^n$. *Then* $\partial f(x) \neq \emptyset$ *whenever* $x \in \mathrm{ri}\,\mathrm{dom}\, f$.

Proof. This is Proposition B.1.2.1.   □

When $\partial f(x) \neq \emptyset$, we obtain a particular relationship at x between f and its convexified version $\overline{\mathrm{co}}\, f$:

Proposition 1.4.3 *For* f *satisfying* (1.1.1), *the following properties hold:*

$$\partial f(x) \neq \emptyset \implies (\overline{\mathrm{co}}\, f)(x) = f(x)\,; \qquad (1.4.3)$$

$$\overline{\mathrm{co}}\, f \leqslant g \leqslant f \text{ and } g(x) = f(x) \implies \partial g(x) = \partial f(x)\,; \qquad (1.4.4)$$

$$s \in \partial f(x) \implies x \in \partial f^*(s)\,. \qquad (1.4.5)$$

Proof. Let s be a subgradient of f at x. From the definition (1.4.1) itself, the function $y \mapsto \ell_s(y) := f(x) + \langle s, y - x \rangle$ is affine and minorizes f, hence $\ell_s \leqslant \overline{\mathrm{co}}\, f \leqslant f$; because $\ell_s(x) = f(x)$, this implies (1.4.3).

Now, $s \in \partial f(x)$ if and only if (1.4.2) holds. From our assumption in (1.4.4), $f^* = g^* = (\overline{\mathrm{co}}\, f)^*$ (Corollary 1.3.6) and $g(x) = f(x)$. Therefore

$$s \in \partial f(x) \iff g^*(s) + g(x) - \langle s, x \rangle = 0\,,$$

which expresses exactly that $s \in \partial g(x)$; (1.4.4) is proved.

Finally, we know that $f^{**} = \overline{\mathrm{co}}\, f \leqslant f$; so, when s satisfies (1.4.2), we have

$$f^*(s) + f^{**}(x) - \langle s, x \rangle = f^*(s) + (\overline{\mathrm{co}}\, f)(x) - \langle s, x \rangle \leqslant 0\,,$$

which means $x \in \partial f^*(s)$: we have just proved (1.4.5).   □

Among the consequences of (1.4.3), we note the following sufficiency condition for convexity: if $\partial f(x)$ is nonempty for all $x \in \mathbb{R}^n$, then f is convex and finite-valued on $\mathbb{R}^n$. Another consequence is important:

Corollary 1.4.4 *If* $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$, *the following equivalences hold:*

$$f(x) + f^*(s) - \langle s, x \rangle = 0 \ (\text{or} \leqslant 0) \iff s \in \partial f(x) \iff x \in \partial f^*(s)\,.$$

Proof. This is a rewriting of Theorem 1.4.1, taking into account (1.4.5) and the symmetric role played by f and f^* when $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$.   □

If $s \in \partial f(x)$ (which in particular implies $x \in \mathrm{dom}\, f$), the property $s \in \mathrm{dom}\, f^*$ comes immediately; beware that, conversely, $\mathrm{dom}\, f^*$ is not entirely covered by $\partial f(\mathbb{R}^n)$: take $f(x) = \exp x$, for which $f^*(0) = 0$; but 0 is only in the closure of $f'(\mathbb{R})$.

Even though it is attached to some designated x, the concept of subdifferential, as defined by (1.4.1), is *global*, in the sense that it uses the values of f on the whole of $\mathbb{R}^n$. For example,

$$\inf \{f(x) \,:\, x \in \mathbb{R}^n\} = -f^*(0)$$

and obviously

$$x \text{ minimizes } f \text{ satisfying } (1.1.1) \quad \Longleftrightarrow \quad 0 \in \partial f(x)\,.$$

Then a consequence of Corollary 1.4.4 is:

$$\operatorname*{Argmin}_{x \in \mathbb{R}^n} f(x) = \partial f^*(0) \quad \text{for} \quad f \in \overline{\text{Conv}}\,\mathbb{R}^n\,. \tag{1.4.6}$$

2 Calculus Rules on the Conjugacy Operation

The function f, whose conjugate is to be computed, is often obtained from some other functions f_i, whose conjugates are known. In this section, we develop a set of calculus rules expressing f^* in terms of the $(f_i)^*$ (some rudimentary such rules were already given in Proposition 1.3.1).

2.1 Image of a Function Under a Linear Mapping

Given a function $g : \mathbb{R}^m \to \mathbb{R} \cup \{+\infty\}$ satisfying (1.1.1), and a linear mapping $A : \mathbb{R}^m \to \mathbb{R}^n$, we recall that the image of g under A is the function defined by

$$\mathbb{R}^n \ni x \mapsto (Ag)(x) := \inf \{g(y) \,:\, Ay = x\}\,. \tag{2.1.1}$$

Let g^* be associated with a scalar product $\langle \cdot, \cdot \rangle_m$ in $\mathbb{R}^m$; we denote by $\langle \cdot, \cdot \rangle_n$ the scalar product in $\mathbb{R}^n$, with the help of which we want to define $(Ag)^*$. To make sure that Ag satisfies (1.1.1), some additional assumption is needed: among all the affine functions minorizing g, there is one with slope in $(\operatorname{Ker} A)^\perp = \operatorname{Im} A^*$:

Theorem 2.1.1 *With the above notation, assume that* $\operatorname{Im} A^* \cap \operatorname{dom} g^* \neq \emptyset$. *Then Ag satisfies* (1.1.1) *and its conjugate is*

$$(Ag)^* = g^* \circ A^*\,.$$

Proof. First, it is clear that $Ag \not\equiv +\infty$ (take $x = Ay$, with $y \in \operatorname{dom} g$). On the other hand, our assumption implies the existence of some $p_0 = A^* s_0$ such that $g^*(p_0) < +\infty$; with Fenchel's inequality (1.1.3), we have for all $y \in \mathbb{R}^m$:

$$g(y) \geqslant \langle A^* s_0, y \rangle_m - g^*(p_0) = \langle s_0, Ay \rangle_n - g^*(p_0)\,.$$

For each $x \in \mathbb{R}^n$, take the infimum over those y satisfying $Ay = x$: the affine function $\langle s_0, \cdot \rangle - g^*(p_0)$ minorizes Ag. Altogether, Ag satisfies (1.1.1).
    Then we have for $s \in \mathbb{R}^n$

$$\begin{aligned}
(Ag)^*(s) &= \sup_{x \in \mathbb{R}^n}[\langle s, x \rangle - \inf_{Ay=x} g(y)] \\
&= \sup_{x \in \mathbb{R}^n, Ay=x}[\langle s, x \rangle - g(y)] \\
&= \sup_{y \in \mathbb{R}^m}[\langle s, Ay \rangle - g(y)] = g^*(A^* s)\,. \qquad \square
\end{aligned}$$

For example, when $m = n$ and A is invertible, Ag reduces to $g \circ A^{-1}$, whose conjugate is therefore $g^* \circ A^*$, given by Proposition 1.3.1(iv).

As a first application of Theorem 2.1.1, it is straightforward to compute the conjugate of a marginal function:

$$f(x) := \inf \{ g(x, z) \; : \; z \in \mathbb{R}^p \}, \tag{2.1.2}$$

where g operates on the product-space $\mathbb{R}^n \times \mathbb{R}^p$. Indeed, just call A the projection onto $\mathbb{R}^n$: $A(x, z) = x \in \mathbb{R}^n$, so f is clearly Ag. We obtain:

Corollary 2.1.2 *With* $g : \mathbb{R}^n \times \mathbb{R}^p =: \mathbb{R}^m \to \mathbb{R} \cup \{+\infty\}$ *not identically* $+\infty$, *let* g^* *be associated with a scalar product preserving the structure of* $\mathbb{R}^m$ *as a product space:* $\langle \cdot, \cdot \rangle_m = \langle \cdot, \cdot \rangle_n + \langle \cdot, \cdot \rangle_p$. *If there exists* $s_0 \in \mathbb{R}^n$ *such that* $(s_0, 0) \in \mathrm{dom}\, g^*$, *then the conjugate of* f *defined by* (2.1.2) *is*

$$f^*(s) = g^*(s, 0) \quad \text{for all } s \in \mathbb{R}^n .$$

Proof. It suffices to observe that, A being the projection defined above, there holds for all $y_1 = (x_1, z_1) \in \mathbb{R}^m$ and $x_2 \in \mathbb{R}^n$,

$$\langle Ay_1, x_2 \rangle_n = \langle x_1, x_2 \rangle_n = \langle x_1, x_2 \rangle_n + \langle z_1, 0 \rangle_p = \langle y_1, (x_2, 0) \rangle_m ,$$

which *defines* the adjoint $A^* x = (x, 0)$ for all $x \in \mathbb{R}^n$. Then apply Theorem 2.1.1. $\square$

More will be said on this operation in §2.4. Here we consider another application: the infimal convolution which, to f_1 and f_2 defined on $\mathbb{R}^n$, associates (see §B.2.3)

$$(f_1 \mathbin{\dot{+}} f_2)(x) := \inf \{ f_1(x_1) + f_2(x_2) \; : \; x_1 + x_2 = x \}. \tag{2.1.3}$$

Then Theorem 2.1.1 can be used, with $m = 2n$, $g(x_1, x_2) := f_1(x_1) + f_2(x_2)$ and $A(x_1, x_2) := x_1 + x_2$:

Corollary 2.1.3 *Let* f_1 *and* f_2 *be two functions from* $\mathbb{R}^n$ *to* $\mathbb{R} \cup \{+\infty\}$, *not identically* $+\infty$, *and satisfying* $\mathrm{dom}\, f_1^* \cap \mathrm{dom}\, f_2^* \neq \emptyset$. *Then their inf-convolution satisfies* (1.1.1), *and* $(f_1 \mathbin{\dot{+}} f_2)^* = f_1^* + f_2^*$.

Proof. Equip $\mathbb{R}^n \times \mathbb{R}^n$ with the scalar product $\langle \cdot, \cdot \rangle + \langle \cdot, \cdot \rangle$. Using the above notation for g and A, we have $g^*(s_1, s_2) = f_1^*(s_1) + f_2^*(s_2)$ (Proposition 1.3.1(ix)) and $A^*(s) = (s, s)$. Then apply the definitions. $\square$

Example 2.1.4 The infimal convolution of a function with the kernel $1/2\, c \| \cdot \|^2$ ($C > 0$) is the so-called *Moreau-Yosida* regularization, which turns out to have a Lipschitzian gradient. With the kernel $c \| \cdot \|$, we obtain the so-called *Lipschitzian* regularization (which can indeed be proved to be Lipschitzian on the whole space). Then Corollary 2.1.3 gives, with the help of Proposition 1.3.1(iii) and Example 1.1.3:

$$(f \mathbin{\dot{+}} \tfrac{1}{2} c \| \cdot \|^2)^* = f^* + \tfrac{1}{2c} \| \cdot \|^2 \quad \text{and} \quad (f \mathbin{\dot{+}} c \| \cdot \|)^* = f^* + \tfrac{1}{c} \mathrm{i}_{B(0,c)} .$$

In particular, if f is the indicator of a nonempty set $C \subset \mathbb{R}^n$, the above formulae yield

$$\left(\tfrac{1}{2} \mathrm{d}_C^2 \right)^* = \sigma_C + \tfrac{1}{2} \| \cdot \|^2, \quad \mathrm{d}_C^* = \sigma_C + \mathrm{i}_{B(0,1)} . \qquad \square$$

2.2 Pre-Composition with an Affine Mapping

In view of the symmetry of the conjugacy operation, Theorem 2.1.1 suggests that the conjugate of $g \circ A$, when A is a linear mapping, is the image-function $A^* g^*$. In particular, a condition was needed in §2.1 to prevent $Ag(x) = -\infty$ for some x. Likewise, a condition will be needed here to prevent $g \circ A \equiv +\infty$. The symmetry is not quite perfect, though: the composition of a closed convex function with a linear mapping is still a closed convex function; but an image-function need not be closed, and therefore cannot be a conjugate function.

We use notation similar to that of §2.1, but we find it convenient to distinguish between the linear and affine cases.

Theorem 2.2.1 *Take* $g \in \overline{\mathrm{Conv}}\,\mathbb{R}^m$, A_0 *linear from* $\mathbb{R}^n$ *to* $\mathbb{R}^m$ *and consider the affine operator* $A(x) := A_0 x + y_0 \in \mathbb{R}^m$. *Suppose that* $A(\mathbb{R}^n) \cap \mathrm{dom}\,g \neq \emptyset$. *Then* $g \circ A \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$ *and its conjugate is the closure of the convex function*

$$\mathbb{R}^n \ni s \mapsto \inf_p \left\{ g^*(p) - \langle y_0, p \rangle_m \ : \ A_0^* p = s \right\}. \tag{2.2.1}$$

Proof. We start with the linear case ($y_0 = 0$): suppose that $h \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$ satisfies $\mathrm{Im}\,A_0 \cap \mathrm{dom}\,h \neq \emptyset$. Then Theorem 2.1.1 applied to $g := h^*$ and $A := A_0^*$ gives $(A_0^* h^*)^* = h \circ A_0$; conjugating both sides, we see that the conjugate of $h \circ A_0$ is the closure of the image-function $A_0^* h^*$.

