

Metric Spaces and Calculus

Nicola Pagani

December 2022

Contents

1	Introduction to metric spaces	5
2	Continuity and convergence in metric spaces	35
3	The topology of metric spaces	51
4	Completeness and compactness	71
5	Spaces of continuous functions	95
6	Introduction to multi-variable calculus	123
7	The inverse function theorem and the implicit function theorem	157

1.

Introduction to metric spaces

Important notions to learn in this Chapter:

1. The definition of a metric space.
2. The notion of subspace metric, and of product of metric spaces.
3. The example of the *discrete* distance on any set.
4. The example of the distance d_p on \mathbb{R}^n , for $p = 1, 2, \dots$ and for $p = \infty$.
5. The examples of the metric spaces (of sequences) (ℓ^p, d_p) for $p = 1, 2, \dots$ and for $p = \infty$.
6. The example of the space of continuous functions $C[0, 1]$ endowed with the distance d_{L^1} , and with the distance d_{L^∞} .
7. The notion of an *isometry* of metric spaces, and that of *isometric* metric spaces.

In this chapter we set the scene for the main characters of this module, the notion of a *distance* and that of a *metric space*. These notions were first singled out in the first decade of the XX century and are by now ubiquitous in Mathematics.

The idea is to have a set enriched with a tool (a distance) to measure how far apart are any two of its elements. This tool needs to satisfy some basic axioms, like the fact that if two points have zero distance, then they must coincide. The main axiom of a distance is the *triangle inequality*, which

in a nutshell says that it always takes more to go from a point A to a point B passing for a third point C , then it takes to go from A to B directly.

The main examples that will be discussed are the different metrics on \mathbb{R}^n , the space $C[0, 1]$ of continuous functions $[0, 1] \rightarrow \mathbb{R}$, and the spaces of sequences ℓ^p . The main technical tool to prove that the distances d_p that we will introduce on the sets \mathbb{R}^n and ℓ^p satisfy the triangle inequality is the so-called *Minkowski inequality*. The case $p = 2$ is the usual Euclidean distance d_2 familiar from Year 1 (the Minkowski inequality in case $p = 2$ takes the name of *Cauchy–Schwarz inequality*).

In terms of the general theory of Metric Spaces, in this chapter we will first see the definition and then discuss how to form metric spaces from subsets and as products. Finally, we will discuss when two metric spaces can be considered equivalent (or isometric).

In later chapters we shall see that most of the real analysis that was introduced in Year 1 (convergence of sequences, continuity of functions, ...) can be recast more generally for arbitrary metric spaces.

1.1. Starting out

Let's begin with a guiding example. Let:

- X = set of cities in Great Britain,
- Point A = Liverpool,
- Point B = London.

Then we have several different notions of distance from A to B , namely:

1. The length of the shortest path from A to B ,
2. The length of the shortest road that connects A to B ,
3. The length of the shortest railway that connects A to B ,
4. The cheapest possible train ticket that allows one to travel from A to B (at a given time).

Depending on the particular context, it may be convenient to consider different notions of distance. For example, if we intend to walk from A to B , then the first distance would be the one that we would matter. On the

1.1 Starting out

other hand, if we are driving from A to B we would only care about the second, etc.

In the example above we highlighted what is a distance of two given points A and B of X. Note that in all examples above the distance is a real number, and that it is never negative. The notion of a distance on X should give the ability to produce one such number for *all* pairs of elements of X, not only A and B (for example, we may be interested in establishing the distance from A to C =Edinburgh, to D =Cardiff, or also to establish the distance between B and C etc.).

The notion of a *distance* (or a metric) on X will then be the assignment of a (non-negative) real number to *all* pairs of elements of X. Here comes the formal definition, which will accompany us throughout this course.

Definition 1.1.1. — Let X be a set. A *distance* on X is a function

$$d: X \times X \rightarrow \mathbb{R}_{\geq 0}$$

(to the non-negative real numbers) that satisfies the following axioms

(M₁) for every $x, y \in X$, $d(x, y) = 0$ if and only if $x = y$,

(M₂) for every $x, y \in X$, $d(x, y) = d(y, x)$ (called *symmetry*),

(M₃) for every $x, y, z \in X$, $d(x, z) \leq d(x, y) + d(y, z)$ (called the *triangle inequality*).

Note that we have initially required a distance to be a *non-negative* real number, this is spelled out as a separate axiom (M₀) in some textbooks.

Definition 1.1.2. — A *metric space* is a pair (X, d) , where X is a set and d is a distance on X.

Aside: Recall that if A and B are two sets, the *cartesian product* $A \times B$ is defined to be the set

$$X \times Y := \{(x, y) : x \in X, y \in Y\}$$

of ordered pairs such that the first element of the pair belongs to X and the second element of the pair belongs to Y.

Example 1.1.3. — If

$$X = \{\text{Dog, Cat, Horse}\}, Y = \{1, 2\},$$

then

$$X \times Y = \{(\text{Dog}, 1), (\text{Cat}, 1), (\text{Horse}, 1), (\text{Dog}, 2), (\text{Cat}, 2), (\text{Horse}, 2)\}.$$

Definition 1.1.4. — The number of elements in a finite set X is denoted by $\#X$. (In fact, the symbol $\#$ can be extended to infinite sets, but this will not be needed in this course).

Note that by definition $\#(X \times Y) = \#X \times \#Y$.

Now let \mathbb{R} be the set of real numbers. We call:

- $\mathbb{R} \times \mathbb{R} = \mathbb{R}^2$ the *real (2-dimensional) plane*. Elements $x \in \mathbb{R}^2$ are vectors of real numbers that we write as $x = (x_1, x_2)$.
- $\mathbb{R} \times \mathbb{R} \times \mathbb{R} = \mathbb{R}^3$ the *real (3-dimensional) space*. Similarly, we write $x = (x_1, x_2, x_3) \in \mathbb{R}^3$.
- $\underbrace{\mathbb{R} \times \dots \times \mathbb{R}}_{n\text{-times}} = \mathbb{R}^n$ the *real n -space*. We write $x = (x_1, \dots, x_n) \in \mathbb{R}^n$.

Let's now see the first example of a distance, which should be familiar from Year 1.

Example 1.1.5. — Let $X = \mathbb{R}$ and define the function

$$d_1(x, y) := |x - y|.$$

(Recall the *absolute value* $|z|$ of some $z \in \mathbb{R}$ is defined by $|z| := \max(z, -z)$).

We now show that d_1 is a distance on \mathbb{R} or, in other words, that (\mathbb{R}, d_1) is a metric space. We start out by observing that for all $x, y \in \mathbb{R}$ we have indeed that $d_1(x, y) = |x - y| \in \mathbb{R}_{\geq 0}$.

(M₁): Suppose $x = y$. Then $|x - y| = |0| = 0$. Similarly, if $d_1(x, y) = 0$, then $|x - y| = 0$ means that $x - y = 0$, which is equivalent to $x = y$.

(M₂): We have $d_1(x, y) = |x - y| = |y - x| = d_1(y, x)$.

(M₃): To prove the triangle inequality, we will use the following lemma:

Lemma 1.1.6. — For all $a, b, c, d \in \mathbb{R}$, we have

$$\max(a + b, c + d) \leq \max(a, c) + \max(b, d)$$

1.1 Starting out

Proof. We have that

$$a \leq \max(a, c), \quad b \leq \max(b, d), \quad c \leq \max(a, c), \quad d \leq \max(b, d)$$

by definition of the function \max .

This implies that

$$a + b \leq \max(a, c) + \max(b, d),$$

and similarly for $c + d$.

Therefore we have that

$$\max(a + b, c + d) \leq \max(a, c) + \max(b, d).$$

□

Using Lemma 1.1.6, we can now complete the proof that the function $d_1: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ of the example is actually a distance.

To prove the triangle inequality, take in the above lemma

$$\begin{aligned} A &= x - y & B &= y - z \\ C &= y - x & D &= z - y \end{aligned}$$

We have then


$$\begin{aligned} d(x, z) &= |x - z| = \max(x - z, z - x) \\ &\leq \max(x - y, y - x) + \max(y - z, z - y) \\ &= |x - y| + |y - z| \\ &= d(x, y) + d(y, z). \end{aligned}$$

This concludes our proof that d_1 is a distance on \mathbb{R} .

Let us move on to more examples, and try to think of distances on \mathbb{R}^2 . Of course the first thing that comes to mind is the usual (or Euclidean) distance, which is the length of the line that connects any two points. This should of course be a distance in the sense of our Definition 1.1.1 (otherwise the theory that we are introducing would be quite strange!). We will analyse this distance later, as proving that it satisfies the axioms of a distance is not that easy.

Example 1.1.7. — For now let's take $X = \mathbb{R}^2$, and define $d_1: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ by the formula $d_1(x, y) := |x_1 - y_1| + |x_2 - y_2|$. We will now see that this function defines a distance on \mathbb{R}^2 that is also known as the *Manhattan distance*. (See Figure below. You should imagine the white squares as

skyscrapers seen from above, and the black lines as streets. The thick black path is one of the shortest paths connecting A to B, and its total length is the distance $d_1(A, B)$.



Proposition 1.1.8. — *The function d_1 is a distance on \mathbb{R}^2 .*

Proof. Before we verify the three properties we observe that

$$d_1(x, y) = |x_1 - y_1| + |x_2 - y_2| \geq 0,$$

so d_1 is indeed a function to $\mathbb{R}_{\geq 0}$ as required by the definition.

(M₁):

$$\begin{aligned} d_1(x, y) = 0 &\iff \underbrace{|x_1 - y_1|}_{\geq 0} + \underbrace{|x_2 - y_2|}_{\geq 0} = 0 \\ &\iff |x_1 - y_1| = |x_2 - y_2| = 0 \\ &\iff x_1 = y_1, x_2 = y_2 \\ &\iff x = y. \end{aligned}$$

(M₂):

$$\begin{aligned} d_1(x, y) &= |x_1 - y_1| + |x_2 - y_2| \\ &= |y_1 - x_1| + |y_2 - x_2| \\ &= d_1(y, x). \end{aligned}$$

(M₃):

$$\begin{aligned} d_1(x, z) &= |x_1 - z_1| + |x_2 - z_2| \\ &\leq |x_1 - y_1| + |y_1 - z_1| + |x_2 - y_2| + |y_2 - z_2| \\ &= d_1(x, y) + d_1(y, z). \end{aligned}$$

□

1.1 Starting out

In fact, the above result generalizes to the case where $x, y \in \mathbb{R}^n$ for every n — define

$$d_1(x, y) := |x_1 - y_1| + |x_2 - y_2| + \dots + |x_n - y_n| = \sum_{i=1}^n |x_i - y_i|. \quad (1.1)$$

The proof is essentially the same as the proof of Proposition 1.1.8, but we omit it here as the result will be later subsumed as a particular case of Theorem 1.1.15.

As promised, it is now about time that we analyse the case of the “usual” distance.

Example 1.1.9. — Consider now the case $X = \mathbb{R}^2$, and let d_2 be defined by

$$d_2(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2}$$

This is the “usual”, “Euclidean” distance, which we can interpret as the length of the segment from x to y in the plane. As we announced before, we now show that this is indeed a distance on \mathbb{R}^2 (but this will be a bit more difficult than the corresponding result for the function d_1 , which we have just discussed).

Proposition 1.1.10. — *The function d_2 is a distance on \mathbb{R}^2 .*

Proof. As usual, we begin by observing that $d_2(x, y) \in \mathbb{R}_{\geq 0}$.

(M₁): (As before — omitted).

(M₂): (As before — omitted).

(M₃): For convenience, define

$$a_1 = x_1 - y_1, \quad a_2 = x_2 - y_2, \quad b_1 = y_1 - z_1, \quad b_2 = y_2 - z_2.$$

With this notation, we have

$$d_2(x, z)^2 = (a_1 + b_1)^2 + (a_2 + b_2)^2 = a_1^2 + a_2^2 + b_1^2 + b_2^2 + 2(a_1b_1 + a_2b_2) \quad (1.2)$$

We will use the Cauchy–Schwarz inequality (with $n = 2$), which we will prove in Lemma 1.1.11, which states

$$a_1b_1 + a_2b_2 \leq \sqrt{a_1^2 + a_2^2} \cdot \sqrt{b_1^2 + b_2^2} \quad (1.3)$$

for all $a_1, a_2, b_1, b_2 \in \mathbb{R}$.

