Set-valued analysis and optimisation

Lecture notes

Tuomo Valkonen
tuomov@iki.fi

Last updated February 4, 2019

Contents

1 Introduction

1.1 Minima of non-smooth functions

We recall from basic optimisation courses and textbooks (e.g., [1]), that if $f : \mathbb{R}^n \to \mathbb{R}$ is differentiable, and $\hat{x}$ is a minimiser of f,

$$f(\hat{x}) = \min_{x \in \mathbb{R}^n} f(x), \tag{1.1}$$

then

$$\nabla f(x) = 0. \tag{1.2}$$

If f is convex, the condition (1.2) is even sufficient to ensure (1.1). But what if f is non-smooth, such as when

$$f(x) = |x|, \quad (x \in \mathbb{R})?$$

It is clear that $\hat{x} = 0$ is a minimiser of this function, but at the same time $\nabla f(0)$ does not exist. If we look at the epigraph of f, defined for general $f : \mathbb{R}^n \to \mathbb{R}$ as the set of points

$$\operatorname{epi} f := \{(x, t) \in \mathbb{R}^n \times \mathbb{R} \mid t \geq f(x),\ x \in \mathbb{R}^n\},$$

as illustrated in Figure 1.1, we can however see that the hyperplane orthogonal to $(0, -1)$ supports f, touching it at $(0, 0)$. It will turn out, as we will see in this course, that this geometric, set-valued approach allows us to differentiate non-smooth functions. In fact, we can even differentiate more general set-valued functions $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$, where for each $x \in \mathbb{R}^n$, the value $F(x) \subset \mathbb{R}^m$ is a set. This will be useful for stability analysis: seeing how the solutions of problems change, as the data changes—we next take a brief look at problems involving data.

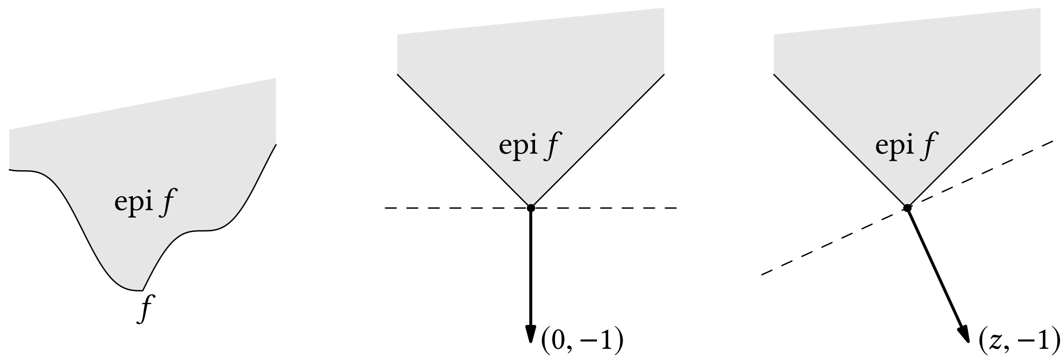

(a) Epigraph of a function.

(b) Epigraph of $f(x) = |x|$ with a supporting hyperplane at $(0, 0)$.

(c) An alternative supporting hyperplane at $(0, 0)$.

Figure 1.1: Epigraphs and supporting hyperplanes. The supporting hyperplane in (b) together with the orthogonal vector, correspond to optimality conditions.

1.2 Applications in image processing

Non-smooth optimisation problems can be found in various fields. In important application area is image processing. The most prototypical problem therein is denoising. This can be done by total variation regularisation,

$$\min_{u \in \mathbb{R}^{n_1 n_2}} \frac{1}{2}\|f - u\|^2 + \alpha \|\nabla_d u\|_{2,1}. \tag{1.1}$$

(a) Noisy image **(b)** Denoised image

Figure 1.2: Demonstration of image denoising with total variation regularisation (1.1). Note how the leaf edges are preserved by the denoising procedure. This is an important feature of total variation type approaches.

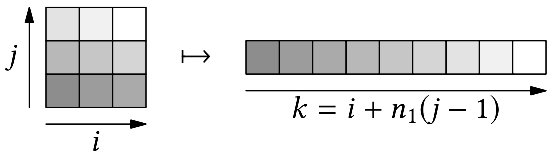

Figure 1.3: Mapping of an $n_1 \times n_2$ pixel grid into a vector of length $n_1 n_2$.

Here we consider images as $n_1 \times n_2$ pixels grids, mapping them for simplicity of overall treatment in these notes, into vectors of length $n_1 n_2$. Thus the pixel at index (i, j) is the element $u_{i+n_1(j-1)}$ of the vector u, as illustrated in Figure 1.3. Here $i \in \{1, \ldots, n_1\}$ and $j \in \{1, \ldots, n_2\}$. The first term in (1.1), the fidelity term, measures the distance of our solution u to the noisy image f. The second regularisation term tells us that the solution should be pretty. The regularisation parameter $\alpha > 0$ balances between these two goals. We illustrate the effect of TV-denoising on a colour image in (1.2).

The matrix

$$\nabla_d = \begin{pmatrix} \nabla_{d,x} \\ \nabla_{d,y} \end{pmatrix} \in \mathbb{R}^{2n_1 n_2 \times n_1 n_2}$$

is a finite-differences approximation of the image gradient. For example, forward differences with Neumann boundary conditions, may be written

$$[\nabla_{d,x} u]_{i+n_1(j-1)} = \begin{cases} u_{i+1+n_1(j-1)} - u_{i+n_1(j-1)}, & 1 \le i < n_1,\ 1 \le j \le n_2 \\ 0, & i = n_1,\ 1 \le j \le n_2 \end{cases}$$

$$[\nabla_{d,y} u]_{i+n_1(j-1)} = \begin{cases} u_{i+n_1 j} - u_{i+n_1(j-1)}, & 1 \le i \le n_1,\ 1 \le j < n_2 \\ 0, & 1 \le i \le n_2,\ j = n_2. \end{cases}$$

We use the 1-2 combination norm

$$\|g\|_{2,1} := \sum_{k=1}^{n_1 n_2} \sqrt{g_k^2 + g_{n_1 n_2 + k}^2},$$

where we take the image-wide 1-norm over the field of 2-norms of the pixelwise gradient approximations. Observe—just try to differentiate!—that this norm is non-smooth: it does not have a conventional gradient if $g_k^2 + g_{n_1 n_2 + k}^2 = 0$. If we replaced $\|\nabla_d u\|_{2,1}$ by the squared norm $\|\nabla_d u\|_{2,1}^2$, we could make the problem smooth. However, the special properties of the image-wide one-norm are important for edge preservation in image processing.

Besides the total variation regulariser $\|\nabla_d u\|$, various more modern higher-order regularisers—that make the image even prettier–also exist. They are also non-smooth. For the purposes of this course, total variation will however suffice.

Various other image processing problems besides denoising can be constructed by replacing the first term by one involving a matrix $T \in \mathbb{R}^{m \times n_1 n_2}$, that is

$$\min_{u \in \mathbb{R}^{n_1 n_2}} \frac{1}{2}\|f - Tu\|^2 + \alpha\|\nabla_d u\|_{2,1}.$$

For deblurring, T can be a convolution operation. For sub-sampled reconstruction from Fourier samples, as is the case with magnetic resonance imaging (MRI) reconstructions, $T = S\mathcal{F}$ for $S \in \mathbb{R}^{m \times n_1 n_2}$ for $m \ll n_1 n_2$ a sub-sampling operator, selecting certain pixels from u and neglecting the rest, and $\mathcal{F}$ the Fourier transform. If simply $T = S$ for a sub-sampling operator, then we are talking about inpainting. This might be used, for example, to hide hairs and scratches in old photographs or films. For a detailed treatment of various image processing tasks, see, for example [2].

1.3 Applications in the data sciences

Problems of similar structure can be found in statistics and machine learning. Various problems therein can be formulated as instances of empirical risk minimisation

$$\min_{x \in \mathbb{R}^m} g(x) + \frac{1}{n}\sum_{i=1}^{n} \phi_i(a_i^T x) \tag{1.1}$$

where $a_i^T x$ is a linear predictor, ϕ_i a convex loss function, and g a *regulariser.*

Example 1.1 (Support vector machines). If a_i is a feature vector associated to a label $b_i = \pm 1$, and we set $\phi_i(z) = \max\{0, 1 - b_i z\}$ to be the hinge loss, and $g(x) = \frac{\lambda}{2}\|x\|_2^2$ for a parameter $\lambda > 0$, then (1.1) becomes a linear support vector machine (SVM). The interpretation here is that ϕ_i does not penalise x if a_i is on the right side of the hyperplane

$$H_x = \{y \in \mathbb{R}^m \mid \langle x, y\rangle = 0\}.$$

This side is determined by the sign of b_i, known as the label or class of the data a_i. If a_i is too far on the wrong side of H_x, meaning

$$a_i \notin H_x + b_i \frac{x}{\|x\|^2}, \tag{1.2}$$

then it is penalised by the amount $1 - b_i a_i^T x$. Observe how the length of x, which is controlled by g and particularly the parameter $\lambda > 0$, controls the thickness of the set in (1.2), also known as the margin of the SVM. Linear support vector machines can thus be used to find—if possible—a separating hyperplane to samples a_i spurning from two different clusters characterised by $b_i = \pm 1$. They are also tolerant to outliers through the penalisation in ϕ_i instead of strict constraints. When the hyperplane is discovered, it can then be used for classifying new samples

Note that the general approach described here only supports hyperplanes H_x containing the origin, but increasing dimensions through replacing a_i by $a_i' = (a_i, 1)$, it is easy to support affine separating hyperplanes.

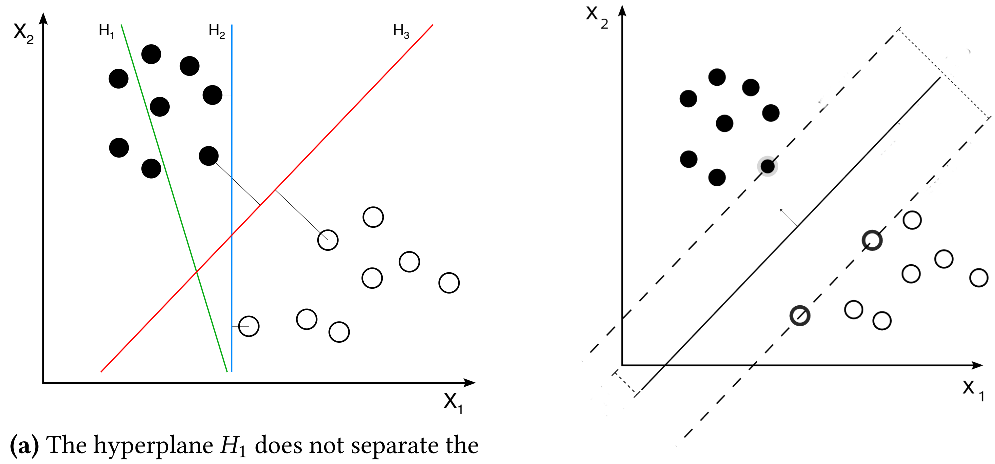

(a) The hyperplane H_1 does not separate the two classes. H_2 does, but only with a small margin. The hyperplane H_3 separates them with the optimum margin.

(b) The margin of the SVM is the distance $2/\|x\|$ between the dashed lines.

Figure 1.4: Illustrations of a linear support vector machines.
(a) is due to user ZackWeinberg on Wikipedia, licensed under Creative Commons BY-SA-3.0. It can be found at https://commons.wikimedia.org/wiki/File:Svm_separating_hyperplanes_(SVG).svg.
(b) is based on an image due to user Peter Buch on Wikipedia, and in the public domain. The original can be found at https://commons.wikimedia.org/wiki/File:Svm_max_sep_hyperplane_with_margin.png.

Non-linear support vector machines, that basically amount to transforming the data into a higher-dimensional space, and then applying the basic linear support vector machine, also exist. These kernel methods are discussed, for example, in [3].

Example 1.2 (Lasso). For linear regression, with each a_i a data vector associated with a dependent variable or measurement $b_i \in \mathbb{R}$, let us set $\phi_i(z) = \frac{1}{2}\|z - b_i\|_2^2$ and $g(x) = \lambda\|x\|_1$. This is the so-called Lasso. It finds sparse least squares solutions to the system $a_i^T x = b_i$, $(i = 1, \ldots, n)$. Sparse here means that the solution x will have many zero components, as enforced by the 1-norm. Thus, to explain the data, Lasso automatically selects more relevant features from the data, ignoring irrelevant ones.

1.4 Saddle-point formulations

The problems we have looked at above, are of the general form

$$\min_{x \in \mathbb{R}^n} g(x) + f(Kx), \tag{1.1}$$

for some $K \in \mathbb{R}^{m \times n}$, $f : \mathbb{R}^m \to \overline{\mathbb{R}}$ non-smooth, and $g : \mathbb{R}^n \to \mathbb{R}$ smooth. If we can write

$$f(z) = \max_{y \in \mathbb{R}^m} \langle y, z \rangle - f^*(y), \tag{1.2}$$

for some conjugate function f^*, then we may write the problem in the saddle-point form

$$\min_{x \in \mathbb{R}^n} \max_{y \in \mathbb{R}^m} g(x) + \langle Kx, y \rangle - f^*(y),$$

It will turn out that this kind of formulations are useful for deriving efficient algorithms. Indeed, if $m = 2M$ for some M, and

$$f(z) = \|z\|_{2,1} = \sum_{k=1}^{M} \sqrt{z_k^2 + z_{M+k}^2},$$

as we would have for the total variation denoising (1.1), we may write

$$f(z) = \sum_{k=1}^{M} \max\{z_k y_k + z_{M+k} y_{M+k} \mid |y_k^2 + y_{M+k}^2| \le 1\} = \max_{y \in \mathbb{R}^m} \langle y, z \rangle - f^*(y).$$

Here

$$f^*(y) = \begin{cases} 0, & \max_{k=1,\ldots,M} |y_k^2 + y_{M+k}^2| \le 1, \\ \infty, & \text{otherwise} \end{cases}$$

is the indicator function of the pointwise (index k) two-dimensional unit ball. This is the most important example of conjugacy for our needs. It turns out that if f is convex, proper, and lower semicontinuous—topics that we will in no time get into—then f^* will always exist, and is defined by

$$f^*(y) = \sup_{z \in \mathbb{R}^m} \langle z, y \rangle - f(y).$$

Exercise 1.1. *What are the conjugate representations of*

(i) $g(x) = \|f - x\|_2^2/2$, *(*$x \in \mathbb{R}^n$*)?*

(ii) $\phi(t) = \max\{0, 1 - bt\}$, *(*$t \in \mathbb{R}$*)?*

Write the support vector machine of Example 1.1 in explicit saddle-point form.

1.5 About the course

As we have already seen, modern approaches to image processing, machine learning, and various big data applications, almost invariably involve the solution of non-smooth optimisation problems. Already at the start, in the characterisation of optimal solutions to these problems, and the development of numerical methods, we run into the most fundamental concept of set-valued analysis: the convex subdifferential, which is the topic of Chapter 2. We then develop some fundamental optimisation methods for convex problems, based the subdifferential and set-valued view in Chapter 3, with an eye to our image processing and data science example applications.

For the understanding of the stability and sensitivity of solutions under perturbations of data and model parameters (Chapter 4), we need to delve further into the differentiation of general set-valued functions–a fascinating concept faced with many challenges. In Chapter 5, we develop general set-valued differentiation and take a look at the central analytical results of this area.

Our main reference on set-valued analysis is [4], and similar to this work, we stay in the finite-dimensional confines. Set-valued analysis in infinite-dimensions is a highly involved affair, studied in detail in [5]. For a less extensive and slightly out-dated treatment, see also [6]. Basic convex analysis, with which we start, may be studied from [7] and [8]. The infinite-dimensional case is treated in the classic [9], and more comprehensively in [10]. For brushing up on basics of numerical optimisation of smooth functions, we point to [1]—such background is however not strictly necessary. All that is required is knowledge of undergraduate calculus and linear algebra, as well as elementary geometry.

2 Convex analysis—subdifferentials

2.1 Convexity

We know intuitively what a convex set is: one can see from any point in the set, to any other point in the set. This is also the proper definition of a convex set.

Definition 2.1. A subset $C \subset \mathbb{R}^n$ is convex if

$$\lambda x + (1-\lambda)y \in C, \quad \text{whenever} \quad x, y \in C, \lambda \in [0,1].$$

Clearly, the intersection of convex sets is a convex set. One way to define a convex function then is, that $\operatorname{epi} f$ is convex. We will however provide a more explicit definition. We work with extended real numbers, $\overline{\mathbb{R}} := [-\infty, \infty]$.

Definition 2.2. We say that $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is convex if

$$f(\lambda x + (1-\lambda)y) \le \lambda f(x) + (1-\lambda)f(y), \quad (x, y \in \mathbb{R}^n;\ \lambda \in [0,1]).$$

Example 2.1. Any norm is convex, indeed $\|\lambda x + (1-\lambda)y\| \le \lambda\|x\| + (1-\lambda)\|y\|$.

For our application purposes, the next exercise covers the most interesting types of convex functions.

Exercise 2.1. *Show that the following functions are convex:*

(i) Any linear function $x \mapsto \langle x, a\rangle$ for some $a \in \mathbb{R}^n$.

(ii) Any linear combination $\sum_{i=1}^n \alpha_i f_i$ of convex functions f_i with $\alpha_i \ge 0$.

(iii) $x \mapsto f(Ax)$, if $A \in \mathbb{R}^{n\times m}$ is a matrix, and $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ convex.

(iv) $t \mapsto |t|^p$ for $t \in \mathbb{R}$ is convex for $p \ge 1$.

(v) $t \mapsto -\log t$ if $t \ge 0$ and ∞ otherwise.

Hint: *For the last two examples, try to write the epigraph as the intersection of affine half-spaces $A_x := \{(z, v) \mid v - f(z) \ge f'(x)(z - x)\}$.*

Example 2.2. For a set $C \subset \mathbb{R}^n$, we define the indicator function

$$\delta_C(x) := \begin{cases} 0, & x \in C, \\ \infty, & x \notin C. \end{cases}$$

Then C is convex if and only if δ_C is convex.

Exercise(Light) 2.2. *For a convex function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$, show that the sub-level sets*

$$\operatorname{lev}_c f := \{x \in \mathbb{R}^n \mid f(x) \le c\}$$

are convex for any $c \in \overline{\mathbb{R}}$.

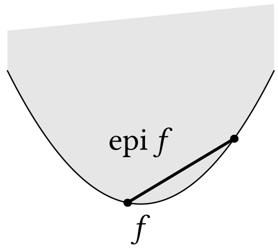

Figure 2.1: The line segment with start points within the epigraph of a convex function f, belongs completely to the epigraph.

We don't work much with general convex sets in this course, but introduce them to give a more geometric flavour to convex functions. We recall the following definition from the introduction.

Definition 2.3. The epigraph of a function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is the set

$$\operatorname{epi} f := \{(x, t) \in \mathbb{R}^n \times \mathbb{R} \mid t \geq f(x),\ x \in \mathbb{R}^n\}.$$

Exercise 2.3. *Show that $f : \Omega \to \overline{\mathbb{R}}$ is a convex function if and only if $\operatorname{epi} f$ is a convex set, cf. Figure 2.1.*

2.2 Properties of (convex) functions

Frequently, we will be making some additional assumptions about our convex function f.

Definition 2.1. Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$. We then say that

(i) f is proper, if $f(x) < \infty$ for some $x \in \mathbb{R}^n$, and $f(x) > -\infty$ for all $x \in \mathbb{R}^n$.

(ii) f is lower semicontinuous at x if for any sequence $\{x^i\}_{i=1}^{\infty} \subset \mathbb{R}^n$, with $x^i \to x$ holds

$$f(x) \leq \liminf_{i \to \infty} f(x^i).$$

(iii) f is lower semicontinuous, if it is lower semicontinuous at every $x \in \mathbb{R}^n$.

(iv) f is closed if $\operatorname{epi} f$ is a closed set.

Exercise 2.4. *Show that f is closed if and only if it is lower semicontinuous, and that $\operatorname{cl} \operatorname{epi} f$ is convex for convex f.*

This exercise motivates the following definition.

Definition 2.2. The closure or lower semicontinuous envelope of $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is the function $\operatorname{cl} f : \mathbb{R}^n \to \overline{\mathbb{R}}$ defined by

$$\operatorname{epi}(\operatorname{cl} f) = \operatorname{cl}(\operatorname{epi} f).$$

All of these properties are important for optimisation problems, as evidence by the next proposition.

Proposition 2.1. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be proper and lower semicontinuous, and $C \subset \mathbb{R}^n$ closed and bounded. Then there exists $\hat{x} \in C$ such that*

$$f(\hat{x}) = \inf_{x \in C} f(x),$$

and this value is finite.

Proof. Let

$$M := \inf_{x \in C} f(x).$$

Suppose $M = -\infty$. Then there exists a sequence $\{x^i\}_{i=1}^{\infty} \subset C$ with $f(x^i) \leq -i$ for each $i \in \mathbb{N}$. Since C is closed and bounded, we can find a limit point $x \in C$ of a subsequence. By lower semicontinuity of f, then

$$f(x) \leq \lim_{i \to \infty}(-i) = -\infty.$$

This is in contradiction to f being proper.

So $M > -\infty$. Since f is proper, there exists a point $x' \in \mathbb{R}^n$ such that $f(x') < \infty$. Therefore also $M < \infty$.

So M is finite. We may then take a minimising sequence $\{x^i\}_{i=1}^{\infty} \subset C$, such that

$$f(x^i) \leq M + 1/i.$$

Again, we may find a limit point x of a subsequence, and see by lower semicontinuity that $f(x) = M$. We have found our $\hat{x} = x$. □

Alternative proof. The set $\widetilde{E} := \operatorname{epi} f \cap (C \times \mathbb{R})$ is closed. Since f is proper, we may find a point x with $f(x) < \infty$. If we let $E := \widetilde{E} \cap ([-\infty, f(x)] \times \mathbb{R})$, then E is non-empty, because f is proper. Now taking $z^i := (x^i, f(x^i)) \in \widetilde{E}$ for a minimising sequence (which eventually and w.log satisfies $f(x^i) \leq f(x)$, we either find that $f(x^i) \searrow -\infty$, a contradiction, or may switch to a compact subset of E, where a subsequence of z^i converges. □

Remark 2.1. Note that we did not yet use convexity for the previous proposition.

For convex sets, a relative definition of the interior is often useful.

Definition 2.3. For a convex set $C \subset \mathbb{R}^n$, we define the relative interior $\operatorname{ri} C$ as the interior of C relative to the smallest affine subspace $V \supset C$.

Example 2.1. For a vector $z \in \mathbb{R}^n$, define the line segment

$$C := \{-\lambda z + (1-\lambda)z \mid 0 \leq \lambda \leq 1\}.$$

This is a one-dimensional set with

$$\operatorname{ri} C = \{-\lambda z + (1-\lambda)z \mid 0 < \lambda < 1\}.$$

Also, if

$$H = \{x \in \mathbb{R}^n \mid x^T z = 0\},$$

is the hyperplane orthogonal to z, a $(n-1)$-dimensional set, we have

$$\operatorname{ri} H = H.$$

Remark 2.2. If $\operatorname{int} A \neq \emptyset$, then $\operatorname{ri} A = \operatorname{int} A$.