In the affine case, consider the function $h := g(\cdot + y_0) \in \overline{\mathrm{Conv}}\,\mathbb{R}^m$; its conjugate is given by Proposition 1.3.1(v): $h^* = g^* - \langle y_0, \cdot \rangle_m$. Furthermore, it is clear that

$$(g \circ A)(x) = g(A_0 x + y_0) = h(A_0 x) = (h \circ A_0)(x),$$

so (2.2.1) follows from the linear case.   □

Thus, the conjugation of a composition by a linear (or affine) mapping is not quite straightforward, as it requires a closure operation. A natural question is therefore: when does (2.2.1) define a closed function of s? In other words, when is an image-function closed? Also, when is the infimum in (2.2.1) attained at some p? We start with a technical result, in which a qualification property comes again into play.

Lemma 2.2.2 *Let* $g \in \overline{\mathrm{Conv}}\,\mathbb{R}^m$ *be such that* $0 \in \mathrm{dom}\,g$ *and let* A_0 *be linear from* $\mathbb{R}^n$ *to* $\mathbb{R}^m$. *Make the following assumption:*

$$\mathrm{Im}\,A_0 \cap \mathrm{ri}\,\mathrm{dom}\,g \neq \emptyset \quad i.e. \quad 0 \in \mathrm{ri}\,\mathrm{dom}\,g - \mathrm{Im}\,A_0 \ \left[= \mathrm{ri}\,(\mathrm{dom}\,g - \mathrm{Im}\,A_0)\right].$$

Then $(g \circ A_0)^* = A_0^* g^*$; *for every* $s \in \mathrm{dom}\,(g \circ A_0)^*$, *the problem*

$$\inf_p \left\{ g^*(p) \ : \ A_0^* p = s \right\} \tag{2.2.2}$$

has at least one optimal solution $\bar{p}$ *and there holds* $(g \circ A_0)^*(s) = A_0^* g^*(s) = g^*(\bar{p})$.

Proof. To prove $(g \circ A_0)^* = A_0^* g^*$, we have to prove that $A_0^* g^*$ is a closed function, i.e. that its sublevel-sets are closed (Definition B.1.2.3).

Thus, for given $r \in \mathbb{R}$, take a sequence (s_k) converging to some s and such that $(A_0^* g^*)(s_k) \leqslant r$. Take also $\delta_k \downarrow 0$; from the definition of the image-function, we can find $p_k \in \mathbb{R}^m$ such that

$$g^*(p_k) \leqslant r + \delta_k \quad \text{and} \quad A_0^* p_k = s_k .$$

Let q_k be the orthogonal projection of p_k onto the subspace $V := \text{lin dom } g - \text{Im } A_0$. Since V contains lin dom g, Proposition 1.3.4 (with $z = 0$) gives $g^*(p_k) = g^*(q_k)$. Furthermore, $V^\perp = (\text{lin dom } g)^\perp \cap \text{Ker } A_0^*$; in particular, $q_k - p_k \in \text{Ker } A_0^*$. In summary, we have singled out $q_k \in V$ such that

$$g^*(q_k) \leqslant r + \delta_k \quad \text{and} \quad A_0^* q_k = s_k \quad \text{for all } k . \tag{2.2.3}$$

Suppose we can bound q_k. Extracting a subsequence if necessary, we will have $q_k \to \bar{q}$ and, passing to the limit, we will obtain (since g^* is l.s.c)

$$g^*(\bar{q}) \leqslant \liminf g^*(q_k) \leqslant r \quad \text{and} \quad A_0^* \bar{q} = s .$$

The required closedness property $A_0^* g^*(\bar{q}) \leqslant r$ will follow by definition. Furthermore, this $\bar{q}$ will be a solution of (2.2.2) in the particular case $s_k \equiv s$ and $r = (A_0^* g^*)(s)$. In this case, (q_k) will be actually a minimizing sequence of (2.2.2).

To prove boundedness of q_k, use the assumption: for some $\varepsilon > 0$, $B_m(0, \varepsilon) \cap V$ is included in dom $g - \text{Im } A$. Thus, for arbitrary $z \in B_m(0, \varepsilon) \cap V$, we can find $y \in \text{dom } g$ and $x \in \mathbb{R}^n$ such that $z = y - A_0 x$. Then

$$\begin{aligned} \langle q_k, z \rangle_m &= \langle q_k, y \rangle_m - \langle A_0^* q_k, x \rangle_n \\ &\leqslant g(y) + g^*(q_k) - \langle A_0^* q_k, x \rangle_n &&\text{[Fenchel (1.1.3)]} \\ &\leqslant g(y) + r + \delta_k - \langle s_k, x \rangle_n . &&\text{[(2.2.3)]} \end{aligned}$$

We conclude that $\sup \{ \langle q_k, z \rangle : k = 1, 2, \ldots \}$ is bounded for any $z \in B_m(0, \varepsilon) \cap V$, which implies that q_k is bounded; this is Proposition V.2.1.3 in the vector space V. □

Adapting this result to the general case is now a matter of translation:

Theorem 2.2.3 *Take* $g \in \overline{\text{Conv}} \, \mathbb{R}^m$, A_0 *linear from* $\mathbb{R}^n$ *to* $\mathbb{R}^m$ *and consider the affine operator* $A(x) := A_0 x + y_0 \in \mathbb{R}^m$. *Make the following assumption:*

$$A(\mathbb{R}^n) \cap \text{ri dom } g \neq \emptyset . \tag{2.2.4}$$

Then, for every $s \in \text{dom } (g \circ A_0)^*$, *the problem*

$$\min_p \{ g^*(p) - \langle p, y_0 \rangle \, : \, A_0^* p = s \} \tag{2.2.5}$$

has at least one optimal solution $\bar{p}$ *and there holds* $(g \circ A)^*(s) = g^*(\bar{p}) - \langle \bar{p}, y_0 \rangle$.

Proof. By assumption, we can choose $\bar{x} \in \mathbb{R}^n$ such that $\bar{y} := A(\bar{x}) \in \operatorname{ri} \operatorname{dom} g$. Consider the function $\bar{g} := g(\bar{y} + \cdot) \in \overline{\operatorname{Conv}} \mathbb{R}^m$. Observing that

$$(g \circ A)(x) = \bar{g}(A(x) - \bar{y}) = (\bar{g} \circ A_0)(x - \bar{x}),$$

we obtain from the calculus rule 1.3.1(v): $(g \circ A)^* = (\bar{g} \circ A_0)^* + \langle \cdot, \bar{x} \rangle$.

Then Lemma 2.2.2 allows the computation of this conjugate. We have 0 in the domain of $\bar{g}$, and even in its relative interior:

$$\operatorname{ri} \operatorname{dom} \bar{g} = \operatorname{ri} \operatorname{dom} g - \{\bar{y}\} \ni 0 \in \operatorname{Im} A_0.$$

We can therefore write: for all $s \in \operatorname{dom} (\bar{g} \circ A_0)^*$ [$= \operatorname{dom} (g \circ A)^*$],

$$(\bar{g} \circ A_0)^*(s) = \min_p \{\bar{g}^*(p) \ : \ A_0^* p = s\},$$

where the minimum is attained at some $\bar{p}$. Using again the calculus rule 1.3.1(v) and various relations from above, we have established

$$\begin{aligned}
(g \circ A)^*(s) - \langle s, \bar{x} \rangle &= \min \{g^*(p) - \langle p, A_0 \bar{x} + y_0 \rangle \ : \ A_0^* p = s\} \\
&= \min \{g^*(p) - \langle p, y_0 \rangle \ : \ A_0^* p = s\} - \langle s, \bar{x} \rangle. \qquad \square
\end{aligned}$$

It can be shown that the calculus rule (2.2.5) does need a qualification assumption such as (2.2.4). First of all, we certainly need to avoid $g \circ A \equiv +\infty$, i.e.

$$A(\mathbb{R}^n) \cap \operatorname{dom} g \neq \emptyset; \tag{2.2.Q.i}$$

but this is not sufficient, unless g is a polyhedral function (a situation which has essentially been treated in §C.3.4).

A "comfortable" sharpening of (2.2.Q.i) is

$$A(\mathbb{R}^n) \cap \operatorname{int} \operatorname{dom} g \neq \emptyset \quad \text{i.e.} \quad 0 \in \operatorname{int} \operatorname{dom} g - A(\mathbb{R}^n), \tag{2.2.Q.ii}$$

but it is fairly restrictive, implying in particular that $\operatorname{dom} g$ is full-dimensional. More tolerant is

$$0 \in \operatorname{int} [\operatorname{dom} g - A(\mathbb{R}^n)], \tag{2.2.Q.iii}$$

which is rather common. Various results from Chap. C, in particular Theorem C.2.2.3(iii), show that (2.2.Q.iii) means

$$\sigma_{\operatorname{dom} g}(p) + \langle p, y_0 \rangle > 0 \quad \text{for all nonzero } p \in \operatorname{Ker} A_0^*.$$

Knowing that $\sigma_{\operatorname{dom} g} = (g^*)'_\infty$ (Proposition 1.2.2), this condition has already been alluded to at the end of §B.3.

Naturally, our assumption (2.2.4) is a further weakening; actually, it is only a slight generalization of (2.2.Q.iii). It is interesting to note that, if (2.2.4) is replaced by (2.2.Q.iii), the solution-set in problem (2.2.5) becomes bounded; to see this, read again the proof of Lemma 2.2.2, in which V becomes $\mathbb{R}^m$.

Finally consider an example with $n = 1$: given a function $g \in \overline{\operatorname{Conv}} \mathbb{R}^m$, fix $x_0 \in \operatorname{dom} g$, $0 \neq d \in \mathbb{R}^m$, and set $\psi(t) := g(x_0 + td)$ for all $t \in \mathbb{R}$. This ψ is the composition of g with the affine mapping which, to $t \in \mathbb{R}$, assigns $x_0 + td \in \mathbb{R}^m$. It is an exercise to apply Theorem 2.2.1 and obtain the conjugate: $\psi^*(\alpha) = \min \{g^*(s) - \langle s, x_0 \rangle : \langle s, d \rangle = \alpha\}$ whenever, for example, x_0 and d are such that $x_0 + td \in \operatorname{ri} \operatorname{dom} g$ for some t – this is (2.2.4). This example can be further particularized to various functions mentioned in §1 (quadratic, indicator, . . .).

2.3 Sum of Two Functions

The formula for conjugating a sum will supplement Proposition 1.3.1(ii) to obtain the conjugate of a positive combination of closed convex functions. Summing two functions is a simple operation (at least it preserves closedness); but Corollary 2.1.3 shows that a sum is the conjugate of something rather involved: an inf-convolution. Likewise, compare the simplicity of the composition $g \circ A$ with the complexity of its dual counterpart Ag; as seen in §2.2, difficulties are therefore encountered when conjugating a function of the type $g \circ A$. The same kind of difficulties must be expected when conjugating a sum, and the development in this section is quite parallel to that of §2.2.

Theorem 2.3.1 *Let g_1, g_2 be in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$ and assume that $\mathrm{dom}\,g_1 \cap \mathrm{dom}\,g_2 \neq \emptyset$. The conjugate $(g_1 + g_2)^*$ of their sum is the closure of the convex function $g_1^* \mathbin{\mkern-1mu\dot{\mkern-3mu\scriptstyle\vee}} g_2^*$.*

Proof. Call $f_i^* := g_i$, for $i = 1, 2$; apply Corollary 2.1.3: $(g_1^* \mathbin{\mkern-1mu\dot{\mkern-3mu\scriptstyle\vee}} g_2^*)^* = g_1 + g_2$; then take the conjugate again.   □

The above calculus rule is very useful in the following framework: suppose we want to compute an inf-convolution, say $h = f \mathbin{\mkern-1mu\dot{\mkern-3mu\scriptstyle\vee}} g$ with f and g in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$. Compute f^* and g^*; if their sum happens to be the conjugate of some known function in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$, this function has to be the closure of h.

Just as in the previous section, it is of interest to ask whether the closure operation is necessary, and whether the inf-convolution is exact.