Plugging Inequality (1.3) into (1.2), we deduce

$$\begin{aligned} d_2(x, z)^2 &\leq a_1^2 + a_2^2 + b_1^2 + b_2^2 + 2\sqrt{a_1^2 + a_2^2} \cdot \sqrt{b_1^2 + b_2^2} \\ &= \left(\sqrt{a_1^2 + a_2^2} + \sqrt{b_1^2 + b_2^2} \right)^2 \\ &= (d_2(x, y) + d_2(y, z))^2. \end{aligned}$$

Now taking the square root of the left hand side and of the right hand side of this inequality, and using that the real numbers $d_2(x, z)$, $d_2(x, y)$ and $d_2(y, z)$ are all non-negative, we find

$$d_2(x, z) \leq d_2(x, y) + d_2(y, z),$$

which concludes our proof. □

In the proof we have used the following lemma:

Lemma 1.1.11 (Cauchy-Schwarz Inequality). — *For every collection of numbers*

$$a_1, \dots, a_n, b_1, \dots, b_n \in \mathbb{R}$$

the following inequality holds

$$\left(\sum_{i=1}^n a_i b_i \right)^2 \leq \sum_{i=1}^n a_i^2 \cdot \sum_{i=1}^n b_i^2. \quad (1.4)$$

(What we used above was this lemma for $n = 2$, but we will soon use this Lemma for arbitrary $n \in \mathbb{N}$).

Proof. Let $F: \mathbb{R} \rightarrow \mathbb{R}$ be defined by

$$t \mapsto F(t) = \sum_{i=1}^n (ta_i + b_i)^2 = At^2 + Ht + B,$$

where we have set


$$A = \sum_{i=1}^n a_i^2, \quad H = 2 \sum_{i=1}^n a_i b_i \quad \text{and} \quad B = \sum_{i=1}^n b_i^2.$$

1.1 Starting out

For fixed A, H and B , $F(t)$ is a polynomial of degree 2 in the variable t .

Because $F(t)$ is a sum of squares, it must be greater than or equal to zero for all $t \in \mathbb{R}$. This means that the polynomial equation $F(t) = 0$ never has more than 1 (real) solution. This implies that its discriminant, which equals $H^2 - 4AB$, must be smaller than or equal to zero. Expanding the inequality $H^2 - 4AB \leq 0$ by replacing the values of H, A, B we precisely obtain Inequality (1.4). \square

Remark 1.1.12. — In the proof of Lemma 1.1.11 we used some properties of a polynomial function F of degree 2. The figure below shows the graph of any such function F defined by $F(s) = As^2 + Hs + B$ in the three possible scenarios: from left to right, when the discriminant is positive, when it is zero, and when it is negative.



Similarly to the case of d_1 , also the distance d_2 can be generalised to the case of $X = \mathbb{R}^n$, where it is defined by

$$d_2(x, y) = \sqrt{(x_1 - y_1)^2 + \dots + (x_n - y_n)^2} = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}. \quad (1.5)$$

That this formula actually defines a distance can be proved as for the case of \mathbb{R}^2 , by using the Cauchy–Schwarz inequality for arbitrary n (exactly the statement of Lemma 1.1.11). This will again be a particular case of Theorem 1.1.15.

Note that in the case where $n = 1$, we have

$$d_2(x, y) = \sqrt{(x - y)^2} = |x - y| = d_1(x, y),$$

whence on $\mathbb{R} = \mathbb{R}^1$ we have that $d_1 = d_2$.

Let's continue and consider more examples of distances.

Example 1.1.13. — Consider the case of \mathbb{R}^2 again, with the distance d_∞ defined by

$$d_\infty(x, y) = \max(|x_1 - y_1|, |x_2 - y_2|).$$

Proposition 1.1.14. — *The function d_∞ is a distance on \mathbb{R}^2 .*

Proof. (M₁): As before.

(M₂): As before.

(M₃):

$$\begin{aligned} d_\infty(x, z) &= \max(|x_1 - z_1|, |x_2 - z_2|) \\ &\leq \underbrace{\max(|x_1 - y_1| + |y_1 - z_1|, |x_2 - y_2| + |y_2 - z_2|)}_{\text{by the triangle inequality for } d_1 \text{ on } \mathbb{R}} \\ &\leq \underbrace{\max(|x_1 - y_1|, |x_2 - y_2|) + \max(|y_1 - z_1|, |y_2 - z_2|)}_{\text{by Lemma 1.1.6}} \\ &= d_\infty(x, y) + d_\infty(y, z) \end{aligned}$$

□

Again, one can prove similarly that d_∞ defined by

$$d_\infty(x, y) = \max_{1 \leq i \leq n} |x_i - y_i|$$

is a distance on \mathbb{R}^n .

We will now recollect the examples given above (except for d_∞) and generalise them to a unique result. Let $n, p \geq 1$ be integers, and consider the distance d_p defined on \mathbb{R}^n by the formula

$$d_p(x, y) = \left(\sum_{i=1}^n |x_i - y_i|^p \right)^{\frac{1}{p}}.$$

Theorem 1.1.15. — *The function d_p defines a distance on \mathbb{R}^n .*

Note that the cases $p = 1, 2$ were discussed in (1.1), (1.5) respectively (and the more specific cases when furthermore $n = 2$ are detailed in Propositions 1.1.8 and 1.1.10 respectively).

Proof. From its very definition, it is evident that $d_p(x, y) \geq 0$ for all $x, y \in \mathbb{R}^n$.

1.1 Starting out

(M₁): As before.

(M₂): As before.

(M₃): Define the vectors in \mathbb{R}^n

$$(a_1, \dots, a_n) = (x_1 - y_1, \dots, x_n - y_n)$$

and

$$(b_1, \dots, b_n) = (y_1 - z_1, \dots, y_n - z_n)$$

We need to prove that

$$\left(\sum_{i=1}^n |a_i + b_i|^p \right)^{\frac{1}{p}} \leq \left(\sum_{i=1}^n |a_i|^p \right)^{\frac{1}{p}} + \left(\sum_{i=1}^n |b_i|^p \right)^{\frac{1}{p}} \quad (1.6)$$

This is the *Minkowski Inequality*, which we will prove later in Theorem 1.3.10.

(Note: Theorem 1.3.10 proves the inequality in the more general case of real *sequences*, when the sum is infinite. Formula (1.6) can be obtained from Theorem 1.3.10 by taking

$$A = (a_1, \dots, a_n, 0, 0, 0, \dots)$$

and

$$B = (b_1, \dots, b_n, 0, 0, 0, \dots),$$

i.e. a n -dimensional vector can be made into a sequence (which in the language of this module is always an infinite list indexed by the natural numbers) by "prolonging" it with infinitely many zeroes.)

□

Next, we observe how the various different distances d_p are related each to the other. We leave this as an exercise.

Exercise 1.1.16. — For fixed $x, y \in \mathbb{R}^n$, prove that

$$d_p(x, y) \leq d_q(x, y)$$

for all $p \geq q$. (Don't worry too much if you cannot do this, it's not that easy! You will be guided into solving this exercise step-by-step in a homework assignment).

Then use the previous paragraph to prove that

$$\lim_{p \rightarrow \infty} d_p(x, y) = d_\infty(x, y).$$

[Note that here we are regarding $(d_p(x, y))_{p \in \mathbb{N}}$ as a sequence of real numbers indexed by $p \in \mathbb{N}$. This is then an exercise of convergence of sequences of real numbers as in Year 1.]

The examples of distances that we have given so far were all on \mathbb{R}^n for some n , and they were all given by some rather complicated formulas. Let's step back for a moment, and try to give the simplest possible example of a distance that satisfies the required axioms, and that we can define on an arbitrary set.

Example 1.1.17. — Let X be any set. Define the *discrete distance* d_{discr} by

$$d_{\text{discr}}(x, y) = \begin{cases} 0 & \text{when } x = y \\ 1 & \text{when } x \neq y \end{cases}$$

for all $x, y \in X$.

We show that d_{discr} is indeed a distance. One sees from the formula that $d(x, y) = 0$ if and only if $x = y$. Moreover, the formula is symmetric in x and y .

For the third condition, if $x = z$, we see that $d_{\text{discr}}(x, z) = 0$ and there is nothing to prove. If $d_{\text{discr}} = 1$, we have that $x \neq z$, so that $y \neq z$ or $y \neq x$. This implies that $d_{\text{discr}}(x, y) + d_{\text{discr}}(y, z)$ equals 1 or 2, in particular $d_{\text{discr}}(x, y) + d_{\text{discr}}(y, z) \geq 1$, so $d_{\text{discr}}(x, y) + d_{\text{discr}}(y, z) \geq d_{\text{discr}}(x, z)$, which is what we wanted to show.

1.2. Subspace distance

In this short section we start addressing the problem of how to create a new metric space from a given metric space. If we start from a metric space and consider a subset of it, then we can just measure distances on the subset using the global distance defined on the larger metric space. This will also work and define a distance on the subset. Let's see this a bit more formally.

Assume that we are given a set X and a subset $Y \subseteq X$, and suppose that we have a distance d_X on X . So d_X is a function

$$d_X: X \times X \rightarrow \mathbb{R}_{\geq 0}.$$

1.3 Spaces of real sequences

The restriction of d_X to the subset $Y \times Y \subseteq X \times X$, is a new function

$$d_Y := d_{X|Y \times Y} : Y \times Y \rightarrow \mathbb{R}_{\geq 0}$$

known as the *subspace distance*. The three axioms that guarantee that d_X is a distance on X imply one by one the three axioms that guarantee that d_Y is a distance on Y .

Example 1.2.1. — The distances d_1, d_2 and d_∞ , which we defined on \mathbb{R}^n , also define distances on subsets of \mathbb{R}^n . For example, it makes sense to speak of the distance d_2 on the subsets

$$Y_1 := \{(x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 < 1\} \subseteq \mathbb{R}^2 = X,$$

$$Y_2 := \{(x_1, x_2) \in \mathbb{R}^2 : x_1 > 0, x_2 < 0\} \subseteq \mathbb{R}^2 = X,$$

$$Y_3 := \{(x_1, x_2) \in \mathbb{R}^2 : x_1 = 0, x_2 < 1\} \subseteq \mathbb{R}^2 = X.$$

1.3. Spaces of real sequences

So far all examples of metric spaces we have given, except for the discrete distance, are on \mathbb{R}^n for some n , or on its subsets. Here we consider some more interesting, “infinite dimensional”, examples. We start by discussing spaces whose *elements* are (infinite) sequences of real numbers. Note that a vector $x \in \mathbb{R}^n$ can be regarded to as a *finite* sequence of real numbers of length n . To avoid confusions, in this module we will *ban* this casual use of the common language word “sequence”, and use that word only for infinite sequences, as in the following definition.

Definition 1.3.1. — A *sequence of real numbers* (or a *real sequence*) is a function

$$A: \mathbb{N} \rightarrow \mathbb{R}$$

We will usually write the sequence as $A = (A_0, A_1, \dots, A_n, \dots) = (A_n)_{n \in \mathbb{N}}$, or more compactly as (A_n) when no confusion is likely to arise on the set of indices.

(Note that we conventionally choose the natural number to contain 0. This is sometimes annoying when considering sequences, for example the sequence $(x_n) = (1/n)_{n \in \mathbb{N}}$ doesn't quite make sense with this notation, because $1/0$ is not defined).

Example 1.3.2. — The following are all sequences of real numbers

$$A = (1, 4, 9, 16, 25, \dots, (n+1)^2, \dots)$$

$$B = (0, 0, 0, \dots, 0, \dots)$$

$$C = (\pi, -\pi, \pi, -\pi, \dots, (-1)^n \pi, \dots)$$

$$D = (1, 1/2, 1/3, \dots, 1/(n+1), \dots)$$

$$E = (1, 1/4, 1/9, \dots, 1/(n+1)^2, \dots)$$

The purpose of this section is to extend our definition of distance d_p to spaces whose elements are sequences. By analogy with the formulas for d_p on \mathbb{R}^n , for two sequences of real numbers $A = (A_n)$ and $B = (B_n)$, we may attempt to define d_p by

$$d_p(A, B) := \left(\sum_{n=0}^{\infty} |A_n - B_n|^p \right)^{\frac{1}{p}} = \lim_{N \rightarrow \infty} \left(\sum_{n=0}^N |A_n - B_n|^p \right)^{\frac{1}{p}} \quad (1.7)$$

and

$$d_{\infty}((A_n), (B_n)) = \sup_{n \in \mathbb{N}} |A_n - B_n|. \quad (1.8)$$

We need some work to prove that these definitions actually produce distances. In fact, none of these functions d_p gives a well-defined distance on the set of *all* sequences. For example, for the sequences A and B of Example 1.3.2, the formula $d_p(A, B)$ gives $+\infty$ for all $p = 1, 2, \dots, \infty$, and $+\infty$ is not even a real number! By definition, a distance is required to be a non-negative real number, which $+\infty$ is not.