Definition 2.4. For a proper function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$, we define the effective domain

$$\operatorname{dom} f := \{x \in \mathbb{R}^n \mid f(x) < \infty\}.$$

Exercise* 2.5. *Show that a convex function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is continuous on $\operatorname{ri} \operatorname{dom} f$. Conclude that $\operatorname{cl} f = f$ on $\operatorname{ri} \operatorname{dom} f$, and that*

$$\operatorname{ri}(\operatorname{epi} f) = \{(x, t) \in \mathbb{R}^{n+1} \mid x \in \operatorname{ri}(\operatorname{dom} f),\ t > f(x)\}.$$

2.3 Subdifferentials

Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be Fréchet-differentiable at $x \in \mathbb{R}$. That is, the gradient $\nabla f(x) := z$ exists, defined by

$$\lim_{h \to 0} \frac{f(x+h) - f(x) - \langle z, h \rangle}{\|h\|} = 0.$$

Note that this can also be written

$$\lim_{h \to 0} \left\langle \begin{pmatrix} x+h \\ f(x+h) \end{pmatrix} - \begin{pmatrix} x \\ f(x) \end{pmatrix}, \begin{pmatrix} z \\ -1 \end{pmatrix} \right\rangle = 0.$$

In terms of epi f, there therefore exists a supporting tangent hyperplane at $(x, f(x))$, orthogonal to $(-1, z)$; see Figure 1.1c.

Indeed, for any $h_i \to 0$, and $t_i \geq f(x + h_i)$, we have

$$\lim_{i \to \infty} \frac{t_i - f(x) - \langle z, h_i \rangle}{\|h_i\|} \geq 0.$$

In particular, as the hardest case

$$\lim_{i \to \infty} \frac{f(x+h_i) - f(x) - \langle z, h_i \rangle}{\|h_i\|} \geq 0. \tag{2.1}$$

Definition 2.1. Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$. If z satisfies (2.1) for any sequence $h_i \to 0$, we say that z is a Fréchet subgradient of f at x. We denote the set of all Fréchet subgradients of f at x by $\partial_F f(x)$.

If f is convex, we have the following simple characterisation.

Lemma 2.1. *If f is convex,* (2.1) *is equivalent to*

$$f(x+h) - f(x) \geq \langle z, h \rangle, \quad (h \in \mathbb{R}^m). \tag{2.2}$$

Proof. Indeed, (2.2) implies

$$\lim_{i \to \infty} f(x+h_i) - f(x) - \langle z, h_i \rangle \geq 0,$$

which implies (2.1).

On the other hand, if (2.2) does not hold, then

$$f(x+h) - f(x) \leq \langle z, h \rangle - \epsilon$$

for some $h \in \mathbb{R}^n \setminus \{0\}$ and $\epsilon > 0$. For any $i \in \mathbb{N}$ it follows

$$f(x+h/2^i) - f(x) = f((x+h)/2^i + (1-1/2^i)x) - f(x) \leq (1/2^i)f(x+h) - (1/2^i)f(x) \leq \langle z, h/2^i \rangle - \epsilon/2^i.$$

Therefore, setting $h_i := h/2^i$, we have

$$\lim_{i \to \infty} \frac{f(x+h_i) - f(x) - \langle z, h_i \rangle}{\|h_i\|} \leq \lim_{i \to \infty} \frac{-\epsilon/2^i}{\|h\|/2^i} = -\epsilon/\|h\|.$$

This violates (2.1). □

This motivates the following definition.

Definition 2.2. Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex, and $x \in \mathbb{R}^n$. If $z \in \mathbb{R}^n$ satisfies

$$f(x') - f(x) \geq \langle z, x' - x \rangle, \quad \text{for all } x' \in \mathbb{R}^n, \tag{2.3}$$

we say that z is a (convex) subgradient of f at x. We denote the set of all convex subgradients of f at x by $\partial f(x)$.

Geometrically, we already know that $(z, -1)$ for any $z \in \partial f(x)$ is normal to a supporting tangent hyperplane

$$H = \{(x', f(x) + \langle z, x' - x\rangle) \in \mathbb{R}^{n+1} \mid x' \in \mathbb{R}^n\}$$

of epi f at $(x, f(x))$; see Figure 1.1c. Therefore the entire set $\partial f(x)$ provides a collection of such. Moreover, each hyperplane supports the whole function globally, not just locally, in the sense that epi f stays on one side of H.

Example 2.1. Let $f(x) = |x|$ for $x \in \mathbb{R}$. Then

$$\partial f(x) = \begin{cases} \{1\}, & x > 0, \\ \{-1\}, & x < 0, \\ [-1, 1], & x = 0. \end{cases}$$

This is illustrated in Figure 2.2.

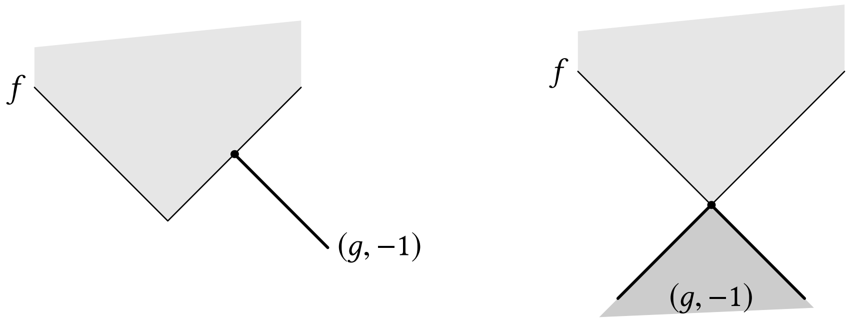

(a) $\partial f(x) = \{\operatorname{sgn} x\}$ at $x \neq 0$ **(b)** $\partial f(x) = [-1, 1]$ at $x = 0$

Figure 2.2: Subdifferentials of $f(x) = |x|$.

Exercise 2.6. *What is the subdifferential of* $\|x\|_2$ *on* $\mathbb{R}^n$*?*

Example 2.2. Let $C \subset \mathbb{R}^n$ be a convex set. Then the subdifferential of the indicator function δ_C is the normal cone

$$\partial \delta_C(x) = N_C(x) := \{z \in \mathbb{R}^n \mid \langle x' - x, z\rangle \leq 0 \text{ for all } x' \in C\}.$$

We illustrate this in Figure 2.3.

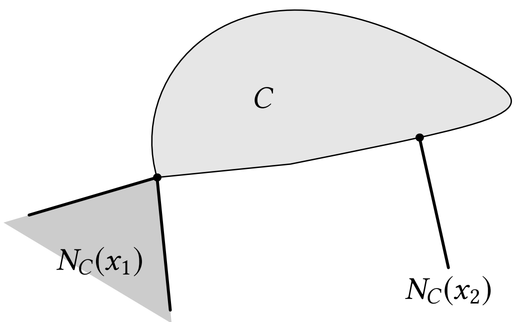

Figure 2.3: Normal cones of $f = \delta_C$ at two points x_1 and x_2.

The subdifferential $\partial f : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is an example of a set-valued map: for each $x \in \mathbb{R}^n$, the value is a subset of $\mathbb{R}^n$, $\partial f(x) \subset \mathbb{R}^n$. For general set-valued functions the equivalent concept of subdifferentiability is given by the next definition.

Definition 2.3. A set-valued function $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is monotone if

$$\langle A(x') - A(x), x' - x \rangle \geq 0, \quad (x', x \in \mathbb{R}^n).$$

(This inequality is to be understood in the sense

$$\langle y' - y, x -' x \rangle \geq 0, \quad (x', x \in \mathbb{R}^n;\ y' \in A(x'),\ y \in A(x))\big).$$

Exercise 2.7. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex. Show that ∂f is monotone.*

Exercise* 2.8. *Show that ∂f is, in fact, maximal monotone. This means that there is no monotone operator $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ such that $\operatorname{Graph} A \supset \operatorname{Graph} \partial f$.* Hint: *Observe that any $z \in \mathbb{R}^n$ can be written as $z = x + y$ for $x \in \mathbb{R}^n$ and $y \in \partial f(x)$ for some convex function f.*

We want to build some calculus rules for the convex subdifferential. For that, we need some additional results and concepts.

Proposition 2.1. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex, proper, and lower semicontinuous. Then the sub-differential mapping ∂f is outer semicontinuous, meaning that for any sequence $x^i \to x$, and $z^i \in \partial f(x^i)$, any limit z of a converging subsequence of $\{z^i\}$, satisfies $z \in \partial f(x)$. We denote*

$$\limsup_{i\to\infty} \partial f(x^i) \subset \partial f(x).$$

Moreover $\partial f(x)$ is a closed set at each $x \in \mathbb{R}^n$.

Proof. Assume, without loss of generality, that $\{z^i\}$ converges to z. Choose arbitrary $x' \in \mathbb{R}^n$. We have by Definition 2.2 that

$$f(x') \geq f(x^i) - \langle z^i, x' - x^i \rangle, \quad (i \in \mathbb{N}).$$

The map $(\widetilde{z}, \widetilde{x}) \mapsto \langle \widetilde{z}, x' - \widetilde{x} \rangle$ is continuous, and by assumption f is lower semicontinuous. Therefore

$$f(x') \geq \liminf_{i\to\infty} \big(f(x^i) - \langle z^i, x' - x^i \rangle\big) \geq f(x) - \langle z, x' - x \rangle.$$

Since this holds for any $x' \in \mathbb{R}^n$, we have proved that $z \in \partial f(x)$.

Finally, the closedness of $\partial f(x)$ is immediate from the definition, or choosing $x^i = x$ above. □

Definition 2.4. For $f : \mathbb{R}^n \to \overline{\mathbb{R}}$, we define the directional differential at $x \in \mathbb{R}^n$ in the direction $h \in \mathbb{R}^n$ by

$$f'(x; h) := \lim_{t \searrow 0} \frac{f(x + th) - f(x)}{t}. \tag{2.4}$$

Lemma 2.2. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex and proper, and $x \in \operatorname{dom} f$. Then*

$$\partial f(x) = \{z \in \mathbb{R}^n \mid \langle z, h \rangle \leq f'(x; h)\ \forall h \in \mathbb{R}^n\}, \tag{2.5}$$

and

$$\operatorname{cl}(h \mapsto f'(x; h)) = \sup_{z \in \partial f(x)} \langle z, h \rangle. \tag{2.6}$$

If $x \in \operatorname{ri} \operatorname{dom} f$, moreover

$$f'(x; h) = \sup_{z \in \partial f(x)} \langle z, h \rangle. \tag{2.7}$$

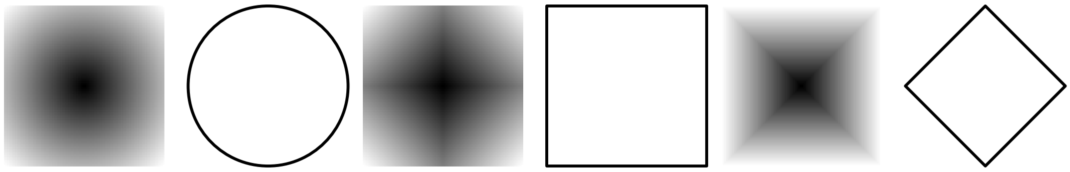

(a) For the two-norm, $\partial\|0\|_2$ is the unit circle. **(b)** For the one-norm, $\partial\|0\|_1$ is the rectangle $[-1,1]^2$. **(c)** For the ∞-norm, $\partial\|0\|_1$ is the diamond.

Figure 2.4: Some support functions on $\mathbb{R}^2$ and their corresponding convex sets.

Proof. Observe that

$$f'(x;h) = \inf_{t>0} \frac{f(x+th)-f(x)}{t}. \tag{2.8}$$

Indeed, for any $0 < s < t$ by convexity

$$\frac{s}{t}f(x+th) + \frac{t-s}{t}f(x) \geq f(x+sh).$$

This gives

$$f(x+th)-f(x) \geq \frac{t}{s}\big(f(x+sh)-f(x)\big).$$

Therefore, the sequence $s \mapsto \frac{f(x+sh)-f(x)}{s}$ is monotonically increasing, proving (2.8).

If we define

$$A := \{z \in \mathbb{R}^n \mid \langle z,h\rangle \leq f'(x;h) \text{ for all } h \in \mathbb{R}^n\},$$

then (2.8) and (2.3) show that $A = \partial f(x)$. This proves (2.5). Observe also from the continuity of $h \mapsto \langle z,h\rangle$ that A is closed (this also follows from Proposition 2.1), and that

$$A = \{z \in \mathbb{R}^n \mid \langle z,h\rangle \leq \mathrm{cl}[f'(x;\,\cdot\,)](h) \text{ for all } h \in \mathbb{R}^n\}. \tag{2.9}$$

Defining the support function of the *closed* convex set A,

$$\sigma_A(h) := \sup_{z\in A} \langle z,h\rangle,$$

we find that σ_A is proper, lower semicontinuous, and sublinear,

$$\sigma_A(s_1h_1 + s_2h_2) \leq s_1\sigma_A(h_1) + s_2\sigma_A(h_2), \quad (h_1,h_2 \in \mathbb{R}^n; s_1,s_2 \geq 0).$$

Also $f'(x;\,\cdot\,)$ is proper and sublinear (although possibly not lower semicontinuous). Proving this is the content of Exercise 2.9. It follows easily that $\mathrm{cl}[f'(x;\,\cdot\,)]$ is proper, sublinear, and lower semicontinuous. Since by (2.9), A is the maximal convex set A' satisfying $\sigma_{A'} \leq \mathrm{cl}[f'(x;\,\cdot\,)]$, the next therefore lemma shows that $\sigma_A = \mathrm{cl}[f'(x;\,\cdot\,)]$.

Finally, if $x \in \mathrm{ri}\,\mathrm{dom}\, f$, we have $\mathrm{cl}[f'(x;\,\cdot\,)] = f'(x;\,\cdot\,)$ by the lower semicontinuity of f on $\mathrm{ri}\,\mathrm{dom}\, f$ (Exercise 2.5). This shows (2.7). □

Exercise 2.9. *For a convex proper function $f : \mathbb{R}^n \to \overline{\mathbb{R}}$, prove that $f'(x;\,\cdot\,)$ is proper and sublinear at $x \in \mathrm{dom}\, f$.*

Lemma 2.3. *Let $\sigma : \mathbb{R}^n \to \mathbb{R}$ be proper, lower semicontinuous, and sub-linear. Then σ is the support function of the convex set*

$$\partial\sigma(0) = \{z \in \mathbb{R}^n \mid \langle z,h\rangle \leq \sigma(h) \textit{ for all } h \in \mathbb{R}^n\}. \tag{2.10}$$

That is

$$\sigma = \sigma_{\partial\sigma(0)},$$

Further, σ_A is sub-linear for any convex set A.

Some very common support functions σ and corresponding "dual balls" $\partial\sigma(0)$ are depicted in Figure 2.4.

Proof. A sub-linear function is convex. For any convex function f, we have

$$f(x) = \sup_{x' \in \mathbb{R}^n,\, z \in \partial f(x')} f(x') + \langle z, x - x' \rangle. \tag{2.11}$$

Indeed, by the definition of the subdifferential, $\geq$ holds here, while choosing $x' = x$ gives equality. Since a sub-linear function is positively homogeneous, meaning

$$\sigma(\lambda x) = \lambda \sigma(x) \quad \text{for} \quad \lambda > 0,$$

we have

$$\partial\sigma(\lambda x) = \partial\sigma(x), \quad \text{for all} \quad x \neq 0,\, \lambda > 0. \tag{2.12}$$

By the outer semicontinuity of $\partial\sigma$ (Proposition 2.1), letting $\lambda \searrow 0$, we see that

$$\partial\sigma(x) \subset \partial\sigma(0), \quad \text{for any} \quad x \in \mathbb{R}^n.$$

Let $x' \in \mathbb{R}^n$, and $z \in \partial\sigma(x')$. Then, since $z \in \partial\sigma(\lambda x')$, we get

$$0 = \sigma(\lambda x') - \sigma(x') \geq \langle z, \lambda x' - x' \rangle \geq \sigma(\lambda x') - \sigma(x').$$

Thus, for any $\lambda > 0$ and $x \in \mathbb{R}^n$, we have

$$\sigma(x') + \langle z, x - x' \rangle = \sigma(\lambda x') + \langle z, x - \lambda x' \rangle.$$

Letting $\lambda \searrow 0$, we have

$$\sigma(x') + \langle z, x - x' \rangle = \langle z, x \rangle.$$

Thus by (2.12), we have

$$\begin{aligned}
\sigma(x) &= \sup_{x' \in \mathbb{R}^n,\, z \in \partial\sigma(x')} (\sigma(x') + \langle z, x - x' \rangle) \\
&= \sup_{x' \in \mathbb{R}^n,\, z \in \partial\sigma(x')} \langle z, x \rangle \\
&= \sup_{z \in \partial\sigma(0)} \langle z, x \rangle.
\end{aligned}$$

This proves (2.10).

Finally, that σ_A is sub-linear for any convex set A, follows immediately from the definition. □

Lemma 2.4. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be a proper convex function. Then*

(i) $\partial f(x) = \emptyset$ for every $x \notin \operatorname{dom} f$.

(ii) $\partial f(x) \neq \emptyset$ for every $x \in \operatorname{ri} \operatorname{dom} f$.

Proof. The first claim is clear from the definition of the subdifferential: if $x \notin \operatorname{dom} f$, (2.3) gives the condition

$$f(x') - \infty \geq \langle z, x' - x \rangle, \quad (x' \in \mathbb{R}^n),$$

which cannot hold.

For the second claim, let $y \in \operatorname{dom} f$, and $x \in \operatorname{ri} \operatorname{dom} f$. If we cannot choose y distinct from x, then $\operatorname{ri} \operatorname{dom} f = \operatorname{dom} f = \{\bar{x}\}$ for some $\bar{x} \in \mathbb{R}^n$. This by the properness of f means that for some constant $c \in \mathbb{R}$ holds

$$f(x) = \begin{cases} c, & x = \bar{x}, \\ \infty & \text{otherwise.} \end{cases}$$

But, as is easily verified, $\partial f(\bar{x}) = \mathbb{R}^n$. Thus (ii) holds in this degenerate case.

For the rest, we may thus assume y distinct from x. With $h := y - x$, writing x as the convex combination

$$x = \frac{t}{1+t}y + \frac{1}{1+t}(x - th),$$

we then deduce

$$\frac{1}{1+t}f(x - th) - f(x) \geq -\frac{t}{1+t}f(y).$$

Thus

$$f(x - th) - f(x) \geq t(f(x) - f(y)).$$

In consequence

$$f'(x; -h) \geq f(x) - f(y) = C' > -\infty.$$

By Lemma 2.2 we observe that $\partial f(x)$ has to be non-empty. □

Remark 2.1. In the context of the proof, it can be that $f'(x; h) = -\infty$. Consider, for example,

$$f(x) = \begin{cases} \infty, & x < 0 \\ 1, & x = 0, \\ 0, & x > 0. \end{cases}$$

With $x = 0$, any $y > 0$, and $h = y - x > 0$, we have $f'(x; h) = -\infty$. The crucial bit is that $f'(x; -h) > -\infty$; in this example $f'(x; -h) = \infty$. This example illustrates how convex functions with non-full domain can exhibit somewhat strange behaviour.

Theorem 2.1. *Suppose $f, g : \mathbb{R}^n \to \overline{\mathbb{R}}$ are convex and proper. Then at any point $x \in \operatorname{dom}(f + g)$ one has*

$$\partial(f + g)(x) \supset \partial f(x) + \partial g(x).$$

If $\operatorname{ri}\operatorname{dom} f \cap \operatorname{ri}\operatorname{dom} g \neq \emptyset$, then this holds as an equality.

Proof. Take first $z \in \partial f(x)$, and $w \in \partial g(x)$. Then (2.3) immediately shows that $z + w \in \partial(f + g)(x)$. This shows the claimed inclusion.

To prove the equality under the additional assumption, we note from Lemma 2.4 for each $q = f, g, f + g$ that $\partial q(x)$ is non-empty. By Lemma 2.2 this implies

$$q'(x; h) > -\infty$$

for any h. Since $q'(x; 0) = 0$, we find that $q'(x; \cdot)$ is proper. Hence we can sum $f'(x; \cdot)$ and $g'(x; \cdot)$. Now

$$\limsup_{t \searrow 0} \frac{(f + g)(x + th) - (f + g)(x)}{t} \leq \limsup_{t \searrow 0} \frac{f(x + th) - f(x)}{t} + \limsup_{t \searrow 0} \frac{g(x + th) - g(x)}{t}.$$

Also

$$\inf_{t>0} \frac{f(x + th) - f(x)}{t} + \inf_{t>0} \frac{g(x + th) - g(x)}{t} \leq \inf_{t>0} \frac{(f + g)(x + th) - (f + g)(x)}{t}.$$

Recalling the equivalence (2.8), and the definition (2.4), therefore

$$(f + g)'(x; h) = f'(x; h) + g'(x; h). \tag{2.13}$$

This would be enough for the application of the formulas provided by Lemma 2.2, if we had $x \in \operatorname{ri}\operatorname{dom} f \cap \operatorname{ri}\operatorname{dom} g \cap \operatorname{ri}\operatorname{dom}(f + g)$. In general, without requiring this condition, using that $q'(x; \cdot)$ is proper for $q = f, g, f + g$, (2.13) implies

$$\liminf_{h' \to h} (f + g)'(x; h') = \liminf_{h' \to h} f'(x; h') + \liminf_{h' \to h} g'(x; h').$$

That is

$$\mathrm{cl}[(f+g)'(x;\cdot)] = \mathrm{cl}[f'(x;\cdot)] + \mathrm{cl}[g'(x;\cdot)].$$

(Note that taking the closure here is only effective if x is not in the relative interior of the domain of one of the functions.) Lemma 2.2 therefore gives

$$\begin{aligned}\sup_{q\in\partial(f+g)(x)} \langle h, q\rangle &= \sup_{z\in\partial f(x)} \langle h, z\rangle + \sup_{w\in\partial g(x)} \langle h, w\rangle \\ &= \sup_{q\in\partial f(x)+\partial g(x)} \langle h, q\rangle.\end{aligned} \tag{2.14}$$

Since this holds for every $h \in \mathbb{R}^n$, and both $\partial(f+g)(x)$ and $\partial f(x) + \partial g(x)$ are *closed* convex sets, we conclude equivalence. Indeed, if there was a point $z \in \big(\partial f(x) + \partial g(x)\big) \setminus \partial(f+g)(x)$, it would be at a positive distance from $\partial(f+g)(x)$, and yield a contradiction to the statement (2.14) on the support functions of these sets. □

Exercise 2.10. *Let $A \in \mathbb{R}^{n\times m}$, and $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex. Show that $\partial(f \circ A)(x) \supset A^T[\partial f](Ax)$ with equality if $\mathcal{R}(A) \cap \mathrm{ri\,dom}\, f \neq \emptyset$.*

2.4 Characterisation of minima

We now concentrate on convex $f : \mathbb{R}^n \to \overline{\mathbb{R}}$. How can we characterise minima of such functions? Going back to (2.3), we see that if $z = 0$, we have

$$f(x') - f(x) \geq 0, \quad \text{for all } x' \in \mathbb{R}^n.$$

This means exactly that x is a minimiser. Since this works both ways, we obtain the following.