Theorem 2.3.2 *Let g_1, g_2 be in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$ and assume that*

$$\left. \begin{array}{c} \textit{the relative interiors of } \mathrm{dom}\,g_1 \textit{ and } \mathrm{dom}\,g_2 \textit{ intersect ,}\\ \textit{or equivalently: } 0 \in \mathrm{ri}\,(\mathrm{dom}\,g_1 - \mathrm{dom}\,g_2)\,. \end{array} \right| \quad (2.3.1)$$

Then $(g_1 + g_2)^* = g_1^* \mathbin{\mkern-1mu\dot{\mkern-3mu\scriptstyle\vee}} g_2^*$ and, for every $s \in \mathrm{dom}\,(g_1 + g_2)^*$, the problem

$$\inf \{g_1^*(p) + g_2^*(q) \ : \ p + q = s\}$$

has at least one optimal solution $(\bar{p}, \bar{q})$, which therefore satisfies

$$g_1^*(\bar{p}) + g_2^*(\bar{q}) = (g_1^* \mathbin{\mkern-1mu\dot{\mkern-3mu\scriptstyle\vee}} g_2^*)(s) = (g_1 + g_2)^*(s)\,.$$

Proof. Define $g \in \overline{\mathrm{Conv}}(\mathbb{R}^n \times \mathbb{R}^n)$ by $g(x_1, x_2) := g_1(x_1) + g_2(x_2)$ and the linear operator $A : \mathbb{R}^n \to \mathbb{R}^n \times \mathbb{R}^n$ by $Ax := (x, x)$. Then $g_1 + g_2 = g \circ A$, and we proceed to use Theorem 2.2.3. As seen in Proposition 1.3.1(ix), $g^*(p, q) = g_1^*(p) + g_2^*(q)$ and straightforward calculation shows that $A^*(p, q) = p + q$. Thus, if we can apply Theorem 2.2.3, we can write

$$\begin{aligned} (g_1 + g_2)^*(s) = (g \circ A)^*(s) &= (A^*g^*)(s)\\ &= \inf_{p,q} \{g_1^*(p) + g_2^*(q) \ : \ p + q = s\} = (g_1^* \mathbin{\mkern-1mu\dot{\mkern-3mu\scriptstyle\vee}} g_2^*)(s) \end{aligned}$$

and the above minimization problem does have an optimal solution.

To check (2.2.4), note that $\operatorname{dom} g = \operatorname{dom} g_1 \times \operatorname{dom} g_2$, and $\operatorname{Im} A$ is the diagonal set $\Delta := \{(s, s) : s \in \mathbb{R}^n\}$. We have

$$(x, x) \in \operatorname{ri} \operatorname{dom} g_1 \times \operatorname{ri} \operatorname{dom} g_2 = \operatorname{ri}(\operatorname{dom} g_1 \times \operatorname{dom} g_2)$$

(Proposition A.2.1.11), and this just means that $\operatorname{Im} A = \Delta$ has a nonempty intersection with $\operatorname{ri} \operatorname{dom} g$. $\qquad\square$

As with Theorem 2.3, a qualification assumption – taken here as (2.3.1) playing the role of (2.2.4) – is necessary to ensure the stated properties; but other assumptions exist that do the same thing. First of all,

$$\operatorname{dom} g_1 \cap \operatorname{dom} g_2 \neq \emptyset, \quad \text{i.e.} \quad 0 \in \operatorname{dom} g_1 - \operatorname{dom} g_2 \qquad (2.3.Q.j)$$

is obviously necessary to have $g_1 + g_2 \not\equiv +\infty$. We saw that this "minimal" condition was sufficient in the *polyhedral* case. Here the results of Theorem 2.3.2 do hold if (2.3.1) is replaced by:

$$g_1 \text{ and } g_2 \text{ are both polyhedral, or also}$$
$$g_1 \text{ is polyhedral and } \operatorname{dom} g_1 \cap \operatorname{ri} \operatorname{dom} g_2 \neq \emptyset.$$

The "comfortable" assumption playing the role of (2.2.Q.ii) is

$$\operatorname{int} \operatorname{dom} g_1 \cap \operatorname{int} \operatorname{dom} g_2 \neq \emptyset, \qquad (2.3.Q.jj)$$

which obviously implies (2.3.1). We mention a non-symmetric assumption, particular to a sum:

$$(\operatorname{int} \operatorname{dom} g_1) \cap \operatorname{dom} g_2 \neq \emptyset. \qquad (2.3.Q.jj')$$

Finally, the weakening (2.2.Q.iii) is

$$0 \in \operatorname{int}(\operatorname{dom} g_1 - \operatorname{dom} g_2), \qquad (2.3.Q.jjj)$$

which means $\sigma_{\operatorname{dom} g_1}(s) + \sigma_{\operatorname{dom} g_2}(-s) > 0$ for all $s \neq 0$; we leave it as an exercise to prove (2.3.Q.jj') $\Rightarrow$ (2.3.Q.jjj).

The following application of Theorem 2.3.2 is important in optimization: take two functions g_1 and g_2 in $\overline{\operatorname{Conv}}\,\mathbb{R}^n$ and assume that

$$\mu := \inf\{g_1(x) + g_2(x) : x \in \mathbb{R}^n\}$$

is a finite number. Under some qualification assumption such as (2.3.1),

$$[(g_1 + g_2)^*(0) =] \quad -\mu = \min_{s \in \mathbb{R}^n}[g_1^*(s) + g_2^*(-s)], \qquad (2.3.2)$$

a relation known as *Fenchel's duality Theorem*. If, in particular, $g_2 = i_C$ is an indicator function, we read (still under some qualification assumption)

$$\inf\{f(x) : x \in C\} = -\min\{f^*(s) + \sigma_C(-s) : s \in \mathbb{R}^n\}.$$

It is appropriate to recall here that conjugate functions are interesting primarily via their arguments and subgradients; thus, it is the *existence* of a minimizing s in the righthand side of (2.3.2) which is useful, rather than a mere equality between two real numbers.

To conclude this study of the sum, we mention one among the possible results concerning its dual operation: the infimal convolution.

Corollary 2.3.3 *Take f_1 and f_2 in $\overline{\mathrm{Conv}}\,\mathbb{R}^n$, with f_1 0-coercive and f_2 bounded from below. Then the inf-convolution problem* (2.1.3) *has a nonempty compact set of solutions; furthermore $f_1 \,\dot{+}\, f_2 \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$.*

Proof. Letting μ denote a lower bound for f_2, we have $f_1(x_1) + f_2(x - x_1) \geqslant f_1(x_1) + \mu$ for all $x_1 \in \mathbb{R}^n$, and the first part of the claim follows.

For closedness of the infimal convolution, we set $g_i := f_i^*$, $i = 1, 2$; because of 0-coercivity of f_1, $0 \in \mathrm{int}\,\mathrm{dom}\,g_1$ (Remark 1.3.10), and $g_2(0) \leqslant -\mu$. Thus, we can apply Theorem 2.3.2 with the qualification assumption (2.3.Q.jj'). $\qquad\square$

2.4 Infima and Suprema

A result of the previous sections is that the conjugacy correspondence establishes a symmetry between the sum and the inf-convolution, and also between the image and the pre-composition with a linear mapping. Here we will see that the operations "supremum" and "closed convex hull of the infimum" are likewise symmetric to each other.

Theorem 2.4.1 *Let $\{f_j\}_{j \in J}$ be a collection of functions satisfying* (1.1.1) *and having a common affine minorant:* $\sup_{j \in J} f_j^*(s) < +\infty$ *for some $s \in \mathbb{R}^n$.*

Then their infimum $f := \inf_{j \in J} f_j$ satisfies (1.1.1), *and its conjugate is the supremum of the f_j^*'s:*

$$(\inf_{j \in J} f_j)^* = \sup_{j \in J} f_j^* \,. \tag{2.4.1}$$

Proof. By definition, for all $s \in \mathbb{R}^n$

$$
\begin{aligned}
f^*(s) &= \sup_x [\langle s, x \rangle - \inf_{j \in J} f_j(x)] \\
&= \sup_x \sup_j [\langle s, x \rangle - f_j(x)] \\
&= \sup_j \sup_x [\langle s, x \rangle - f_j(x)] = \sup_{j \in J} f_j^*(s) \,. \qquad\square
\end{aligned}
$$

This result should be compared to Corollary 2.1.2, which after all expresses the conjugate of the same function, and which is proved in just the same way; compare the above proof with that of Theorem 2.1.1. The only difference (apart from the notation j instead of z) is that the space $\mathbb{R}^p$ is now replaced by an arbitrary set J, with no special structure, in particular no scalar product. In effect, Theorem 2.4.1 supersedes Corollary 2.1.2: the latter just says that, for g defined on a product-space, the easy-to-prove formula $g^*(s, 0) = \sup_z g_x^*(s, z)$ holds, where the subscript x of g_x^* indicates conjugacy with respect to the first variable only. In a word, the conjugate with respect to the *couple* (x, z) is of no use for computing the conjugate at $(\cdot, 0)$. Beware that this is true only when the scalar product considers x and z as two separate variables (i.e. is compatible with the product-space). Otherwise, trouble may occur.

Example 2.4.2 The Euclidean distance to a nonempty set $C \subset \mathbb{R}^n$ is an inf-function:

$$x \mapsto \mathrm{d}_C(x) = \inf\{f_y(x) : y \in C\} \quad \text{with} \quad f_y(x) := \|x - y\| \,.$$

Remembering that the conjugate of the norm $\|\cdot\|$ is the indicator of the unit ball (combine (C.2.3.1) and Example 1.1.5), the calculus rule 1.3.1(v) gives

$$(f_y)^*(s) = \mathrm{i}_{B(0,1)}(s) + \langle s, y \rangle \tag{2.4.2}$$

and (2.4.1) may be written as

$$\mathrm{d}_C^*(s) = \mathrm{i}_{B(0,1)}(s) + \sup_{y \in C} \langle s, y \rangle = \mathrm{i}_{B(0,1)}(s) + \sigma_C(s) \,,$$

a formula already given in Example 2.1.4. The same exercise can be carried out with the squared distance.   □

Example 2.4.3 (Directional Derivative) With $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$, let x_0 be a point where $\partial f(x_0)$ is nonempty and consider for $t > 0$ the functions

$$\mathbb{R}^n \ni d \mapsto f_t(d) := \frac{f(x_0 + td) - f(x_0)}{t} \,. \tag{2.4.3}$$

They are all minorized by $\langle s_0, \cdot \rangle$, with $s_0 \in \partial f(x_0)$; also, the difference quotient is an increasing function of the stepsize t; their infimum is therefore obtained for $t \downarrow 0$. This infimum is $f'(x_0, \cdot)$, the directional derivative of f at x_0 (already encountered in Chap. D, in a finite-valued setting). The conjugate of f_t can be computed directly or using Proposition 1.3.1:

$$s \mapsto (f_t)^*(s) = \frac{f^*(s) + f(x_0) - \langle s, x_0 \rangle}{t} \,, \tag{2.4.4}$$

so we obtain from (2.4.1)

$$[f'(x_0, \cdot)]^*(s) = \sup_{t>0} \frac{f^*(s) + f(x_0) - \langle s, x_0 \rangle}{t} \,.$$

In view of Theorem 1.4.1, the supremand is always nonnegative and is 0 if and only if $s \in \partial f(x_0)$. In a word: $[f'(x_0, \cdot)]^* = \mathrm{i}_{\partial f(x_0)}$. Taking the conjugate again, we obtain

$$[f'(x_0, \cdot)]^{**}(d) = \sigma_{\partial f(x_0)}(d) = \sup_{s \in \partial f(x_0)} \langle s, d \rangle \quad \text{for all } d \,,$$

i.e. the support function of the subdifferential is *the closure of* the directional derivative; a result which must be compared to §D.1.1.

At this point, it may be worth mentioning that the closure operation may be necessary: indeed take $f = \mathrm{i}_{B(0,1)}$, where $B(0,1)$ is the unit-ball of $\mathbb{R}^2$. Taking $x_0 = (0, -1)$ and $d = (\gamma, \delta)$, we clearly obtain as directional derivative

$$f'(x, d) = \begin{cases} 0 & \text{if } \delta > 0, \\ +\infty & \text{if } \delta \leqslant 0; \end{cases}$$

thus, $f'(x_0, \cdot)$ is not closed, and is therefore not a support function. However, its closure is indeed the support function of $\{0\} \times \mathbb{R}_-$ $(= \partial f(x_0))$.   □

A logical supplement to (2.4.1) is the conjugate of a supremum.

Theorem 2.4.4 *Let $\{g_j\}_{j \in J}$ be a collection of functions in $\overline{\text{Conv}}\,\mathbb{R}^n$. If their supre-mum $g := \sup_{j \in J} g_j$ is not identically $+\infty$, it is in $\overline{\text{Conv}}\,\mathbb{R}^n$, and its conjugate is the closed convex hull of the g_i^*'s:*

$$\left(\sup_{j \in J} g_i\right)^* = \overline{\text{co}}\left(\inf_{j \in J} g_j^*\right). \tag{2.4.5}$$

Proof. Call $f_j := g_j^*$, hence $f_j^* = g_j$, and g is nothing but the f^* of (2.4.1). Taking the conjugate of both sides, the result follows from (1.3.4). □

Example 2.4.5 Given, as in Example 2.4.2, a nonempty bounded set C, let

$$\mathbb{R}^n \ni x \mapsto \Delta_C(x) := \sup\{\|x - y\| \ : \ y \in C\}$$

be the distance from x to the most remote point of $\text{cl}\,C$. Using (2.4.2),

$$\inf_{y \in C}(f_y)^*(s) = i_{B(0,1)}(s) + \inf_{y \in C}\langle s, y\rangle,$$

so (2.4.5) gives $\Delta_C^* = \overline{\text{co}}\left(i_{B(0,1)} - \sigma_{-C}\right)$. □

Example 2.4.6 (Asymptotic Function) With $f \in \overline{\text{Conv}}\,\mathbb{R}^n$, $x_0 \in \text{dom}\,f$ and f_t as defined by (2.4.3), consider the asymptotic function of §B.3.2;

$$f'_\infty(d) := \sup\{f_t(d) \ : \ t > 0\} = \lim\{f_t(d) \ : \ t \to +\infty\}.$$

Clearly, $f_t \in \overline{\text{Conv}}\,\mathbb{R}^n$ and $f_t(0) = 0$ for all $t > 0$, so we can apply Theorem 2.4.4: $f'_\infty \in \overline{\text{Conv}}\,\mathbb{R}^n$. In view of (2.4.4), the conjugate of f'_∞ is therefore the closed convex hull of the function

$$s \mapsto \inf_{t > 0} \frac{f^*(s) + f(x_0) - \langle s, x_0\rangle}{t}.$$

Since the infimand lies in $[0, +\infty]$, the infimum is 0 if $s \in \text{dom}\,f^*$, $+\infty$ if not. In summary, $(f'_\infty)^*$ is the indicator of $\text{cl}\,\text{dom}\,f^*$. Conjugating again, we obtain $f'_\infty = \sigma_{\text{dom}\,f^*}$, a formula already seen in (1.2.3).