Definition 1.3.3. — For $p \geq 1$, we let

$$\ell^p := \left\{ (A_n) \text{ real sequence such that } \sum_{n=0}^{\infty} |A_n|^p < \infty \right\}$$

For $p = \infty$ we set

$$\ell^{\infty} := \{(A_n) \text{ real sequence such that } (A_n) \text{ is bounded}\}$$

(We say a sequence (A_n) is bounded if there exists $M \in \mathbb{R}$ such that $|A_n| \leq M$ for all $n \in \mathbb{N}$.)

1.3 Spaces of real sequences

For example, it is well-known from Year 1 that the (harmonic) sequence $(D_n) = \frac{1}{n+1}$ is *not* an element of ℓ^1 as the sum

$$\sum \frac{1}{n+1} = \infty.$$

On the other hand, $(E_n) = \frac{1}{(n+1)^2}$ is in ℓ^1 , and for the same reason (D_n) is in ℓ^2 . Both sequences are in ℓ^∞ .

(As for the other examples mentioned in 1.3.2, the sequence A is not in ℓ^p for any p nor it is in ℓ^∞ , the sequence B is in ℓ^p for all p including $p = \infty$, and C is in ℓ^p only for $p = \infty$).

Theorem 1.3.4. — For all $p \geq 1$ (and also for $p = \infty$) the function d_p defines a distance on the set ℓ^p .

Proof. (Of Theorem 1.3.4) We start by proving that $d_p(A, B)$ is a nonnegative real number for all $A, B \in \ell^p$. Observe that, by its very definition, $d_p(A, B) \in \mathbb{R}_{\geq 0} \cup \{+\infty\}$, so we only have to rule out the possibility that $d_p(A, B)$ equals $+\infty$.

Case $p < \infty$. The Minkowski inequality, which we will prove in Theorem 1.3.10, reads

$$\left(\sum_{n=0}^{\infty} |A_n + B_n|^p \right)^{\frac{1}{p}} \leq \left(\sum_{n=0}^{\infty} |A_n|^p \right)^{\frac{1}{p}} + \left(\sum_{n=0}^{\infty} |B_n|^p \right)^{\frac{1}{p}}. \quad (1.9)$$

The right hand side is $< \infty$ because $(A_n) \in \ell^p$ and because $(B_n) \in \ell^p$. Therefore the left hand side must also be $< \infty$.

Case $p = \infty$. Here $d_\infty(A, B)$ is not $+\infty$ for $(A_n), (B_n) \in \ell^\infty$, because if both sequences (A_n) and (B_n) are bounded, then so is $(A_n - B_n)$.

Properties M1 and M2 are proven as usual, and we omit the details here. As for M3, let's take $C = (C_n), D = (D_n)$ and $E = (E_n)$ in ℓ^p .

Consider first the case $p < \infty$. Define the two real sequences

$$(A_n) = (C_n - D_n) \text{ and } (B_n) = (D_n - E_n).$$

Then the Minkowski inequality, Equation 1.9, can be rewritten as

$$d_p(C, E) \leq d_p(C, D) + d_p(D, E),$$

which is the triangle inequality for d_p on ℓ^p .

When $p = \infty$, the triangle inequality for the distance d_∞ on \mathbb{R}^n , applied to the first N entries of the sequences C, D and E , reads

$$\max_{0 \leq n \leq N-1} |C_n - E_n| \leq \max_{0 \leq n \leq N-1} |C_n - D_n| + \max_{0 \leq n \leq N-1} |D_n - E_n|.$$

Taking the limit for $N \mapsto \infty$ of the left hand side and of the right hand side (such limits exist because both the left hand side and the right hand side are non-decreasing real sequences indexed by the natural number N), gives

$$\sup_{n \in \mathbb{N}} |C_n - E_n| \leq \sup_{n \in \mathbb{N}} |C_n - D_n| + \sup_{n \in \mathbb{N}} |D_n - E_n|.$$

This is the triangle inequality for d_∞ on ℓ^∞ , which concludes our proof. \square

We now show that the set ℓ^p can also be endowed with the metric d_q for all $q \geq p$, including the case when $q = \infty$.

Proposition 1.3.5. — *For $p \leq q$, the inclusion $\ell^p \subseteq \ell^q$ holds (and this inclusion also holds when $q = \infty$).*

Proof. Suppose $(A_n) \in \ell^p$. This means that $\sum_{n=0}^{\infty} |A_n|^p < \infty$, which implies that $\lim_{n \rightarrow \infty} |A_n| = 0$. Therefore there is some N such that for all $n > N$ we have $|A_n| < 1$. It follows that $|A_n|^q \leq |A_n|^p$ whenever $n > N$. Taking the sum over $n > N$ gives us

$$\sum_{n=N+1}^{\infty} |A_n|^q \leq \sum_{n=N+1}^{\infty} |A_n|^p < \infty,$$

which implies that $\sum_{n=0}^{\infty} |A_n|^q < \infty$.

For the case when $q = \infty$, recall from Year 1 the implication

$$\sum_{n=0}^{\infty} |A_n|^p < \infty \Rightarrow \lim_{n \rightarrow \infty} A_n = 0,$$

which, in particular, implies that the sequence (A_n) is bounded. \square

From the construction of subspace distance that we discussed in Section 1.2 we deduce:

Corollary 1.3.6. — *(ℓ^p, d_q) is a metric space (a subspace metric of (ℓ^q, d_q)) whenever $p \leq q$.*

1.3 Spaces of real sequences

Example 1.3.7. — Let $x \in \mathbb{R}$. The geometric sequence $(x^n)_{n \in \mathbb{N}}$ of ratio x belongs to ℓ^p for all $p \geq 1$ if and only if $|x| < 1$.

Indeed, for $x \in \mathbb{R}$, we have that

$$\sum_{n=0}^{\infty} |x|^n = \frac{1}{1 - |x|}, \quad \text{when } |x| < 1.$$

Therefore the sequence $(A_n)_{n \in \mathbb{N}}$ where each element is defined by $A_n = x^n$ belongs to $\ell^1, \ell^2, \dots, \ell^\infty$ when $|x| < 1$, it does not belong to ℓ^1, ℓ^2, \dots and ℓ^∞ when $|x| > 1$. Finally, for $|x| = 1$ the sequence (A_n) does belong to ℓ^∞ but it does not belong to ℓ^p for any other $p < \infty$.

Exercise 1.3.8. — Let $(A_n) = (4^{-n}) = (1, \frac{1}{4}, \frac{1}{16}, \dots)$, $(B_n) = (0, 1, 0, 0, \dots)$. Prove that $(A_n), (B_n) \in \ell^1, \ell^2, \ell^\infty$. Compute the distances d_1, d_2, d_∞ between (A_n) and (B_n) .

1.3.9. The Minkowski Inequality — (The essential inequality that we used to prove that d_p satisfies the triangle inequality on ℓ^p (and, in particular, also on \mathbb{R}^n)).

Theorem 1.3.10 (The Minkowski Inequality). — Let $(A_n)_{n \in \mathbb{N}}$ and $(B_n)_{n \in \mathbb{N}}$ be two real sequences and let $p \geq 1$. Then the inequality

$$\left(\sum_{n=0}^{\infty} |A_n + B_n|^p \right)^{\frac{1}{p}} \leq \left(\sum_{n=0}^{\infty} |A_n|^p \right)^{\frac{1}{p}} + \left(\sum_{n=0}^{\infty} |B_n|^p \right)^{\frac{1}{p}}. \quad (1.10)$$

holds.

In the following example we discuss the fact that the Minkowski Inequality can be seen as a generalisation of the Cauchy-Schwarz inequality.

Example 1.3.11. — For $p = 2$, the Minkowski Inequality (1.10) becomes:

$$\sqrt{\sum_{n=0}^{\infty} (A_n)^2 + \sum_{n=0}^{\infty} (B_n)^2 + 2 \sum_{n=0}^{\infty} A_n B_n} \leq \sqrt{\sum_{n=0}^{\infty} (A_n)^2} + \sqrt{\sum_{n=0}^{\infty} (B_n)^2}$$

After taking squares and simplifying the expression, this reduces to:

$$\sum_{n=0}^{\infty} A_n B_n \leq \sqrt{\sum_{n=0}^{\infty} (A_n)^2 \sum_{n=0}^{\infty} (B_n)^2}$$

which could also have been obtained from the Cauchy-Schwarz inequality (Theorem 1.1.11) by taking the limit for $N \rightarrow \infty$.

We will now discuss a proof of the Minkowski inequality.

Proof. (of Theorem 1.3.10)

Define $F: [0, \infty) \rightarrow [0, \infty)$ by $F(t) := t^p$. We will use the fact that the function F is both increasing and convex, two facts that should be well-known from Year 1. (The first is proven by showing that, for all $x \in [0, \infty)$, the first derivative of F is > 0 and the second by showing the same for the second derivative of F).

Since F is convex on $[0, \infty)$, for every $t, t' \geq 0$ we have

$$F(\alpha t + \beta t') \leq \alpha F(t) + \beta F(t')$$

for all $0 \leq \alpha, \beta \leq 1$ such that $\alpha + \beta = 1$.

Let $A := (\sum_{n=0}^{\infty} |A_n|^p)^{\frac{1}{p}}$, $B := (\sum_{n=0}^{\infty} |B_n|^p)^{\frac{1}{p}}$. We will assume that both $A, B < \infty$ for otherwise Inequality (1.10) holds trivially. We begin by normalising the two sequences (A_n) and (B_n) by setting

$$\hat{A}_n := \frac{A_n}{A}, \quad \hat{B}_n := \frac{B_n}{B}$$

for all n , so that

$$\sum_{n=0}^{\infty} |\hat{A}_n|^p = \sum_{n=0}^{\infty} |\hat{B}_n|^p = 1.$$

Then we have, for all $n \in \mathbb{N}$,

$$\begin{aligned} |A_n + B_n|^p &\leq (|A_n| + |B_n|)^p \\ &= (A |\hat{A}_n| + B |\hat{B}_n|)^p \\ &= (A + B)^p \left(\frac{A}{A+B} |\hat{A}_n| + \frac{B}{A+B} |\hat{B}_n| \right)^p \\ &\leq (A + B)^p \underbrace{\left(\frac{A}{A+B} |\hat{A}_n|^p + \frac{B}{A+B} |\hat{B}_n|^p \right)}_{\text{by convexity}} \end{aligned}$$

since $\frac{A}{A+B} + \frac{B}{A+B} = 1$. (The first inequality occurs because F is increasing).

1.4 Spaces of functions

Summing over all n , we obtain

$$\begin{aligned}
 \sum_{n=0}^{\infty} |A_n + B_n|^p &\leq (A + B)^p \sum_{n=0}^{\infty} \left(\frac{A}{A + B} |\hat{A}_n|^p + \frac{B}{A + B} |\hat{B}_n|^p \right) \\
 &= (A + B)^p \left(\frac{A}{A + B} \underbrace{\sum_{n=0}^{\infty} |\hat{A}_n|^p}_{=1} + \frac{B}{A + B} \underbrace{\sum_{n=0}^{\infty} |\hat{B}_n|^p}_{=1} \right) \\
 &= (A + B)^p \underbrace{\left(\frac{A}{A + B} + \frac{B}{A + B} \right)}_{=1} \\
 &= (A + B)^p
 \end{aligned}$$

Taking the p -th root in the last inequality we deduce the desired result. \square

Note that the Minkowski inequality is valid for any *real* $p \in [1, \infty)$, the number p does not need to necessarily be an integer! For this reason, we could have as well defined distances d_p on \mathbb{R}^n for an arbitrary $p \in \mathbb{R}_{\geq 1} \cup \{+\infty\}$ (not just for p a natural number), and we could have done the same for the spaces of sequences (ℓ^p, d_p) .

1.4. Spaces of functions

In this section we introduce more examples of “infinite dimensional” metric spaces. We will discuss examples of metric spaces whose elements are functions. We are just scratching the surface of an important and beautiful theory, called *functional analysis*, which develops the tools of real analysis (limits, derivatives etc.) for sets whose elements are functions (as opposed to sets of numbers, such as \mathbb{R} or vectors, such as \mathbb{R}^n). The main motivation is solving ordinary and partial differential equations. We will see a little bit of this in a later chapter of these notes.