Theorem 2.1. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ by convex. Then $x \in \mathbb{R}^n$ is a minimiser of f,*

$$f(x) = \min_{x'\in\mathbb{R}^n} f(x'),$$

if and only if

$$0 \in \partial f(x). \tag{2.1}$$

Example 2.1. Let $C \subset \mathbb{R}^n$ be a convex set, and $f : \mathbb{R}^n \to \mathbb{R}$ convex and differentiable. Then, by Theorem 2.1 and Example 2.2, we have

$$\hat{x} \in \arg\min_{x\in C} f(x)$$

if and only if

$$0 \in \nabla f(\hat{x}) + N_C(\hat{x}).$$

In particular, let

$$C = \{x \in \mathbb{R}^n \mid g(x) \leq 0\},$$

for some convex, differentiable, constraint function g satisfying

$$\inf_{x\in\mathbb{R}^n} g(x) < 0. \tag{2.2}$$

Then, as we will shortly see

$$N_C(x) = \begin{cases} \emptyset, & g(x) > 0, \\ \{0\}, & g(x) < 0, \\ [0,\infty)\nabla g(x), & g(x) = 0. \end{cases} \tag{2.3}$$

Therefore, we recover the usual Karush-Kuhn-Tucker conditions

$$\nabla f(\hat{x}) + \lambda \nabla g(\hat{x}) = 0 \quad \text{with} \quad \lambda \geq 0,\ \lambda g(\hat{x}) = 0,\ g(\hat{x}) \leq 0.$$

To see the expression (2.3), we first of all recall that if $x \notin C = \operatorname{dom} \delta_C$, then $N_C(x) = \partial \delta_C(x)$ is empty. Otherwise, $z \in N_C(x)$ for $x \in C$ is defined by

$$0 \geq \langle z, x' - x \rangle, \quad (\text{for all } x',\ g(x') \leq 0). \tag{2.4}$$

If $g(x) < 0$, we can find $\delta > 0$ such that $g(x') < 0$ for $\|x' - x\| < \delta$. Therefore, we see that the only possibility is $z = 0$, that is, $N_C(x) = \{0\}$. The case $g(x) = 0$ remains. Since g is convex and (2.2) holds, we deduce

$$C = \operatorname{cl}\{x' \in \mathbb{R}^n \mid g(x') < 0\}.$$

Indeed, by convexity of g, if $g(x') < 0$, then $g(\lambda x' + (1-\lambda) < \lambda g(x') < 0$ for any $\lambda \in (0,1)$. We now note from (2.8) that $g(x') < 0$ if and only if $x' = x + \lambda h$ for some $\lambda > 0$ and $h \in \mathbb{R}^n$ with $g'(x; h) < 0$. Therefore

$$\begin{aligned} C &= \operatorname{cl}\{x + \lambda h \mid \lambda > 0,\ h \in \mathbb{R}^n,\ g'(x;h) < 0\} \\ &= \operatorname{cl}\{x + \lambda h \mid \lambda > 0,\ h \in \mathbb{R}^n,\ \operatorname{cl}[g'(x;\,\cdot\,)](h) \leq 0\}. \end{aligned}$$

Since the the normal cone of an open set agrees with the normal cone of the closure, we deduce that $z \in N_C(x)$ if and only if

$$0 \geq \langle z, h \rangle, \quad (\text{for all } h,\ \operatorname{cl}[g'(x;\,\cdot\,)](h) \leq 0).$$

By Lemma 2.2, this is the same as

$$0 \geq \langle z, h \rangle, \quad (\text{for all } h, \langle \nabla g(x), h \rangle \leq 0).$$

This shows that $z = \lambda \nabla g(x)$ for some $\lambda \geq 0$.

2.5 Strong convexity and smoothness

Definition 2.1. Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex. We say that f is

(i) strictly convex if (2.3) holds strictly, that is

$$f(x') - f(x) > \langle \partial f(x), x' - x \rangle, \quad (x' \neq x \in \mathbb{R}^n).$$

(ii) γ-strongly-convex for $\gamma > 0$ if

$$f(x') - f(x) \geq \langle \partial f(x), x' - x \rangle + \frac{\gamma}{2}\|x' - x\|^2, \quad (x', x \in \mathbb{R}^n).$$

Obviously, strong convexity implies strict convexity.

Lemma 2.1. *Suppose $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is strictly convex. Then it has at most one minimiser.*

Proof. Let $\hat{x}$ be a minimiser. By Theorem 2.1, $0 \in \partial f(\hat{x})$. By strict convexity then

$$f(x') > f(x), \quad (x' \in \mathbb{R}^n).$$

Definition 2.2. Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex. We say that f is L-smooth if it is differentiable and

$$f(x') \leq f(x) + \langle \nabla f(x), x' - x \rangle + \frac{L}{2}\|x' - x\|^2, \quad (x', x \in \mathbb{R}^n). \tag{2.1}$$

One could, in principle, not require differentiability in Definition 2.2, and replace ∇f by ∂f in (2.1). Exercise 2.11 shows that this would lead nowhere.

For the next chapter, on optimisation methods, the following consequence is important. It introduces a stronger version of monotonicity of ∇f.

Lemma 2.2. *Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and L-smooth. Then ∇f is L^{-1}-co-coercive, that is*

$$L^{-1}\|\nabla f(x) - \nabla f(y)\|^2 \leq \langle \nabla f(x) - \nabla f(y), x - y\rangle, \quad (x, y \in \mathbb{R}^n). \tag{2.2}$$

Proof. We have

$$f(x') \leq f(x) + \langle \nabla f(x), x' - x\rangle + \frac{L}{2}\|x' - x\|^2. \tag{2.3}$$

Thus, adding $\langle \nabla f(y), x - x'\rangle$ on both sides, we get

$$f(x') - \langle \nabla f(y), x'\rangle \leq f(x) - \langle \nabla f(y), x\rangle - \langle \nabla f(x) - \nabla f(y), x' - x\rangle + \frac{L}{2}\|x' - x\|^2.$$

The left hand side is minimised by $x' = y$. Using $x' = x + L^{-1}(\nabla f(x) - \nabla f(y))$ on the right-hand side gives

$$f(y) - \langle \nabla f(y), y\rangle \leq f(x) - \langle \nabla f(y), x\rangle - \frac{1}{2L}\|\nabla f(x) - \nabla f(y)\|^2.$$

A fully analogous argument, starting from (2.3) with roles of x and y exchanged, gives

$$f(x) - \langle \nabla f(x), x\rangle \leq f(y) - \langle \nabla f(x), y\rangle - \frac{1}{2L}\|\nabla f(x) - \nabla f(y)\|^2.$$

Summing these two estimates, we obtain (2.2). □

Exercise 2.11. *Show that the following are equivalent:*

(i) L-smoothness of f,

(ii) L^{-1}-co-coercivity of ∇f.

(iii) Lipschitz continuity of ∇f with factor L.

2.6 Convex conjugates and duality

Let us briefly get back into the conjugate functions introduced in Section 1.4. We now make this precise.

Definition 2.1. Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be a general, possibly non-convex function. We then define the (convex) conjugate

$$f^*(y) := \sup_{x \in \mathbb{R}^n} (\langle x, y\rangle - f(x)).$$

We also denote the second conjugate $f^{**} := (f^*)^*$.

Example 2.1. The support function σ_A equals δ_A^* for any set $A \subset \mathbb{R}^n$. In Theorem 2.1 below we will see that if $A \neq \emptyset$ is convex and closed, then the opposite also holds, $\delta_A = \sigma_A^*$. In particular, the norms in Figure 2.4 are in one-to-one correspondence with the corresponding unit balls $B_q = \partial\|\cdot\|_p(0)$ also through $\delta_{B_q} = (\|\cdot\|_p)^*$ for q the conjugate exponent of p. This is defined through $1/p + 1/q = 1$.

The next exercise and proposition list some basic properties of f^* for arbitrary f.

Exercise 2.12. *Show that the function f^* is convex and lower semicontinuous for any $f : \mathbb{R}^n \to \overline{\mathbb{R}}$. Also show that f^* is proper if f is proper, lower semicontinuous, and level-bounded. The latter means that all of the level sets $\operatorname{lev}_c f$ are bounded.*

Proposition 2.1. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$. Then*

(i) $f \geq f^{**}$.

(ii) (Fenchel–Young) $f(x) + f^*(y) \geq \langle x, y\rangle$ for all $x, y \in \mathbb{R}^n$.

Proof. We first of all note that by defininition of f^* holds

$$f^*(y) \geq \langle y, x\rangle - f(x), \quad (y, x \in \mathbb{R}^n). \tag{2.2}$$

Since f is proper, we cannot have $f(x) = -\infty$, so simple rearrangements quickly yield (ii).

To prove (i), we note that if $f^{**}(x) < \infty$, then for every $\epsilon > 0$ we can find y with

$$f^{**}(x) \leq \langle x, y\rangle - f^*(y) + \epsilon.$$

Combining this with (2.2) yields

$$f^{**}(x) \leq f(x) + \epsilon.$$

Since $\epsilon > 0$ was arbitrary, we get $f^{**} \leq f$.

If $f^{**}(x) = \infty$, we can likewise for any $k \geq 0$ find y such that

$$f^{**}(x) \geq \langle x, y\rangle - f^*(y) \geq k.$$

This shows for any $x' \in \mathbb{R}^n$ that

$$\langle x, y\rangle - \big(\langle x', y\rangle - f(x')\big) \geq k.$$

Choosing $x' = x$ shows that $f(x) \geq k$. Since $k \geq 0$ was arbitrary, $f(x) = \infty$. This finishes the proof of (i). □

For convex f, we have the following stronger relationships. In particular, (i) justifies the expression (1.2) in the introduction.

Theorem 2.1. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex, proper, and lower semicontinuous. Then*

(i) (Fenchel–Moreau) $f = f^{**}$.

(ii) $f(x) + f^*(y) = \langle x, y\rangle$ if and only if $y \in \partial f(x)$.

(iii) $y \in \partial f(x)$ if and only if $x \in \partial f^*(y)$.

Proof. We already know from Proposition 2.1(i) that $f \geq f^{**}$. If $f^{**}(x) = \infty$, then this already shows that $f(x) = f^{**}(x)$. We may therefore suppose that $f^{**}(x) < \infty$. By Exercise 2.12, we know that f^{**} is proper, so also $f^{**}(x) > -\infty$. If there exists some $y \in \partial f(x) \neq \emptyset$, then by Theorem 2.1, $f^*(y) = \langle y, x\rangle - f(x)$. This shows that

$$f^{**}(x) \geq \langle x, y\rangle - f^*(y) \geq f(x).$$

This establishes that $f^{**} = f$ on $\operatorname{dom} \partial f = \{x \in \mathbb{R}^n \mid \partial f(x) \neq \emptyset\}$.

We then observe that

$$y \in \partial f(x) \implies x \in \partial f^*(y). \tag{2.3}$$

Indeed, suppose $y \in \partial f(x)$. By Theorem 2.1, this holds of and only if

$$f^*(y) = \langle y, x\rangle - f(x). \tag{2.4}$$

In particular, (ii) holds. By Proposition 2.1(ii), moreover

$$f(x) + f^*(y') \geq \langle x, y' \rangle. \tag{2.5}$$

The inequality (2.5) and equality (2.4) imply

$$f^*(y') - f^*(y) \geq \langle y' - y, x \rangle.$$

Thus $x \in \partial f^*(y')$, so (2.3) holds.

The same argument naturally also establishes

$$x \in \partial f^*(y) \implies y \in \partial f^{**}(x). \tag{2.6}$$

Thus $\partial f^{**}(x) \supset \partial f(x)$ for all $x \in \mathbb{R}^n$. We recall from (2.11) that

$$f(x) = \sup_{x' \in \mathbb{R}^n,\ y \in \partial f(x')} f(x') + \langle y, x - x' \rangle.$$

Since the sup-expression can be limited to $x' \in \operatorname{dom} \partial f$, having established that $f^{**} = f$ and $\partial f^{**}(x) \supset \partial f(x)$ on $\operatorname{dom} \partial f$, we see that $f^{**} \geq f$ on $\mathbb{R}^n$. This is what we needed to prove (i).

To prove (iii), we simply use (i) in (2.6), and combine this with (2.3). □

The next theorem provides a very useful duality correspondence.

Theorem 2.2 (Fenchel–Rockafellar "lite"). *Let $f : \mathbb{R}^m \to \overline{\mathbb{R}}$ and $g : \mathbb{R}^m \to \overline{\mathbb{R}}$ be convex, proper, and lower semicontinuous, and $K \in \mathbb{R}^{m \times n}$. Then we have weak duality*

$$\inf_{x \in \mathbb{R}^n} (g(x) + f(Kx)) + \inf_{y \in \mathbb{R}^m} \left(g^*(-K^T y) + f^*(y) \right) \geq 0. \tag{2.7}$$

Suppose

$$K(\operatorname{ri} \operatorname{dom} g) \cap \operatorname{int} \operatorname{dom} f \neq \emptyset, \tag{2.8}$$

and that $x \mapsto g(x) + f(Kx)$ has a minimiser $\hat{x}$. Then we have strong duality

$$\min_{x \in \mathbb{R}^n} (g(x) + f(Kx)) + \min_{y \in \mathbb{R}^m} \left(g^*(-K^T y) + f^*(y) \right) = 0. \tag{2.9}$$

Proof. For any $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$, we have by Proposition 2.1(ii) that

$$g(x) + g^*(-K^T y) \geq -\langle x, K^T y \rangle \quad \text{and} \quad f(Kx) + f^*(y) \geq \langle Kx, y \rangle. \tag{2.10}$$

Summing these expressions shows (2.7).

For (2.9) we plan to use Theorem 2.1(ii). By Theorem 2.1 and Theorem 2.1, whose conditions are verified by (2.8), we have

$$0 \in \partial g(\hat{x}) + (f \circ K)(\hat{x}).$$

The condition (2.8) also implies $\mathcal{R}(K) \cap \operatorname{ri} \operatorname{dom} f \neq \emptyset$ by (2.8). Exercise 2.10 therefore shows that $\partial(f \circ K)(\hat{x}) = K^T \partial f(K\hat{x})$. Thus there exists

$$\hat{y} \in \partial f(K\hat{x}) \tag{2.11a}$$

such that $0 \in \partial g(\hat{x}) + K^T \hat{y}$. In other words,

$$-K^T \hat{y} \in \partial g(\hat{x}). \tag{2.11b}$$

The "primal-dual" optimality conditions (2.11) and Theorem 2.1(ii) now make the inequalities in (2.10) to hold as equalities for $x = \hat{x}$ and $y = \hat{y}$. Thus (2.9) holds, with the first "min" in place of "inf" in (2.9) justified by our assumption of $\hat{x}$ being a minimiser. The second "min" is also justified. Indeed by (2.11), Theorem 2.1(iii), Exercise 2.10, and Theorem 2.1, we see that $0 \in \partial h(\hat{y})$ for $h(y) := g^*(-K^T y) + f^*(y)$. Therefore $\hat{y}$ is a minimiser by Theorem 2.1. □

Remark 2.1. The condition $K(\operatorname{ri}\operatorname{dom} g)\cap \operatorname{int}\operatorname{dom} f \neq \emptyset$ is enough for strong duality, without requiring the existence of a minimiser, albeit with the first "min" remaining an "inf" in (2.9). Even more relaxed conditions exist [11]. We stick to our stronger requirements, as the relaxed ones demand a little bit more machinery than we have time for, and in practise we are interested in the case when (2.11) is satisfied.

Remark 2.2. Note that (2.11) holding implies that $\hat{x}$ is the minimiser required for the theorem. Moreover, under (2.11), it is not necessary to assume (2.8), which was only used to prove (2.11).

Due to the relationships (2.7) and (2.9), we call

$$\min_{y\in\mathbb{R}^m} g^*(-K^T y) + f^*(y) \tag{D}$$

the dual problem of the primal problem

$$\min_{x\in\mathbb{R}^n} g(x) + f(Kx). \tag{P}$$

We denote by

$$\mathcal{G}(x, y) := g(x) + f(Kx) + g^*(-K^T y) + f^*(y) \geq 0$$

the duality gap. It is only zero when x solves (P), and y solves (D), hence $\mathcal{G}(x, y) \leq \epsilon$ for a suitable level $\epsilon > 0$ forms a good stopping criterion, independent of any knowledge of the optimal solution, for primal-dual algorithms. These simultaneously look for x and y by working on the saddle point problem

$$\min_{x\in\mathbb{R}^n} \max_{y\in\mathbb{R}^m} g(x) + \langle y, Kx\rangle - f^*(y). \tag{S}$$

Under the conditions of Theorem 2.2, this problem can be derived from (P) by writing $f(Kx) = \sup_y \big(\langle y, Kx\rangle - f^*(y)\big)$. The supremum isn't a maximum for all $x \in \mathbb{R}^n$, but under (2.11), it is for $x = \hat{x}$; hence the "max" in (S).

Exercise 2.13. *Assuming* (2.11) *to hold, show that*

$$\min_{x\in\mathbb{R}^n} \sup_{y\in\mathbb{R}^m} g(x) + \langle y, Kx\rangle - f^*(y) = \max_{y\in\mathbb{R}^m} \inf_{x\in\mathbb{R}^n} g(x) + \langle y, Kx\rangle - f^*(y). \tag{2.12}$$

Hence, show that (S) *can alternatively be derived from* (D) *by writing*

$$g^*(-K^T y) = \sup_x \big(-\langle x, K^T y\rangle - g(x)\big).$$

Remark 2.3. The property (2.12) is why (S) is called a *saddle point* problem. Without that, it should merely be called a min-max problem. Without conditions such as (2.11) holding, it is indeed generally not possible to swap the order of "min" and "max" in (S) without changing the problem. Generally $\inf_x \sup_y L(x, y) \leq \sup_x \inf_y L(x, y)$ for any function L. A saddle point $(\hat{x}, \hat{y})$ satisfies $L(\hat{x}, y) \leq L(\hat{x}, \hat{y}) \leq L(x, \hat{y})$ for all x and y.

Exercise 2.14. *Show that a solution* $(\hat{x}, \hat{y})$ *of* (2.11) *is a saddle point of the Lagrangian*

$$L(x, y) := g(x) + \langle y, Kx\rangle - f^*(y).$$

Example 2.2. Consider the empirical risk minimisation problem (1.1), that is

$$\min_{x\in\mathbb{R}^m} g(x) + \frac{1}{n}\sum_{i=1}^{n} \phi_i(a_i^T x).$$

We can also write this as

$$\min_{x \in \mathbb{R}^m} g(x) + \phi(A^T x) \quad \text{for} \quad A := \begin{pmatrix} a_1 & \dots a_n \end{pmatrix} \in \mathbb{R}^{n \times m} \text{ and } \phi(z) := \frac{1}{n} \sum_{i=1}^{n} \phi_i(z_i).$$

The dual problem is

$$\min_{y \in \mathbb{R}^n} g^*(-Ay) + \phi^*(y),$$

which we can also write as

$$\min_{y \in \mathbb{R}^n} \frac{1}{n} g^*(-Ay) + \sum_{i=1}^{n} \phi_i^*(y_i).$$

(You can easily observe that since each ϕ_i only depends on z_i, the conjugate of ϕ is the sum of the conjugates ϕ_i^* acting on y_i.)

For the linear SVM,

$$g(x) = \frac{\alpha}{2}\|x\|^2, \quad \text{and} \quad \phi_i(t) := \max\{0, 1 - b_i t\}.$$

These have the conjugates

$$g^*(z) = \frac{1}{2\alpha}\|z\|^2, \quad \text{and} \quad \phi_i^*(y_i) := \begin{cases} y_i/b, & y_i \in [-b, 0], \\ \infty, & \text{otherwise}, \end{cases}$$

where we denote $[-b, 0] := [0, -b]$ if $b < 0$. In this dual formulation, the non-smooth function ϕ^* therefore nicely splits into componentwise functions, with the “mixing” of the different coordinates of y_i by A moved into the smooth part $g^*(-Ay)$. This dual form of the problem will be easy to solve with the forward–backward splitting method that we introduce in the next section, while the original form is less trivial. This dual form also forms the computationally tractable basis of non-linear support vector machines, which might have $A \in \mathbb{R}^{n \times m}$ for extremely large n stemming from a non-linear transformation of the data. This is however completely hidden in the dual formulation that only operates on variables of dimension m.

Exercise 2.15. *What is the dual problem of the Lasso? Is this likely to be useful? How about the saddle point problem?*

3 Non-smooth optimisation methods

3.1 Surrogate objectives and gradient descent

Let $f : \mathbb{R}^n \to \mathbb{R}$ be convex and differentiable. We want to find a point $\hat{x}$ such that

$$f(\hat{x}) = \min_{x \in \mathbb{R}^n} f(x). \tag{P}$$

As we have learned, this is of course characterised by

$$\nabla f(\hat{x}) = 0.$$

This system is, however, in most interesting cases difficult to solve analytically. So let us try to derive a numerical methods. One way of deriving numerical methods is to replace the original difficult objective with a simpler one whose minimisation provides improvement to the original objective.

Definition 3.1. A function $\widetilde{f}_{\bar{x}} : \mathbb{R}^n \to \overline{\mathbb{R}}$ is a surrogate objective for $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ at $\bar{x}$ if $\widetilde{f}_{\bar{x}} \geq f$, and $\widetilde{f}_{\bar{x}}(\bar{x}) = f(\bar{x})$.

Starting with a point x^0, we would then minimise $\widetilde{f}_{x^0}$ to obtain a new point x^{i+1}. Through the properties of the surrogate objective, this will not increase the value of f. Hopefully it will provide significant improvement! Then we repeat the process, minimising $\widetilde{f}_{x^1}$, and so on.

What options are there for surrogate objectives, and what would be a good one? If f is differentiable, one possibility is

$$\min_{x \in \mathbb{R}^n} \widetilde{f}_{\bar{x}}(x) := f(\bar{x}) + \langle \nabla f(\bar{x}), x - \bar{x} \rangle + \frac{1}{2\tau} \|x - \bar{x}\|^2. \tag{3.1}$$

Here $\tau > 0$ is a suitable factor. In general $f(\bar{x}) = \widetilde{f}_{\bar{x}}(\bar{x})$. If f is L-smooth per Definition 2.2, and $L\tau \leq 1$, then also $f \leq \widetilde{f}_i$. Therefore, in this case, $\widetilde{f}_{\bar{x}}$ is a valid surrogate objective, and minimising $\widetilde{f}_{\bar{x}}$ will provide improvement to f as well.

The optimality condition $0 \in \partial \widetilde{f}_{x^i}(x)$ becomes

$$\nabla f(x^i) + \tau^{-1}(x - x^i) = 0. \tag{3.2}$$

This holds if $x^i = \hat{x}$ by taking also $x = \hat{x}$. Therefore, there is a direct correspondence between the solutions of the surrogate objective and the original. If $x^i \neq \hat{x}$, solving (3.2) for $x = x^{i+1}$, we get the rule

$$x^{i+1} = x^i - \tau \nabla f(x^i). \tag{GD}$$

This is known as the gradient descent method. In this context the quadratic term in (3.1) can be seen as a step length condition.

Will sequentially minimising $\widetilde{f}_{x^i}$ provide sufficient decrease in f such that we obtain convergence of $\{x^i\}$ to a minimiser $\hat{x}$ of f? This is what we study next.

3.2 Fixed point theorems

Convergence of optimisation methods can often by proved through various fixed point theorems applied to the operator $T : x^i \mapsto x^{i+1}$, mapping one iterate to the next one. We will in particular use the following result from [12].

Theorem 3.1 (Browder fixed point theorem, version 1). *Let $T : \mathbb{R}^n \to \mathbb{R}^n$ be firmly non-expansive, that is*

$$\|T(x) - T(y)\|^2 \le \langle T(x) - T(y), x - y\rangle, \quad (x, y \in \mathbb{R}^n).$$

Suppose T admits some fixed point $x^* = T(x^*)$. Then, for any starting point $x^0 \in \mathbb{R}^n$, the iteration sequence $x^{i+1} := T(x^i)$ satisfies $x^i \to \widetilde{x}$ for some *fixed point $\widetilde{x} = T(\widetilde{x})$.*

Remark 3.1. Firm non-expansivity is the co-coercivity of (2.2) with constant $L = 1$.