The comparison with Example 2.4.3 is interesting. In both cases we extremize the same function, namely the difference quotient (2.4.3); and in both cases a support function is obtained (neglecting the closure operation). Naturally, the supported set is larger in the present case of maximization. Indeed, $\text{dom}\,f^*$ and the union of all the subdifferentials of f have the same closure. □

2.5 Post-Composition with an Increasing Convex Function

For $f \in \overline{\text{Conv}}\,\mathbb{R}^n$ and $g \in \overline{\text{Conv}}\,\mathbb{R}$, the function $g \circ f$ is in $\overline{\text{Conv}}\,\mathbb{R}^n$ if g is increasing (Theorem B.2.1.7; we assume $f(\mathbb{R}^n) \cap \text{dom}\,g \neq \emptyset$ and we set $g(+\infty) = +\infty$). A relevant question is then how to express its conjugate, in terms of f^* and g^*.

We start with some preliminary observations, based on the particular nature of g. The domain of g is an interval unbounded from the left, whose relative interior is $\text{int}\,\text{dom}\,g \neq \emptyset$. Also, $\text{dom}\,g^* \subset \mathbb{R}^+$ and (remember Proposition B.1.2.5)

$$g(y) = g^{**}(y) = \sup\{py - g^*(p) \ : \ p \in \text{ri}\,\text{dom}\,g^*\}.$$

Theorem 2.5.1 *With f and g defined as above, assume that $f(\mathbb{R}^n) \cap \text{int dom } g \neq \emptyset$. For all $s \in \text{dom } (g \circ f)^*$, define the function $\psi_s \in \overline{\text{Conv}} \, \mathbb{R}$ by*

$$\mathbb{R} \ni \alpha \mapsto \psi_s(\alpha) := \begin{cases} \alpha f^*\left(\frac{1}{\alpha}s\right) + g^*(\alpha) & \text{if } \alpha > 0, \\ \sigma_{\text{dom } f}(s) + g^*(0) & \text{if } \alpha = 0, \\ +\infty & \text{if } \alpha < 0. \end{cases}$$

Then $(g \circ f)^*(s) = \min_{\alpha \in \mathbb{R}} \psi_s(\alpha)$.

Proof. By definition,

$$\begin{aligned} -(g \circ f)^*(s) &= \inf_x [g(f(x)) - \langle s, x \rangle] \\ &= \inf_{x,r} \{g(r) - \langle s, x \rangle \,:\, f(x) \leqslant r\} \qquad [g \text{ is increasing}] \\ &= \inf_{x,r} [g(r) - \langle s, x \rangle + i_{\text{epi } f}(x, r)]. \end{aligned}$$

We must compute the conjugate of the sum of the two functions $f_1(x, r) := g(r) - \langle s, x \rangle$ and $f_2 := i_{\text{epi } f}$, both with argument $(x, r) \in \mathbb{R}^n \times \mathbb{R}$. We have $\text{dom } f_1 = \mathbb{R}^n \times \text{dom } g$, so $\text{int dom } f_1 = \mathbb{R}^n \times \text{int dom } g$; hence, by assumption:

$$\text{int dom } f_1 \cap \text{dom } f_2 = (\mathbb{R}^n \times \text{int dom } g) \cap \text{epi } f \neq \emptyset.$$

Theorem 2.3.2, more precisely Fenchel's duality theorem (2.3.2), can be applied with the qualification condition (2.3.Q.jj'):

$$(g \circ f)^*(s) = \min \{f_1^*(-p, \alpha) + f_2^*(p, -\alpha) \,:\, (p, \alpha) \in \mathbb{R}^n \times \mathbb{R}\}.$$

The computation of the above two conjugates is straightforward and gives

$$(g \circ f)^*(s) = \min_{p, \alpha} \left[g^*(\alpha) + i_{\{-s\}}(-p) + \sigma_{\text{epi } f}(p, -\alpha)\right] = \min_\alpha \psi_s(\alpha),$$

where the second equality comes from (1.2.2). $\qquad\qquad\square$

Let us take two examples, using the *ball-pen function* defined on its domain $B(0, 1)$ by $f(x) = 1 - \sqrt{1 - \|x\|^2}$: here $f(\mathbb{R}^n) = [0, 1] \cup \{+\infty\}$.

– With $g(r) = \max\{0, r\}$, we leave it to the reader to compute the relevant conjugates: $\psi_s(\alpha)$ has the unique minimum-point $\bar{\alpha} = 0$ for all s.

– With $g = i_{]-\infty, 0]}$, the qualification assumption in Theorem 2.5.1 is no longer satisfied. We have $(g \circ f)^*(s) = 0$ for all s, while $\min_\alpha \psi_s(\alpha) = \|s\| - 1$.

Remark 2.5.2 Under the qualification assumption, ψ_s always attains its minimum at some $\bar{\alpha} \geqslant 0$. As a result, if for example $\sigma_{\text{dom } f}(s) = (f^*)'_\infty(s) = +\infty$, or if $g^*(0) = +\infty$, we are sure that $\bar{\alpha} > 0$. $\qquad\qquad\square$

Theorem 2.5.1 takes an interesting form in case of positive homogeneity, which can apply to f or g. For the two examples below, we recall that a closed sublinear function is the support of its subdifferential at the origin.

Example 2.5.3 If g is positively homogeneous, say $g = \sigma_\ell$ for some closed interval ℓ, the domain of ψ_s is contained in ℓ, on which the term $g^*(\alpha)$ disappears. The interval ℓ supported by g has to be contained in $\mathbb{R}^+$ (to preserve monotonicity) and the following example motivates the case $\ell = \mathbb{R}^+$.

Consider the sublevel set $C := \{x \in \mathbb{R}^n \,:\, f(x) \leqslant 0\}$, with $f : \mathbb{R}^n \to \mathbb{R}$ convex; its support function can be expressed in terms of f^*. Indeed, write the indicator of C as the composed function $i_C = i_{]-\infty, 0]} \circ f$. This places us in the present framework: g is the support

of $\mathbb{R}^+$. To satisfy the hypotheses of Theorem 2.5.1, we have only to assume the existence of an x_0 with $f(x_0) < 0$; then the result is: for $0 \neq s \in \operatorname{dom} \sigma_C$, there exists $\bar{\alpha} > 0$ such that

$$\sigma_C(s) = \min_{\alpha > 0} \alpha f^*(s/\alpha) = \bar{\alpha} f^*(s/\bar{\alpha}).$$

Indeed, we are in the case of Remark 2.5.2 because $\operatorname{dom} f = \mathbb{R}^n$.   □

Example 2.5.4 In our second example, it is f that is positively homogeneous. Then, for all $s \in \operatorname{dom}(g \circ f)^*$,

$$\psi_s(\alpha) = \begin{cases} g^*(\alpha) & \text{if } \alpha > 0 \text{ and } \frac{1}{\alpha}s \in \partial f(0), \\ \sigma_{\operatorname{dom} f}(s) + g^*(0) & \text{if } \alpha = 0, \\ +\infty & \text{if } \alpha < 0. \end{cases}$$

If f is finite everywhere, i.e. $\partial f(0)$ is bounded, we have for all $s \neq 0$

$$(g \circ f)^*(s) = \min\left\{ g^*(\alpha) : \alpha > 0 \text{ and } \tfrac{1}{\alpha}s \in \partial f(0) \right\}.$$

As an application, let Q be symmetric positive definite and define the function

$$\mathbb{R}^n \ni x \mapsto f(x) := \sqrt{\langle Q^{-1}x, x \rangle},$$

which is the support function of the elliptic set (see Example C.2.3.4)

$$E_Q := \partial f(0) = \{ s \in \mathbb{R}^n : \langle s, Qs \rangle \leqslant 1 \}.$$

The composition of f with the function $r \mapsto g(r) := 1/2\, r^2$ is the quadratic form associated with Q^{-1}, whose conjugate is known (see Example 1.1.3):

$$(g \circ f)(x) = \tfrac{1}{2}\langle Q^{-1}x, x \rangle \quad \text{and} \quad (g \circ f)^*(s) = \tfrac{1}{2}\langle s, Qs \rangle.$$

Then

$$\tfrac{1}{2}\langle s, Qs \rangle = \min\left\{ \tfrac{1}{2}\alpha^2 : \alpha \geqslant 0 \text{ and } s \in \alpha E_Q \right\}$$

is the half-square of the gauge function γE_Q (Example C.2.3.4).   □

3 Various Examples

The examples listed below illustrate some applications of the conjugacy operation. It is an obviously useful tool thanks to its convexification ability; note also that the knowledge of f^* fully characterizes $\overline{\operatorname{co}} f$, which is f in case the latter is already in $\overline{\operatorname{Conv}}\,\mathbb{R}^n$; another fundamental property is the characterization (1.4.2) of the subdifferential. All this concerns convex analysis more or less directly, but the conjugacy operation is instrumental in several areas of mathematics.

3.1 The Cramer Transformation

Let P be a probability measure on $\mathbb{R}^n$ and define its *Laplace transform* $\mathrm{L}_P : \mathbb{R}^n \to {]0, +\infty]}$ by

$$\mathbb{R}^n \ni s \mapsto \mathrm{L}_P(s) := \int_{\mathbb{R}^n} \exp \langle s, z \rangle \, \mathrm{d}P(z) \, .$$

The so-called *Cramer transform* of P is the function

$$\mathbb{R}^n \ni x \mapsto \mathrm{C}_P(x) := \sup \left\{ \langle s, x \rangle - \log \mathrm{L}_P(s) \; : \; s \in \mathbb{R}^n \right\} .$$

It is used for example in statistics. We recognize in C_P the conjugate of the function $\log \mathrm{L}_P$, which can be shown to be convex (as a consequence of Hölder's inequality) and lower semi-continuous (from Fatou's lemma). We conclude that C_P is a closed convex function, whose conjugate is $\log \mathrm{L}_P$.

3.2 The Conjugate of Convex Partially Quadratic Functions

Given a linear symmetric positive semi-definite operator B and a subspace H, let us consider

$$\mathbb{R}^n \ni x \mapsto g(x) := \begin{cases} \frac{1}{2} \langle Bx, x \rangle & \text{if } x \in H \, , \\ +\infty & \text{otherwise} \, . \end{cases} \tag{3.2.1}$$

Such a function (closed, convex, homogeneous of degree 2) is said to be *partially quadratic*. Its conjugate turns out to have just the same form, so the set of convex partially quadratic functions is stable with respect to the conjugacy operation (cf. Example 1.1.4).

Proposition 3.2.1 *The function g of* (3.2.1) *has the conjugate*

$$g^*(s) = \begin{cases} \frac{1}{2} \langle s, (\mathrm{p}_H \circ B \circ \mathrm{p}_H)^- s \rangle & \text{if } s \in \mathrm{Im}\, B + H^\perp \, , \\ +\infty & \text{otherwise} \, , \end{cases} \tag{3.2.2}$$

where p_H *is the operator of orthogonal projection onto H and* $(\cdot)^-$ *is the Moore-Penrose pseudo-inverse.*

Proof. Set $f := 1/2 \langle B\cdot, \cdot \rangle$, so that $g = f + \mathrm{i}_H$ and $g^* = (f \circ \mathrm{p}_H)^* \circ \mathrm{p}_H$ (Proposition 1.3.2). Knowing that $(f \circ \mathrm{p}_H)(x) = 1/2 \langle (\mathrm{p}_H \circ B \circ \mathrm{p}_H)x, x \rangle$, we obtain from Example 1.1.4 the conjugate $g^*(s)$ under the form

$$(f \circ \mathrm{p}_H)^*(\mathrm{p}_H s) = \begin{cases} \frac{1}{2} \langle s, (\mathrm{p}_H \circ B \circ \mathrm{p}_H)^- s \rangle & \text{if } \mathrm{p}_H s \in \mathrm{Im}\,(\mathrm{p}_H \circ B \circ \mathrm{p}_H) \, , \\ +\infty & \text{otherwise} \, . \end{cases}$$

It could be checked directly that $\mathrm{Im}\,(\mathrm{p}_H \circ B \circ \mathrm{p}_H) + H^\perp = \mathrm{Im}\, B + H^\perp$. A simpler argument, however, is obtained via Theorem 2.3.2, which can be applied since $\mathrm{dom}\, f = \mathbb{R}^n$. Thus,

$$g^*(s) = (f^* \mathbin{\dot{\triangledown}} \mathrm{i}_{H^\perp})(s) = \min \left\{ \tfrac{1}{2} \langle p, B^- p \rangle \; : \; p \in \mathrm{Im}\, B, \; s - p \in H^\perp \right\} ,$$

which shows that $\mathrm{dom}\, g^* = \mathrm{Im}\, B + H^\perp$. $\qquad\square$

For example, suppose $H = \operatorname{Im} B$. Calling $A := B^-$ the pseudo-inverse of B (hence $A^- = (B^-)^- = B$), we see that g of (3.2.1) is just f^* of Example 1.1.4 with $b = 0$. Since $\operatorname{Im} B + H^\perp = \mathbb{R}^n$, and using the relations $\mathrm{p}_H \circ B = B \circ \mathrm{p}_H = B$, we obtain finally that g^* is the f of (1.1.4) – and this is perfectly normal: $g^* = f^{**} = f$.