Our prototype example of a space of functions will be the following.

Definition 1.4.1. — We define $C[0, 1]$ to be the set

$$C[0, 1] := \{f: [0, 1] \rightarrow \mathbb{R} \text{ such that } f \text{ is continuous}\}.$$

(The notion of continuity for a function $f: [0, 1] \rightarrow \mathbb{R}$ was defined in Year 1. We will discuss that notion more thoroughly and generalise it in Chapter 2 of these notes).

What should a distance $d(f, g)$ on $C[0, 1]$ look like? We will take inspiration from the cases of the distances d_p that we introduced in the previous sections. We will only describe the cases of $p = 1, \infty$ (but one could similarly discuss the cases of other $p \in \mathbb{R}_{\geq 1}$).

Definition 1.4.2. — We define a distance d_{L^1} on $C[0, 1]$ by

$$d_{L^1}(f, g) = \int_0^1 |f(x) - g(x)| \, dx.$$


Similarly we define a distance d_{L^∞} by

$$d_{L^\infty}(f, g) = \max_{x \in [0, 1]} |f(x) - g(x)|.$$

Remark 1.4.3. — We know from Year 1 that a continuous function

$$h: [0, 1] \rightarrow \mathbb{R}$$

always has a maximum and that its integral exists and it is a real number. This, together with the fact that the function $|f - g|$ is both continuous and nonnegative, implies that d_{L^1} and d_{L^∞} are both nonnegative real numbers.



Theorem 1.4.4. — d_{L^∞} is a distance on $C[0, 1]$.

Proof. (M₁)

$$\begin{aligned} d_{L^\infty}(f, g) = 0 &\iff \max_{x \in [0, 1]} |f(x) - g(x)| = 0 \\ &\iff |f(x) - g(x)| = 0 \, \forall x \in [0, 1] \\ &\iff f(x) = g(x) \, \forall x \in [0, 1] \\ &\iff f = g \end{aligned}$$

(M₂) $|f(x) - g(x)| = |g(x) - f(x)|$ is true for every x .

1.4 Spaces of functions

(M₃) For all functions $f, g, h \in C[0, 1]$, we have

$$\begin{aligned}
 d_{L^\infty}(f, h) &= \max_{x \in [0, 1]} |f(x) - h(x)| \\
 &= |f(x') - h(x')| \text{ for some } x' \in [0, 1] \\
 &\leq |f(x') - g(x')| + |g(x') - h(x')| \\
 &\leq \max_{x \in [0, 1]} |f(x) - g(x)| + \max_{x \in [0, 1]} |g(x) - h(x)| \\
 &= d_{L^\infty}(f, g) + d_{L^\infty}(g, h)
 \end{aligned}$$

□

For the case of d_{L^1} , we will need a lemma:

Lemma 1.4.5. — Let $h: [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ be a continuous function. Then:

$$\int_0^1 h(x) \, dx = 0 \Rightarrow h = 0.$$

For the first time here we are directly using the notion of *continuity*, as learned in MATH101 - Calculus I. The notion will be reviewed and discussed extensively later in Chapter 2 of these notes.

Proof. If $h \neq 0$ then there exists an $x' \in [0, 1]$ such that $h(x') > 0$. The function h is continuous at x' . By MATH101, this means that for all sequences (x_n) that converge to x' , we have that the sequence $(h(x_n))$ converges to $h(x')$.

Claim: there exists an interval $I = (x' - \delta, x' + \delta)$ centered at x' for some $\delta > 0$ such that $h(x) > \frac{h(x')}{2}$ for all $x \in I$. **Proof of the claim:** if this wasn't the case, then for all $n > 0$ there would exist $x_n \in (x' - 1/n, x' + 1/n)$ with $h(x_n) \leq h(x')/2$. This would imply that $x_n \rightarrow x'$, but $h(x_n)$ doesn't converge to $h(x')$ (because $h(x_n) \leq h(x')/2$ for all $n \in \mathbb{N}$ and $h(x') > 0$), which contradicts the assumption that h is continuous at x' .

Using the claim, we deduce:

$$\begin{aligned}
 \int_0^1 h(x) \, dx &\geq \int_{x' - \delta}^{x' + \delta} h(x) \, dx \\
 &\geq \int_{x' - \delta}^{x' + \delta} \frac{h(x')}{2} \, dx \\
 &= 2\delta \frac{h(x')}{2} = \delta h(x') > 0.
 \end{aligned}$$

(The first inequality follows because $h(x) \geq 0$ for all x , the second from the claim, and the last step is the just integration of a constant function). □

The above Lemma is used to prove the first part of the following result.

Theorem 1.4.6. — d_{L^1} is a distance on $C[0, 1]$.

Proof. We start by pointing out that the integral of a nonnegative *continuous* function on $[0, 1]$ is always a nonnegative real number.

(M₁) First, for all $f \in C[0, 1]$ we have

$$d_{L^1}(f, f) = \int_0^1 |f(x) - f(x)| \, dx = \int_0^1 0 \, dx = 0.$$

Moreover, for all $f, g \in C[0, 1]$, we have

$$d_{L^1}(f, g) = \int_0^1 |f(x) - g(x)| \, dx = 0 \quad \Rightarrow \quad f - g = 0$$

by Lemma 1.4.5, which means that $f = g$.

(M₂) As usual (the formula for d_{L^1} is symmetric in the two entries).

(M₃) For all $f, g, h \in C[0, 1]$ we have:


$$\begin{aligned} d_{L^1}(f, h) &= \int_0^1 |f(x) - h(x)| \, dx \\ &\leq \int_0^1 |f(x) - g(x)| + |g(x) - h(x)| \, dx \\ &= \int_0^1 |f(x) - g(x)| \, dx + \int_0^1 |g(x) - h(x)| \, dx \\ &= d_{L^1}(f, g) + d_{L^1}(g, h). \end{aligned}$$

□

Example 1.4.7. — We compute the distances d_{L^1} and d_{L^∞} of f and g where f and g are defined by $f(x) = x$ and $g(x) = x^2$. The calculation is simplified in this case by the fact that $f(x) \geq g(x)$ for all x , so that $|f - g| = f - g$ on $[0, 1]$.

$$\begin{aligned} d_{L^\infty}(f, g) &= \max_{x \in [0, 1]} |x - x^2| \\ &= \max_{x \in [0, 1]} (x - x^2) \end{aligned}$$

1.5 Product metrics



Define $h = f - g$.

The maximum of $h = f - g$ on the interval $[0, 1]$, is either attained at the values x where the derivative $h'(x)$ equals zero and $h''(x) < 0$, or at the two extremes of the interval (0 and 1).

We calculate:

$$h'(x) = 1 - 2x. \text{ Therefore}$$

$$h'(x) = 0 \iff x = \frac{1}{2}$$

and $h''(x) = -2$ which is always negative. On the other hand, we have $h(0) = h(1) = 0$. We conclude that

$$d_{L^\infty}(f, g) = \max(h(0), h(1), h(1/2)) = \frac{1}{4}.$$

On the other hand,

$$\begin{aligned} d_{L^1}(f, g) &= \int_0^1 |x - x^2| dx \\ &= \int_0^1 (x - x^2) dx \\ &= \left[\frac{x^2}{2} - \frac{x^3}{3} \right]_0^1 \\ &= \frac{1}{2} - \frac{1}{3} = \frac{1}{6} \end{aligned}$$

Exercise 1.4.8. — Compute d_{L^1} and d_{L^∞} of f and g , where $f(x) = 1$ and $g(x) = 2 \sin(\pi x)$.

1.5. Product metrics

This section is an analogue of Section 1.2. Our aim is to expand on the theory of how one can define new metric spaces from existing ones.

Let X and Y be metric spaces with respective distances d_X and d_Y . In this section we define distances on the product $X \times Y$. Taking inspiration from the example $X = Y = \mathbb{R}$, we have already seen that there isn't a unique way to "naturally" define one such distance.

For $(x_1, y_1), (x_2, y_2) \in X \times Y$ we define distances D_p for $p = 1, 2, \infty$ on $X \times Y$ by the formulas:

$$D_1((x_1, y_1), (x_2, y_2)) := d_X(x_1, x_2) + d_Y(y_1, y_2)$$

$$D_2((x_1, y_1), (x_2, y_2)) := \sqrt{d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2}$$

$$D_\infty((x_1, y_1), (x_2, y_2)) := \max(d_X(x_1, x_2), d_Y(y_1, y_2))$$

As a homework exercise, you will be asked to prove that D_1 and D_∞ are distances on $X \times Y$. Here we will discuss the more difficult case of D_2 . The proof is formally very similar to the proof that d_2 is a distance on \mathbb{R}^2 .

Theorem 1.5.1. — *The function D_2 is a distance on $X \times Y$.*

Proof. (M₁)

$$\begin{aligned} D_2((x_1, y_1), (x_2, y_2)) = 0 & \\ \iff \sqrt{d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2} = 0 & \\ \iff d_X(x_1, x_2) = 0 \text{ and } d_Y(y_1, y_2) = 0 & \\ \iff (x_1, y_1) = (x_2, y_2) & \end{aligned}$$

(M₂)

$$\begin{aligned} D_2((x_1, y_1), (x_2, y_2)) &= \sqrt{d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2} \\ &= \sqrt{d_X(x_2, x_1)^2 + d_Y(y_2, y_1)^2} \\ &= D_2((x_2, y_2), (x_1, y_1)) \end{aligned}$$

(M₃) Take three arbitrary elements $(x_1, y_1), (x_2, y_2), (x_3, y_3)$ of $X \times Y$. We want to prove that the number

$$D_2((x_1, y_1), (x_3, y_3)) = \sqrt{d_X(x_1, x_3)^2 + d_Y(y_1, y_3)^2}$$

is smaller than or equal to the number

$$\begin{aligned} &D_2((x_1, y_1), (x_2, y_2)) + D_2((x_2, y_2), (x_3, y_3)) = \\ &= \sqrt{d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2} + \sqrt{d_X(x_2, x_3)^2 + d_Y(y_2, y_3)^2}. \end{aligned}$$

1.5 Product metrics

After taking squares (all distances are nonnegative), this is equivalent to proving the inequality:

$$d_X(x_1, x_3)^2 + d_Y(y_1, y_3)^2 \leq d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2 + d_X(x_2, x_3)^2 + d_Y(y_2, y_3)^2 + 2\sqrt{d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2} \sqrt{d_X(x_2, x_3)^2 + d_Y(y_2, y_3)^2}$$

Call RHS (for Right Hand Side) the right hand side of the last inequality. By the Cauchy-Schwarz inequality (1.1.11) with $n = 2$:

$$\sqrt{a_1^2 + a_2^2} \cdot \sqrt{b_1^2 + b_2^2} \geq a_1 b_1 + a_2 b_2,$$

we have:

$$\begin{aligned} \text{RHS} &\geq d_X(x_1, x_2)^2 + d_Y(y_1, y_2)^2 + d_X(x_2, x_3)^2 + d_Y(y_2, y_3)^2 \\ &\quad + 2d_X(x_1, x_2)d_X(x_2, x_3) + 2d_Y(y_1, y_2)d_Y(y_2, y_3) \\ &= (d_X(x_1, x_2) + d_X(x_2, x_3))^2 + (d_Y(y_1, y_2) + d_Y(y_2, y_3))^2 \\ &\geq d_X(x_1, x_3)^2 + d_Y(y_1, y_3)^2 \end{aligned}$$

which concludes our proof. □

In fact, the cases $p = 1, 2, \infty$ have nothing special, and one could have as well taken any $p \in \mathbb{N} \cup \{+\infty\}$ (or even $p \in \mathbb{R}_{\geq 1} \cup \{+\infty\}$).

Example 1.5.2. — Let $(X, d_X) = (\mathbb{R}^2, d_2)$ and $(Y, d_Y) = (\mathbb{R}, d_{\text{discr}})$. Let $x = ((1, 2), 3)$ and $y = ((4, 5), 6)$ be elements of $\mathbb{R}^2 \times \mathbb{R}$. Then

$$\begin{aligned} D_\infty(d_2, d_{\text{discr}})(x, y) &= \max(d_2((1, 2), (4, 5)), d_{\text{discr}}(3, 6)) \\ &= \max(\sqrt{3^2 + 3^2}, 1) \\ &= 3\sqrt{2} \end{aligned}$$

The fact that, unlike for the case of subspace distances (see Section 1.2), the definition of a product distance isn't natural (it depends on a choice of p) may seem disappointing. We will later comment on the fact that all these choices define "equivalent metrics" in some precise sense that will be introduced and explored in Chapter 3.