Th above variant of Browder's fixed point theorem follows from a more general one for averaging operators.

Definition 3.1. A map $T : \mathbb{R}^n \to \mathbb{R}^n$ is non-expansive, if

$$\|T(x) - T(y)\| \le \|x - y\|, \quad (x, y \in \mathbb{R}^n).$$

It is α-averaging, if $T = (1 - \alpha)I + \alpha J$ for some non-expansive $J : \mathbb{R}^n \to \mathbb{R}^n$, and $\alpha \in (0, 1)$.

Theorem 3.2 (Browder fixed point theorem, version 2). *Let $T : \mathbb{R}^n \to \mathbb{R}^n$ be averaging, and suppose T admits some fixed point $x^* = T(x^*)$. Then, for any starting point $x^0 \in \mathbb{R}^n$, the iteration sequence $x^{i+1} := T(x^i)$ satisfies $x^i \to \widetilde{x}$ for* some *fixed point $\widetilde{x} = T(\widetilde{x})$.*

Theorem 3.1 now follows from Theorem 3.2 and the following lemma.

Lemma 3.1. *$T : \mathbb{R}^n \to \mathbb{R}^n$ is firmly non-expansive if and only if it is $(1/2)$-averaging.*

Proof. Suppose T is $(1/2)$-averaging. Then $T = (I + J)/2$ for some non-expansive J. We compute

$$\begin{aligned}
\|T(x) - T(y)\|^2 &= \frac{1}{4}\left(\|J(x) - J(y)\|^2 + 2\langle J(x) - J(y), x - y\rangle + \|x - y\|^2\right) \\
&\le \frac{1}{2}\left(\langle J(x) - J(y), x - y\rangle + \|x - y\|^2\right) \\
&= \langle T(x) - T(y), x - y\rangle.
\end{aligned}$$

Thus T is firmly non-expansive.

Suppose then that T is firmly non-expansive. If we show that $J := 2T - I$ is non-expansive, it follows that T is $(1/2)$-averaging. This is established by the simple calculations

$$\begin{aligned}
\|J(x) - J(y)\|^2 &= 4\|T(x) - T(y)\|^2 - 4\langle T(x) - T(y), x - y\rangle + \|x - y\|^2 \\
&\le \|x - y\|^2.
\end{aligned}$$

This completes the proof. □

Browder's fixed point theorem is a practical improvement over the classical Banach fixed point theorem.

Theorem 3.3 (Banach fixed point theorem). *Let $T : \mathbb{R}^n \to \mathbb{R}^n$ be a contraction mapping, that is for some $\kappa \in [0, 1)$ holds*

$$\|T(x) - T(y)\| \le \kappa\|x - y\|, \quad (x, y \in \mathbb{R}^n). \tag{3.1}$$

Then T admits a unique fixed point $x^* = T(x^*)$. This can be moreover discovered as the limit of the iteration sequence $x^{i+1} := T(x^i)$ for any starting point x^0.

Note that firm non-expansivity implies non-expansivity, that is (3.1) with $\kappa = 1$, motivating the choice of the term. While non-expansivity is enough to show the existence of a fixed point of T in some cases (T maps a bounded convex set C into itself [13]), it is not enough to show the convergence of the sequence $x^{i+1} := T(x^i)$ to a fixed point. So we need one of the stronger conditions: firm non-expansivity, the averaging property, or contractivity with $\kappa < 1$.

Theorem 3.4. *Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is convex and L-smooth. If the step length $\tau \le L^{-1}$, then, for any starting point $x^0 \in \mathbb{R}^n$, the iterates $\{x^i\}_{i=0}^\infty$ of the gradient descent method (GD) converge to a minimiser $\hat{x}$ of f.*

Proof. By Lemma 2.2, we have

$$L^{-1}\|\nabla f(x) - \nabla f(y)\|^2 \le \langle \nabla f(x) - \nabla f(y), x - y\rangle, \quad (x, y \in \mathbb{R}^n). \tag{3.2}$$

The iteration (GD) may be written in terms of the operator

$$T(x) := x - \tau\nabla f(x).$$

Now

$$\begin{aligned}\|T(x) - T(y)\|^2 &= \langle T(x) - T(y), x - y\rangle - \tau\langle T(x) - T(y), \nabla f(x) - \nabla f(y)\rangle \\ &= \langle T(x) - T(y), x - y\rangle + \tau^2\|\nabla f(x) - \nabla f(y)\|^2 - \tau\langle \nabla f(x) - \nabla f(y), x - y\rangle \\ &\le \langle T(x) - T(y), x - y\rangle.\end{aligned}$$

In the final step we have used (3.2) and $\tau \le L^{-1}$. Thus T is firmly non-expansive. Theorem 3.1 now proves the claim. □

3.3 Variational inclusions and the proximal point method

The gradient descent method is very basic, but often not very good. In particular, subgradient extensions of (GD) can have very slow convergence. Therefore we need alternative methods.

We now allow for general (possibly non-differentiable) convex functions $f : \mathbb{R}^n \to \overline{\mathbb{R}}$, and replace the surrogate objective in (3.1) by another surrogate

$$\min_{x\in\mathbb{R}^n} \bar{f}_{\bar{x}}(x) := f(x) + \frac{1}{2\tau}\|x - \bar{x}\|^2. \tag{3.1}$$

In other words, we remove the linearisation, and try to minimise f directly with a step length condition. Again $\bar{f}_{\bar{x}}(\bar{x}) = f(\bar{x})$, and clearly $f_{\bar{x}} \ge f$. Therefore $\bar{f}_{\bar{x}}$ is a valid surrogate objective for f at $\bar{x}$. This time the optimality conditions for x minimising $\bar{f}_{x^i}$ are

$$0 \in \partial f(x) + \tau^{-1}(x - x^i). \tag{3.2}$$

If $x^i = \hat{x}$ for $\hat{x}$ a minimiser of the original objective f, then (3.2) is solved by $x = \hat{x}$, so again there is a direct correspondence between the solutions of the surrogate objective and the original.

The method based on solving (3.2) resp. (3.1) is known as the proximal point method. The step is the backward step, or the implicit step, since we cannot in general derive an explicit solution $x = x^{i+1}$, and try to go "back to x^i from x^{i+1}". However *often, and especially in context of splitting algorithms,* (3.2) *is easy to solve.* We will get back to this. By contrast, the gradient descent step (GD) is also known as the forward step or the explicit step, because we calculate $\nabla f(x^i)$ already at the current iterate, going "forward" from it.

Re-ordering as

$$x^i \in \tau\partial f(x^{i+1}) + x^{i+1},$$

the iteration resulting from the condition (3.2) may also be written as

$$x^{i+1} := (I + \tau\partial f)^{-1}(x^i), \tag{PP}$$

where the proximal mapping

$$\operatorname{prox}_{\tau\partial f} := (I + \tau\partial f)^{-1}$$

is the inverse of the set-valued map $A := I + \tau\partial f$, defined simply as

$$A^{-1}y := \{x \mid y \in Ax\}.$$

(Thus $y \in Ax \iff x \in A^{-1}y$.) As is evident from the expression

$$\operatorname{prox}_{\tau\partial f}(x) = \arg\min_{x'} f(x') + \frac{1}{2\tau}\|x' - x\|^2,$$

the proximal mapping is, in fact, single-valued.

Remark 3.1. Let $f_\tau := \min_{x'} f(x') + \frac{1}{2\tau}\|x' - x\|^2$. This is known as the Moreau–Yosida regularisation of f—a type of smoothing. In this way, the proximal step also corresponds to solving a sequence of smoothed problems.

Exercise 3.1. *Calculate* $\operatorname{prox}_{\tau\partial f}$ *on* $\mathbb{R}^n$ *for*

(i) $f(x) = \|f - x\|_2^2/2$.

(ii) $f(x) = \delta_{\alpha B}(x)$, *where* B *is the unit ball and* $\alpha > 0$.

(iii) $f(x) = \alpha\|x\|_2$.

Hint: *For (iii) you may find the next Exercise 3.2 useful.*

Exercise 3.2. *Suppose the convex function* $f(x) = \sup_{y\in\mathbb{R}^m} (\langle y, x\rangle - f^*(y))$ *for another proper convex lower semicontinuous function* f^*. *Prove Moreau's identity*

$$y = \operatorname{prox}_{\tau\partial f^*}(y) + \tau \operatorname{prox}_{\tau^{-1}\partial f}(\tau^{-1}y). \tag{3.3}$$

Hint: *Use Theorem 2.1.*

The proximal point method (PP) readily generalises to solving for monotone operators $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ the monotone variational inclusion

$$0 \in A(x). \tag{MVI}$$

The method is simply

$$x^{i+1} := \operatorname{prox}_{\tau A}(x^i) = (I + \tau A)^{-1}(x^i). \tag{MPP}$$

Theorem 3.1. *Let* $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ *be monotone, and suppose there exists a solution* $\hat{x}$ *to* (MVI). *Then for any starting point* $x^0 \in \mathbb{R}^n$, *and any* $\tau > 0$, *the iterates* $\{x^i\}_{i=0}^\infty$ *of the proximal point method* (MPP) *converge to a solution of* (MVI).

Proof. We again use the Browder fixed point theorem, writing the iteration (MPP) in terms of the mapping $T := \operatorname{prox}_{\tau A}$. We have

$$Tx \in x - \tau A(Tx).$$

Thus

$$\|Tx - Ty\|^2 \in \langle Tx - Ty, x - y\rangle - \tau\langle Tx - Ty, A(Tx) - A(Ty)\rangle \le \langle Tx - Ty, x - y\rangle.$$

In the latter step we have used the Cauchy–Schwarz inequality and the monotonicity of A. Thus T is non-expansive, and the rest follows from Theorem 3.1. □

Corollary 3.1. *Suppose* $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ *is convex and proper, and there exists a solution* $\hat{x}$ *to* (P). *Then for any starting point* $x^0 \in \mathbb{R}^n$, *and any* $\tau > 0$, *the iterates* $\{x^i\}_{i=0}^\infty$ *of the proximal point method* (PP) *converge to a solution* $\hat{x}$ *of* (P).

Example 3.1. Let us return to the saddle point problems of Section 1.4 and Section 2.6. That is, let us try to solve

$$\min_x \max_y \; g(x) + \langle Kx, y\rangle - f^*(y), \tag{3.4}$$

for some convex and proper $g : \mathbb{R}^n \to \overline{\mathbb{R}}$, and $f^* : \mathbb{R}^m \to \overline{\mathbb{R}}$, and some matrix $K \in \mathbb{R}^{m\times n}$. As we have seen in Section 2.6, the optimality conditions for this system are

$$-K^T\hat{y} \in \partial g(\hat{x}), \quad \text{and} \quad K\hat{x} \in \partial f^*(\hat{y}).$$

This may be encoded as $0 \in H(x, y)$ in terms of the monotone operator

$$H(x, y) := \begin{pmatrix} \partial g(x) + K^T y \\ \partial f^*(y) - Kx \end{pmatrix}. \tag{3.5}$$

In principle, we may therefore apply (MPP) to solve the saddle point problem (3.4). In practise we however need to work a little bit more, as the step (MPP) can rarely be given an explicit, easily solvable form.

3.4 Forward–backward splitting

Let us consider the minimisation of the composite objective

$$\min_{x\in\mathbb{R}^n} h(x) := g(x) + f(x), \tag{3.1}$$

where g is smooth, but f possibly non-smooth. By Theorem 2.1, we may write the optimality conditions as

$$0 \in \nabla g(x) + \partial f(x).$$

We can rewrite this as

$$\tau^{-1}x - \nabla g(x) \in \tau^{-1}x + \partial f(x),$$

or

$$x = (I + \tau\partial f)^{-1}(x - \tau\nabla g(x)).$$

This gives the iteration

$$x^{i+1} = \operatorname{prox}_{\tau\partial f}(x^i - \tau\nabla g(x^i)). \tag{FB}$$

In other words, we do a gradient/forward step with respect to g, and a proximal/backward step with respect to f. The resulting method is known as forward–backward splitting. Particular instances include the so-called iterative soft-thresholding (IST) algorithm for Lasso.

Exercise 3.3. *When does the method* (FB) *converge to a solution of* (3.1)*?* Hint: *You will need to use the second version of Browder's fixed point theorem.*

Exercise 3.4. *Implement* (FB) *for the Lasso problem of Example 1.2. With your implementation, find the two most relevant physicochemical attributes for the quality of Portuguese* vinho verde*, according to the Wine Quality data set from the UCI machine learning repository at* `http://archive.ics.uci.edu/ml/datasets/Wine+Quality`. Note: *you will need to choose a stopping criterion for the algorithm. For the purposes of this exercise, it is sufficient to take a fixed number of iterations, let's say 1000.*

Exercise(Light) 3.5. *Express forward–backward splitting in terms of a surrogate objective.*

Exercise 3.6. *The total variation denoising problem* (1.1) *may be written in a dual form (cf. Section* 2.6*)*

$$\min_{\phi\in\mathbb{R}^{2n_1n_2}} \frac{1}{2}\|f+\nabla_d^T\phi\|^2, \quad \text{s.t.} \quad \sqrt{\phi_k^2+\phi_{n_1n_2+k}^2} \le \alpha \ \forall k=1,\dots,n_1n_2.$$

Implement (FB) *for this problem. The solution of the original primal problem, the desired image, is* $\hat{x} = f + \nabla_d^T\hat{\phi}$ *for* $\hat{\phi}$ *the solution of the dual problem.*

Let $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ be a general (set-valued) monotone operator, and $B : \mathbb{R}^n \to \mathbb{R}^n$ a single-valued monotone operator. Completely analogously to (3.1) and (FB), we can derive for the inclusion

$$B(x) + A(x) \ni 0 \tag{3.2}$$

the iteration

$$x^{i+1} = \operatorname{prox}_{\tau A}(x^i - \tau B(x^i)).$$

3.5 Douglas–Rachford splitting

Let us try to derive an improved algorithm for (3.2), now both $A, B : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ general (set-valued) monotone operators. This will of course give an algorithm for (3.1) as well, through the choice $A = \partial f$ and $B = \partial g$. Picking $\lambda > 0$, let us set $z \in (I + \lambda B)(x)$. Then $\operatorname{prox}_{\lambda B}(z) = x$. Multiplying (3.2) by λ, and inserting this, we obtain

$$z + \lambda A(\operatorname{prox}_{\lambda B}(z)) \ni \operatorname{prox}_{\lambda B}(z).$$

This reorganises into

$$\operatorname{prox}_{\lambda B}(z) + \lambda A(\operatorname{prox}_{\lambda B}(z)) \in (2\operatorname{prox}_{\lambda B} - I)(z),$$

and further into

$$\operatorname{prox}_{\lambda B}(z) = \operatorname{prox}_{\lambda A}((2\operatorname{prox}_{\lambda B} - I)(z)).$$

This gives the fixed point equation

$$z = \operatorname{prox}_{\lambda A}((2\operatorname{prox}_{\lambda B} - I)(z)) + (I - \operatorname{prox}_{\lambda B})(z).$$

Consequently, we derive the algorithm

$$z^{i+1} := \operatorname{prox}_{\lambda A}((2\operatorname{prox}_{\lambda B} - I)(z^i)) + (I - \operatorname{prox}_{\lambda B})(z^i). \tag{3.1}$$

Note that this is for the transformed variable z, not our variable of interest x. To get a useful result, after the final step i, we therefore need to set

$$x^{i+1} := \operatorname{prox}_{\lambda B}(z^i). \tag{3.2}$$

Performing this at each step, and employing the result in (3.1), we may divide the algorithm into two distinct steps that are called the Douglas–Rachford splitting algorithm

$$x^{i+1} := \operatorname{prox}_{\lambda B}(z^i), \tag{DRS-0}$$
$$z^{i+1} := z^i + \operatorname{prox}_{\lambda A}(2x^{i+1} - z^i) - x^{i+1}. \tag{DRS-1}$$

Theorem 3.1 ([14, 15]). *Let* $A, B : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ *be maximal monotone operators, and suppose there exists a solution* $\hat{x}$ *to* $0 \in A(\hat{x}) + B(\hat{x})$*. Then, for any* $\lambda > 0$*, and any starting point* z^0*, the iterates* $\{x^i\}_{i=1}^\infty$ *of the method* (DRS-0)–(DRS-1) *converge to a point* $\widetilde{x}$ *satisfying* $0 \in A(\widetilde{x}) + B(\widetilde{x})$.

In particular, since the convex subdifferential can be shown to be a maximal monotone operator, we have the following.

Corollary 3.1. *Let $f, g : \mathbb{R}^n \to \overline{\mathbb{R}}$ be convex, and suppose there exists a solution to the composite minimisation problem (3.1). Then, for any $\lambda > 0$, and any starting point z^0, the iterates $\{x^i\}_{i=1}^{\infty}$ of the method*

$$x^{i+1} := \operatorname{prox}_{\lambda \partial g}(z^i), \tag{DRS'-0}$$
$$z^{i+1} := z^i + \operatorname{prox}_{\lambda \partial f}(2x^{i+1} - z^i) - x^{i+1} \tag{DRS'-1}$$

converge to a solution of (3.1).

Exercise 3.7. *Implement the Douglas–Rachford splitting algorithm for dual form of total variation denoising, described in Exercise 3.6. How does the performance compare to basic forward–backward splitting?*
Note: *You will need to invert $I + \nabla_d^T \nabla_d$. For small images, you can simply employ sparse matrices and the slash operator in Matlab, but for bigger images it is beneficial use Fourier transform techniques, familiar from basic numerical analysis courses.*

Remark 3.1. The Douglas–Rachford splitting method (DRS-0)–(DRS-1), when applied to $A = \partial[g^*(-K^T \cdot)]$, and $B = \partial f^*$, is also known as the Alternating Direction Method of Multipliers (ADMM) for the solution of (1.1). In Exercise 3.7 we have, in fact, already implemented the ADMM for the TV denoising problem (1.1). Since a solution of (1.1) corresponds the condition $0 \in H(x, y)$ for H as in (3.5), we have therefore finally, through splitting, found a practical variant for solving the latter problem.

3.6 The Chambolle–Pock method

Let us study a very effective primal–dual method for the saddle point problem (3.1). For parameters $\tau, \sigma > 0$, the primal variable x, and the dual variable y, we define the iteration

$$x^{i+1} := (I + \tau\partial g)^{-1}(x^i - \tau K^T y^i), \tag{CP-0}$$
$$\bar{x}^{i+1} := 2x^{i+1} - x^i, \tag{CP-1}$$
$$y^{i+1} := (I + \sigma\partial f^*)^{-1}(y^i + \sigma K \bar{x}^{i+1}). \tag{CP-2}$$

The step (CP-0) is simply a proximal step for x in (3.4), keeping $y = y^i$ fixed. The step (CP-2) is likewise a proximal step for y in (3.4), keeping x fixed, not to x^i or x^{i+1} but to the inertial variable $\bar{x}^{i+1}$ defined in (CP-1). This may be visualised as a "heavy ball" version of x^{i+1} that has enough inertia to not get stuck in small bumps in the landscape.

With the general notation

$$u = (x, y),$$

the steps (CP-0)–(CP-2) may also be written in the preconditioned proximal point form

$$H(u^{i+1}) + M(u^{i+1} - u^i) \ni 0, \tag{3.1}$$

for the monotone operator H as in (3.5), and the preconditioning matrix

$$M := \begin{pmatrix} I/\tau & -K^T \\ -K & I/\sigma \end{pmatrix}.$$

Through the replacement of I by M in the basic proximal point iteration $u^{i+1} := (I + H)^{-1}(u^i)$, we thus have in (CP-0)–(CP-1) a proximal point method for which the steps can often be solved explicitly.

Theorem 3.1. *Let $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ and $g : \mathbb{R}^m \to \overline{\mathbb{R}}$ be convex, proper, and lower semicontinuous, and $K \in \mathbb{R}^{m \times n}$. Choose $\tau, \sigma > 0$ such that $\tau\sigma\|K\|^2 < 1$. Then the iterates of (CP-0)–(CP-2) convergece for any starting point $u^0 = (x^0, y^0)$ to a saddle point $u^* = (x^*, y^*)$ of (3.4).*

Proof. A saddle point $\widehat{u}$ satisfies $0 \in H(\widehat{u})$. Therefore

$$\langle H(u^{i+1}) - H(\widehat{u}), u^{i+1} - \widehat{u}\rangle \geq 0.$$

Thus (3.1) gives

$$\langle M(u^{i+1} - u^i), u^{i+1} - \widehat{u}\rangle \leq 0. \tag{3.2}$$

With the notation

$$\langle x, y\rangle_M := \langle Mx, y\rangle, \quad \text{and} \quad \|x\|_M := \sqrt{\langle x, x\rangle_M}, \tag{3.3}$$

we calculate

$$\begin{aligned}\langle u^{i+1} - u^i, u^{i+1} - \widehat{u}\rangle_M &= \|u^{i+1} - u^i\|_M^2 + \langle u^i - \widehat{u}, u^{i+1} - u^i\rangle_M \\ &= \|u^{i+1} - u^i\|_M^2 - \|u^i - \widehat{u}\|_M^2 + \langle u^i - \widehat{u}, u^{i+1} - \widehat{u}\rangle_M \\ &= \|u^{i+1} - u^i\|_M^2 - \|u^i - \widehat{u}\|_M^2 + \|u^{i+1} - \widehat{u}\|_M^2 \\ &\quad + \langle u^i - u^{i+1}, u^{i+1} - \widehat{u}\rangle_M.\end{aligned}$$

That is

$$\langle u^{i+1} - u^i, u^{i+1} - \widehat{u}\rangle_M = \frac{1}{2}\|u^{i+1} - u^i\|_M^2 - \frac{1}{2}\|u^i - \widehat{u}\|_M^2 + \frac{1}{2}\|u^{i+1} - \widehat{u}\|_M^2.$$

Now (3.2) shows that

$$\frac{1}{2}\|u^{i+1} - \widehat{u}\|_M^2 + \frac{1}{2}\|u^{i+1} - u^i\|_M^2 \leq \frac{1}{2}\|u^i - \widehat{u}\|_M^2. \tag{3.4}$$

Summing (3.4) over $i = 0, \ldots, N-1$ shows that

$$\frac{1}{2}\|u^N - \widehat{u}\|_M^2 + \sum_{i=0}^{N-1} \frac{1}{2}\|u^{i+1} - u^i\|_M^2 \leq \frac{1}{2}\|u^0 - \widehat{u}\|_M^2. \tag{3.5}$$

Now, the condition $\tau\sigma\|K\|^2 < 1$ ensures that $\|u\|_M^2 \geq \theta\|u\|^2$ for some $\theta > 0$. Therefore (3.5) shows that $\|u^{i+1} - u^i\| \to 0$, and that $\{u^i\}_{i\in\mathbb{N}}$ is bounded. It follows that the whole sequence $\{u^i\}_{i\in\mathbb{N}}$ converges to some u^*. Using this and $\|u^{i+1} - u^i\| \to 0$ in (3.1) shows that $0 \in H(u^*)$, so we have found a saddle point. (It might be that $u^* \neq \widehat{u}$.) □

Exercise 3.8. *Implement the Chambolle–Pock method for total variation denoising, described in Exercise 3.6. What is the effect of the choice of τ and σ? How does the performance compare to forward–backward splitting and the ADMM of Exercise 3.7?*

A few pointers for the aficionados

We won't go deeper into optimisation methods in this course, concentrating next on sensitivity analysis. Various further splitting algorithms exist in the literature, many of which are closely linked to each other. The Chambolle–Pock method and forward–backward splitting can also be accelerated, to obtain fast convergence rates on strongly convex problems [16–18]. We refer in particular to [19–21] as starting points for further studies.