As another example, take the identity for B. Then $\operatorname{Im} B = \mathbb{R}^n$ and Proposition 3.2.1 gives immediately the formula already seen in Example 1.1.4:

$$\left[\tfrac{1}{2}\|\cdot\|^2 + \mathrm{i}_H\right]^*(s) = \tfrac{1}{2}\langle s, (\mathrm{p}_H)^- s\rangle = \tfrac{1}{2}\langle s, \mathrm{p}_H s\rangle = \tfrac{1}{2}\|\mathrm{p}_H s\|^2\,, \quad \text{for all } s \in \mathbb{R}^n\,.$$

3.3 Polyhedral Functions

For given $(s_i, b_i) \in \mathbb{R}^n \times \mathbb{R}$, $i = 1, \ldots, k$, set $f_i(x) := \langle s_i, x\rangle - b_i$ and define the piecewise affine function

$$\mathbb{R}^n \ni x \mapsto f(x) := \max\{f_i(x) : i = 1, \ldots, k\}\,. \tag{3.3.1}$$

Proposition 3.3.1 *At each* $s \in \operatorname{co}\{s_1, \ldots, s_k\} = \operatorname{dom} f^*$, *the conjugate of f has the value (Δ_k is the unit simplex)*

$$f^*(s) = \min\left\{\sum_{i=1}^k \alpha_i b_i : \alpha \in \Delta_k,\ \sum_{i=1}^k \alpha_i s_i = s\right\}\,. \tag{3.3.2}$$

Proof. Set $g_i(s) := \mathrm{i}_{\{s_i\}} + b_i$ and

$$g(s) := (\inf g_i)(s) = \begin{cases} b_i & \text{if } s = s_i \text{ for some } i = 1, \ldots, k\,, \\ +\infty & \text{otherwise}\,. \end{cases}$$

Apply Proposition B.2.5.4 to see that $\overline{\operatorname{co}}\, g = f^*$ of (3.3.2). The rest follows from Theorem 2.4.1 or 2.4.4, with a notational flip of f and g. $\qquad\square$

In $\mathbb{R}^n \times \mathbb{R}$ (considered here as the dual of the graph-space), each of the k vertical half-lines $\{(s_i, r) : r \geqslant b_i\}$, $i = 1, \ldots, k$ is the epigraph of $g_i = f_i^*$, and these half-lines are the "needles" of Fig. B.2.5.1. Now, the convex hull of these half-lines is a closed convex polyhedron, which is just epi f^*. Needless to say, (3.3.1) is obtained by conjugating again (3.3.2).

Example 3.3.2 Suppose that the only available information about a function $f \in \overline{\operatorname{Conv}}\, \mathbb{R}^n$ is a finite sampling of function- and subgradient-values, say $f(x_i)$ and $s_i \in \partial f(x_i)$ for $i = 1, \ldots, k$. By convexity, we therefore know that $\varphi_1 \leqslant f \leqslant \varphi_2$, where

$$\varphi_1(y) := \max\{f(x_i) + \langle s_i, y - x_i\rangle : i = 1, \ldots, k\}\,,$$
$$\varphi_2 := \operatorname{co}(\min_i g_i) \quad \text{with} \quad g_i(y) := \mathrm{i}_{\{x_i\}}(y) + f(x_i) \text{ for } i = 1, \ldots, k\,.$$

The resulting bracket on f is drawn in the left-part of Fig. 3.3.1. It has a counterpart in the dual space: $\varphi_2^* \leqslant f^* \leqslant \varphi_1^*$, illustrated in the right-part of Fig. 3.3.1, where
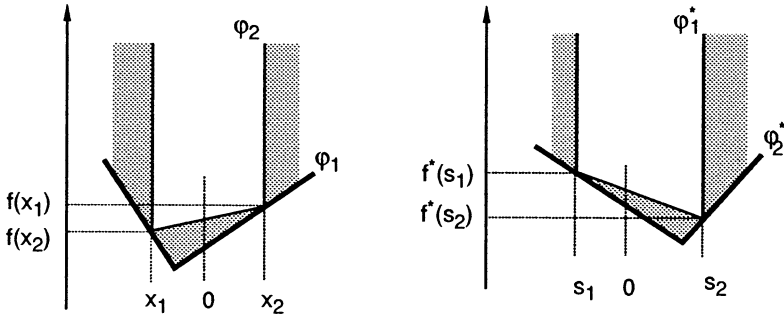
Fig. 3.3.1. Sandwiching a function and its conjugate

$$\varphi_1^*(s) = \text{co} \left[\min \{ i_{\{s_i\}}(s) + \langle s_i, x_i \rangle - f(x_i) \} \right],$$
$$\varphi_2^*(s) = \max \{ \langle s, x_i \rangle - f(x_i) \ : \ i = 1, \ldots, k \}.$$

Note: these formulae can be made more symmetric, with the help of the relations $\langle s_i, x_i \rangle - f(x_i) = f^*(s_i)$.

The function φ_1 above is often called the *cutting-plane* approximation of f. The relation $\varphi_1 \leqslant f$ is certainly richer than the relation $f \leqslant \varphi_2$, which ignores all the first-order information contained in the subgradients s_i. The interesting point is that the situation is reversed in the dual space: the cutting-plane approximation φ_2^* of f^* is obtained from the poor primal approximation φ_2 and vice versa.   □

Because the f of (3.3.1) increases at infinity no faster than linearly, its conjugate f^* has a bounded domain: it is not piecewise affine, but rather polyhedral. A natural idea is then to develop a full calculus for polyhedral functions, just as was done in §3.3 for the quadratic case.

As seen in §B.1.3(b), a polyhedral function has the general form $g = f + i_P$, where P is a closed convex polyhedron and f is defined by (3.3.1). The conjugate of g is given by Theorem 2.3.2: $g^* = f^* \mathbin{\dot\triangledown} \sigma_P$, i.e.

$$g^*(s) = \min_{\alpha \in \Delta_k} \left[\sum_{i=1}^{k} \alpha_i b_i + \sigma_P \left(s - \sum_{i=1}^{k} \alpha_i s_i \right) \right]. \tag{3.3.3}$$

This formula may take different forms, depending on the particular description of P. When $P = \text{co} \{ p_1, \ldots, p_m \}$ is (compact and) described as a convex hull, σ_P is the maximum of just as many linear functions $\langle p_j, \cdot \rangle$ and (3.3.3) becomes an explicit formula. Let us consider the case of a dual description of P.

Example 3.3.3 With the notation above and f of (3.3.1), suppose that P is described as an intersection (assumed nonempty) of half-spaces:

$$H_j := \{ x \in \mathbb{R}^n : \langle c_j, x \rangle \leqslant d_j \} \quad \text{for } j = 1, \ldots, m$$
$$\text{and} \quad P := \cap H_j = \{ x \in \mathbb{R}^n : Cx \leqslant d \}.$$

It is convenient to introduce the following notation: $\mathbb{R}^k$ and $\mathbb{R}^m$ are equipped with their standard dot-products; A [resp. C] is the linear operator which, to $x \in \mathbb{R}^n$, associates the vector whose coordinates are $\langle s_i, x \rangle$ in $\mathbb{R}^k$ [resp. $\langle c_j, x \rangle \in \mathbb{R}^m$]. It is not difficult to compute the adjoints of A and C: $A^* \alpha = \sum_{i=1}^{k} \alpha_i s_i$ for $\alpha \in \mathbb{R}^k$ and $C^* \gamma = \sum_{j=1}^{m} \gamma_j c_j$ for $\gamma \in \mathbb{R}^m$. Then we want to compute the conjugate of the polyhedral function

$$x \mapsto g(x) := \begin{cases} \max_{j=1,\dots,k} \langle s_j, x \rangle - b_j & \text{if } Cx \leqslant d, \\ +\infty & \text{if not,} \end{cases} \tag{3.3.4}$$

and (3.3.3) becomes $g^*(s) = \min \{ b^\mathsf{T} \alpha + \sigma_P(s - A^* \alpha) : \alpha \in \Delta_k \}$.

To compute σ_P, we have its conjugate i_P as the sum of the indicators of the H_j's; using Theorem 2.3.2 (with the qualification assumption (2.3.Q.j), since all the functions involved are polyhedral),

$$\sigma_P(p) = \min \left\{ \sum_{j=1}^{m} \sigma_{H_j}(s_j) : \sum_{j=1}^{m} s_j = p \right\}.$$

Using Example C.3.4.4 for the support function of H_j, it comes

$$\sigma_P(p) = \min \{ d^\mathsf{T} \gamma : C^* \gamma = p \}.$$

Piecing together, we finally obtain $g^*(s)$ as the optimal value of the problem in the pair of variables (α, γ):

$$\left| \begin{array}{l} \min (b^\mathsf{T} \alpha + d^\mathsf{T} \gamma), \quad \alpha \in \Delta_k, \ \gamma \in (\mathbb{R}^+)^m, \\ A^* \alpha + C^* \gamma = s. \end{array} \right. \tag{3.3.5}$$

Observe the nice image-function exhibited by the last constraint, characterized by the linear operator $[A^* | C^*]$. The only difference between the variables α (involving the piecewise affine f) and γ (involving the primal constraints that define P) is that the latter do not have to sum up to 1; they have no a priori bound. Also, note that the constraint of (3.3.4) can be more generally expressed as $Cx - d \in K$ (a closed convex polyhedral cone), which induces in (3.3.5) the constraint $\gamma \in K^\circ$ (another closed convex polyhedral cone, standing for the nonnegative orthant).

Finally, g^* is a closed convex function. Therefore, there is some $s \in \mathbb{R}^n$ such that the feasible domain of (3.3.5) is nonempty, and the minimal value is never $-\infty$. These properties hold provided that g in (3.3.4) satisfies (1.1.1), i.e. that the domain $Cx \leqslant d$ is nonempty. $\qquad \square$

4 Differentiability of a Conjugate Function

For $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$, we know from Corollary 1.4.4 that $s \in \partial f(x)$ if and only if $x \in \partial f^*(s)$; and this actually was our very motivation for defining f^*, see the introduction to this chapter. Geometrically, the graph of ∂f and of ∂f^* in $\mathbb{R}^n \times \mathbb{R}^n$ are images of each other under the mapping $(x, s) \mapsto (s, x)$. Knowing that a convex function is differentiable when its subdifferential is a singleton, smoothness properties of f^* correspond to monotonicity properties of ∂f, in the sense of §D.6.1.

4.1 First-Order Differentiability

Theorem 4.1.1 *Let $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$ be strictly convex. Then* $\mathrm{int\,dom}\,f^* \neq \emptyset$ *and* f^* *is continuously differentiable on* $\mathrm{int\,dom}\,f^*$.

Proof. For arbitrary $x_0 \in \mathrm{dom}\,f$ and nonzero $d \in \mathbb{R}^n$, consider Example 2.4.6. Strict convexity of f implies that

$$0 < \frac{f(x_0 - td) - f(x_0)}{t} + \frac{f(x_0 + td) - f(x_0)}{t} \quad \text{for all } t > 0\,,$$

and this inequality extends to the suprema: $0 < f'_\infty(-d) + f'_\infty(d)$. Remembering that $f'_\infty = \sigma_{\mathrm{dom}\,f^*}$ (Proposition 1.2.2), this means

$$\sigma_{\mathrm{dom}\,f^*}(d) + \sigma_{\mathrm{dom}\,f^*}(-d) > 0\,,$$

i.e. $\mathrm{dom}\,f^*$ has a positive breadth in every nonzero direction d: its interior is nonempty – Theorem C2.2.3(iii).