1.6. Isometries

In this section we will introduce the notion of an isometry between metric spaces (X, d_X) and (Y, d_Y) . Two isometric metric spaces are two spaces that are in some sense indistinguishable from the metric viewpoint.

Definition 1.6.1. — An *isometry* from (X, d_X) to (Y, d_Y) is a function $\phi: X \rightarrow Y$ such that

(I₁) for every $x_1, x_2 \in X$, $d_Y(\phi(x_1), \phi(x_2)) = d_X(x_1, x_2)$, and

(I₂) ϕ is surjective.

(Note: some textbooks have a more relaxed notion of an isometry and do not require it to be surjective).

Lemma 1.6.2. — *An isometry is injective.*

Proof. With the same notation as in the above definition, assume $x_1, x_2 \in X$. Then

$$\begin{aligned} \phi(x_1) = \phi(x_2) &\iff d_Y(\phi(x_1), \phi(x_2)) = 0 \\ &\iff d_X(x_1, x_2) = 0 \\ &\iff x_1 = x_2. \end{aligned}$$

□

Definition 1.6.3. — Two metric spaces $(X, d_X), (Y, d_Y)$ are *isometric* if there exists an isometry $\phi: (X, d_X) \rightarrow (Y, d_Y)$.

Note that if two metric spaces are isometric, then typically there is more than 1 isometry between them.

The rest of this section is to give examples of isometries and of isometric metric spaces. We will do so first in the most familiar situations, and then in the less familiar ones.

Example 1.6.4. — The metric spaces $([0, 1], d_1)$ and $([2, 3], d_1)$ are isometric. Define ϕ by $\phi(x) = x + 2$. Then

(I₁) $d_1(\phi(x_1), \phi(x_2)) = |(x_1 + 2) - (x_2 + 2)| = |x_1 - x_2| = d_1(x_1, x_2)$.

(I₂) for every $y \in [2, 3]$, we can take x to be $y - 2$. This is in $[0, 1]$, and $\phi(x) = y$, so ϕ is surjective.

1.6 Isometries

Therefore the function ϕ is an isometry according to our definition. Note that there is another isometry $\phi': ([0, 1], d_1) \rightarrow ([2, 3], d_1)$, defined by $\phi'(x) = 3 - x$ (we will leave it to the reader to check that this is also an isometry).

Example 1.6.5. — The metric spaces $([0, 2], d_1)$ and $([0, 1], d_1)$ are *not* isometric. If there existed an isometry ϕ between them, then we would have

$$d_1(\phi(0), \phi(2)) = d_1(0, 2) = 2.$$

But no two points in $[0, 1]$ have a distance between them greater than 1.

We now briefly discuss translations, rotations and reflections.

Exercise 1.6.6. — Show that $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by

$$(x_1, \dots, x_n) \rightarrow (x_1 + v_1, \dots, x_n + v_n)$$

for some $(v_1, \dots, v_n) \in \mathbb{R}^n$ (a translation) defines an isometry $(\mathbb{R}^n, d_p) \rightarrow (\mathbb{R}^n, d_p)$ for all $p \geq 1$ (including $p = \infty$).

Exercise 1.6.7. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be given by

$$f(x_1, x_2) = (\cos(\theta)x_1 - \sin(\theta)x_2, \sin(\theta)x_1 + \cos(\theta)x_2)$$

for some $\theta \in \mathbb{R}$. From year 1 linear algebra you know that f is the counter-clockwise rotation of angle θ centered at the origin.

Show that f is an isometry from \mathbb{R}^2 endowed with the Euclidean distance d_2 to \mathbb{R}^2 with the same distance.

Are these isometries when \mathbb{R}^2 is endowed with any other distance d_p ? (Hint: think first about the cases of $p = 1$ or of $p = \infty$. What are the points that have distance equal to 1 from the origin?)

Exercise 1.6.8. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be given by


$$f(x_1, x_2) = (\cos(\theta)x_1 + \sin(\theta)x_2, \sin(\theta)x_1 - \cos(\theta)x_2)$$

for some $\theta \in \mathbb{R}$. This f is the reflection along the line centered at the origin and of slope $\theta/2$.

Show that f is an isometry from \mathbb{R}^2 endowed with the Euclidean distance d_2 to \mathbb{R}^2 with the same distance.

Remark 1.6.9. — It is a theorem in elementary Euclidean plane geometry that every isometry $\phi: (\mathbb{R}^2, d_2) \rightarrow (\mathbb{R}^2, d_2)$ is the composition of a translation, a rotation and a reflection. (A similar statement is valid for all isometries in \mathbb{R}^n , but this is considerably more complicated to state and prove rigorously, and it doesn't belong in this course.)

Example 1.6.10. — The closed square S of edge length 2 and the closed unit disk D are not isometric with the standard Euclidean metric d_2 . To prove this, observe that the distance of two opposite vertices of S equals $2\sqrt{2}$, but the maximum distance of two points of D equals 2.



Exercise 1.6.11. — 1. Prove that if $\phi: (X, d_X) \rightarrow (Y, d_Y)$ is an isometry, $p \in X$ and $S_1^{d_X}(p)$ is the set of points of X that have distance 1 from p , and $S_1^{d_Y}(\phi(p))$ is the set of points of Y that have distance 1 from $\phi(p)$, then the restriction of ϕ to $S_1^{d_X}(p)$ defines an isometry of the latter with $S_1^{d_Y}(\phi(p))$ (both with the induced metrics from X and Y respectively).

2. Prove that (\mathbb{R}^2, d_2) and $(\{(x_1, x_2) \in \mathbb{R}^2 : x_1 = 0\}, d_2)$ are not isometric.
3. Prove that (\mathbb{R}^n, d_2) and (\mathbb{R}^n, d_∞) are not isometric.
4. Let $\phi: X \rightarrow Y$ be a bijection and let d_Y be a distance on Y . Define a distance d_X on X such that ϕ is an isometry $(X, d_X) \rightarrow (Y, d_Y)$.

(One could prove, using the ideas of the previous exercise, that two distances d_p and d_q on \mathbb{R}^n are never isometric unless of course when p equals q).

The next example shows what are the isometries when one considers *discrete* metric spaces.

1.6 Isometries

Example 1.6.12. — Any bijection $\phi: X \rightarrow Y$ is an isometry of (X, d_{discr}) to (Y, d_{discr}) . Indeed, the following equivalences hold

$$\begin{aligned} d_{\text{discr}}(\phi(x_1), \phi(x_2)) = 0 &\iff \phi(x_1) = \phi(x_2) \\ &\iff x_1 = x_2 \\ &\iff d_{\text{discr}}(x_1, x_2) = 0 \\ d_{\text{discr}}(\phi(x_1), \phi(x_2)) = 1 &\iff \phi(x_1) \neq \phi(x_2) \\ &\iff x_1 \neq x_2 \\ &\iff d_{\text{discr}}(x_1, x_2) = 1 \end{aligned}$$

To conclude, we give a couple of worked examples of isometries first between spaces of functions, and then between spaces of real sequences. We are giving these examples as a further way to illustrate how to play with those metric spaces, we are not attempting a classification of isometries in those spaces.

Example 1.6.13. — Let $A, B \subseteq C[0, 1]$ be sets defined by

$$A = \{f: [0, 1] \rightarrow \mathbb{R} \text{ such that } f(1/2) = 0\},$$

and

$$B = \{f: [0, 1] \rightarrow \mathbb{R} \text{ such that } f(1/2) = 1\}.$$

We claim that (A, d_{L^∞}) and (B, d_{L^∞}) are isometric. Let the proposed isometry be given by $\phi: A \rightarrow B, f \mapsto f + 1$, where $f + 1$ is the function

$$f + 1: x \mapsto f(x) + 1 \quad \text{for all } x.$$

We show that this is indeed an isometry:

$$\begin{aligned} \text{(I}_1\text{)} \quad d_{L^\infty}(\phi(f_1)\phi(f_2)) &= \max_{x \in [0,1]} |f_1(x) + 1 - (f_2(x) + 1)| \\ &= \max_{x \in [0,1]} |f_1(x) - f_2(x)| \\ &= d_{L^\infty}(f_1, f_2) \end{aligned}$$

(I₂) ϕ is also surjective. Let $g \in B$. Define $f := g - 1$. Since

$$g(1/2) = 1 \iff f(1/2) = 0$$

and g is continuous if and only if f is continuous, we deduce that f is in A . Finally, $\phi(f) = f + 1 = g$ so ϕ is surjective.

Example 1.6.14. — There are functions $L, R: \ell^p \rightarrow \ell^p$ (shift to the Left and shift to the Right respectively) defined by

$$L(a_0, a_1, \dots, a_n, \dots) = (a_1, a_2, \dots),$$

and

$$R(a_0, a_1, \dots, a_n, \dots) = (0, a_0, a_1, \dots).$$

Neither L nor R is an isometry. The first is not injective, and the second is not surjective. (For those who know a bit of linear algebra, the maps L and R are linear maps from the vector space ℓ^p to itself. Note that a linear map from a *finite dimensional* vector space to itself is injective if and only if it is surjective).

However, L does define an isometry

$$L: \{(A_n) \in \ell^p : A_0 = 0\} \rightarrow \ell^p$$

(source and target endowed with the same distance d_p), and R does define an isometry

$$R: \ell^p \rightarrow \{(A_n) \in \ell^p : A_0 = 0\}.$$

(source and target endowed with the same distance d_p)

Example 1.6.15. — Let

$$X = \{(A_n) \in \ell^\infty \text{ such that } A_0 = 1, A_2 < 3\}$$

and

$$Y = \{(A_n) \in \ell^\infty \text{ such that } A_0 = 2, A_1 > 0\}.$$

Then there is an isometry $\phi: X \rightarrow Y$ given by

$$\phi(A_0, A_1, A_2, \dots) = (A_0 + 1, -A_2 + 3, A_1, A_3, A_4, \dots).$$

We check the first condition of an isometry, and leave the proof of surjectivity as an easy exercise to the reader.

$$\begin{aligned} d_\infty(\phi(A), \phi(B)) &= \sup(|A_0 + 1 - (B_0 + 1)|, |-A_2 + 3 - (-B_2 + 3)|, |A_1 - B_1|, |A_3 - B_3|, \dots) \\ &= \sup(|A_0 - B_0|, |A_2 - B_2|, |A_1 - B_1|, |A_3 - B_3|, \dots) \\ &= d_\infty(A, B) \end{aligned}$$

In general it is an interesting, and difficult, question to characterise all isometries (if any) between two given metric spaces, or even all auto-isometries (isometries from a metric space to itself). This topic could be an advanced course of its own, and we will not discuss any of it. In MATH241 we will only use isometries in an elementary way.

2.

Continuity and convergence in metric spaces

Important notions to learn from this section:

1. The notion of convergence for a sequence of elements of a metric space.
2. The notion of continuity for a function between metric spaces.
3. A characterization of continuity by means of convergence of sequences (Lemma 2.2.7).
4. The notion of open/closed balls.
5. A characterization of continuity using open balls (Lemma 2.2.4).

In this chapter we introduce the notion of convergence for a sequence of elements of a metric space, and the notion of continuity for a function between metric spaces. We also discuss how these two notions are related. The key notion is that of an open (and that of a closed) ball in a metric space.

2.1. Convergence in metric spaces

Let's first review the notion of convergence for real sequences that is familiar from Year 1. You can find more details on the Calculus I (MATH101) notes.

Definition 2.1.1. — A sequence $(x_n)_{n \in \mathbb{N}}$ of real numbers *converges* to $\ell \in \mathbb{R}$ if for all $\epsilon > 0$ there exists $N \in \mathbb{N}$ such that $|\ell - x_n| < \epsilon$ for every $n > N$.

A sequence (x_n) of real numbers *converges* if there exists $\ell \in \mathbb{R}$ such that (x_n) converges to ℓ , according to the previous definition.

This notion of convergence captures the notion of a limit of a real sequence (from Year 1), except in the case when the limit is $+\infty$ or $-\infty$, where the definition is slightly different.

Example 2.1.2. — Let $(x_n)_{n \in \mathbb{N}} = (\frac{1}{n+1})$, then (x_n) converges to $\ell = 0$. Indeed, for every $\epsilon > 0$, take N to be any natural number larger than $\frac{1}{\epsilon}$. Then

$$|x_n| = \left| \frac{1}{n+1} \right| < \frac{1}{N+1} < \frac{1}{N} \leq \frac{1}{1/\epsilon} = \epsilon \text{ whenever } n > N.$$

We have already encountered and worked with the notion of a *real* sequence. There is nothing special with the real number, we can consider sequences of an arbitrary set, as defined in the following.