4 Set-valued maps and sensitivity analysis

In practise, interesting optimisation problems incorporate a model, parameters, and data. We are therefore solving problems of the type

$$\min_{x \in \mathbb{R}^n} f(x; p), \tag{4.1}$$

for some parameter $p \in \mathbb{R}^m$, which could also be our data. To be able to rely on the solution $\hat{x}$ under noise and other corruptions to the data, or under very approximate parameter choices, we need to know how much $\hat{x}$ can vary as the data or parameters varies. This is the topic of the rest of the course.

Our study is centred around the set-valued solution map

$$S : \mathbb{R}^m \rightrightarrows \mathbb{R}^n, \quad S(p) := \{x \in \mathbb{R}^n \mid 0 \in \partial_x f(x; p)\},$$

where we assume f to be convex, so that the minima of (4.1) for fixed $p \in \mathbb{R}^m$ are encoded in terms of the convex subdifferential $\partial_x f(x; p) := \partial[f(\,\cdot\,; p)](x)$. If we have a solution $\hat{x}$ for some parameter $\hat{p}$—that is, $\hat{x} \in S(\hat{p})$—we would then like to obtain Lipschitz-style estimates

$$\inf_{x \in S(p)} \|x - \hat{x}\| \le \ell \|p - \hat{p}\| \tag{4.2}$$

for the parameter p close to $\hat{p}$. This says that there exists solutions x for p close to $\hat{x}$ in a Lipschitz sense. Of course, even mere continuity of S would be encouraging.

4.1 Basic properties of set-valued maps

We have already learned the concept of outer semicontinuity of the subdifferential in the context of Proposition 2.1. Set-valued maps can also be inner semicontinuous.

Definition 4.1. Let $\{A^i\}_{i=1}^\infty$ be a sequence of subsets of $\mathbb{R}^n$. We define the

1. outer limit as the set

$$\limsup_{i \to \infty} A^i := \{x \in \mathbb{R}^n \mid \forall j \in \mathbb{N}\, \exists i_{j+1} > i_j, x^{i_j} \in A_{i_j} : x^{i_j} \to x\}.$$

2. inner limit as the set

$$\liminf_{i \to \infty} A^i := \{x \in \mathbb{R}^n \mid \forall i \in \mathbb{N}\, \exists x^i \in A^i : x^i \to x\}.$$

The vast difference between inner and outer limits is illustrated by the extreme example Figure 4.1. We also extend these definitions to functions, i.e., an uncountable index set.

Definition 4.2. Let $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ be a set-valued map. Then we set $\limsup_{x' \to x} A(x') := \bigcup_{x^i \to x} \limsup_{i \to \infty} A(x^i)$, and $\liminf_{x' \to x} A(x') := \bigcap_{x^i \to x} \liminf_{i \to \infty} A(x^i)$

Exercise 4.1. *Show that both $\limsup_{x' \to x} A(x')$ and $\liminf_{x' \to x} A(x')$ are always (possibly empty) closed sets.*

Definition 4.3. Let $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ be a set-valued map. We say that

Figure 4.1: Illustration of a sequence of sets $\{A^i\}_{i\in\mathbb{N}}$, $A^i \subset [0,1]$, with vastly different inner and outer limits. In this case, the outer limit $\limsup_{i\to\infty} A^i = [0,1]$, while the inner limit $\liminf_{i\to\infty} A^i = \emptyset$. The outer limit consists of all points approximable through some subsequence of the sets A^i, while the inner limit has to be approximable via every subsequence. In this case, we can for any $x \in [0,1]$, find a subsequence of the somewhat "random" $\{A^i\}$ that will not contain points approaching x.

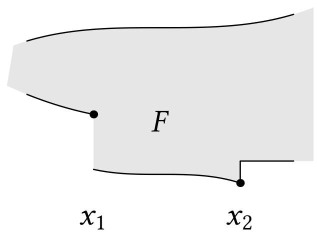

Figure 4.2: Illustration of outer and inner semicontinuity. The black line indicates the bounds on the boundary of Graph F that belong to the graph. The set-valued map F is not outer semicontinuous at x_1, because $F(x_1)$ does not include all limits from the right. It is outer semicontinuous at the "discontinuous" point x_2, as $F(x_2)$ includes all limits from both sides. The map F is not inner semicontinuous at x_2, because at this point, $F(x)$ cannot be approximated from both sides. It is inner semicontinuous at every other point x, including x_1, as at this points $F(x)$ can be approximated from both sides.

(i) A is outer-semicontinuous at x if $\limsup_{x' \to x} A(x') \subset A(x)$.

(ii) A is inner-semicontinuous at x if $\liminf_{x' \to x} A(x') \supset A(x)$.

(iii) continuous (at x), if it is both inner- and outer-semicontinuous (at x).

We illustrate these concepts in Figure 4.2.

In fact, outer semicontinuity can be reduced to a simple property on $\operatorname{Graph} A$, as the next easy exercise shows.

Exercise(Light) 4.2. *Show that $A : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is outer semicontinuous if and only if $\operatorname{Graph} A$ is closed.*

The outer semicontinuity of the solution map is crucial for approximation: If $p^i \to p$, and we have solution $x^i \in S(p^i)$, it would be highly desirable that any accumulation point x of $\{x^i\}$ satisfies $x \in S(p)$. When can this hold?

Proposition 4.1. *Let $f : \mathbb{R}^n \times \mathbb{R}^m \to \overline{\mathbb{R}}$, and suppose $\partial_x f$ is outer semicontinuous as a function of (x, p). Then S is also outer semicontinuous.*

Proof. Let us set $F(x,p) = \partial_x f(x;p)$, and $\widetilde{S}(p,z) := \{x \in \mathbb{R}^n \mid z \in F(x,p)\}$. Then $S(p) = \widetilde{S}(p,0)$. Clearly the outer semicontinuity of S follows from that of $\widetilde{S}$. But, now

$$\operatorname{Graph} F = \{(x,p,z) \in \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R} \mid z \in F(x,p)\},$$

while

$$\operatorname{Graph} \widetilde{S} = \{(p,z,x) \in \mathbb{R}^n \times \mathbb{R} \times \mathbb{R}^m \mid z \in F(x,p)\}.$$

In other words $\operatorname{Graph} \widetilde{S} = \mathcal{P} \operatorname{Graph} F$ for the permutation $\mathcal{P}(x,p,z) := (p,z,x)$. By Exercise 4.2, $\operatorname{Graph} F$ is closed. Clearly then also $\operatorname{Graph} \widetilde{S}$ is closed, which again by Exercise 4.2 is equivalent to $\widetilde{S}$ being outer semicontinuous. □

Example 4.1. Let $f(x;p) := g(x-p) + h(x)$, where $f : \mathbb{R}^n \to \overline{\mathbb{R}}$ is convex, proper, and lower semicontinuous, and $g : \mathbb{R}^n \to \mathbb{R}$ is convex and continuously differentiable. Then

$$\partial_x f(x,p) = \nabla g(x-p) + \partial h(x)$$

is outer semicontinuous as a function of (x,p). In particular, the solutions to Lasso and total variation denoising are outer semicontinuous for varying data. By analogous calculations, we can verify the same property for the SVM.

4.2 The Aubin property

So far, we have defined basic properties of continuity of set-valued maps. But how about more strict forms of continuity, such as Lipschitz continuity? One successful approach of extending the definition of Lipschitz continuity to set-valued maps is given by the next definition; for further approaches we refer to [4]. Here and throughout, we denote the closed ball at $x \in \mathbb{R}^n$, for suitable n, of radius $\rho > 0$ by

$$\mathbb{B}(x,\rho) := \{x' \in \mathbb{R}^n \mid \|x' - x\| \le \rho\}.$$

Definition 4.1. The set-valued mapping $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ has the Aubin property at $\bar{x}$ for $\bar{y} \in F(\bar{x})$ if $\operatorname{Graph} F$ is locally closed (see below) at $(\bar{x}, \bar{y})$ and for some $\ell > 0$ there exist $\rho, \delta > 0$ such that

$$F(x') \cap \mathbb{B}(\bar{y}, \rho) \subset F(x) + \mathbb{B}(0, \ell\|x' - x\|), \quad (x', x \in \mathbb{B}(\bar{x}, \delta)). \tag{4.1}$$

The infimum of all possible factors ℓ over all $\rho, \delta > 0$ is denoted $\operatorname{lip} F(\bar{x}|\bar{y})$, and called the graphical modulus of F at $\bar{x}$ for $\bar{y}$.

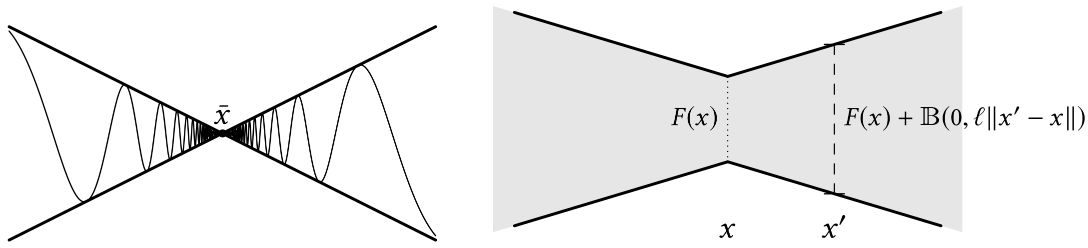

(a) An oscillating single-valued function **(b)** The graph of $x \mapsto F(x) + \mathbb{B}(0, \ell\|x' - x\|)$

Figure 4.3: The single-valued example in (a) illustrates why Lipschitz properties $\|f(x') - f(x)\| \leq \ell\|x'-x\|$ at a point $\bar{x}$ need to be based on two points $x, x' \in \mathbb{B}(\bar{x}, \delta)$. If we fixed $x = \bar{x}$, this highly oscillatory example would be Lipschitz at $\bar{x}$. Indeed, the graph lies between the two thick lines, demonstrating the bounds $\|f(x') - f(\bar{x})\| \leq \ell\|x' - \bar{x}\|$. If we do not fix $x = \bar{x}$, the function clearly will not be Lipschitz, and will also not satisfy the Aubin property. In (b) we illustrate the "fat cone" structure $\operatorname{Graph}(x' \mapsto F(x) + \mathbb{B}(0, \ell\|x' - x\|)$ appearing on the right-hand-side in (4.1), and varying with the second base point x around $\bar{x}$. This is to be contrasted with the leaner cone $\operatorname{Graph}(x' \mapsto f(\bar{x}) + \mathbb{B}(0, \ell\|x' - \bar{x}\|))$ bounding the function in (a).

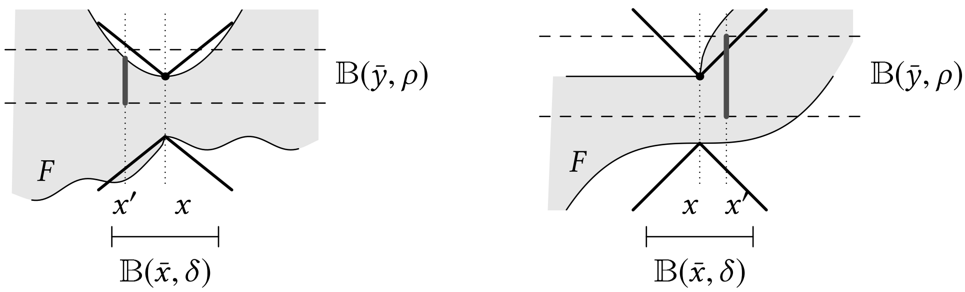

(a) Property is satisfied **(b)** Property is not satisfied

Figure 4.4: Illustration of satisfaction and dissatisfaction of the Aubin property for $x = \bar{x}$. The dashed lines indicate $\mathbb{B}(\bar{y}, \rho)$, and the dot marks $(\bar{x}, \bar{y})$, while the dark gray thick line indicates $F(x') \cap \mathbb{B}(\bar{y}, \rho)$. It should remain within the bounds the bounds of the black thick lines indicating $F(x) + \mathbb{B}(0, \ell\|x' - x\|)$. The violation of the bounds at the bottom in (a) does not matter, because we are only interested in the area between the dashed lines.

We required the following concept on $A = \operatorname{Graph} F$.

Definition 4.2. A set A is locally closed at x if there exists $\rho > 0$ such that $\mathbb{B}(x, \rho) \cap A$ is closed.

Naturally, in the case of $A = \operatorname{Graph} F$, we can replace the ball in the definition by the product neighbourhood $\mathbb{B}(\bar{x}, \delta) \times \mathbb{B}(\bar{y}, \rho)$.

The single-valued example in Figure 4.3a illustrates why we need two points x' and x in (4.1), instead of fixing one of them to equal $\bar{x}$. Figure 4.3b illustrates the "fat cone" structure on the right hand side of 4.1. It should locally at each base point x around $\bar{x}$ bound F for the Aubin property to be satisfied. The satisfaction and dissatisfaction of the Aubin property we illustrate in Figure 4.4.

In fact, we do not need to restrict x' into a tiny neighbourhood of $\bar{x}$ in the Aubin property.

Lemma 4.1. *The condition* (4.1) *is equivalent to the existence of (possibly different)* $\rho, \delta > 0$ *satisfying*

$$F(x') \cap \mathbb{B}(\bar{y}, \rho) \subset F(x) + \mathbb{B}(0, \ell\|x' - x\|), \quad (x \in \mathbb{B}(\bar{x}, \delta);\ x' \in \mathbb{R}^n). \tag{4.2}$$

Proof. Clearly (4.2) implies (4.1). To show the implication in the other direction, we start by applying

(4.1) to $x' = \bar{x}$. For the moment, we also fix $x \in \mathbb{B}(\bar{x}, \delta')$. This gives for any $\delta' \in (0, \delta]$ the estimate

$$\bar{y} \in F(x) + \mathbb{B}(0, \ell\|\bar{x} - x\|).$$

Thus

$$\bar{y} \in F(x) + \mathbb{B}(0, \ell\delta').$$

In particular, for any $\epsilon' > 0$, we have

$$\mathbb{B}(\bar{y}, \epsilon') \subset F(x) + \mathbb{B}(0, \ell\delta' + \epsilon'). \tag{4.3}$$

For $x' \in \mathbb{B}(\bar{x}, \delta)$, (4.2) is clear, so we may concentrate on $x' \in \mathbb{R}^n \setminus \mathbb{B}(\bar{x}, \delta)$. Then

$$\|x' - x\| \geq \|x' - \bar{x}\| - \|\bar{x} - x\| \geq \delta - \delta'.$$

If we pick $\epsilon', \delta' > 0$ such that $\ell\delta' + \epsilon' \leq \ell(\delta - \delta')$, it follows

$$\ell\delta' + \epsilon' \leq \ell\|x' - x\|.$$

Thus (4.3) gives

$$\begin{aligned} F(x') \cap \mathbb{B}(\bar{y}, \epsilon') &\subset \mathbb{B}(\bar{y}, \epsilon') \\ &\subset F(x) + \mathbb{B}(0, \ell\delta' + \epsilon') \\ &\subset F(x) + \ell\mathbb{B}(0, \|x' - x\|), \end{aligned}$$

as illustrated in Figure 4.5. □

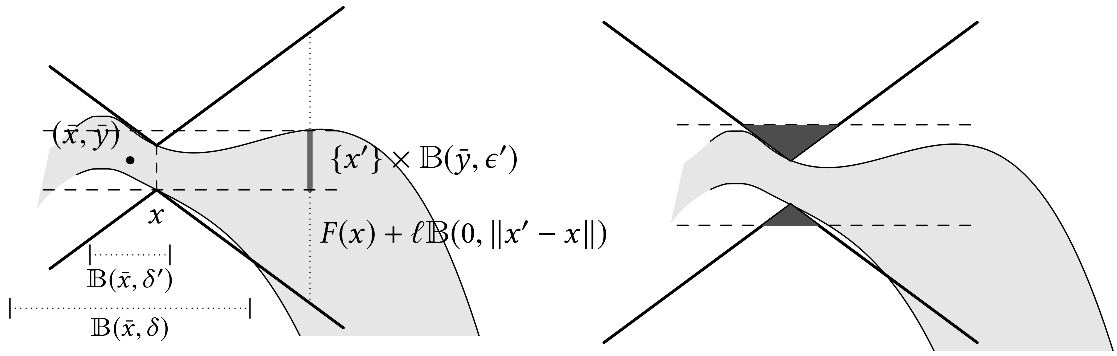

(a) Illustration of the technique **(b)** The critical areas

Figure 4.5: (a) Illustration of the technique in Lemma 4.1. For x' outside the ball $\mathbb{B}(\bar{x}, \delta)$, the set $\mathbb{B}(\bar{y}, \epsilon')$ indicated by the thick dark grey line, is completely contained in the fat-cone structure $F(x) + \ell\mathbb{B}(0, \|x' - x\|)$ of Figure 4.3b, indicated by the thick black and dotted lines. Closer to x, within $\mathbb{B}(\bar{x}, \delta)$, this is not the case, although $F(x') \cap \mathbb{B}(\bar{y}, \epsilon')$ itself is still contained in the structure. (b) highlights in dark grey the critical areas for the Aubin property.

With the help of the above equivalent formulation, to facilitate development of the Mordukhovich criterion and stability analysis later on, we still introduce a few further equivalent formulations of the Aubin property.

Proposition 4.1. *Suppose $\bar{y} \in F(\bar{x})$. We may equivalently state (4.1) as*

(i) The property

$$d(y, F(x)) \leq \ell d(F^{-1}(y), x), \quad (x \in \mathbb{B}(\bar{x}, \delta),\ y \in \mathbb{B}(\bar{y}, \rho)), \tag{4.4}$$

where

$$d(A, y) := d(y, A) := \inf_{y' \in A} \|y' - y\|.$$

(ii) For some $\rho', \delta', \ell' > 0$, the Lipschitz continuity with uniform factor $\ell' > 0$ over $y \in \mathbb{B}(\bar{y}, \rho')$, of the maps $x \mapsto d(y, F(x))$, where $x \in \mathbb{B}(\bar{x}, \delta')$.

Proof. We begin by showing (i). By Lemma 4.1, it suffices to prove the equivalence of (4.4) to (4.2). The latter may be expanded as

$$\{y'\} \cap \mathbb{B}(\bar{y}, \rho) \subset F(x) + \mathbb{B}(0, \ell\|x' - x\|), \quad (y' \in F(x'); x \in \mathbb{B}(\bar{x}, \delta);\ x' \in \mathbb{R}^n).$$

In other words

$$\inf_{y \in F(x)} \|y' - y\| \le \ell\|x' - x\|, \quad (y' \in F(x') \cap \mathbb{B}(\bar{y}, \rho);\ x \in \mathbb{B}(\bar{x}, \delta);\ x' \in \mathbb{R}^n).$$

This may further be rewritten as

$$\inf_{y \in F(x)} \|y' - y\| \le \inf_{x' \in F^{-1}(y')} \ell\|x' - x\|, \quad (x \in \mathbb{B}(\bar{x}, \delta); y' \cap \mathbb{B}(\bar{y}, \rho)).$$

Thus (4.4) is equivalent to (4.1).

Regarding (ii), let us begin by defining the two distance functions

$$\mathbb{d}_{\rho'}(A, B) := \max_{\|y\| \le \rho'} |d(A, y) - d(B, y)|.$$

and

$$\hat{\mathbb{d}}_\rho(A, B) := \inf\{t \ge 0 \mid A \cap \mathbb{B}(0, \rho) \subset B + \mathbb{B}(0, t),\ B \cap \mathbb{B}(0, \rho) \subset A + \mathbb{B}(0, t)\}.$$

Then, it is easily observed that, (ii) is equivalent to

$$\mathbb{d}_{\rho'}(F(x), F(x')) \le \ell'\|x - x'\|, \quad (x, x' \in \mathbb{B}(\bar{x}, \delta')),$$

while (4.1) is equivalent to

$$\hat{\mathbb{d}}_\rho(F(x), F(x')) \le \ell\|x - x'\|, \quad (x, x' \in \mathbb{B}(\bar{x}, \delta)).$$

Observe also that (ii) guarantees for some constant $C > 0$ the bound

$$\begin{aligned} d(0, F(x)) &\le \|\bar{y}\| + d(\bar{y}, F(x)) \\ &\le \|\bar{y}\| + |d(\bar{y}, F(x)) - d(\bar{y}, F(\bar{x}))| + d(\bar{y}, \bar{F}(\bar{x})) \le C, \quad (x \in \mathbb{B}(\bar{x}, \delta')). \end{aligned}$$

Therefore, we can find a constant ρ' satisfying

$$\rho' \ge 2\rho + \sup\{d(0, F(x)) \mid x \in \mathbb{B}(\bar{x}, \delta')\}.$$

To prove the equivalence of (ii) to (4.1), it therefore suffices to show that

$$\hat{\mathbb{d}}_\rho(A, B) \le \mathbb{d}_\rho(A, B) \le \hat{\mathbb{d}}_{\rho'}(A, B)$$

whenever $\rho' \ge 2\rho + \max\{d(0, A), d(0, B)\}$. This follows from the next little lemma. □

Lemma 4.2. *Let $A, B \subset \mathbb{R}^n$ be closed. Then we have the following implications.*

(i) $d(\cdot, A) \le d(\cdot, B) + \epsilon$ on $\mathbb{B}(0, \rho) \implies B \cap \mathbb{B}(0, \rho) \subset A + \mathbb{B}(0, \epsilon)$.

(ii) $B \cap \mathbb{B}(0, \rho') \subset A + \mathbb{B}(0, \epsilon)$ for some $\rho' \ge 2\rho + d(0, B) \implies d(\cdot, A) \le d(\cdot, B) + \epsilon$ on $\mathbb{B}(0, \rho)$.

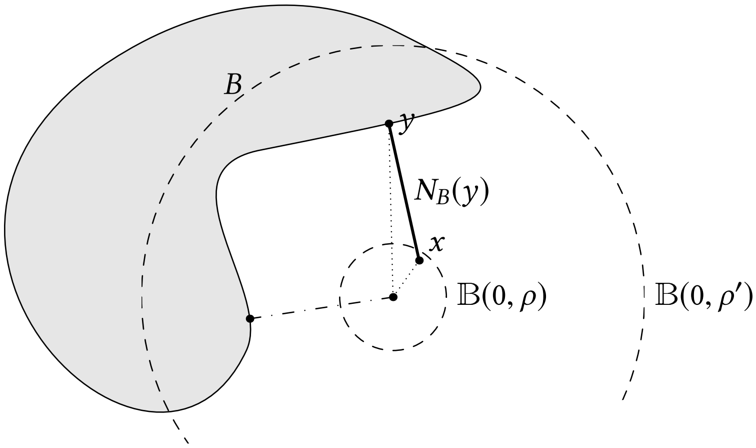

Figure 4.6: Illustration for the proof of (4.5) in Lemma 4.2. The thick solid line depicts y as a projection of x. The dotted lines illustrate the argument in (4.6), while the dash-dotted line depicts $d(0, B)$.

Proof. Regarding (i), take $x \in B \cap \mathbb{B}(0,\rho)$. Then $d(x, B) = 0$, so $d(x, A) \le \epsilon$. That is, $x \in A + \mathbb{B}(0,\epsilon)$.