Now suppose that there is some $s \in \mathrm{int\,dom}\,f^*$ such that $\partial f^*(s)$ contains two distinct points x_1 and x_2. Then $s \in \partial f(x_1) \cap \partial f(x_2)$; by convex combination of the relations

$$f^*(s) + f(x_i) = \langle s, x_i \rangle \quad \text{for } i = 1, 2$$

we deduce, using Fenchel's inequality (1.1.3):

$$f^*(s) + \sum_{i=1}^{2} \alpha_i f(x_i) = \left\langle s, \sum_{i=1}^{2} \alpha_i x_i \right\rangle \leqslant f^*(s) + f\left(\sum_{i=1}^{2} \alpha_i x_i \right),$$

which implies that f is affine on $[x_1, x_2]$, a contradiction. In other words, ∂f^* is single-valued on $\mathrm{int\,dom}\,f^*$, and this means that f^* is continuously differentiable there (Remark 6.2.6). □

For an illustration, consider the function $\mathbb{R}^n \ni x \mapsto f(x) := \sqrt{1 + \|x\|^2}$, whose conjugate is

$$\mathbb{R}^n \ni s \mapsto f^*(s) = \begin{cases} -\sqrt{1 - \|s\|^2} & \text{if } \|s\| \leqslant 1\,, \\ +\infty & \text{otherwise.} \end{cases}$$

Here, f is strictly convex (compute $\nabla^2 f$ to check this), $\mathrm{dom}\,f^*$ is the unit ball, on the interior of which f^* is differentiable, but ∂f^* is empty on the boundary.

Incidentally, observe also that f^* is strictly convex; as a result, f is differentiable. Such is not the case in our next example: with $n = 1$, take

$$f(x) := |x| + \tfrac{1}{2}x^2 = \max\left\{ \tfrac{1}{2}x^2 + x, \tfrac{1}{2}x^2 - x \right\}. \tag{4.1.1}$$

Use direct calculations or calculus rules from §2 to compute

$$f^*(s) = \begin{cases} \tfrac{1}{2}(s+1)^2 & \text{for } s \leqslant -1\,, \\ 0 & \text{for } -1 \leqslant s \leqslant 1\,, \\ \tfrac{1}{2}(s-1)^2 & \text{for } s \geqslant 1\,. \end{cases}$$

This example is illustrated by the instructive Fig. 4.1.1. The left part displays gr f, made up of two parabolas; $f^*(s)$ is obtained by leaning onto the relevant parabola a straight line with slope s. The right part illustrates Theorem 2.4.4: it displays $1/2\, s^2 + s$ and $1/2\, s^2 - s$, the conjugates of the two functions making up f; epi f^* is then the convex hull of the union of their epigraphs.
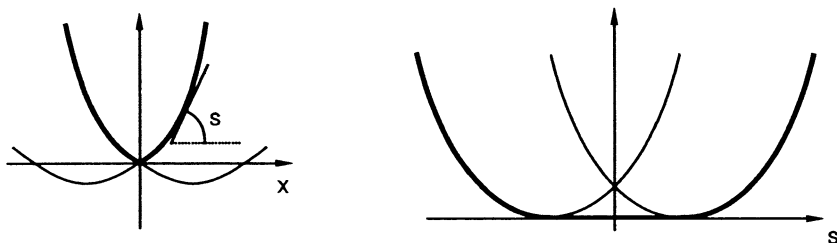
Fig. 4.1.1. Strict convexity corresponds to differentiability of the conjugate

Theorem 4.1.1 gives a sufficient condition for smoothness of a closed convex function (namely f^*). Apart from possible side-effects on the boundary, this condition is also necessary:

Theorem 4.1.2 *Let $f \in \overline{\mathrm{Conv}}\,\mathbb{R}^n$ be differentiable on the set $\Omega := \mathrm{int\,dom}\, f$. Then f^* is strictly convex on each convex subset $C \subset \nabla f(\Omega)$.*

Proof. Let C be a convex set as stated. Suppose that there are two distinct points s_1 and s_2 in C such that f^* is affine on the line-segment $[s_1, s_2]$. Then, setting $s := 1/2\,(s_1 + s_2) \in C \subset \nabla f(\Omega)$, there is $x \in \Omega$ such that $\nabla f(x) = s$, i.e. $x \in \partial f^*(s)$. Using the affine character of f^*, we have

$$0 = f(x) + f^*(s) - \langle s, x \rangle = \tfrac{1}{2} \sum_{i=1}^{2} [f(x) + f^*(s_i) - \langle s_i, x \rangle]$$

and, in view of Fenchel's inequality (1.1.3), this implies that each term in the bracket is 0: $x \in \partial f^*(s_1) \cap \partial f^*(s_2)$, i.e. $\partial f(x)$ contains the two points s_1 and s_2, a contradiction to the existence of $\nabla f(x)$. $\qquad\square$

The strict convexity of f^* cannot in general be extended outside $\nabla f(\Omega)$; and be aware that this set may be substantially smaller than one might initially guess.

We gather the results of this section, with a characterization of the "ideal situation", in which the Legendre transform alluded to in the introduction is well-defined.

Corollary 4.1.3 *Let $f : \mathbb{R}^n \to \mathbb{R}$ be strictly convex, differentiable and 1-coercive. Then*

  (i) *f^* is also finite-valued on $\mathbb{R}^n$, strictly convex, differentiable and 1-coercive;*
 (ii) *the continuous mapping ∇f is one-to-one from $\mathbb{R}^n$ onto $\mathbb{R}^n$, and its inverse is continuous;*

(iii) $f^*(s) = \langle s, (\nabla f)^{-1}(s) \rangle - f\big((\nabla f)^{-1}(s)\big)$ *for all* $s \in \mathbb{R}^n$. $\square$

The simplest such situation occurs when f is a strictly convex quadratic function, as in Example 1.1.3, corresponding to an affine Legendre transform. Another example is $f(x) := \exp\big(\tfrac{1}{2}\|x\|^2\big)$.

4.2 Lipschitz Continuity of the Gradient Mapping

Along the same lines as in §4.1, better than strict convexity of f and better than differentiability of f^* correspond to each other.

The next result is stated for a finite-valued f, mainly because the functions considered in Chap. D were such; but this assumption is actually useless.

Theorem 4.2.1 *Assume that* $f : \mathbb{R}^n \to \mathbb{R}$ *is strongly convex with modulus* $c > 0$ *on* $\mathbb{R}^n$ *: for all* $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$ *and* $\alpha \in \,]0, 1[$,

$$f(\alpha x_1 + (1-\alpha)x_2) \leqslant \alpha f(x_1) + (1-\alpha)f(x_2) - \tfrac{1}{2}c\alpha(1-\alpha)\|x_1 - x_2\|^2 . \quad (4.2.1)$$

Then $\operatorname{dom} f^* = \mathbb{R}^n$ *and* ∇f^* *is Lipschitzian with constant* $1/c$ *on* $\mathbb{R}^n$:

$$\|\nabla f^*(s_1) - \nabla f^*(s_2)\| \leqslant \tfrac{1}{c}\|s_1 - s_2\| \quad \text{for all } (s_1, s_2) \in \mathbb{R}^n \times \mathbb{R}^n .$$

Proof. We use the various equivalent definitions of strong convexity (see Theorem D.6.1.2). Fix x_0 and $s_0 \in \partial f(x_0)$: for all $0 \neq d \in \mathbb{R}^n$ and $t \geqslant 0$

$$f(x_0 + td) \geqslant f(x_0) + t\langle s_0, d \rangle + \tfrac{1}{2}ct^2\|d\|^2 ,$$

hence $f'_\infty(d) = \sigma_{\operatorname{dom} f^*}(d) = +\infty$, i.e. $\operatorname{dom} f^* = \mathbb{R}^n$. Also, f is in particular strictly convex, so we know from Theorem 4.1.1 that f^* is differentiable (on $\mathbb{R}^n$). Finally, strong convexity of f can also be written $\langle s_1 - s_2, x_1 - x_2 \rangle \geqslant c\|x_1 - x_2\|^2$, in which we have $s_i \in \partial f(x_i)$, i.e. $x_i = \nabla f^*(s_i)$, for $i = 1, 2$. The rest follows from the Cauchy-Schwarz inequality. $\square$

This result is quite parallel to Theorem 4.1.1: improving the convexity of f from "strict" to "strong" amounts to improving ∇f^* from "continuous" to "Lipschitzian". The analogy can even be extended to Theorem 4.1.2:

Theorem 4.2.2 *Let* $f : \mathbb{R}^n \to \mathbb{R}$ *be convex and have a gradient-mapping Lipschitzian with constant* $L > 0$ *on* $\mathbb{R}^n$ *: for all* $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$,

$$\|\nabla f(x_1) - \nabla f(x_2)\| \leqslant L\|x_1 - x_2\| .$$

Then f^* *is strongly convex with modulus* $1/L$ *on each convex subset* $C \subset \operatorname{dom} \partial f^*$. *In particular, there holds for all* $(x_1, x_2) \in \mathbb{R}^n \times \mathbb{R}^n$

$$\langle \nabla f(x_1) - \nabla f(x_2), x_1 - x_2 \rangle \geqslant \tfrac{1}{L}\|\nabla f(x_1) - \nabla f(x_2)\|^2 . \quad (4.2.2)$$

Proof. Let s_1 and s_2 be arbitrary in $\mathrm{dom}\,\partial f^* \subset \mathrm{dom}\,f^*$; take s and s' on the segment $[s_1, s_2]$. To establish the strong convexity of f^*, we need to minorize the remainder term $f^*(s') - f^*(s) - \langle x, s' - s \rangle$, with $x \in \partial f^*(s)$. For this, we minorize $f^*(s') = \sup_y[\langle s', y \rangle - f(y)]$, i.e. we majorize $f(y)$:

$$
\begin{aligned}
f(y) &= f(x) + \langle \nabla f(x), y - x \rangle + \int_0^1 \langle \nabla f(x + t(y - x)) - \nabla f(x), y - x \rangle \mathrm{d}t \\
&\leqslant f(x) + \langle \nabla f(x), y - x \rangle + \tfrac{1}{2} L \|y - x\|^2 = \\
&= -f^*(s) + \langle s, y \rangle + \tfrac{1}{2} L \|y - x\|^2
\end{aligned}
$$

(we have used the property $\int_0^1 t\,\mathrm{d}t = 1/2$, as well as $x \in \partial f^*(s)$, i.e. $\nabla f(x) = s$). In summary, we have

$$
f^*(s') \geqslant f^*(s) + \sup_y \left[\langle s' - s, y \rangle - \tfrac{1}{2} L \|y - x\|^2 \right].
$$

Observe that the last supremum is nothing but the value at $s' - s$ of the conjugate of $1/2\,L\| \cdot -x\|^2$. Using the calculus rule 1.3.1, we have therefore proved

$$
f^*(s') \geqslant f^*(s) + \langle s' - s, x \rangle + \tfrac{1}{2L} \|s' - s\|^2 \tag{4.2.3}
$$

for all s, s' in $[s_1, s_2]$ and all $x \in \partial f^*(s)$. Replacing s' in (4.2.3) by s_1 and by s_2, and setting $s = \alpha s_1 + (1 - \alpha)s_2$, the strong convexity (4.2.1) for f^* is established by convex combination.

On the other hand, replacing (s, s') by (s_1, s_2) in (4.2.3):

$$
f^*(s_2) \geqslant f^*(s_1) + \langle s_2 - s_1, x_1 \rangle + \tfrac{1}{2L} \|s_2 - s_1\|^2 \quad \text{for all } x_1 \in \partial f^*(s_1).
$$

Then, replacing (s, s') by (s_2, s_1) and summing: $\langle x_1 - x_2, s_1 - s_2 \rangle \geqslant 1/L \|s_1 - s_2\|^2$. In view of the differentiability of f, this is just (4.2.2), which has to hold for all (x_1, x_2) simply because $\mathrm{Im}\,\partial f^* = \mathrm{dom}\,\nabla f = \mathbb{R}^n$. $\qquad\square$

For a convex function, the Lipschitz property of the gradient mapping thus appears as equivalently expressed by (4.2.2); a result which is of interest in itself. Naturally, Corollary 4.1.3 has also its equivalent, namely: if f is strongly convex and has a Lipschitzian gradient-mapping on $\mathbb{R}^n$, then f^* enjoys the same properties. These properties do not leave much room, though: f (and f^*) must be "sandwiched" between two positive definite quadratic functions.

Exercises

1. Let C be a closed convex set.

(i) Show that the subdifferential of its indicator function is the normal cone: $\partial i_C(x) = \mathrm{N}_C(x)$ for all $x \in C$.

(ii) In particular, compute $\partial i_C(0)$, where C is the two-dimensional disk of center $e_1 := (1,0)$ and radius 1:

$$C := \{x \in \mathbb{R}^2 \; : \; \|x - e_1\| \leqslant 1\} = \{x = (\xi, \eta) \; : \; (\xi - 1)^2 + \eta^2 \leqslant 1\}.$$

(iii) With f_1, f_2 closed convex, prove the formula $\partial(f_1+f_2)(x) \supset \partial f_1(x)+\partial f_2(x)$ for all x.

(iv) With C of (ii), use the example $(i_C + i_{-C})$ to demonstrate that equality need not hold in general.

(v) Realize that equality holew ir C is shifted to the left.

(vi) Meditate this exercise with Theorems 1.4.2 and A.2.1.10 in mind.

2*. In what follows, g is a function from $\mathbb{R}^n$ to $\mathbb{R} \cup \{+\infty\}$ and $h : \mathbb{R}^n \to \mathbb{R}$ is convex; set $f := g - h$.

(i) Show that $f^*(s) = \sup\limits_{p \in \operatorname{dom} h^*} [g^*(s + p) - h^*(p)]$ for all $s \in \mathbb{R}^n$.