Definition 2.1.3. — A *sequence of elements* of a set X is a function $x: \mathbb{N} \rightarrow X$, which we will write as $x = (x_n)_{n \in \mathbb{N}}$ or simply as $x = (x_n)$.

We now generalise the notion of convergence of a real sequence to the case of a sequence of an arbitrary metric space, by replacing $d_1(\ell, x_n) = |\ell - x_n|$ with an arbitrary distance.

Definition 2.1.4. — Let (X, d) be a metric space and (x_n) be a sequence of elements of X and $\ell \in X$. Then we say a sequence (x_n) *converges to ℓ in X* if for every $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that $d(x_n, \ell) < \epsilon$ for all $n > N$.

Definition 2.1.5. — Let (X, d) be a metric space and (x_n) be a sequence of elements of X . Then we say that (x_n) *converges*, if there exists $\ell \in X$ such that (x_n) converges to ℓ in X (according to Definition 2.1.4).

Notation. — In a general metric space (X, d) , if (x_n) is a sequence of X and $\ell \in X$ we will write

$$\lim_{n \rightarrow \infty} x_n = \ell$$

to mean that the sequence (x_n) converges to the limit ℓ . Alternatively, if we want to emphasise that the sequence converges with respect to the given distance d , we write

$$x_n \xrightarrow{d} \ell.$$

2.1 Convergence in metric spaces

Exercise 2.1.6. — Let (X, d) be a metric space and (x_n) be a sequence of X and $\ell \in X$. Set $y_n = d(x_n, \ell)$. Then $(y_n)_{n \in \mathbb{N}}$ is a sequence of real numbers. Show that

$$x_n \xrightarrow{d} \ell \iff y_n \xrightarrow{d_1} 0.$$

Solution: The fact that the sequence (x_n) converges to ℓ means that for all $\epsilon > 0$ there is N such that $d(x_n, \ell) < \epsilon$ for all $n > N$. Because

$$d_1(y_n, 0) = |y_n - 0| = d(x_n, \ell),$$

we have that the previous sentence is the same as saying that (y_n) converges to 0 in (\mathbb{R}, d_1) .

This notion of convergence can be restated in a more geometric manner by introducing the notion of a *ball* in an arbitrary metric space.

Definition 2.1.7. — Let (X, d) be a metric space, p be an element of X and $R \in \mathbb{R}_{>0}$ be a positive real number. Then the *open ball of radius R centered at the point $p \in X$* is defined by:

$$B_R(p) = \{x \in X : d(x, p) < R\}.$$

Similarly the closed ball of radius R centered at the point $p \in X$ is defined by:


$$\overline{B}_R(p) = \{x \in X : d(x, p) \leq R\}.$$

We will sometimes write $B_R^d(p)$ and $\overline{B}_R^d(p)$ when we want to emphasise that the balls are defined using the distance d .

We will now look at examples of balls in the metric spaces that we have been considering.


Example 2.1.8. — Examples of balls in \mathbb{R}^2 :

1. $B_R^{d_2}((0, 0)) = \{(x_1, x_2) \in \mathbb{R}^2 : d_2((x_1, x_2), (0, 0)) < R\}$




Continuity and convergence in metric spaces

2. $B_{\mathbb{R}}^{d_{\infty}}((0,0)) = \{(x_1, x_2) \in \mathbb{R}^2 : d_{\infty}((x_1, x_2), (0,0)) < R\}$



3. $B_{\mathbb{R}}^{d_1}((0,0)) = \{(x_1, x_2) \in \mathbb{R}^2 : d_1((x_1, x_2), (0,0)) < R\}$



Exercise 2.1.9. — Draw the balls $B_{\mathbb{R}}^{d_p}((0,0))$ in \mathbb{R}^2 for other values of $p \geq 1$.

Example 2.1.10. — Balls with the *discrete* distance. Let X be a set and consider the metric space (X, d_{discr}) . Let $x_0 \in X$ and $R > 0$. Then

$$B_{\mathbb{R}}^{d_{\text{discr}}}(x_0) = \begin{cases} \{x_0\} & \text{if } R \leq 1 \\ X & \text{if } R > 1 \end{cases}$$

and

$$\bar{B}_{\mathbb{R}}^{d_{\text{discr}}}(x_0) = \begin{cases} \{x_0\} & \text{if } R < 1 \\ X & \text{if } R \geq 1. \end{cases}$$


Example 2.1.11. — Balls in $(C[0,1], d_{L^{\infty}})$. Fix a continuous function $f \in C[0,1]$ and make this the centre of a ball. Let $R \in \mathbb{R}_{\geq 0}$. Then the open (resp. closed) balls in this metric space are

$$B_{\mathbb{R}}^{d_{L^{\infty}}}(f) = \{g \in C[0,1] : \max |f(x) - g(x)| < R\}$$

$$\bar{B}_{\mathbb{R}}^{d_{L^{\infty}}}(f) = \{g \in C[0,1] : \max |f(x) - g(x)| \leq R\}.$$

In other words, a function g is in the ball if and only if its graph is a subset of the shaded area in the picture below.

2.1 Convergence in metric spaces



(If the graph of g touches the edges of the shaded area then g is in the closed ball but it is not in the open ball).

Example 2.1.12. — Balls in (ℓ^∞, d_∞) .

Let $A = (A_n) = (0, 0, \dots)$ be the zero sequence and $R = 1$. What are the sequences $B = (B_n)$ in the open ball $B_1(A)$ and in the closed ball $\bar{B}_1(A)$? The second question is easier.

Indeed we have

$$B \in \bar{B}_1^{d_\infty}(A) \iff \sup_n |B_n| \leq 1 \iff |B_n| \leq 1 \text{ for all } n \in \mathbb{N}.$$

At this point it might be tempting to guess that $B \in B_1(A)$ if and only if $|B_n| < 1$ for all $n \in \mathbb{N}$, but this is wrong! Indeed, consider the sequence $(B_n) = (1 - 1/n)$. All of its elements are nonnegative and < 1 , but the supremum of B_n **equals** 1, therefore $B \notin B_1(A)$! This example shows that it is possible for the supremum to equal 1 even when each individual element of the sequence satisfies $|B_n| < 1$.

We have therefore

$$B \in B_1^{d_\infty}(A) \iff \sup_n |B_n| < 1 \iff \exists \epsilon > 0 : |B_n| \leq 1 - \epsilon \quad \forall n \in \mathbb{N}.$$

After having seen all these examples, we are now ready to restate the definition of convergence of a sequence using the notion of a ball.

Definition 2.1.13. — Let (X, d) be a metric space. Then we say that (x_n) *converges* to $\ell \in X$ if for every $\epsilon > 0$ there exists an $N \in \mathbb{N}$ such that $x_n \in B_\epsilon(\ell)$ for all $n > N$.

We will see later that the notion of convergence really depends on the choice of the distance on X : we will produce examples of sequences that converge for some distance d but that do not converge for another distance d' on the same set X .

Remark 2.1.14. — What does it mean to say that (x_n) does *not* converge in (X, d) ? In logical notation we can write this condition as

$$\neg(\exists \ell \in X : \forall \epsilon > 0 \exists N \in \mathbb{N} : x_n \in B_\epsilon(\ell) \quad \forall n > N)$$

Through the use of De Morgan's laws, we can rewrite this as the following:

$$\forall \ell \in X \exists \epsilon > 0 : \forall N \in \mathbb{N}, \exists n > N : x_n \notin B_\epsilon(\ell)$$

There are also some very special sequences of X , that converge for *any* notion of a distance d that we could define on X .

Definition 2.1.15. — We say a sequence (x_n) in X is *eventually constant* if there exists an $\ell \in X$ and a natural number N such that $x_n = \ell$ for every $n > N$.

Exercise 2.1.16. — Let (X, d) be a metric space and (x_n) be a sequence of X . Show that if (x_n) is eventually constant, it converges (and it converges precisely to the value that it is eventually constant to).

Exercise 2.1.17. — Let d be a distance on X . Prove that a sequence (x_n) is eventually constant and equal to ℓ if and only if there exists an N such that for every $\epsilon > 0$, we have $x_n \in B_\epsilon(\ell)$ for every $n > N$. [Note that, unlike in Definition 2.1.4, here we are requiring N to be chosen **independently** of ϵ .]

Exercise 2.1.18. — We have already observed that if a sequence (x_n) is eventually constant, then it converges. The converse is also true in the very special case when X is endowed with the *discrete* metric.

Proof. Since (x_n) is convergent, there exists an ℓ such that for all $\epsilon > 0$ there exists $N \in \mathbb{N}$ such that $x_n \in B_\epsilon(\ell)$ for every $n > N$. To show that in fact (x_n) is eventually constant, take any $\epsilon < 1$. Then there exists an N such that $x_n \in \{\ell\}$ for every $n > N$, which is true if and only if $x_n = \ell$ for every $n > N$. □

Example 2.1.19. — Let $(X, d) = (C[0, 1], d_{L^\infty})$. Let $f_n(x) = \frac{x}{n+1}$ and $f(x) = 0$. We claim that (f_n) converges to f in X . Note that


$$\max_{x \in [0,1]} |f_n(x) - f(x)| = \max_{x \in [0,1]} f_n(x) = \frac{1}{n+1},$$

so proving that f_n converges to f is equivalent to proving that $\frac{1}{n+1}$ converges to 0 in (\mathbb{R}, d_1) , which is what we did in Example 2.1.2.

2.1 Convergence in metric spaces

Example 2.1.20. — Let $(X, d) = (C[0, 1], d_{L^\infty})$, $f_n(x) = x^n$ and $f(x) = 0$ for all $x \in [0, 1]$. We claim that f_n does not converge to f in d_{L^∞} . But this is clear - if we take $\epsilon = \frac{1}{2}$, then for every N ,

$$\max_{x \in [0, 1]} |f_N(x) - f(x)| = |f_N(1) - f(1)| = 1 > \frac{1}{2}.$$



We now want to produce some example of convergence in the spaces of sequences ℓ^p . For this, we will need to introduce some notation and clarifications first.

Notation. — Let's now try and write a *sequence* of elements in ℓ^p and see if it converges or not. This might at first sight be a bit confusing (and we have to be extra careful with language), because an *element* of ℓ^p is already a sequence itself! A sequence of elements in ℓ^p is therefore a *sequence of sequences*.

Now a sequence $(A_k)_{k \in \mathbb{N}}$ of elements of ℓ^p is the datum of a function $A_k: \mathbb{N} \rightarrow \ell^p$ for every $k \in \mathbb{N}$, that we can therefore write as

$$(A_{k,0}, A_{k,1}, \dots, A_{k,n}, \dots) = ((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}}$$

where each $A_{k,n}$ is a real number. Then we have that each of

$$(A_{0,n})_{n \in \mathbb{N}}, (A_{1,n})_{n \in \mathbb{N}}, (A_{2,n})_{n \in \mathbb{N}}, \dots, (A_{k,n})_{n \in \mathbb{N}}, \dots$$

is an element of ℓ^p . (In fact this notation would work for any sequence of sequences of real numbers, the fact that we are considering elements of the space ℓ^p has played no role here).

If we write each element of the sequence (of sequences) on a subsequent row, then what we obtain looks like a matrix with infinite rows and

columns

$$(A_{0,n})_{n \in \mathbb{N}} = (A_{0,0}, A_{0,1}, \dots, A_{0,n}, \dots)$$

$$(A_{1,n})_{n \in \mathbb{N}} = (A_{1,0}, A_{1,1}, \dots, A_{1,n}, \dots)$$

\vdots

$$(A_{k,n})_{n \in \mathbb{N}} = (A_{k,0}, A_{k,1}, \dots, A_{k,n}, \dots)$$

\vdots

with the only peculiarity that the indices for row and columns start from zero rather than from one.