The proof of (ii) is somewhat longer. To start with, let x be arbitrary. Then

$$d(x, B \cap \mathbb{B}(0,\rho')) \ge d(x, A + \mathbb{B}(0,\epsilon)) \ge d(x, A) - \epsilon.$$

If

$$d(x, B) \ge d(x, B \cap \mathbb{B}(0,\rho')) \quad \text{when} \quad x \in \mathbb{B}(0,\rho), \tag{4.5}$$

we obtain (ii). To verify (4.5), let $x \in \mathbb{B}(0,\rho)$, and let $y \in B$ satisfy $\|x - y\| = d(x, B)$ (such a point exists, because B is closed). Now

$$\|y\| \le \|x\| + \|y - x\| \le \rho + d(x, B) \le 2\rho + d(0, B) \le \rho'. \tag{4.6}$$

Therefore $y \in \mathbb{B}(0,\rho')$. This is illustrated in Figure 4.6. We have thus discovered a point y such that

$$d(x, B) = \|y - x\| = d(x, B \cap \mathbb{B}(0,\rho')).$$

This immediately proves (4.5) and consequently (ii). □

Let us try to make sense of (4.2). If F has the Aubin property, then by Proposition 4.1, picking $y = \bar{y}$, we have

$$d(\bar{y}, F(x)) \le \ell d(\bar{x}, x), \quad (x \in \mathbb{B}(\bar{x},\delta)).$$

or in other words

$$\inf_{y \in F(x)} \|\bar{y} - y\| \le \ell\|x - \bar{x}\|, \quad (x \in \mathbb{B}(\bar{x},\delta)).$$

This is what we need: *the Aubin property of the solution map S at $\hat{p}$ for $\hat{x}$ implies the stability of solutions $x = S(p)$ under perturbations p to the parameter $\hat{p}$.*

Unfortunately, in practise the direct calculation of the graphical modulus $\ell = \operatorname{lip} S(\hat{p}|\hat{x})$ is infeasible. The rest of this course concentrates on calculating the factor in special cases.

4.3 Tilt stability and metric regularity

Let us study

$$f(x; p) = g(x) - \langle p, x\rangle,$$

for some $g : \mathbb{R}^n \to \mathbb{R}$, and a tilt parameter p. Stability of solutions $x \in S(p)$ with respect to p is then called tilt stability. Now

$$S(p) = \{x \in \mathbb{R}^n \mid p \in \partial g(x)\} = (\partial g)^{-1}(p).$$

The Aubin property 4.4 of S may now be written

$$d(x, (\partial g)^{-1}(p)) \leq \ell d(\partial g(x), p), \quad (p \in \mathbb{B}(\hat{p}, \delta);\ x \in \mathbb{B}(\hat{x}, \rho)).$$

In other words, S is stable, if $(\partial g)^{-1}$ has the Aubin property. This concept has a special name.

Definition 4.1. A set-valued map $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ is metrically regular at $\bar{x}$ for $\bar{y} \in F(\bar{x})$ if F^{-1} has the Aubin property at $\bar{y}$ for $\bar{x}$, that is, when there exists $\ell > 0$ such that for some $\rho, \delta > 0$ holds

$$d(x, F^{-1}(y)) \leq \ell d(F(x), y), \quad (x \in \mathbb{B}(\bar{x}, \delta);\ y \in \mathbb{B}(\bar{y}, \rho)).$$

The infimum of all possible factors ℓ over all $\rho, \delta > 0$ is denoted $\operatorname{lip} F^{-1}(\bar{y}|\bar{x})$, and called the modulus of metric regularity of F at $\bar{x}$ for $\bar{y}$.

Exercise 4.3. *Study the tilt-stability of $g(x) = \|a - x\|^2/2$. Observe how the tilt parameter p corresponds to change in the data a. How about $g(x) = \|a - Tx\|^2/2$ for some matrix $T \in \mathbb{R}^{m \times n}$?*

Exercise 4.4. *Study the tilt-stability of $g(x) = \|a - x\|$.*

Exercise 4.5. *Is the Lasso (Example 1.2) with $a_i = e_i$ for the unit coordinate vectors e_i stable with respect to the data? (We will return to the general case after introducing the Mordukhovich criterion.)*

5 Graphical derivatives and coderivatives

Recall how the subdifferential of a convex function f can be defined in terms of normal or tangent cones to the epigraph. This idea forms the basis of differentiating general set-valued maps $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$, where instead of taking the tangent or normal cone at $(x, f(x))$ to $\operatorname{epi} f$, we can do this at any point (x, y) for $y \in F(x)$. Since we are generally not in the nice convex setting—even for a convex function f, $\operatorname{Graph} \partial f$ is not convex unless f is linear—there are however some complications, which result in various definitions.

5.1 Tangent and normal cones

Our first type of tangent cone is based on roughly the same limiting process on difference quotients, as the definition of directional derivatives.

Definition 5.1. A vector $z \in \mathbb{R}^k$ is tangent to a set $A \subset \mathbb{R}^k$ at $x \in \mathbb{R}^k$, if

$$z = \lim_{i\to\infty} \frac{x^i - x}{\tau^i}, \quad \text{for some} \quad A \ni x^i \to x,\ \tau^i \searrow 0.$$

The set $T_A(x)$ of all such z is the tangent cone or contingent cone.

In other words

$$T_A(x) = \limsup_{\tau \searrow 0} \frac{A - x}{\tau}.$$

Thus the tangent cone consists of all possible limits of the "blown up" sets $A_\tau := (A - x)/\tau$ for $\tau \searrow 0$.

The tangent cone is closely related to the Fréchet normal cone, which is based on the same limiting process as the Fréchet subdifferential in (2.1).

Definition 5.2. A vector $z \in \mathbb{R}^k$ is a Fréchet normal to a set $A \subset \mathbb{R}^k$ at $x \in \mathbb{R}^k$ if,

$$\limsup_{A \ni x' \to x} \frac{\langle z, x' - x\rangle}{\|x' - x\|} \leq 0.$$

The set $\widehat{N}_A(x)$ of all such z is the Fréchet normal cone.

One difficulty with the Fréchet normal cone is that it is not outer semicontinuous. By taking all outer limits, we obtain a more "regular" normal cone.

Definition 5.3. The normal cone of a set $A \subset \mathbb{R}^k$ at $x \in \mathbb{R}^k$ is the set

$$N_A(x) := \limsup_{A \ni x' \to x} \widehat{N}_A(x').$$

Remark 5.1. Despite N_A obtained by the regularisation of $\widehat{N}_A$, the *latter* is sometimes in the literature called the regular normal cone. We stick to the convention of calling it the Fréchet normal cone, and N_A simply the normal cone.

This normal cone is closely related to a another type of tangent cone, smaller than $T_A(x)$.

Definition 5.4. The regular (or Clarke) tangent cone of a set $A \subset \mathbb{R}^k$ at $x \in \mathbb{R}^k$ is the set

$$\widehat{T}_A(x) := \liminf_{\substack{A \ni x' \to x,\\ \tau \searrow 0}} \frac{A - x'}{\tau}.$$

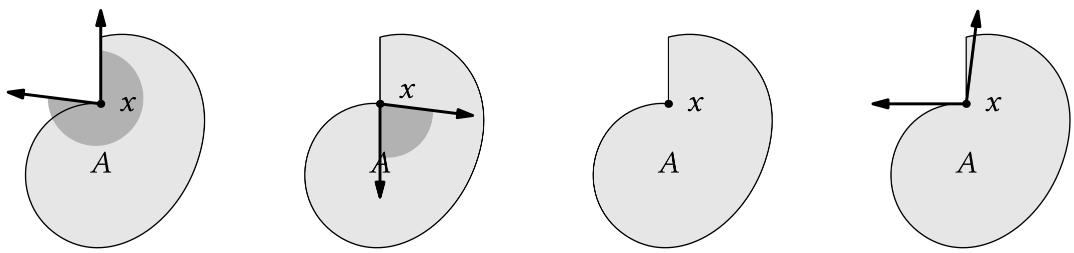

(a) $T_A(x)$ **(b)** $\widehat{T}_A(x)$ **(c)** $\widehat{N}_A(x) = \{0\}$ **(d)** $N_A(x)$

Figure 5.1: Illustration of the different normal and tangent cones at a non-regular point of a set A. The dot indicates the base point x. The thick arrows and dark grey areas the directions included in the cones.

We will later see that for closed A, in fact, $\widehat{T}_A(x) = \liminf_{A\ni x'\to x} T_A(x')$.

Exercise(Light) 5.1. *Compute all the different tangent and normal cones at all points $x \in A \subset \mathbb{R}^2$ for $A = [0,1]^2$, $A = \mathbb{B}(0,1)$, and $A = [0,1]^2 \setminus [1/2,1]^2$. How about the tangent and normal cones to ∂A for each of these sets?*

We illustrate the differences of the different tangent and normal cones in Figure 5.1. The next proposition lists some of their most basic properties.

Proposition 5.1. *Let $A \subset \mathbb{R}^k$. Then at every $x \in \mathbb{R}^k$, each of the sets $T_A(x)$, $\widehat{T}_A(x)$, $N_A(x)$, and $\widehat{N}_A(x)$ is a closed cone. The sets $\widehat{N}_A(x)$ and $\widehat{T}_A(x)$ are moreover convex.*

Proof. By the definition of $\widehat{N}_A(x)$, if z belongs to this set, then any multiple of z also belongs. Thus $\widehat{N}_A(x)$ is a cone. Let then $z^i \in \widehat{N}_A(x)$ converge to some $z \in \mathbb{R}^k$. Let $A \ni x^j \to x$. Then, for any $i \in \mathbb{N}$, we have

$$\frac{\langle z, x^j - x\rangle}{\|x^j - x\|} \le \frac{\langle z^i, x^j - x\rangle}{\|x^j - x\|} + \|z^i - z\|.$$

Thus

$$\limsup_{j\to\infty} \frac{\langle z, x^j - x\rangle}{\|x^j - x\|} \le \|z^i - z\|.$$

Since $i \in \mathbb{N}$ was arbitrary, and $z^i \to z$, we see that $z \in \widehat{N}_A(x)$, and may conclude that $\widehat{N}_A(x)$ is closed. Let then $\bar z, \widetilde z \in \widehat{N}_A(x)$. To show convexity, since $\widehat{N}_A(x)$ is a cone, it suffices to show that $z := \bar z + \widetilde z \in \widehat{N}_A(x)$. We have

$$\frac{\langle z, x^j - x\rangle}{\|x^j - x\|} = \frac{\langle \bar z, x^j - x\rangle}{\|x^j - x\|} + \frac{\langle \widetilde z, x^j - x\rangle}{\|x^j - x\|}.$$

This immediately allows us to deduce $z \in \widehat{N}_A(x)$, and conclude the convexity of this normal cone.

The set $N_A(x)$ is a closed cone as the outer limit of the (closed) cones $\widehat{N}_A(x^i)$ as $x^i \to x$.

Similarly, we see that $T_A(x)$ is closed as the outer limit of the sets $A_\tau := (A - x)/\tau$, as $\tau \searrow 0$; see Exercise 4.1. To see that it is a cone, suppose $z \in T_A(x)$. Then there exists a sequence $z^i \in A_{\tau^i}$, with $\tau^i \searrow 0$ and $z^i \to z$ as $i \to \infty$. Now, for any $\lambda > 0$ holds $\lambda z^i \in A_{\tau^i/\lambda}$, so clearly $\lambda z \in \limsup_{i\to\infty} A_{\tau^i/\lambda} \subset T_A(x)$. This proves that $T_A(x)$ is a cone.

Finally, $\widehat{T}_A(x)$ is a closed set as an inner limit, cf. Exercise 4.1, and seen to be a cone analogously to the proof for $T_A(x)$. To see that it is convex, take $\bar z, \widetilde z \in \widehat{T}_A(X)$. Since $\widehat{T}_A(x)$ is a cone, we need to show that $z := \bar z + \widetilde z \in \widehat{T}_A(x)$. By the definition of $\widehat{T}_A(x)$, we therefore have to show that for any sequence $\tau^i \searrow 0$ and $x^i \to x$, there exist $z^i \in (A - x^i)/\tau^i$ such that $z^i \to z$. Since $\bar z \in \widehat{T}_A(x)$, by the definition of $\widehat{T}_A(x)$ again, we can find points $\bar w^i \in A$ with $(\bar w^i - x^i)/\tau^i \to \bar z$. This implies $\bar w^i \to x$. Thus, since $\widetilde z \in \widehat{T}_A(x)$, we can find $\widetilde w^i \in A$ such that $(\widetilde w^i - \bar w^i)/\tau^i \to \widetilde z$. Now

$$\frac{\widetilde w^i - x^i}{\tau^i} = \frac{\widetilde w^i - \bar w^i}{\tau^i} + \frac{\bar w^i - x^i}{\tau^i} \to \bar z + \widetilde z = z.$$

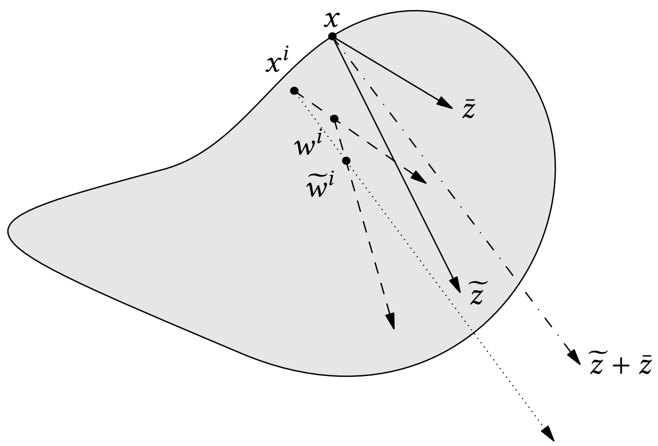

Figure 5.2: Illustration of the "bridging" argument in the proof of Proposition 5.1. As x^i converges to x, the dashed arrows converge to the solid arrows, while the dotted arrow converges to the dash-dotted one, which depicts the point $\widetilde{z} + \bar{z}$ we are trying to prove to belong to $\widehat{T}_A(x)$.

Thus $z \in \widehat{T}_A(x)$, so the cone of regular tangents is convex. Observe how we have used the varying base point in the definition of $\widehat{T}_A(x)$ to "bridge" between the two tangent cones. This is illustrated in Figure 5.2. □

Exercise(Light) 5.2. *Show that the map $x \mapsto N_A(x)$ is outer semicontinuous.*

To prove further relationships of the different cones, we need the following lemma on projections.

Lemma 5.1. *Let $A \subset \mathbb{R}^k$ be closed, and for $x \in \mathbb{R}^k$, define the (possibly multi-valued) projection*

$$P_A(x) := \operatorname*{arg\,min}_{x' \in A} \|x' - x\|.$$

Then, any $\bar{x} \in P_A(x)$ satisfies

$$x \in \bar{x} + \widehat{N}_A(\bar{x}).$$

Proof. Since A is closed, a solution $\bar{x} \in P_A(x)$ exists. We claim that $x - \bar{x} \in N_A(\bar{x})$. Indeed, for any other $x' \in A$, we have

$$\|x - \bar{x}\|^2 \le \|x - x'\|^2,$$

which, after reorganisation, is to say that

$$\langle x - \bar{x}, x' - \bar{x} \rangle \le \langle x' - x, x' - \bar{x} \rangle.$$

This implies

$$\lim_{A \ni x' \to x} \langle x - \bar{x}, \frac{x' - \bar{x}}{\|x' - \bar{x}\|} \rangle \le \lim_{A \ni x' \to x} \langle x' - x, \frac{x' - \bar{x}}{\|x' - \bar{x}\|} \rangle \le \lim_{A \ni x' \to x} \|x - x'\| = 0.$$

Thus $x - \bar{x} \in \widehat{N}_A(\bar{x})$. □

Remark 5.2. We have $P_A = \operatorname{prox}_{\delta_A}$ for δ_A the indicator function of A.

Remark 5.3. Let $\epsilon > 0$, and suppose $\bar{x} \in \partial \mathbb{B}(x, \epsilon) \cap A$, with $\operatorname{int} \mathbb{B}(x, \epsilon) \cap A = \emptyset$. It is not difficult to see that this is equivalent to $\bar{x} \in P_A(x)$. Consequently $x - \bar{x} \in \widehat{N}_A(\bar{x}) \subset N_A(\bar{x})$. We will frequently use this property, illustrated in Figure 5.3, in the proof of Proposition 5.2 below.

We also need the following notion.

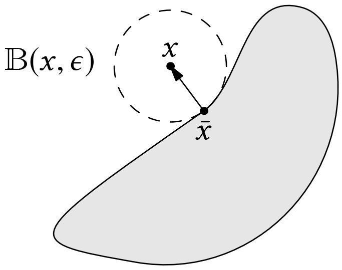

Figure 5.3: Illustration of Remark 5.3 and Lemma 5.1. Only the boundary of the ball $\mathbb{B}(x,\epsilon)$ intersects A at $\bar{x}$, so the vector $x-\bar{x}$ is normal to A at $\bar{x}$.

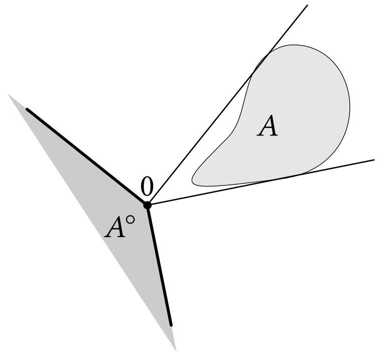

Figure 5.4: The polar cone is the normal cone at zero, of the smallest cone containing A.

Definition 5.5. For a set $A \subset \mathbb{R}^k$, we define the polar cone as

$$A^\circ := \{v \in \mathbb{R}^k \mid \langle v, z\rangle \le 0 \text{ for all } z \in A\}.$$

The polar cone is illustrated in Figure 5.4.

Exercise* 5.3. *For a set A, show that A° is closed and convex, and that $A \subset A^{\circ\circ}$. Also show that if K is a closed and convex cone, then $K = K^{\circ\circ}$.*

Proposition 5.2. *We have the inclusions*

$$\widehat{T}_A(x) \subset T_A(x), \tag{5.1}$$

$$\widehat{N}_A(x) \subset N_A(x), \tag{5.2}$$

and the relationship

$$\widehat{N}_A(x) = T_A(x)^\circ \tag{5.3}$$

If A is locally closed at x, then also

$$\widehat{T}_A(x) = N_A(x)^\circ. \tag{5.4}$$

Proof. If we fix the base point x' as x in the definition of $\widehat{T}_A(x)$, the tangent inclusion (5.1) is clear from the definition of $T_A(x)$ as an outer limit, and of $\widehat{T}_A(x)$ as an an inner limit. The normal inclusion (5.2) likewise follows from the definition of $N_A(x)$ as the outer limit of $\widehat{N}_A(x')$.

To see the relationship (5.3) between Frechét normals and tangents, take first $z \in T_A(x)$, and $w \in \widehat{N}_A(x)$. Then there exist $\tau^i \searrow 0$, and $A \ni x^i \to x$ such that $(x^i - x)/\tau^i \to z$. Since $x^i, x \in A$, we have

$$\langle z, w\rangle = \limsup_{i\to\infty} \langle x^i - x, w\rangle/\tau^i \le 0.$$

Since this holds for every $z \in T_A(x)$, we see that $w \in T_A(x)^\circ$. This shows $\widehat{N}_A(x) \subset T_A(x)^\circ$. To prove the inclusion in the other direction, we take $w \notin \widehat{N}_A(x)$. Then there must exist $A \ni x^i \to x$ with

$$\lim_{i\to\infty} \frac{\langle w, x^i - x\rangle}{\|x^i - x\|} > 0.$$

Taking $z^i := (x^i - x)/\|x^i - x\|$, and passing to a subsequence, we may assume that $z^i \to z$ for some z. Necessarily $z \in T_A(x)$ by the definition of the latter. However $\langle z, w\rangle > 0$, so $w \notin T_A(x)^\circ$. This shows $T_A(x)^\circ \subset \widehat{N}_A(x)$, concluding the proof of (5.3).

For the relationship (5.4) between regular normals and tangents, we need to work quite a bit. First of all, we show that

$$\widehat{T}_A(x) \subset K := \liminf_{A\ni x'\to x} T_A(x'). \tag{5.5}$$

Indeed, if $z \notin K$, then there exists $\epsilon > 0$ and a sequence $x^i \to x$ such that

$$\inf_{z^i \in T_A(x^i)} \|z^i - z\| \geq 2\epsilon.$$

This implies

$$\lim_{\tau \searrow 0} \inf_{\widetilde{x} \in A} \left\| \frac{\widetilde{x} - x^i}{\tau} - z \right\| \geq 2\epsilon.$$

In other words, we can find $\tau^i > 0$, with $\tau^i \searrow 0$, satisfying for each $i \in \mathbb{N}$ the inequality

$$\inf_{\widetilde{x} \in A} \left\| \frac{\widetilde{x}^i - x^i}{\tau^i} - z \right\| \geq \epsilon.$$

This says exactly that $z \notin \widehat{T}_A(x)$. Therefore (5.5) holds.

Now, to see that $\widehat{T}_A(x) \subset N_A(x)^\circ$, we take $w \in N_A(x)$ and $z \in \widehat{T}_A(x)$. This gives by the definition of $N_A(x)$, a sequence $x^i \to x$ and $w^i \to w$ with $w^i \in \widehat{N}_A(x^i)$. By (5.5), we can for each $i \in \mathbb{N}$ find $z^i \in T_A(x^i)$ with $z^i \to z$. Now the polarity relationship (5.3) gives

$$\langle w, z \rangle = \lim_{i \to \infty} \langle w^i, z^i \rangle \leq 0.$$

Therefore $\widehat{T}_A(x) \subset N_A(x)^\circ$,

In the other direction, let us take $z \notin \widehat{T}_A(x)$. Then we need to show that $z \notin N_A(x)^\circ$, or that there exists $w \in N_A(x)$ such that $\langle z, w \rangle > 0$. By Definition 5.4, $z \notin \widehat{T}_A(x)$ means the existence of $\epsilon > 0$, $x^i \to x$, and $\tau^i \searrow 0$ such that

$$\inf_{x' \in A} \left\| \frac{x' - x^i}{\tau^i} - z \right\| > \epsilon, \quad (i \in \mathbb{N}).$$

With

$$d_A(x') := \min_{x'' \in A} \|x' - x''\|,$$

this says

$$d_A(x^i + \tau^i z) \geq \tau^i \epsilon. \tag{5.6}$$

Since A is locally closed, choosing i large enough that τ^i is small enough that we may assume A closed. Lemma 5.1 then gives a vector z^i such that

$$P_A(x^i + \tau^i z) = x^i + z^i \in A, \quad \text{and} \quad \tau^i z - z^i \in \widehat{N}_A(x^i + z^i) \subset N_A(x^i + z^i).$$

Let us set

$$w^i := \frac{\tau^i z - z^i}{\|\tau^i z - z^i\|} \in \widehat{N}_A(x^i + z^i),$$

as well as $\epsilon^i := \|\tau^i z - z^i\| \geq \epsilon$. We claim that we can take $\langle z, w^i \rangle \geq \epsilon^i$. Suppose this does not hold. By Remark 5.3, we have

$$A \cap \operatorname{int} \mathbb{B}(x^i + \tau^i z, \epsilon^i) = \emptyset \quad \text{and} \quad x^i + z^i \in A \cap \partial \mathbb{B}(x^i + \tau^i z, \epsilon^i).$$

Therefore $\langle z, w^i \rangle \geq \epsilon^i$ means that $x^i + z^i = x^i + \tau^i z - \epsilon^i w^i \in L(\tau^i; x^i, \epsilon^i)$ for

$$L(\tau; x^i, \epsilon^i) := \{x^i + \tau(z - \epsilon^i v/\tau^i) \in \mathbb{R}^n \mid \|v\| = 1, \ \langle z, v \rangle \geq \epsilon^i\}$$

the "leading arc" of the ball

$$U(\tau; x^i, \epsilon^i) := \mathbb{B}(x^i + \tau z, \tau\epsilon^i/\tau^i).$$

The leading arc is illustrated by the thick line in Figure 5.5a.