   Deduce

$$\inf_{x \in \mathbb{R}^n} [g(x) - h(x)] = \inf_{p \in \operatorname{dom} h^*} [h^*(p) - g^*(p)].$$

(ii) Now we consider the problem of maximizing a finite-valued function h over a set $S \subset \mathbb{R}^n$. Let $\alpha := \sup_{x \in S} h(x)$.

   – Write α as the infimum over $\mathbb{R}^n$ of a function f as described in (i).
   – Show that $-\alpha = \inf\limits_{s \in \operatorname{dom} h^*} [h^*(s) - \sigma_S(s)]$.
   – Let A be a symmetric positive definite operator. Using the above method, show that

$$-\lambda_{\max}(A) = \inf_{s \in \mathbb{R}^n} \left(2\|s\| - \langle A^{-1}s, s \rangle\right) = \inf_{s \in \mathbb{R}^n} \left(\|s\|^2 - 2\sqrt{\langle As, s \rangle}\right).$$

3. A function $f \in \overline{\operatorname{Conv}} \, \mathbb{R}^n$ is said to be invariant under an orthogonal transformation U if $f(Ux) = f(x)$ for all $x \in \mathbb{R}^n$. Show that f is invariant under U if and only if so is f^*.

4. Let $f \in \overline{\operatorname{Conv}} \, \mathbb{R}^n$; take a set $G \subset \mathbb{R}^n$ and a function $g : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ satisfying

$$\operatorname{ri} \operatorname{dom} f \subset G \subset \operatorname{dom} f \,,$$
$$g(x) \geq f(x) \quad \text{for all } x \in \mathbb{R}^n \,,$$
$$g(x) = f(x) \quad \text{for all } x \in \operatorname{ri} \operatorname{dom} f \,.$$

Then establish the relations

$$\inf_{x \in G} g(x) = \inf_{x \in \mathbb{R}^n} g(x) = \inf_{x \in \mathbb{R}^n} f(x) \,.$$

5*. This exercise is based on Proposition 3.3.1 giving the conjugate of a polyhedral function. Given $(s_i, b_i) \in \mathbb{R}^n \times \mathbb{R}$ for $i = 1, \ldots, k$, consider the closed convex polyhedron $P := \{x \in \mathbb{R}^n : \langle s_i, x \rangle \leqslant b_i, i = 1, \ldots, k\}$. Show

$$P \neq \emptyset \quad \Longleftrightarrow \quad \sum_{i=1}^{k} \alpha_i b_i \geqslant 0 \text{ for all } \alpha \in (\mathbb{R}_+)^k \text{ such that } \sum_{i=1}^{k} \alpha_i s_i = 0 \,,$$

or equivalently

$$P = \emptyset \quad \Longleftrightarrow \quad \exists \alpha \in (\mathbb{R}^+)^k \text{ such that } \sum_{i=1}^{k} \alpha_i s_i = 0 \text{ and } \sum_{i=1}^{k} \alpha_i b_i < 0 \,.$$

6*. Given $p > 0$, consider the functions f_p, g_p, h_p defined on $\mathbb{R}$ by:

$$f_p(x) = g_p(x) = h_p(x) = \frac{1}{p+1} x^{p+1} \quad \text{for } x \geqslant 0 \,,$$

while, for $x < 0$:

$$f_p(x) = +\infty \,, \quad g_p(x) = 0 \,, \quad h_p(x) = \frac{1}{p+1}(-x)^{p+1} \,.$$

Show that $f_p^* = g_{1/p}$, $g_p^* = f_{1/p}$, $h_p^* = h_{1/p}$.

7*. Let $p; r_1, \ldots, r_m; a_1, \ldots, a_m$ be real numbers, with p and the r_i's positive. For $i = 1, \ldots, m$, set

$$\mathbb{R} \ni x \mapsto f_i(x) := \frac{1}{(p+1)r_i^p} |x - a_i|^{p+1} \,.$$

Show that $(f_1 \mathbin{\dot\vee} \cdots \mathbin{\dot\vee} f_m)(x) = \frac{1}{(p+1)r^p} |x - a|^{p+1}$ for all $x \in \mathbb{R}$, where $r := \sum_i r_i$ and $a := \sum_i a_i$.

8*. Let $a_1, \ldots, a_m$ be positive real numbers. Using the inf-convolution of the functions $\mathbb{R} \ni x \mapsto 1/2 \, a_i x^2$ and of their conjugates, establish the formula

$$\left(\sum_{i=1}^{m} \frac{1}{a_i} \right)^{-1} = \min \left\{ \sum_{i=1}^{m} a_i x_i^2 : \sum_{i=1}^{m} x_i = 1 \right\} \,.$$

9. Let K be a closed convex cone of $\mathbb{R}^n$ and $f \in \overline{\mathrm{Conv}}\, \mathbb{R}^n$. Assume f is continuous at some $x_0 \in K$. Then show $\inf_{x \in K} f(x) = -\min_{s \in K^\circ} f^*(s)$.

10. Compute the conjugate of the function ch (hyperbolic cosine) on $\mathbb{R}$. Deduce an antiderivative of $(\mathrm{sh})^{-1}$ (inverse of the hyperbolic sine).

11*. For real m and σ, consider the function $q_{m,\sigma} : \mathbb{R} \to \mathbb{R}$ defined by

$$q_{m,\sigma}(x) := \frac{(x-m)^2}{2\sigma^2} \quad \text{if } \sigma \neq 0 \,, \quad \text{and } q_{m,0} = \mathrm{i}_{\{m\}} \,.$$

Compute $(q_{m,\sigma})^*$. Deduce $q_{m,\sigma} \mathbin{\dot\vee} q_{m',\sigma'} = q_{m+m', \sqrt{\sigma^2 + \sigma'^2}}$.

12*. In the space $S_n(\mathbb{R})$ equipped with its standard scalar product, consider

$$f(X) := \frac{1}{n}\|X\|^2 - \left(\frac{\operatorname{tr} X}{n}\right)^2.$$

Show that f is convex and that $f^*(S) = \begin{cases} \dfrac{n}{4}\operatorname{tr} S^2 & \text{if } \operatorname{tr} S = 0, \\ +\infty & \text{if } \operatorname{tr} S \neq 0 \end{cases}.$

13*. As in §B.1.3(f), consider the function

$$S_n(\mathbb{R}) \ni A \mapsto f(A) := \begin{cases} -\log(\det A) - n/2 & \text{if } A \text{ is positive definite}, \\ +\infty, & \text{otherwise}. \end{cases}$$

– Knowing that f is differentiable on its domain, what can be said about strict convexity of f^*?
– Assume A positive definite. Develop $\det A$ along the i^{th} column of A to obtain $\nabla(\det)(A) = (\det A)A^{-1}$, so that $\nabla f(A) = -A^{-1}$.
– Compute f^*. Show directly that f is convex, and even strictly convex.
– Deduce the inequality $\operatorname{tr}(A^{-1}B) + \operatorname{tr}(B^{-1}A) > 2n$, valid for any two different positive definite matrices A and B in $S_n(\mathbb{R})$.

14. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex, C^∞ and 1-coercive on $\mathbb{R}^n$. Show that f^* is C^∞.

15. Let f and g be in $\overline{\operatorname{Conv}}\,\mathbb{R}^n$.

For $s \in \mathbb{R}^n$, show the equivalence of the following statements:

(i) $f^*(s) + g^*(-s) \leqslant 0$;
(ii) there exists $r \in \mathbb{R}$ such that

$$f(x) \geqslant \langle s, x \rangle + r \geqslant -g(x), \quad \text{for all } x \in \mathbb{R}^n.$$

Assuming the existence of some $\bar{x} \in \mathbb{R}^n$ such that $f(\bar{x}) = -g(\bar{x})$, establish the relation

$$\{s \in \mathbb{R}^n : f^*(s) + g^*(-s) \leqslant 0\} = \partial f(\bar{x}) \cap -\partial g(\bar{x}).$$

Bibliographical Comments

Some notion of convexity appeared already in Archimedes' works. Closer to us, modern convexity theory resulted in several branches, using often the same tools or basic concepts, but with problems to solve of diverse nature: geometric convexity is one example. In the present book, it is the *variational* aspect, or the relationship with *continuous optimization* that we have stressed.

The development of convexity during the last half-century owes much to W. Fenchel (1905-1988), J.-J. Moreau (1923-), R.T. Rockafellar (1935-). Fenchel was very "geometrical"; Moreau, according to his own words, did applied mechanics: he "applied Mechanics to Mathematics"; while the concept of "dual problem" is a constant leading thread for Rockafellar. Besides mechanics, one should not forget that convexity comes naturally into play in another branch of science: thermodynamics. There, "convexifying" a function (passing from f to $\overline{\mathrm{co}}\, f$) is a common operation. The works of the physicist J.W. Gibbs (1839-1903) were a benchmark in this respect: read A.S. Wightman's introduction to [24]: "Convexity and the notion of equilibrium state in thermodynamics and statistical mechanics".

For the reader wanting to start a library, here are some suggested books:

- First of all, *the* reference book to keep on the shelf, as far as convex analysis in finite dimension is concerned, is [42].
- Convex sets and functions in infinite dimension: [34], [14, Chaps I–III], [43, 7], [5, Part I], [11, 21, 38], [28, Chaps VI–VII], [9, Chap 9].
- Convexity and mathematical economics: [35, 6, 5].
- Convexity in variational problems: [1, 14, 23, 48, 50]
- and, concerning more particularly optimization algorithms: [18, Vol. 2], [36, 47].
- Convexity and approximation theory: [20, 28].
- Convexity in statistics, in statistical mechanics: [8, 15].
- Use of convexity in nonsmooth analysis: [44, 12, 45, 9].

Let us also mention [17] for a load of additional exercises and references.

We now give some comments for more detail on subjects treated in the present volume.

Chapter A. The first systematic study of convexity (in finite dimension) is due to H. Minkowski (1864-1909); most ideas on the subject can be found in his works, at least in seminal form. Theorem 1.3.6 of Carathéodory (1873-1950) goes back also to the very beginning of the XX$^{\mathrm{th}}$ Century. A proof of the Fenchel-Bunt Theorem 1.3.7

can be found in [13, Thm 18(ii)], or [25, Lemma B.2.2]. Along the lines of these results, we mention the following theorem (of Shapley-Folkman). Let $S_1, \ldots, S_k$ be subsets of $\mathbb{R}^n$ and $S := S_1 + \cdots + S_k$ (we know that $\operatorname{co} S = \operatorname{co} S_1 + \cdots + \operatorname{co} S_k$). Any $x \in \operatorname{co} S$ can be expressed as $x_1 + \cdots + x_k$, with $x_i \in \operatorname{co} S_i$, and the set of i such that $x_i \notin S_i$ having at most n elements. This theorem and Carathéodory's can be viewed as particular cases of a more general result ([3, Lemma I]), relating the dimension of a face of C exposed by s and that of $A(C)$ (A being affine) exposed by As. Minkowski's Theorem 2.3.4 has a generalization to infinite dimension, due to Krein and Milman: a compact convex set C in a Hausdorff locally convex vector space is the closed convex hull of its extreme points. This explains that Minkowski's theorem often appears under the banner "theorem of Krein-Milman".

Moreau's Theorem 3.2.5 goes back to 1962: [33]. For developments around the Minkowski-Farkas lemma and the associated historical context, consult for example [40, 46]. The use of separation theorems to obtain multipliers in nonlinear programming, and their development through the ages, are recorded in [39].

We have totally neglected here combinatorial aspects in the study of convex sets, and convex geometry. These subjects are treated in [29, 46, 16]. To know more on closed convex polyhedra, we recommend [10].

Chapter B. The variational formulation of the sum of the m largest eigenvalues of a symmetric matrix A, as the support function of $\Omega = \{Q : Q^\top Q = \mathrm{i}_m\}$ (§1.3(e)), is due to Ky Fan (1949). Incidentally, it can be shown that $\operatorname{co} \Omega$ is nothing more than the set of positive semi-definite matrices A satisfying $\operatorname{tr} A = m$ and $\lambda_1(A) \leqslant 1$.

The operator $(A_1^{-1} + A_2^{-1})^{-1}$, constructed in Example 2.3.9, is called the parallel sum of A_1 and A_2; this parallel addition appeared for the first time in [2]. For more on the variational approach of this operation, see [32].

A proof of Theorem 4.2.3 can be found for example in [41, §44]. This last work contains interesting complements on convex functions.

Chapter C. The concept of support function has its roots in the works of Minkowski, who considered the three-dimensional case. However, the note [22] of L. Hörmander, written in a general context of locally convex topological spaces, was extremely influential in modern developments.

Let $\mathcal{H}$ be the set of positively homogeneous functions from $\mathbb{R}^n$ to $\mathbb{R}$. This is a vector space, which can be equipped with the norm (assumed finite for simplicity)

$$\mathcal{H} \ni h \mapsto \|h\| := \sup_{\|x\|=1} |h(x)|.$$

In $\mathcal{H}$, we find the convex cone $\mathcal{K}$ of (finite-valued) sublinear functions $\sigma : \mathbb{R}^n \to \mathbb{R}$. Besides, consider the set $\mathcal{C}$ of nonempty compact convex sets in $\mathbb{R}^n$. The isomorphism of §3

$$\mathcal{C} \ni C \mapsto \sigma_C \in \mathcal{K} \subset \mathcal{H}$$

is sometimes called Rådström's embedding.