Example 2.1.21. — Consider the sequence of sequences defined by

$$((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}} = \left(\frac{(-1)^n}{k+1} \right).$$

Set $A_k = (A_{k,n})_{n \in \mathbb{N}}$ for all $k \in \mathbb{N}$. We claim that each A_k is an element of ℓ^∞ . This is formally obvious because

$$\sup_{n \in \mathbb{N}} |A_{k,n}| = \frac{1}{k+1} < \infty$$

for all $k \in \mathbb{N}$. To avoid any confusion, let's write out the first few elements of this sequence:

$$(A_{0,n})_{n \in \mathbb{N}} = (1, -1, 1, -1, 1, -1, \dots)$$

$$(A_{1,n})_{n \in \mathbb{N}} = \left(\frac{1}{2}, -\frac{1}{2}, \frac{1}{2}, -\frac{1}{2}, \frac{1}{2}, -\frac{1}{2}, \dots \right)$$

$$(A_{2,n})_{n \in \mathbb{N}} = \left(\frac{1}{3}, -\frac{1}{3}, \frac{1}{3}, -\frac{1}{3}, \frac{1}{3}, -\frac{1}{3}, \dots \right)$$

\vdots

After having written it out, it should be evident that the sequence $(A_k)_{k \in \mathbb{N}}$ is converging, and that it is converging to the zero sequence $B = (0, 0, \dots)$. Here is a proof. For every $\epsilon > 0$, choose an integer N greater than $\frac{1}{\epsilon}$. Then

$$\sup_{n \in \mathbb{N}} |A_{k,n} - B_n| = \frac{1}{k+1} \leq \frac{1}{N} < \epsilon.$$

for all $k > N$.

As usual, we write

$$A_k \xrightarrow{d_\infty} B$$

or equivalently

$$(A_{k,n})_{n \in \mathbb{N}} \xrightarrow{d_\infty} (B_n)_{n \in \mathbb{N}}$$

2.1 Convergence in metric spaces

to mean that the sequence $(A_k)_{k \in \mathbb{N}}$ converges to B for $k \rightarrow \infty$.

(For the record, it should be clear that $((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}}$ does not define a sequence in ℓ^p for any $p < \infty$).

Let's now consider a similar example, where we swap the roles of the two indices k and n .

Example 2.1.22. — Consider now the sequence of sequences

$$((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}} = \left(\frac{(-1)^k}{n+1} \right),$$

obtained from the previous example by swapping the roles of the indices k and n on the right hand side. Writing out the first few sequences gives:

$$\begin{aligned} (A_{0,n})_{n \in \mathbb{N}} &= \left(1, \frac{1}{2}, \frac{1}{3}, \dots \right) \\ (A_{1,n})_{n \in \mathbb{N}} &= \left(-1, -\frac{1}{2}, -\frac{1}{3}, \dots \right) \\ (A_{2,n})_{n \in \mathbb{N}} &= \left(1, \frac{1}{2}, \frac{1}{3}, \dots \right) \\ &\vdots \end{aligned}$$

It should be quite clear that, for each fixed k , the sequence $(A_{k,n})_{n \in \mathbb{N}}$ is in ℓ^p for all $p > 1$, including $p = \infty$.

We claim that the sequence $(A_k)_{k \in \mathbb{N}}$ does not converge in (ℓ^∞, d_∞) (and in fact it also does not converge in any (ℓ^p, d_p) for any $p > 1$). [The idea of the proof is that if (A_k) did converge for $k \rightarrow \infty$ then, in particular, the real sequence $(A_{k,0})$ that consists of the first coordinates of each A_k would also converge in (\mathbb{R}, d_1) . That sequence is $((-1)^k)_{k \in \mathbb{N}}$, which does not converge.]

Proof. Assume that $(A_k)_{k \in \mathbb{N}}$ converges to some sequence B . We fix $\epsilon = \frac{1}{2}$. Then by definition of convergence, there exists $K \in \mathbb{N}$ such that for all $k > K$

$$\frac{1}{2} > \sup_n |A_{k,n} - B_n| \geq |A_{k,0} - B_0|.$$

For $k > K$ even, this constrains B_0 to satisfy $\frac{1}{2} < B_0 < \frac{3}{2}$. On the other hand, for $k > K$ odd, this constrains B_0 to satisfy $-\frac{3}{2} < B_0 < -\frac{1}{2}$. It is not possible for B_0 to satisfy both constraints, and this shows that the sequence $(A_k)_{k \in \mathbb{N}}$ does not converge in ℓ^∞ . \square

Continuity and convergence in metric spaces

(Try to prove yourself that the same sequence does not converge as a sequence of ℓ^p for any $p \geq 1$. It is essentially the same argument!)

After having seen some examples, let's go back to the general theory of convergence of sequences in metric spaces. You have seen in Year 1 that a real sequence cannot possibly converge to two different limits. The same is true of a sequence of an arbitrary metric space.

Lemma 2.1.23 (Uniqueness of the limit). — *Let (X, d) be a metric space, and let (x_n) be a sequence of X . Then, if we have $\ell_1, \ell_2 \in X$ such that $\lim_n x_n = \ell_1$ and $\lim_n x_n = \ell_2$, then $\ell_1 = \ell_2$.*

Proof. Suppose $\ell_1 \neq \ell_2$. Take $\epsilon = \frac{d(\ell_1, \ell_2)}{3}$. Then $\epsilon > 0$ because $d(\ell_1, \ell_2) > 0$ since $\ell_1 \neq \ell_2$.

Using the definition of convergence with this ϵ , we find N_1 such that $d(x_n, \ell_1) < \epsilon$ for every $n > N_1$, and similarly an N_2 such that $d(x_n, \ell_2) < \epsilon$ for every $n > N_2$. Take now $N' > \max(N_1, N_2)$. Then by the triangle inequality for d we have:

$$d(\ell_1, \ell_2) \leq d(\ell_1, x_{N'}) + d(\ell_2, x_{N'}) < \frac{2}{3}d(\ell_1, \ell_2).$$

This is a contradiction, because no positive real number can be smaller than or equal to its two-thirds! This concludes our proof. \square

2.2. Continuity in metric spaces

We now review the notion of continuity of a real function of 1 variable and then generalise it to the case of a function between two abstract metric spaces. We take a slightly different approach from the previous chapter. We start from a definition of continuity that is *different* from the one that was given in MATH101, but we will show the equivalence with the Year 1 notion in Lemma 2.2.7.

Definition 2.2.1. — Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a function. We say that f is continuous at x_0 if for every $\epsilon > 0$ there exists $\delta > 0$ such that

$$|f(x) - f(x_0)| = d_1(f(x), f(x_0)) < \epsilon$$

whenever $|x - x_0| = d_1(x, x_0) < \delta$. We say that f is *continuous* if f is continuous at every $x_0 \in \mathbb{R}$.

2.2 Continuity in metric spaces

We say that the *limit of f at x_0* is equal to $\ell \in \mathbb{R}$ (written $\lim_{x \rightarrow x_0} f(x) = \ell$) if for every $\epsilon > 0$, $\exists \delta > 0$ such that

$$|f(x) - \ell| = d_1(f(x), \ell) < \epsilon$$

whenever $0 < |x - x_0| < \delta$.

We now extend these definitions to the general case of functions between metric spaces.

Definition 2.2.2. — Let $(X, d_X), (Y, d_Y)$ be metric spaces. Let $f: X \rightarrow Y$ be a function, and let $x_0 \in X$. Then we say f is *continuous at x_0* if for every $\epsilon > 0$ there exists a $\delta > 0$ such that

$$d_Y(f(x), f(x_0)) < \epsilon$$

whenever $d_X(x, x_0) < \delta$. We say that f is continuous if it is continuous at x_0 for every $x_0 \in X$.

Let now $\ell \in Y$. We say the *limit of f at x_0* is equal to ℓ (and we write $\lim_{x \rightarrow x_0} f(x) = \ell$) if for every $\epsilon > 0$, $\exists \delta > 0$ such that

$$d_Y(f(x), \ell) < \epsilon$$

whenever $0 < d_X(x, x_0) < \delta$.

The notion of limit and of continuity are intimately related. Indeed, it immediately follows from the definition that f is continuous at x_0 if and only if $\lim_{x \rightarrow x_0} f(x)$ exists and it equals $f(x_0)$.

We will now describe the notion of continuity in other equivalent manners. First we recall some notation.

Remark 2.2.3 (Direct and Inverse images of a function). — Here we recall some notation.

Let $f: X \rightarrow Y$ be a function. Then we have that $f(x) \in Y$.

We can define, for a subset $A \subseteq X$, its *direct image*

$$f(A) := \{y \in Y : \exists x \in A \text{ such that } f(x) = y\}.$$

(Note that if $A = \{x\}$ consists of one element, then $f(A) = \{f(x)\}$, the set that contains precisely one element: $f(x)$).

Similarly, for a subset $B \subseteq Y$ we define its *inverse image* (or *preimage*)

$$f^{-1}(B) := \{x \in X \text{ such that } f(x) \in B\}.$$

Continuity and convergence in metric spaces

Using this notation, we now rewrite the definition of continuity using only the notion of balls in the source and in the target metric space. Let (X, d_X) and (Y, d_Y) be metric spaces, and $f: X \rightarrow Y$. Then

$$\begin{aligned} f \text{ is continuous at } x_0 &\iff \forall \epsilon > 0 \exists \delta > 0 \text{ such that } d_X(x, x_0) < \delta \implies d_Y(f(x), f(x_0)) < \epsilon \\ &\iff \forall \epsilon > 0 \exists \delta > 0 \text{ such that } f(B_\delta^{d_X}(x_0)) \subseteq B_\epsilon^{d_Y}(f(x_0)) \end{aligned}$$

where we have emphasized the distance that defines each ball by denoting it with a superscript. We will write what we have just observed in a corollary.

Corollary 2.2.4. — *Let $(X, d_X), (Y, d_Y)$ be metric spaces. Let $f: X \rightarrow Y$ be a function, and let $x_0 \in X$. Then f is continuous at x_0 if and only if*

$$\forall \epsilon > 0 \exists \delta > 0 \text{ such that } f(B_\delta^{d_X}(x_0)) \subseteq B_\epsilon^{d_Y}(f(x_0))$$

This characterization is quite useful, and permits for example to give a slick proof of the following result.

Lemma 2.2.5. — *Let $(X, d_X), (Y, d_Y), (Z, d_Z)$ be metric spaces. Suppose that $x_0 \in X, f: X \rightarrow Y, g: Y \rightarrow Z$. If f is continuous at x_0 and g continuous at $f(x_0)$, then $g \circ f$ is continuous at x_0 .*

Corollary 2.2.6. — *If f and g are both continuous (i.e. they are continuous at all points where they are defined), then so is $g \circ f$.*

Proof. (Of Lemma 2.2.5) Let $\epsilon > 0$. Since g is continuous at $f(x_0)$ we can find a δ_1 such that $g(B_{\delta_1}(f(x_0))) \subseteq B_\epsilon(g(f(x_0)))$. Similarly since f is continuous at x_0 we can find a $\delta > 0$ such that $f(B_\delta(x_0)) \subseteq B_{\delta_1}(f(x_0))$. Applying the function g to both sides of this second inclusion gives us

$$\begin{aligned} g(f(B_\delta(x_0))) &\subseteq g(B_{\delta_1}(f(x_0))) \subseteq B_\epsilon(g(f(x_0))) \\ &\implies g(f(B_\delta(x_0))) \subseteq B_\epsilon(g(f(x_0))) \end{aligned}$$

and so using Corollary 2.2.4, we see that $g \circ f$ is indeed continuous at x_0 as claimed. □

We will now see how the notion of continuity relates to the one of convergence. More precisely, the notion of continuity can be recast completely in terms of convergence of sequences in the source and in the target. This equivalent formulation is the one that was used as the definition of continuity in Year 1.

2.2 Continuity in metric spaces

Lemma 2.2.7. — Let (X, d_X) and (Y, d_Y) be metric spaces, $f: X \rightarrow Y$ a function and $p \in X$. The following are equivalent:

1. f is continuous at p .
2. for every sequence (x_n) in (X, d_X) such that $x_n \xrightarrow{d_X} p$, $f(x_n) \xrightarrow{d_Y} f(p)$.

Remark 2.2.8. — If f is continuous at p , by the above we have that

$$f\left(\lim_n x_n\right) = \lim_n f(x_n).$$

i.e. limits can be moved in and out of a parenthesis, when the corresponding function is continuous.

The above lemma is often useful when trying to show that certain functions are *not* continuous, as we see in the following example.

Example 2.2.9. — Let $(X, d) = (C[0, 1], d_{L^1})$. Define

$$\phi: C[0, 1] \rightarrow \mathbb{R}$$

be defined by $f \mapsto \max_{x \in [0, 1]} f(x)$. Then we claim that ϕ is not continuous at

$g = 0$. To see this, let $(f_n)(x) = x^n$. Then in X we have $f_n \xrightarrow{d_{L^1}} g$, since

$$\int_0^1 f_n(x) dx = \frac{1}{n+1} \rightarrow 0.$$

However, $\phi(f_n) = 1$ for every n , whilst $\phi(g) = 0$. Therefore

$$\underbrace{\phi(\lim f_n)}_{=0} \neq \underbrace{\lim \phi(f_n)}_{=1}$$

and so ϕ is not continuous at g .