We claim that the leading arc $L(\tau^i; x^i, \epsilon^i)$ is the only part of $U(\tau^i; x^i, \epsilon^i)$ covered by the balls $U(\tau; x^i, \epsilon^i)$ for $\tau \in [0, \tau^i)$. Indeed, the point $x' := x^i + \tau^i(z - \epsilon^i v/\tau^i) \in \partial U(\tau^i; x^i, \epsilon^i)$ satisfies $x' \in U(\tau; x^i, \epsilon^i)$ for $\tau < \tau^i$ when

$$\|x' - (x^i + \tau z)\| \leq \epsilon^i \tau/\tau^i.$$

In other words

$$\|(\tau^i - \tau)z - \epsilon^i v\| \leq \epsilon^i \tau/\tau^i.$$

Squaring and taking $\|v\| = 1$ as in the definition of the leading arc, we may rewrite this as

$$\langle z, v\rangle \geq \frac{(\epsilon^i)^2(\tau^i + \tau) + (\tau^i - \tau)(\tau^i\|z\|)^2}{2\tau^i\epsilon^i}.$$

The right hand side is minimised in the limit $\tau = \tau^i$, goving condition $\langle z, v\rangle \geq \epsilon^i$. So we have proved our claim that $L(\tau^i; x^i, \epsilon^i)$ is the only part of $U(\tau^i; x^i, \epsilon^i)$ covered by the balls $U(\tau; x^i, \epsilon^i))$ for $\tau \in [0, \tau^i)$. The union of all these balls is illustrated by ice cream cone structure in Figure 5.5a.

Now, if $\langle z, w^i\rangle < \epsilon^i$, we have $x^i + z^i \notin L(\tau^i; x^i, \epsilon^i)$. Therefore, by the above covering property, we may decrease τ until there exists a point $x^i + \widetilde{z}^i \in L(\tau; x^i, \epsilon^i)$, while maintaining $A \cap \operatorname{int} U(\tau; x^i, \epsilon^i) = \emptyset$. Indeed, since smaller balls corresponding to smaller τ only cover the leading arc of the boundary of the bigger balls, $U(\widetilde{\tau}^i; x^i, \epsilon^i)$ for decreasing τ will first intersect A on $L(\widetilde{\tau}^i; x^i, \epsilon^i)$ for some $\widetilde{\tau}^i \in [0, \tau^i)$. Now, if $\widetilde{\tau}^i > 0$, by Remark 5.3, we have

$$\widetilde{w} := \frac{\widetilde{\tau}^i z - \widetilde{z}^i}{\|\widetilde{\tau}^i z - \widetilde{z}^i\|} \in \widehat{N}_A(x^i + \widetilde{z}), \tag{5.7}$$

and, by the definition of $L(\widetilde{\tau}^i; x^i, \epsilon^i)$, also $\langle \widetilde{w}^i, z\rangle \geq \epsilon^i$. Replacing z^i by $\widetilde{z}^i$, we may therefore assume that $\langle w^i, z\rangle \geq \epsilon^i$.

If $\widetilde{\tau}^i = 0$, the first point of contact between A and the balls $U(\tau; x^i, \epsilon^i)$ for $\tau \geq 0$ was x^i itself. In this case, we proceed to alter x^i as in Figure 5.5b. Specifically, we replace x^i by $\bar{x}^i := x^i - \delta^i z$ for some small $\delta^i \searrow 0$. This is doable while maintaining $\operatorname{int} U(\tau^i; \bar{x}^i, \epsilon') \cap A = \emptyset$, when we also replace ϵ^i by some fixed $\epsilon' < \epsilon$ guaranteeing $U(\tau^i; \bar{x}^i, \epsilon') \subset U(\tau^i; \bar{x}^i, \epsilon^i)$. It may be that $\bar{x}^i \notin A$, but this is inconsequential. The important fact of this modification is that now $\widetilde{\tau}^i > 0$, because for some $\hat{\tau}^i > 0$ holds $x^i \in U(\hat{\tau}^i; \bar{x}^i, \epsilon')$. If $\widetilde{\tau}^i > \hat{\tau}^i$, then we have again found our $\widetilde{z}^i$ and $\widetilde{w}^i \in \widehat{N}_A(\bar{x}^i + \widetilde{z}^i)$ as in (5.7). This time, $\langle \widetilde{w}^i, z\rangle \geq \epsilon'$, which is enough.

If we can find a subsequence of indices $i \in \mathbb{N}$, unrelabelled, such that we never go all the way to the case $\widetilde{\tau}^i = \hat{\tau}^i$, then we have discovered a sequence of vectors $w^i \in \widehat{N}_A(x^i + z^i)$, with $x^i + z^i \to x$, as well as $\|w^i\| = 1$ and $\langle w^i, z\rangle \geq \epsilon'$. By going to a further subsequence we can assume that $w^i \to w \in N_A(x)$. But then $\langle w, z\rangle \geq \epsilon' > 0$. This says that $z \notin N_A(x)^\circ$. Therefore $\widehat{T}_A(x) \supset N_A(x)^\circ$.

If, on the other hand, only the case $\widetilde{\tau}^i = \hat{\tau}^i$ occurs infinitely often, we proceed as follows. We deduce by referral to Remark 5.3 that x^i is a closest point in A to $\bar{x}^i + \hat{\tau}z = x^i + (\hat{\tau} - \delta^i)z$. Therefore $z \in N_A(x^i)$. Since this happens for all large enough $i \in \mathbb{N}$, we find that in the limit $z \in N_A(x)$. But then most definitely $z \notin N_A(x)^\circ$. This shows that again $\widehat{T}_A(x) \supset N_A(x)^\circ$. ◻

Corollary 5.1. *Let $A \subset \mathbb{R}^k$ be locally closed at x. Then*

$$\widehat{T}_A(x) = \liminf_{A \ni x' \to x} T_A(x).$$

Proof. We have already proved one direction of the claim in (5.5). For the other direction, we note that by Proposition 5.2, using the local closedness in its application, we have to prove

$$N_A(x)^\circ \supset K := \liminf_{A \ni x' \to x} T_A(x). \tag{5.8}$$

So suppose $z \notin N_A(x)^\circ$. Then there exists $w \in N_A(x)$ satisfying $\langle z, w\rangle > 0$. We have again $w = \lim_{i\to\infty} w^i$ for some $x^i \to x$ and $w^i \in \widehat{N}_A(x^i)$. But if $z \in K$, we can find $z^i \in T_A(x^i)$ with $z^i \to z$. By the tangent-normal polarity (5.3), we then obtain $\langle z, w\rangle = \lim_{i\to\infty}\langle z^i, w^i\rangle \leq 0$, which is a contradiction. Therefore $z \notin K$, so (5.8) holds, finishing the proof of the claim. ◻

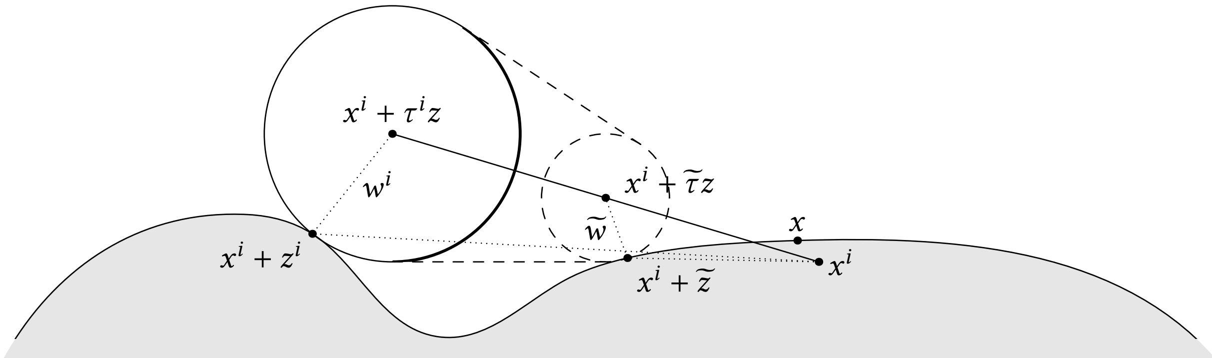

(a) Illustration of the τ^i reduction argument. Replacement of τ^i and z^i by $\widetilde{\tau}$ and $\widetilde{z}$, hence w^i by $\widetilde{w}$, ensures $\langle z, w^i \rangle \geq \epsilon$. The thick line indicates the "leading arc" $L^i(\tau^i)$.

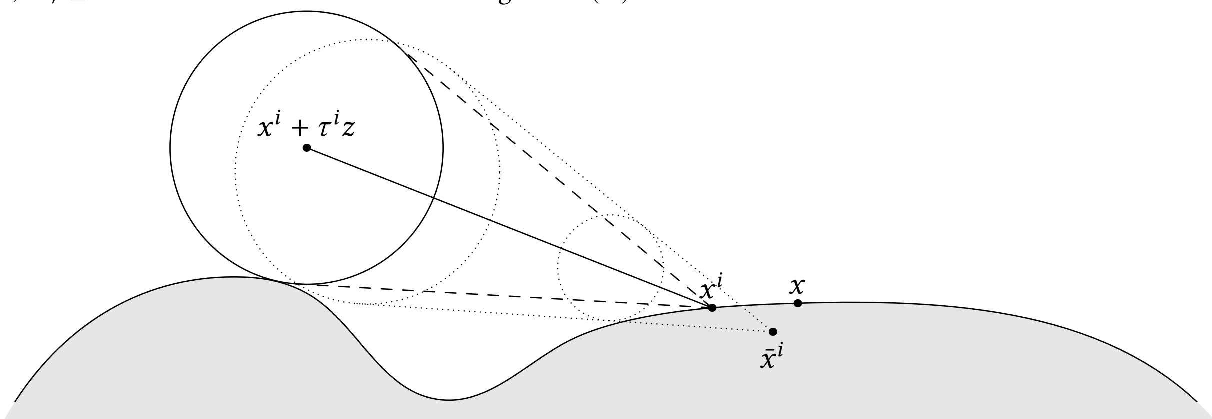

(b) Alteration of the base point x^i to $\bar{x}^i$, in case $\widetilde{\tau}$ cannot be found in (a).

Figure 5.5: Illustrations for the proof of $\widehat{T}_A(x) \supset N_A(x)^\circ$ in Proposition 5.2.

Theorem 5.1. *At any point x where $A \subset \mathbb{R}^k$ is locally closed, the following conditions are equivalent, and when they hold, we say that A is regular at x.*

(i) $N_A(x) = \widehat{N}_A(x)$.

(ii) $T_A(x) = \widehat{T}_A(x)$.

Proof. Suppose (i) holds. Then by Exercise 5.3, and Proposition 5.2 holds

$$T_A(x) \subset T_A(x)^{\circ\circ} = \widehat{N}_A(x)^{\circ} = N_A(x)^{\circ} = \widehat{T}_A(x) \subset T_A(x)$$

This shows (ii). The other direction is completely analogous, exchanging the roles of "N" and "T". □

Corollary 5.2. *If A is locally closed and regular at x, then both $T_A(x)$ and $N_A(x)$ are convex and polar to each other,*

$$N_A(x) = T_A(x)^{\circ}, \quad T_A(x) = N_A(x)^{\circ}.$$

Theorem 5.2. *Every convex set $A \subset \mathbb{R}^k$ is regular at every point $x \in \mathbb{R}^k$.*

Proof. We have already seen that

$$\widehat{N}_A(x) = \partial_F \delta_A(x) = \partial \delta_A(x),$$

and that the convex subdifferential is outer semicontinuous. Thus

$$N_A(x) := \limsup_{x' \to x} \widehat{N}_A(x') \subset \widehat{N}_A(x) \subset N_A(x).$$

This proves $N_A(x) = \widehat{N}_A(x)$, that is, A is regular at x. □

5.2 Derivatives and coderivatives

We are finally ready to differentiate set-valued mappings. Although there's a multitude different definitions, they are actually pretty straightforward now with the preparations of the previous section.

Definition 5.1. Let $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$, and take $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$. We then define

(i) The Fréchet coderivative of F at x for y as the map $\widehat{D}^*F(x|y) : \mathbb{R}^m \rightrightarrows \mathbb{R}^n$ given by

$$\widehat{D}^*F(x|y)(\Delta y) := \{\Delta x \mid (\Delta x, -\Delta y) \in \widehat{N}_{\operatorname{Graph} F}(x, y)\}.$$

(ii) The coderivative of F at x for y as the map $D^*F(x|y) : \mathbb{R}^m \rightrightarrows \mathbb{R}^n$ given by

$$D^*F(x|y)(\Delta y) := \{\Delta x \mid (\Delta x, -\Delta y) \in N_{\operatorname{Graph} F}(x, y)\}.$$

(iii) The regular derivative of F at x for y as the map $\widehat{D}F(x|y) : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ given by

$$\widehat{D}F(x|y)(\Delta x) := \{\Delta y \mid (\Delta x, \Delta y) \in \widehat{T}_{\operatorname{Graph} F}(x, y)\}.$$

(iv) The graphical derivative of F at x for y as as the map $DF(x|y) : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ given by

$$DF(x|y)(\Delta x) := \{\Delta y \mid (\Delta x, \Delta y) \in T_{\operatorname{Graph} F}(x, y)\}.$$

Regarding the minus sign in the definitions of the coderivatives, recall that $(\nabla f(x), -1)$ is normal to epi f for a smooth function f; see also Figure 5.6. Also observe how the coderivatives operate from $\mathbb{R}^m$ to $\mathbb{R}^n$, while the derivatives operate from $\mathbb{R}^n$ to $\mathbb{R}^m$. We illustrate the different graphical derivatives and coderivatives on $\partial \delta_{[-1,1]}$ in Figure 5.7.

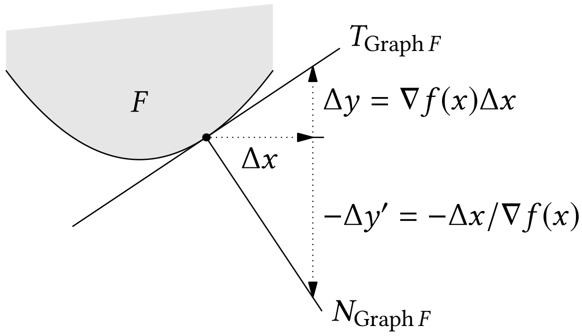

Figure 5.6: Illustration, why the coderivatives negate Δy in comparison to the normal cone.

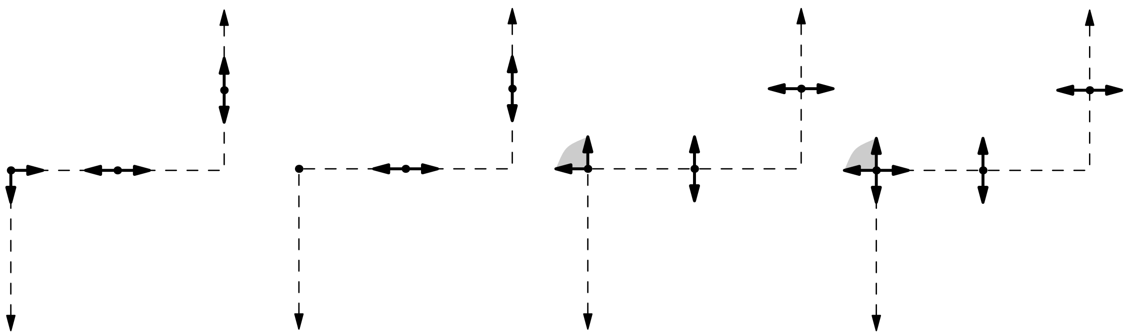

(a) $D(\partial f)$ **(b)** $\widehat{D}(\partial f)$ **(c)** $\widehat{D}^*(\partial f)$ **(d)** $D^*(\partial f)$

Figure 5.7: Illustration of the different graphical derivatives and coderivatives for ∂f with $f = \delta_{[-1,1]}$. The dashed line is Graph ∂f. The dots indicate the base points (x, y) where $D(\partial f)(x|y)$ is calculated, and the thick arrows and grey areas the directions of $(\Delta x, \Delta y)$ relative to the base point. ($(\Delta x, -\Delta y)$ for the coderivatives.) Observe that there is no graphical regularity at $(x, y) = (-1, 0)$, and, analogously, $(x, y) = (1, 0)$. Everywhere else, ∂f is graphically regular.

Exercise(Light) 5.4. *Show that*

$$\Delta y \in DF(x|y)(\Delta x) \iff \Delta x \in DF^{-1}(y|x)(\Delta y),$$

and analogously for $\widehat{D}^*$, $\widehat{D}$, *and* D^*. *What can you say about other basic relationships that the different graphical derivatives and coderivatives have?*

The following result is immediate from Theorem 5.1.

Theorem 5.1. *If* Graph F *is closed at* (x, y), *the following conditions are equivalent, and when they hold, we say that* F *is* graphically regular *at* x *for* y.

(i) $D^*F(x|y) = \widehat{D}^*F(x|y)$.

(ii) $DF(x|y) = \widehat{D}F(x|y)$.

We will generally only work with graphically regular maps. Then, with the help of the upper and lower adjoints of a set-valued function $H : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$, defined as

$$H^{\circ+}(y) := \{x \mid \langle x, x' \rangle \le \langle y, y' \rangle \text{ when } y' \in H(x')\},$$

and

$$H^{\circ-}(y) := \{x \mid \langle x, x' \rangle \ge \langle y, y' \rangle \text{ when } y' \in H(x')\},$$

Proposition 5.2 gives the following.

Corollary 5.1. *We have the relationships*

$$\widehat{D}^* F(x|y) = DF(x|y)^{\circ +} \tag{5.1}$$

If Graph F *is locally closed at* (x, y)*, then also*

$$\widehat{D}F(x|y) = D^* F(x|y)^{\circ -}. \tag{5.2}$$

Proof. By definition

$$DF(x|y)(\Delta x) := \{\Delta y \mid (\Delta x, \Delta y) \in T_{\operatorname{Graph} F}(x, y)\},$$

and

$$\widehat{D}^* F(x|y)(\Delta y) := \{\Delta x \mid (\Delta x, -\Delta y) \in \widehat{N}_{\operatorname{Graph} F}(x, y)\}.$$

Using (5.3) from Proposition 5.2 We thus calculate

$$\begin{aligned}
\Delta x \in DF(x|y)^{\circ +}(\Delta y) &\iff \langle \Delta x, \Delta x' \rangle \le \langle \Delta y, \Delta y' \rangle \text{ when } \Delta y' \in DF(x|y)(\Delta x') \\
&\iff \langle \Delta x, \Delta x' \rangle + \langle -\Delta y, \Delta y' \rangle \le 0 \text{ when } (\Delta x', \Delta y') \in T_{\operatorname{Graph} F}(x, y) \\
&\iff (\Delta x, -\Delta y) \in T_{\operatorname{Graph} F}(x, y)^{\circ} \\
&\iff (\Delta x, -\Delta y) \in \widehat{N}_{\operatorname{Graph} F}(x, y) \\
&\iff \Delta x \in \widehat{D}^* F(x|y)(\Delta y).
\end{aligned}$$

This proves (5.1). The expression (5.2) is calculated completely analogously using (5.4). Going back from the coderivatives involving a negated Δy in the relation to normal cones, we now need the lower adjoint instead of the upper adjoint in the other direction. □

Exercise(Light) 5.5. *Write down the upper and lower adjoints for a linear map* $H : \mathbb{R}^n \to \mathbb{R}^m$.

Exercise 5.6. *Let* $F : \mathbb{R}^n \to \mathbb{R}^m$ *be single-valued and differentiable. Show with* $y = F(x)$ *that*

$$DF(x|y)(\Delta x) = \nabla F(x)\Delta x \quad \text{and} \quad \widehat{D}^* F(x|y)(\Delta y) = [\nabla F(x)]^* \Delta y.$$

(Here $[\nabla F(x)]^*$ *stands for the adjoint of* $\nabla F(x)$ *as a linear operator; a reader unfamiliar with functional analysis may consider it a matrix transpose.) Also show that if* F *is graphically regular at* (x, y)*, if it is strictly differentiable at* x*, meaning that there exists* $L \in \mathbb{R}^{m \times n}$ *such that*

$$\lim_{h \to 0,\, x' \to x} \frac{F(x +' h) - F(x') - Lx}{\|h\|} = 0.$$

What does this mean for linear F*?*

Exercise 5.7. *Let* $f(x) = |x|$ *on* $\mathbb{R}$*. Calculate the derivatives and coderivatives of* ∂f*. When is* ∂f *graphically regular?*

5.3 Basic calculus

What is the coderivative of the a sum $H = F + G$ of set-valued maps $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ and $G : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$? How about a composition $F \circ J$ for $J : \mathbb{R}^k \rightrightarrows \mathbb{R}^n$? Generally, such results are very limited, and the most we can derive are the coderivative inclusions

$$D^* H(x|y) \subset \bigcup_{\substack{y = y_1 + y_2, \\ y_1 \in F(x), \\ y_2 \in G(x)}} D^* F(x|y_1) + D^* G(x|y_2),$$

and

$$D^*(F \circ J)(x|y) \subset \bigcup_{w \in J(x) \cap F^{-1}(y)} D^*J(x|w) \circ D^*F(w|y),$$

with the inclusions in terms of graphs.

We will not go into these general results in detail. Instead, in this course, we concentrate on practically useful cases, where we do have sharp results. These generally involve sums and compositions with single-valued and affine mappings. Also recall that we are mainly interested in the cases $F = \partial f$. In particular, if $g(x) = \|Tx - f\|^2/2$, then $G = \partial g$ is the affine mapping $G(x) = T^*(Tx - f)$, and $\nabla G = T^*T$. This allows f to be non-smooth when calculating a second-order graphical differential for the sum $f + g$. For a similarly exact theory in the infinite-dimensional space $L^2(\Omega)$, see [22].

Theorem 5.1 (Addition of a single-valued differentiable map). *Let $H = F + G$ for $F : \mathbb{R}^m \rightrightarrows \mathbb{R}^n$, and a differentiable single-valued function $G : \mathbb{R}^m \to \mathbb{R}^n$. With $y - G(x) \in F(x)$, we have*

$$DH(x|y)(\Delta x) = DF(x|y - G(x))(\Delta x) + \nabla G(x)\Delta x. \tag{5.1}$$

Proof. Observe that, applying the definition of $T_{\operatorname{Graph} F}$, for any set-valued map F, we may write

$$DF(x|y)(\Delta x) = \limsup_{\substack{\tau \searrow 0 \\ \Delta x' \to \Delta x}} \frac{F(x + \tau\Delta x') - y}{\tau}. \tag{5.2}$$

Indeed

$$(\Delta x, \Delta y) \in \limsup_{\tau \searrow 0} \frac{\operatorname{Graph} F - (x, y)}{\tau}$$

if and only if there exist $\tau^i \searrow 0$ and x^i such that

$$\Delta x = \lim_{i \to \infty} \frac{x^i - x}{\tau^i},$$

and

$$\Delta y \in \limsup_{i \to \infty} \frac{F(x^i) - y}{\tau^i}.$$

The former forces $x^i = x - \tau^i \Delta x^i$ for $\Delta x^i \to \Delta x$, so the latter gives (5.2).