Formula (3.3.5), expressing the Hausdorff distance Δ_H between two sets S and S' with the help of their support functions, appears in [22]. We mention that distance functions can also be used: in fact,

$$\Delta_H(S, S') = \sup_{x \in \mathbb{R}^n} |\mathrm{d}_S(x) - \mathrm{d}_{S'}(x)| .$$

Both formulae are applicable to some extent to unbounded sets S and S'.

Chapter D. Various names appear in 1963 to denote a vector s satisfying (1.2.1): R.T. Rockafellar in his thesis (1963) calls s "a differential of f at x"; it is J.-J. Moreau who, in a Note aux Comptes-Rendus de l'Académie des Sciences (Paris, 1963), introduces for s the word "sous-gradient".

In the years 1965-70, various calculus rules concerning sup-functions (§4.4) started to emerge. The time was ripe and several authors contributed to the subject: B.N. Pschenichnyi, V.L. Levin, R.T. Rockafellar, A. Sotskov, ... who worked in the field and used various assumptions. However, the most elaborated results are due to M. Valadier. Concerning counter-examples, K.C. Kiwiel found (4.4.7) and Remark 4.5.4 comes from [26]. Generally speaking, our assumptions in this Section 4.4 (finite-dimensional context, finite-valued functions) are more restrictive than those used by most of the above-mentioned authors; however, they allow more refined statements and less technical proofs.

The work presented in §5.1 on the maximal eigenvalue has a complete extension to the sum $f_m(M)$ of the m largest eigenvalues of a symmetric matrix M. As indicated in our preceding comments on Chap. IV, $\partial f_m(0)$ is the set of positive semi-definite matrices A such that $\mathrm{tr}\, A = m$ and $\lambda_1(A) \leqslant 1$. A general formula for the subdifferential is then:

$$\partial f_m(M) = \{A \in \partial f_m(0) \: : \: \langle\!\langle A, M \rangle\!\rangle = f_m(M)\} ,$$

i.e. $\partial f_m(M)$ is the face of $\partial f_m(0)$ exposed by M. More explicit formulae can be given; see [19, 37] and the references therein.

Chapter E. As suggested from the introduction to this chapter, the transformation $f \mapsto f^*$ has its origins in a publication of A. Legendre (1752-1833), dated from 1787. Since then, this transformation has received a number of names in the literature: conjugate, polar, maximum transformation, etc. However, it is now generally agreed that an appropriate terminology is *Legendre-Fenchel transform*, which we precisely adopted in this chapter. Let us mention that W. Fenchel (1905-1988) wrote in a letter to C. Kiselman, dated March 7, 1977: "I do not want to add a new name, but if I had to propose one now, I would let myself be guided by analogy and the relation with polarity between convex sets (in dual spaces) and I would call it for example *parabolic polarity*". Fenchel was influenced by his geometric (projective) approach, and also by the fact that the "parabolic" function $f : x \mapsto f(x) = 1/2 \, \|x\|^2$ is the only one that satisfies $f^* = f$. In our present finite-dimensional context, it is mainly Fenchel who studied "his" transformation, in papers published between 1949 and 1953.

The close-convexification – biconjugacy – of a function arises naturally in variational problems: minimizing an objective function like $\int_a^b \ell(t, x(t), \dot{x}(t))\mathrm{d}t$ is related to the minimization of the "relaxed" form $\int_a^b \overline{\mathrm{co}}\, \ell(t, x(t), \dot{x}(t))\mathrm{d}t$, where $\overline{\mathrm{co}}\, \ell$

denotes the closed convex hull of the function $\ell(t, x, \cdot)$. For references on these questions, see [14, 23].

The calculus rules of §2 are all rather classical. Among the proofs giving the conjugate of a sum, see for example [4]; the conjugate of an infimal convolution can then be seen as a corollary, and this reverts our present path §2.2 – §2.3. Post-composition with an increasing convex function is treated in [27] for vector-valued functions.

Proposition 3.2.1 was published in [32, §II.3]. As demonstrated in [30], conjugating quadratic functions turns out to be a fundamental operation for the so-called *SDP relaxation*, which has many important applications; see an overview in [49].

Section 4.2 opens the way to second-order studies of convex functions; see [31] and the references therein.

The Founding Fathers of the Discipline

Hermann Minkowski
(1864–1909)
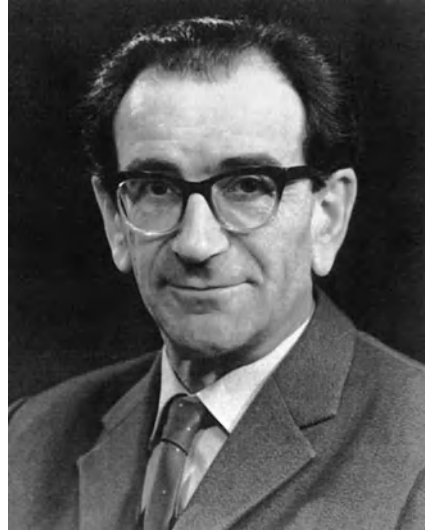
Werner Fenchel
(1905–1988)
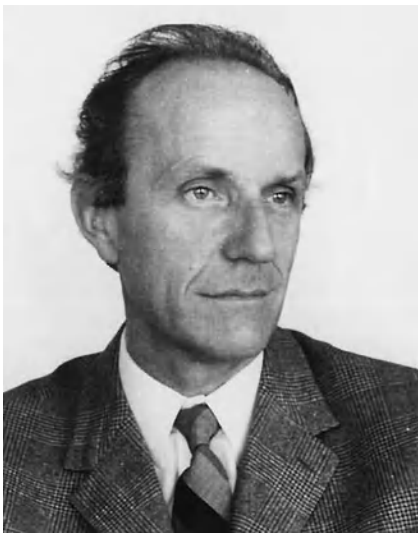
Jean-Jacques Moreau
(* 1923)

R. Tyrrell Rockafellar
(* 1935)
and research collaborator
exploring new pathways

References

1. V. Alexeev, V.M. Tikhomirov, and S. Fomine. *Commande Optimale*. Mir, Moscow, 1982.
2. W.N. Anderson Jr and R.J. Duffin. Series and parallel addition of matrices. *Journal of Mathematical Analysis and Applications*, 26:576–594, 1969.
3. Z. Artstein. Discrete and continuous bang-bang and facial spaces or: look for the extreme points. *SIAM Review*, 22(2):172–185, 1980.
4. H. Attouch and H. Brezis. Duality for the sum of convex functions in general Banach spaces. In J.A. Barroso, editor, *Aspects of Mathematics and its Applications*, pages 125–133. North Holland, Amsterdan, 1986.
5. J.-P. Aubin. *Mathematical Methods of Games and Economic Theory*. North Holland, Amsterdam, 1982.
6. J.-P. Aubin. *Optima and Equilibria: an Introduction to Nonlinear Analysis*. Springer Verlag, Heidelberg, 1993.
7. V. Barbu and T. Precupanu. *Convexity and Optimization in Banach Spaces*. Sijthoff and Noordhoff, 1992.
8. O. Barndorff-Nielsen. *Information and Exponential Families in Statistical Theory*. Wiley & Sons, 1978.
9. J. Borwein and A.S. Lewis. *Convex Analysis and Nonlinear Optimization*. Springer Verlag, New York, 2000.
10. A. Brøndsted. *An Introduction to Convex Polytopes*. Springer Verlag, New York, 1983.
11. C. Castaing and M. Valadier. *Convex Analysis and Measurable Multifunctions*. Lecture Notes in Mathematics 580. Springer Verlag, Heidelberg, 1977.
12. F.H. Clarke. *Optimization and Nonsmooth Analysis*. Wiley, 1983; reprinted by SIAM, 1983.
13. H.G. Eggleston. *Convexity*. Cambridge University Press, London, 1958.
14. I. Ekeland and R. Temam. *Convex Analysis and Variational Problems*. North Holland, 1976; reprinted by SIAM, 1999.
15. R.S. Ellis. *Entropy, Large Deviations and Statistical Mechanics*. Springer Verlag, New York, 1985.
16. P. Gritzmann and V. Klee. Mathematical programming and convex geometry. In P.M. Gruber and J.M. Wills, editors, *Handbook of Convex Geometry*, pages 627–674. North Holland, 1993.
17. J.-B. Hiriart-Urruty. *Optimisation et Analyse Convexe*. Presses Universitaires de France, Paris, 1998.
18. J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms*. Springer Verlag, Heidelberg, 1993. Two volumes.
19. J.-B. Hiriart-Urruty and D. Ye. Sensitivity analysis of all eigenvalues of a symmetric matrix. *Numerische Mathematik*, 70(1):45–72, 1995.
20. R.B. Holmes. *A Course on Optimization and Best Approximation*. Lecture Notes in Mathematics 257. Springer Verlag, Heidelberg, 1972.
21. R.B. Holmes. *Geometrical Functional Analysis and its Applications*. Springer Verlag, New York, 1975.

22. L. Hörmander. Sur la fonction d'appui des ensembles convexes dans un espace locale-
    ment convexe. *Ark. Mat.*, 3(12):181–186, 1954.
23. A.D. Ioffe and V.M. Tikhomirov. *Theory of Extremal Problems*. North Holland, 1979.
24. R.B. Israel. *Convexity in the Theory of Lattice Gases*. Princeton University Press, 1979.
25. S. Karlin. *Mathematical Methods and Theory in Games*. Mc Graw-Hill, New York,
    1960.
26. C.O. Kiselman. How smooth is the shadow of a smooth convex body? *J. London Math.
    Soc.*, 22(2):101–109, 1986.
27. S.S. Kutateladze. Changes of variables in the Young tranformation. *Soviet Mathematics
    Doklady*, 18(2):545–548, 1977.
28. P.-J. Laurent. *Approximation et Optimisation*. Hermann, Paris, 1972.
29. S.R. Lay. *Convex Sets and their Applications*. Wiley & Sons, 1982.
30. C. Lemaréchal and F. Oustry. Semi-definite relaxations and Lagrangian duality with
    application to combinatorial optimization. Rapport de Recherche 3710, Inria, 1999.
    http://www.inria.fr/rrrt/rr-3710.html.
31. C. Lemaréchal, F. Oustry, and C. Sagastizábal. The $\mathcal{U}$-Lagrangian of a convex function.
    Transactions of the AMS, 352(2):711–729, 2000.
32. M.-L. Mazure. L'addition parallèle d'opérateurs interprétée comme inf-convolution
    de formes quadratiques convexes. *Modélisation Mathématique et Analyse Numérique*,
    20:497–515, 1986.
33. J.-J. Moreau. Décomposition orthogonale d'un espace hilbertien selon deux cônes
    mutuellement polaires. *C. R. Acad. Sci. Paris*, 255:238–240, 1962.
34. J.-J. Moreau. *Fonctionnelles convexes*. Séminaire "Équations aux dérivées partielles".
    Collège de France, Paris, 1966.
35. H. Moulin and F. Fogelman-Soulié. *La Convexité dans les Mathématiques de la
    Décision*. Hermann, Paris, 1979.
36. Yu.E. Nesterov and A.S. Nemirovskii. *Interior-Point Polynomial Algorithms in Convex
    Programming*. SIAM Studies in Applied Mathematics 13. SIAM, Philadelphia, 1994.
37. M.L. Overton and R.S. Womersley. Optimality conditions and duality theory for mini-
    mizing sums of the largest eigenvalues of symmetric matrices. *Mathematical Program-
    ming*, 62(2):321–358, 1993.
38. R.R. Phelps. *Convex Functions, Monotone Operators and Differentiability*. Lecture
    Notes in Mathematics 1364. Springer Verlag, Heidelberg, 1993.
39. B.H. Pourciau. Modern multiplier rules. *American Mathematical Monthly*, 87:433–452,
    1980.
40. A. Prékopa. On the development of optimization theory. *American Mathematical Month-
    ly*, 87:527–542, 1980.
41. A.W. Roberts and D.E. Varberg. *Convex Functions*. Academic Press, 1973.
42. R.T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
43. R.T. Rockafellar. *Conjugate Duality and Optimization*. Regional Conference Series in
    Applied Mathematics 16. SIAM Publications, 1974.
44. R.T. Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal
    on Control and Optimization*, 14:877–898, 1976.
45. R.T. Rockafellar and R.J.-B. Wets. *Variational Analysis*. Springer Verlag, Heidelberg,
    1998.
46. A. Schrijver. *Theory of Linear and Integer Programming*. Wiley-Interscience, 1986.
47. V.M. Tikhomirov. The evolutions of methods of convex optimization. *American Mathe-
    matical Monthly*, 103(1):65–71, 1996.
48. J.L. Troutman. *Variational Calculus with Elementary Convexity*. Springer Verlag, New
    York, 1983.
49. L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95,
    1996.
50. M. Willem. *Analyse Convexe et Optimisation*. Éditions CIACO, Louvain-la-Neuve,
    1989.

Index