Returning to (5.1), we just use (5.2) to compute

$$\begin{aligned}
\Delta y \in DH(x|y)(\Delta x) &\iff \Delta y = \lim_{i \to \infty} \frac{y^i - y}{\tau^i} \quad (\text{for some } \tau^i \searrow 0,\ y^i \in H(x + \tau^i \Delta x^i)) \\
&\iff \Delta y = \lim_{i \to \infty} \frac{\tilde{y}^i + G(x + \tau^i \Delta x^i) - y}{\tau^i} \quad (\ldots \text{ for some } \tilde{y}^i \in F(x + \tau^i \Delta x^i)) \\
&\iff \Delta y = \lim_{i \to \infty} \frac{\tilde{y}^i - (y - G(x))}{\tau^i} + \nabla G(x)\Delta x \\
&\iff \Delta y - \nabla G(x)\Delta x \in DF(x|y - G(x))(\Delta x).
\end{aligned}$$

□

Exercise 5.8 (Outer composition with a linear map). *Let $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ be a set-valued map, and $A : \mathbb{R}^m \to \mathbb{R}^k$ a linear map satisfying*

$$Az = 0 \implies z = 0. \tag{5.3}$$

Show that

$$D(A \circ F)(x|y)(\Delta x) = \bigcup_{z : Az = y} A\,DF(x|z)(\Delta x). \tag{5.4}$$

Remark 5.1. Note that (5.3) implies that for each y there is a unique z satisfying $Az = y$. Therefore the union in (5.4) can be avoided.

Theorem 5.2 (Inner composition with a linear map). *Let $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ be a set-valued map, and $A : \mathbb{R}^k \to \mathbb{R}^n$ a linear map satisfying*

$$A^T z = 0 \implies z = 0. \tag{5.5}$$

Then

$$D(F \circ A)(x|y)(\Delta x) = DF(Ax|y)(A\Delta x).$$

Remark 5.2. The above result can be extended to more general single-valued maps under graphical regularity [4].

Proof. Let us set $z := Ax$ and $\Delta z := A\Delta x$. We use the expression (5.2) to derive

$$\begin{aligned} &\Delta y \in D(F \circ A)(x|y)(\Delta x) \\ &\quad \iff \Delta y = \lim_{i\to\infty} \frac{y^i - y}{\tau^i} \quad (\text{for some } \tau^i \searrow 0,\ \Delta x^i \to \Delta x,\ y^i \in F(Ax + \tau^i A\Delta x^i)) \\ &\quad \iff \Delta y = \lim_{i\to\infty} \frac{y^i - y}{\tau^i} \quad (\ldots \text{ for } \Delta z^i := A\Delta x^i, y^i \in F(z + \tau^i \Delta z^i)) \\ &\quad \iff \Delta y = \lim_{i\to\infty} \frac{y^i - y}{\tau^i} \quad (\text{for some } \Delta z^i \to \Delta z, y^i \in F(z + \tau^i \Delta z^i)) \\ &\quad \iff \Delta y \in DF(z|y)(\Delta z). \end{aligned}$$

In the semifinal step we have used the fact that the range of A is full, which is equivalent to the condition (5.5). This implies that every Δz^i satisfies $\Delta z^i = \Delta x^i$ for some Δx^i. □

Corollary 5.1. *Let $f(x) = f_0(Ax)$ for some convex function f_0 and linear map A satisfying (5.5) and $\mathcal{R}(A) \cap \operatorname{ri} \operatorname{dom} f_0 \neq \emptyset$. Then*

$$D(\partial f)(x|y)(\Delta x) = \bigcup_{z: A^T z = y} A^T D(\partial f_0)(Ax|z)(A\Delta x).$$

5.4 The Mordukhovich criterion

We finally have all the tools necessary to express the following simple criterion for the Aubin property to be satisfied.

Theorem 5.1. *Let $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$. Suppose $\operatorname{Graph} F$ is locally closed at (x, y) and*

$$D^* F(x|y)(0) = \{0\}. \tag{5.1}$$

Then F has the Aubin property at x for y with

$$\operatorname{lip} F(x|y) = |D^* F(x|y)|^+, \tag{5.2}$$

where the outer norm

$$|H|^+ := \sup_{\|w\|\le 1} \sup_{z\in H(w)} \|z\|.$$

We illustrate the outer norm in Figure 5.8. Of particular intrest to us is the H whose graph is a cone in Figure 5.8b. As we recall from Proposition 5.1 and Definition 5.1, $\operatorname{Graph} D^* F(x|y)$ and is a cone related to $N_{\operatorname{Graph} F}(x, y)$. For cones, we see that the outer norm is, in a sense, the normalised maximum opening of the cone away from the line $z = 0$.

Using Definition 5.1 and the fact that we only care about the magnitudes of the vectors Δx and Δy, not their exact directions, we may write

$$\begin{aligned} |D^* F(x|y)|^+ &= \sup\{\|\Delta x\| \mid (\Delta x, -\Delta y) \in N_{\operatorname{Graph} F}(x, y),\ \|\Delta y\| \le 1\} \\ &= \sup\{\|\Delta x\| \mid (\Delta x, \Delta y) \in N_{\operatorname{Graph} F}(x, y),\ \|\Delta y\| \le 1\}. \end{aligned}$$

With this in mind, using the examples from Figure 4.4, we illustrate in Figure 5.9 how the outer norm of the coderivative relates to the Aubin property.

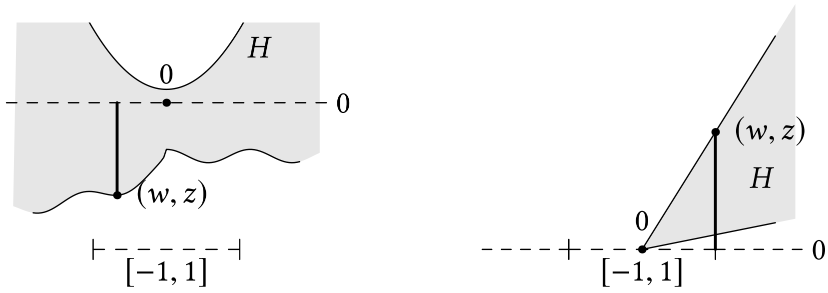

(a) A general function **(b)** Graph H is a cone

Figure 5.8: Points (w, z) achieving the supremum in the expression of the outer norm $|H|^+$.

(a) Property is satisfied **(b)** Property is not satisfied

Figure 5.9: Illustration of outer norm $|D^*F(x|y)|^+ = \sup\{\|\Delta x\| \mid (\Delta x, \Delta y) \in N_{\operatorname{Graph} F}(x, y),\ \|\Delta y\| \leq 1\}$. The arrows illustrate the directions contained in the normal cone. In (a), $\Delta y \in [0, \infty)$, but $\Delta x = 0$, so $|D^*F(x|y)|^+ = 0$, and the Aubin property is satisfied. In (b), we can for $\Delta y = 0$ take any $\Delta x \in (-\infty, 0]$, so $|D^*F(x|y)|^+ = \infty$, and the Aubin property is not satisfied.

Corollary 5.1 (Inverse function theorem). *Let $F : \mathbb{R}^m \to \mathbb{R}^k$, ($k \leq m$), be single-valued and differentiable. Let $x \in \mathbb{R}^m$. If the Jacobian at x satisfies*

$$\nabla F(x)^T z = 0 \implies z = 0, \tag{5.3}$$

then there exist $\ell > 0$ and $\rho > 0$, and for all $y' \in \mathbb{B}(F(x), \rho)$ a single-valued selection $J(y') \in F^{-1}(y')$, satisfying

$$\|x - J(y')\| \leq \ell \|F(x) - y'\|. \tag{5.4}$$

Proof. We let $y := F(x)$. Since F is continuous, its graph is closed. Since $D^*F^{-1}(y|x) = [D^*F(x|y)]^{-1}$, cf. Exercise 5.4, the condition (5.1), applied to the inverse, can also be written

$$D^*F(x|y)(z) = 0 \implies z = 0.$$

For single-valued F, this is exactly (5.3), as we can see from Exercise 5.6. The conditions of Theorem 5.1 are therefore satisfied for F^{-1}. Consequently the latter satisfies the Aubin property at $y = F(x)$ for x. (In other words, F is metrically regular at x for y.) By Proposition 4.1, this can be written simply as the existence for any $\ell' > \operatorname{lip} F^{-1}(y|x)$ of $\rho, \delta > 0$ such that

$$\inf_{x'' \in F^{-1}(y')} \|x' - x''\| \leq \ell' \|F(x') - y'\|, \quad (x' \in \mathbb{B}(x, \delta), y' \in \mathbb{B}(y, \rho)).$$

Taking $x' = x$, we have in particular

$$\inf_{x'' \in F^{-1}(y')} \|x - x''\| \leq \ell' \|F(x) - y'\|, \quad (y' \in \mathbb{B}(F(x), \rho)).$$

The infimum might not be reached, but choosing $\ell > \ell'$, we can for $y' \in \mathbb{B}(F(x), \rho)$ find $x'' = J(y') \in F^{-1}(y')$ satisfying

$$\|x - x''\| \leq \ell \|F(x) - y'\|.$$

This is exactly (5.4). □

5.5 Application to sensitivity analysis

The following proposition applied to $F(x,p) = \partial_x f(x;p)$ provides a general tool for analysing the stability of solution maps.

Proposition 5.1. *Let $F : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^k$, and define the implicit mapping*

$$S(p) := \{x \in \mathbb{R}^n \mid 0 \in F(x,p)\}.$$

Then S has the Aubin property at $\bar{p}$ for $\bar{x} \in S(\bar{p})$ if

$$(0, \Delta p) \in D^* F(\bar{x}, \bar{p}|0)(\Delta u) \implies \Delta u = 0, \Delta p = 0. \tag{5.1}$$

Proof. Let us define the map

$$Q(u,p) := \{x \in \mathbb{R}^n \mid u \in F(x,p)\}.$$

Then $S(p) = Q(0,p)$. Moreover, the Aubin property of Q at $(0,\bar{p})$ for $\bar{x}$ will imply the Aubin property of S at $\bar{p}$ for $\bar{x}$ by simple limitation of x, x' in (4.1) to the subspace $\{0\} \times \mathbb{R}^m$.

For Q to have the Aubin property at $(0,\bar{p})$ for $\bar{x}$, the Mordukhovich criterion, Theorem 5.1 requires

$$D^* Q(0, \bar{p}|\bar{x})(0) = \{0\},$$

which is to say that

$$(\Delta u, \Delta p, 0) \in N_{\operatorname{Graph} Q}(0, \bar{p}, \bar{x}) \implies \Delta u = 0,\ \Delta p = 0. \tag{5.2}$$

Now

$$\operatorname{Graph} Q = \{(u,p,x) \mid u \in F(x,p)\} = \mathcal{P} \operatorname{Graph} F,$$

for the permutation $\mathcal{P}(x,p,u) := (u,p,x)$. Therefore also $N_{\operatorname{Graph} Q}(u,p,x) = \mathcal{P} N_{\operatorname{Graph} F}(\mathcal{P}(u,p,x))$. In particular, (5.2) becomes

$$(0, \Delta p, \Delta u) \in N_{\operatorname{Graph} F}(\bar{x}, \bar{p}, 0) \implies \Delta u = 0,\ \Delta p = 0. \tag{5.3}$$

This is exactly (5.1). ∎

Remark 5.1. Proposition 5.1 is related to the classical implicit function theorem. If F is graphically regular at $(\bar{x}, \bar{p}, 0)$, it also possible derive formulas for DS, such as

$$DS(\bar{p}|\bar{x})(\Delta p) = \{\Delta x \in \mathbb{R}^n \mid DF(\bar{x}, \bar{p}|0)(\Delta x, \Delta p) \ni 0\}.$$

For details we refer to [4, Theorem 9.56 & Proposition 8.41]. Using Lemma ??, we can then define a more wieldy expression for the graphical modulus, namely

$$\begin{aligned}
\operatorname{lip} S(\bar{p}|\bar{x}) &= \sup_{\|w\| \le 1} \sup\{\|z\| \mid \langle z, \Delta p \rangle + \langle w, \Delta x \rangle \le 0 \text{ when } \Delta x \in DS(\bar{p}|\bar{x})(\Delta p)\} \\
&= \sup_{\|w\| \le 1} \sup\{\|z\| \mid \langle z, \Delta p \rangle + \langle w, \Delta x \rangle \le 0 \text{ when } DF(\bar{x}, \bar{p}|0)(\Delta x, \Delta p) \ni 0\}.
\end{aligned}$$

This can be further estimated similarly to Corollary ??.

5.6 Tilt-stability revisited

Let us recall the tilt-stability of Section 4.3. Namely, we studied

$$f(x;p) = g(x) - \langle p, x\rangle,$$

for some $g : \mathbb{R}^n \to \overline{\mathbb{R}}$, and a tilt parameter p. Then the stability of the solution map

$$S(p) = \{x \in \mathbb{R}^n \mid p \in \partial g(x)\} = (\partial g)^{-1}(p)$$

was reduced to the metric regularity of ∂g. We now have a criterion to study it. Indeed, by Corollary **??**, we have the estimate

$$\operatorname{lip} [\partial g]^{-1}(p|x) \le \sup\{\|z\| \mid \langle z, \Delta p\rangle \le \|\Delta x\| \text{ when } \Delta p \in D[\partial g](x|p)(\Delta x)\}. \tag{5.1}$$

This can also be deduced from Remark 5.1 applied to $F(x,p) := \partial g(x) - p$. To estimate the modulus of metric regularity of ∂g, we just need to calculate the graphical "second derivative" of g.

Let us study this for Lasso-type problems, only analysing the stability with respect to the dependent variables or measurements b_i. We do not analyse stability with respect to the features $\{a_i\}_{i=1}^n$, encoded into the matrix

$$A = \begin{pmatrix} a_1^T \\ \vdots \\ a_n^T \end{pmatrix} \in \mathbb{R}^{n\times m}.$$

In principle this can, however, be done through Proposition 5.1.

Now

$$g(x) = \frac{1}{2}\|b - Ax\|_2^2 + \lambda\psi(x)$$

for $\psi(x) = \|x\|_1$. Of course, Theorem 2.1 gives

$$\partial g(x) = A^T(Ax - b) + \lambda\partial\psi(x).$$

Since the first part of $\partial g(x)$ is single-valued linear, using Theorem 5.1, and Exercise 5.8 on the linear map $q \mapsto \lambda q$, we deduce for any $p \in \partial g(x)$ that

$$\begin{aligned} D[\partial g](x|p)(\Delta x) &= A^T A\Delta x + D[\lambda\partial\psi](x|p - A^T(Ax - b))(\Delta x) \\ &= A^T A\Delta x + \lambda D[\partial\psi](x|\lambda^{-1}(p - A^T(Ax - b)))(\Delta x) \end{aligned} \tag{5.2}$$

For tilt-stability, we only need to consider $p = 0$ and an optimal solution $\hat{x}$ to $0 \in \partial g(\hat{x})$. We therefore set

$$\widetilde{p} := -\lambda^{-1}A^T(A\hat{x} - b).$$

For Lasso, in Exercise 5.7, we have already calculated the graphical derivative of $\partial\psi$; namely

$$\partial\psi(\hat{x}) = \prod_{i=1}^m \begin{cases} \operatorname{sgn}\hat{x}_i, & \hat{x}_i \neq 0 \\ [-1, 1], & \hat{x}_i = 0. \end{cases}$$

Recalling the expressions for $D[\partial\psi]$ and $\widehat{D}[\partial\psi]$ from Exercise 5.7, we see that graphical regularity corresponds to the strict complementarity condition

$$\text{either } \hat{x}_i \neq 0 \text{ or } |\widetilde{p}_i| < 1, \quad (i = 1, \dots, n). \tag{5.3}$$

This says that graphical regularity fails if $\hat{x}_i = 0$ and $[A^T(A\hat{x} - b)]_i = \pm\lambda$. This means that the optimality condition $0 \in \partial g(\hat{x})$ is not satisfied "strictly" at a point of non-smoothness, because at the coordinate i it becomes

$$[A^T(A\hat{x} - b)]_i \in [-\lambda\partial\psi(\hat{x})]_i = [-\lambda, \lambda].$$

The only way to satisfy the optimality condition $p \in \partial g(x)$ for small $p \neq 0$ might therefore be to make $\hat{x}_i = 0$ an active variable $x_i \neq 0$. This might directly give x_i a large value.

If (5.3) however holds, then

$$\widehat{D}[\partial\psi](\hat{x}|\widetilde{p})(\Delta x) = D[\partial\psi](\hat{x}|\widetilde{p})(\Delta x) = \prod_{i=1}^{n} \begin{cases} \{0\}, & \hat{x}_i \neq 0,\ \widetilde{p}_i = \operatorname{sgn} \hat{x}_i, \\ \mathbb{R}, & \hat{x}_i = 0,\ |\widetilde{p}_i| < 1, \Delta x = 0 \\ \emptyset, & \text{otherwise}. \end{cases}$$

Therefore, if

$$\mathcal{I} := \{i \in \{1, \dots, n\} \mid \hat{x}_i = 0\}$$

is the set of "inactive" indices, and we set

$$V = \prod_{i=1}^{n} \begin{cases} \mathbb{R}, & i \notin \mathcal{I}, \\ \{0\}, & i \in \mathcal{I}, \end{cases}$$

then referring back to (5.2) we obtain

$$D[\partial g](\hat{x}|0)(\Delta x) = \begin{cases} A^T A \Delta x + V^\perp, & \Delta x \in V, \\ \emptyset, & \Delta x \notin V. \end{cases}$$

Note how λ disappears fro the expression, as V and $V^\perp$ are, as subspaces, invariant with respect to multiplication by λ. Using this in (5.1), we obtain

$$\begin{aligned} \operatorname{lip} [\partial g]^{-1}(0|\hat{x}) &\leq \sup\{\|z\| \mid \langle z, A^T A\Delta x + V^\perp\rangle \leq \|\Delta x\| \text{ when } \Delta x \in V\} \\ &= \sup\{\|z\| \mid z \in V,\ \langle z, A^T A \Delta x\rangle \leq \|\Delta x\| \text{ when } \Delta x \in V\} \\ &= \sup\{\|P_V z\| \mid \langle z, P_V^T A^T A P_V \Delta x\rangle \leq \|\Delta x\| \text{ for all } \Delta x\} \\ &= \sup\{\|P_V z\| \mid \|P_V A^T A P_V z\| \leq 1\}. \end{aligned}$$

Here P_V denotes the orthogonal projection into V. This says that for S to be tilt-stable, the matrix $P_V A^T A P_V$ has to be non-singular on the subspace V. In other words, if

$$A = (\widetilde{a}_1 \quad \dots, \widetilde{a}_n)$$

for some column vectors $\widetilde{a}_i \in \mathbb{R}^m$ corresponding to different features (not to be confused with the sample-specific feature vectors a_j), if the complement $\mathcal{I}^c = \{i_1, \dots, i_k\}$, and we build

$$A_{\mathcal{I}^c} := (\widetilde{a}_{i_1} \quad \dots \quad \widetilde{a}_{i_k}) \in \mathbb{R}^{n \times k},$$

then the matrix

$$A_{\mathcal{I}^c}^T A_{\mathcal{I}^c}$$

corresponding to the "active features" $\hat{x}_i \neq 0$ has to be non-singular. The inverse of its smallest eigenvalue will be $\operatorname{lip} [\partial g]^{-1}(0|\hat{x})$. Thus the Lasso is stable with respect to the dependent variables b_i, as long there are at least as much of them as active features x, that is $n \geq k$, and the corresponding feature-specific measurement vectors $\widetilde{a}_{i_1}, \dots \widetilde{a}_{i_k}$ are linearly independent. In our Exercise 3.4 on wine quality, the Lasso is stable with respect to the rankings of the wines, as long as the found two most significant features provide two linearly independent vectors of measurements of the different physicochemical properties. Linear independence here an be seen as the properties of the wine actually being descriptive of the classifications; the Lasso in this case may be unstable if the discovered properties were not truly descriptive of the classifications, while it is stable if they are.

Remark 5.1. Similar techniques can be used to study other problems. In case of total variation image reconstruction the general structure is the same as above, except ψ will involve the gradient operator ∇_d. This significantly complicates the analysis, which is, in fact, more easily performed for the saddle point problem (3.4) and the corresponding variational inclusion $0 \in H(x, y)$ for H the monotone operator in (3.5). This will not generally be stable, unless we introduce Moreau–Yosida regularisation—often known in image processing as Huber regularisation—replacing at each pixel i the function $f_i(g) = \|g\|$ by the "smooth around zero" function

$$f_{i,\tau}(g) := \min_{g'} f_i(g') + \frac{1}{2\tau}\|g' - g\|^2.$$

This corresponds to replacing in the saddle point formulation f_i^* by

$$f_{i,\tau}^*(y) = f_i^*(y) + \frac{\tau}{2}\|y\|^2.$$

The squared norm adds a level of strong convexity, which is helpful for stability.

Bibliography

[1] J. Nocedal & S. Wright, *Numerical Optimization*, Springer, 2006.

[2] T. Chan & J. Shen, *Image Processing and Analysis: Variational, PDE, Wavelet, and Stochastic Methods*, Society for Industrial & Applied Mathematics (SIAM), 2005.

[3] B. Schölkopf & A. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, 2002.

[4] R. T. Rockafellar & R. J.-B. Wets, *Variational Analysis*, Springer, 1998.

[5] B. S. Mordukhovich, *Variational Analysis and Generalized Differentiation I: Basic Theory*, Springer Verlag, 2006.

[6] J. Aubin & H. Frankowska, *Set-Valued Analysis*, Springer, 1990.

[7] J.-B. Hiriart-Urruty & C. Lemaréchal, *Convex analysis and minimization algorithms I-II*, Springer, 1993.

[8] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, 1972.

[9] I. Ekeland & R. Temam, *Convex analysis and variational problems*, SIAM, 1999.

[10] H. Bauschke & P. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, 2011.

[11] H. Attouch & H. Brezis, *Duality for the Sum of Convex Functions in General Banach Spaces*, in: *Aspects of Mathematics and its Applications*, vol. 34, North-Holland, Elsevier, 1986, 125–133.

[12] F. Browder & W. Petryshyn, *Construction of fixed points of nonlinear mappings in Hilbert space*, Journal of Mathematical Analysis and Applications **20** (1967).

[13] F. E. Browder, *Nonexpansive nonlinear operators in a Banach space*, Proceedings of the National Academy of Sciences of the United States of America **54** (1965).

[14] P. L. Lions & B. Mercier, *Splitting Algorithms for the Sum of Two Nonlinear Operators*, SIAM Journal on Numerical Analysis **16** (1979).

[15] J. Eckstein & D. Bertsekas, *On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators*, Mathematical Programming **55** (1992).

[16] A. Chambolle & T. Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, Journal of Mathematical Imaging and Vision **40** (2011).

[17] A. Beck & M. Teboulle, *A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems*, SIAM Journal on Imaging Sciences **2** (2009).

[18] T. Valkonen & T. Pock, *Acceleration of the PDHGM on partially strongly convex functions*, submitted, 2015, arXiv:1511.06566.

[19] T. Valkonen, *Optimising Big Images*, in: *Big Data Optimization: Recent Developments and Challenges*, Studies in Big Data, Springer-Verlag, 2016, 97–131.

[20] S. Setzer, *Operator Splittings, Bregman Methods and Frame Shrinkage in Image Processing*, International Journal of Computer Vision **92** (2011).

[21] E. Esser, X. Zhang & T. F. Chan, *A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science*, SIAM Journal on Imaging Sciences **3** (2010).

[22] C. Clason & T. Valkonen, *Stability of saddle points via explicit coderivatives of pointwise subdifferentials*, Set-valued and Variational Analysis (2016).