Exercise 2.2.10. — Prove that the function ϕ of the previous example is continuous if $C[0, 1]$ is endowed with the d_{L^∞} distance (instead of d_{L^1}).

Solution: The difficult part of this exercise is proving the inequality

$$\left| \max_{x \in [0, 1]} f(x) - \max_{x \in [0, 1]} g(x) \right| \leq \max_{x \in [0, 1]} |f(x) - g(x)| \quad (2.1)$$

for all $f, g \in C[0, 1]$. Assuming (2.1), for $g \in C[0, 1]$ we prove continuity of Φ at g by defining $\delta = \epsilon$ for all $\epsilon > 0$. Then by (2.1) we deduce

$$d_{L^\infty}(f, g) < \delta \implies |\Phi(f) - \Phi(g)| \leq d_{L^\infty}(f, g) < \delta = \epsilon.$$

Continuity and convergence in metric spaces

In order to prove (2.1) we observe that, by definition of the absolute value, it is equivalent to showing the two inequalities

$$-\max_{x \in [0,1]} |f(x) - g(x)| \leq \max_{x \in [0,1]} f(x) - \max_{x \in [0,1]} g(x) \leq \max_{x \in [0,1]} |f(x) - g(x)|$$

We only prove the second inequality and leave the first to the reader.

Let $t \in [0, 1]$ be one value where f attains its maximum in $[0, 1]$, so $\max_{x \in [0,1]} f(x) = f(t)$. Then

$$f(t) - g(t) \leq \max_{x \in [0,1]} (f(x) - g(x)) \leq \max_{x \in [0,1]} |f(x) - g(x)| \quad (2.2)$$

On the other hand, $-g(t) \geq -\max_{x \in [0,1]} g(x)$, therefore

$$f(t) - \max_{x \in [0,1]} g(x) \leq f(t) - g(t). \quad (2.3)$$

Combining (2.2) and (2.3) we deduce

$$\max_{x \in [0,1]} f(x) - \max_{x \in [0,1]} g(x) \leq \max_{x \in [0,1]} |f(x) - g(x)|,$$

which is what we wanted to prove.

We now proceed with the proof of Lemma 2.2.7.

Proof of Lemma 2.2.7. Assume that f is continuous at p and suppose we have a sequence (x_n) of (X, d_X) such that

$$\lim_{n \rightarrow \infty} x_n = p.$$

Since f is continuous, for every $\epsilon > 0$ we can find a $\delta > 0$ such that $f(B_\delta(p)) \subseteq B_\epsilon(f(p))$. Since $(x_n) \rightarrow p$, we can find an $N \in \mathbb{N}$ such that for every $n > N$, $x_n \in B_\delta(p)$. Therefore $f(x_n) \in f(B_\delta(p)) \subseteq B_\epsilon(f(p))$ and so we have convergence in Y too.

Conversely, suppose f is not continuous at p . Then there exists an $\epsilon > 0$ such that for every $\delta > 0$, $f(B_\delta(p)) \not\subseteq B_\epsilon(f(p))$. In other words, for all $\delta > 0$, there exists a $y \in f(B_\delta(p))$ such that $y \notin B_\epsilon(f(p))$. Therefore take

$$\begin{aligned} \delta = 1 &\longrightarrow \text{find a } y_1 \in f(B_1(p)) : y_1 \notin B_\epsilon(f(p)) \\ \delta = \frac{1}{2} &\longrightarrow \text{find a } y_2 \in f(B_{\frac{1}{2}}(p)) : y_2 \notin B_\epsilon(f(p)) \\ &\vdots \\ \delta = \frac{1}{n} &\longrightarrow \text{find a } y_n \in f(B_{\frac{1}{n}}(p)) : y_n \notin B_\epsilon(f(p)) \\ &\vdots \end{aligned}$$

2.2 Continuity in metric spaces

For all $n \in \mathbb{N}$, because $y_n \in f(B_{\frac{1}{n}}(p))$, there is $x_n \in B_{\frac{1}{n}}(p)$ such that $f(x_n) = y_n$. Therefore we have constructed a sequence (x_n) in X such that, for all $n \in \mathbb{N}$,

$$\underbrace{x_n \in B_{\frac{1}{n}}(p)}_{\text{so } x_n \rightarrow p} \text{ but } \underbrace{f(x_n) = y_n \notin B_\epsilon(f(p))}_{\text{so } f(x_n) \not\rightarrow f(p)}$$

This gives the required contradiction. □

3.

The topology of metric spaces

Important notions to learn from this section:

1. The notion of open and closed subset of a metric space.
2. Formal properties of open subsets.
3. A characterisation of closed sets by convergence of sequences.
4. A characterisation of convergence/continuity using only open subsets.
5. The notion of *equivalence* for two distances on the same set.
6. The notions of homeomorphism and of homeomorphic metric spaces.

The central notion of this chapter is that of open (and that of closed) subset of a metric space. We will see how the notions of convergence and continuity can be rephrased in alternative definitions that only refer to the open subsets. The two core notions of convergence and continuity are therefore insensitive of the actual distance, as they only sense it via the underlying “topology” of the metric space (by topology we mean the “ability to distinguish” the open subsets within the collection of all subsets). This leads us to define a second notion of equivalence for metric spaces, called “homeomorphism”, which is less rigid than the notion of an isometry that we encountered in the first chapter. The main point is that two homeomorphic metric spaces are, in some sense, indistinguishable from the point of view of convergence and continuity.

3.1. Open and closed sets

In this section we introduce and discuss the notion of *open* and that of *closed* subsets of a metric space. These notions should already be somehow familiar from Year 1 for the subsets of the real line \mathbb{R} endowed with the standard (Euclidean) distance. The standard examples to keep in mind are that of the open interval $(0, 1) \subset \mathbb{R}$ and that of the closed interval $[0, 1] \subset \mathbb{R}$.

Definition 3.1.1. — Let (X, d) be a metric space, and let $A \subseteq X$ be a subset of X . Then we say that A is *open in* (X, d) if for every $p \in A$, there exists an $\epsilon > 0$ such that $B_\epsilon(p) \subseteq A$. If $B \subseteq X$ then we say that B is *closed in* (X, d) if $X \setminus B$ is open.

Note that in common language the word “open” means the opposite of the word “closed”. This is not so in technical mathematical language: a subset may fail to be open and closed, and a subset may be simultaneously open AND closed.

Example 3.1.2. — 1. The open interval $(0, 1)$ is open in (\mathbb{R}, d_1) . Indeed, for all $p \in (0, 1)$, take $0 < \epsilon < \min(p, 1 - p)$. Then

$$B_\epsilon(p) = (p - \epsilon, p + \epsilon) \subseteq (0, 1).$$

2. The closed interval $[0, 1]$ is closed in (\mathbb{R}, d_1) . Indeed,

$$\mathbb{R} \setminus [0, 1] = (-\infty, 0) \cup (1, \infty)$$

is open because for all $p \in (-\infty, 0) \cup (1, \infty)$ we will choose ϵ any positive real number that satisfies

$$\begin{cases} \epsilon < p - 1 & \text{if } p > 1 \\ \epsilon < -p & \text{if } p < 0. \end{cases}$$

With this choice of ϵ , we have

$$B_\epsilon(p) = (p - \epsilon, p + \epsilon) \subseteq (-\infty, 0) \cup (1, \infty).$$

3. There is nothing special about the values 0 and 1 in the previous two examples. With a similar reasoning, one can easily see that all open intervals of the form (a, b) for $a, b \in \mathbb{R}$ are indeed open subsets of (\mathbb{R}, d_1) and all closed intervals of the form $[a, b] \in \mathbb{R}$ are closed subsets of (\mathbb{R}, d_1) .

3.1 Open and closed sets

4. A subset can be neither open nor closed! Let $A = [0, 1)$. Then A is neither open nor closed in $(X, d) = (\mathbb{R}, d_1)$.

To see it isn't open, let $p = 0$. Then for every $\epsilon > 0$,

$$B_\epsilon(p) = (p - \epsilon, p + \epsilon) \not\subseteq [0, 1).$$

Likewise to see that A isn't closed, take $p = 1$. Then for every $\epsilon > 0$,

$$B_\epsilon(1) = (1 - \epsilon, 1 + \epsilon) \not\subseteq \mathbb{R} \setminus [0, 1).$$

5. In fact a subset may even be simultaneously open *and* closed! Take $A = \mathbb{R}$, then we claim that A is open and closed in (\mathbb{R}, d_1) .

It is clear that A is open. Indeed, for all $p \in A = \mathbb{R}$, we have that $B_\epsilon(p) \subseteq A = \mathbb{R}$ for all possible choices of ϵ !

To see that A is also closed will require a small effort of logic. Indeed, the complement $\mathbb{R} \setminus \mathbb{R}$ is the empty set \emptyset . How do we check that the empty set is open in (\mathbb{R}, d_1) ? For all $p \in \emptyset$ we have to verify something, but \emptyset has no elements, so there is nothing to verify!

Incidentally, because $A = \mathbb{R}$ is open and closed in (\mathbb{R}, d_1) , then the same applies to $\mathbb{R} \setminus A = \emptyset$, so we find that \emptyset is also open and closed in (\mathbb{R}, d_1) .

6. Arguing as in the previous point, we see that the empty set and X are open and closed subsets of any metric space (X, d) . This apparently bizarre observation will become especially relevant later.

Let's see now an example that is not a subset of the real line.

Example 3.1.3. — Let $(X, d) = (\mathbb{R}^2, d_2)$ and $R > 0$. Then the open ball $B_R(0)$ is actually an open subset. To show this, let p be a point in $B_R(0)$. Then we can write $p = (r \cos \theta, r \sin \theta)$ where $r < R$ and $\theta \in [0, 2\pi)$. For $\epsilon = R - r$, then $B_\epsilon(p) \subseteq B_R(0)$.

On the other hand, we can also see that $B_R(0)$ is not closed. Consider now the point $p = (R, 0) \in \mathbb{R}^2 \setminus B_R(0)$. Then for every $\epsilon > 0$,

$$B_\epsilon(p) \not\subseteq \mathbb{R}^2 \setminus B_R(0).$$

(There is nothing special about the origin in this example, one could have considered the open ball centered at an arbitrary point).

In the same metric space, it is not difficult to explicitly verify that the closed ball $\overline{B}_R(0)$ is closed and not open.

The notion of being “open” is relative rather than absolute, in the sense that it also depends on the ambient space (X, d) . Let us illustrate this important subtlety with some more examples.

7. Let $(X, d) = ([0, \infty), d_1)$, $A = [0, 1)$. Then A is open in X . For $p \in [0, 1)$, take $\epsilon = 1 - p > 0$. Then

$$B_\epsilon(p) = (p - \epsilon, p + \epsilon) \cap [0, \infty) \subseteq [0, 1),$$

because $p + \epsilon = p + 1 - p = 1$.

8. Let $(X, d) = ([0, 1] \cup [2, 3], d_1)$ and $A = [0, 1]$. Then A is both open and closed in X . Indeed, as seen in the previous example, for ϵ sufficiently small one has $B_\epsilon(0) = [0, \epsilon)$ and $B_\epsilon(1) = (1 - \epsilon, 1]$. Similarly, one shows that $[2, 3]$ is also open in X , hence that $[0, 1]$ is closed.

9. Let $(X, d) = (C[0, 1], d_{L^\infty})$ and

$$A = \{f \in X : f(1/3) > 1\}.$$

We claim that A is open in X . To see this, take $f_0 \in A$. Then let $\epsilon = f_0(1/3) - 1$. Then

$$B_\epsilon(f_0) = \{f : \max |f(x) - f_0(x)| < \epsilon\}.$$

Then ϵ is positive because $f_0 \in A$ and

$$\begin{aligned} f \in B_\epsilon(f_0) &\implies -\epsilon < f(1/3) - f_0(1/3) < \epsilon \\ &\implies f(1/3) > f_0(1/3) - \epsilon = 1 \end{aligned}$$

i.e. $B_\epsilon(f_0) \subseteq A$.

Similarly, one could prove that A is not closed in (X, d) .

10. Let $(X, d) = (X, d_{\text{discr}})$ and A be any subset of X . Then A is both open and closed. To see this, for every $p \in A$, take $\epsilon = \frac{1}{2}$. The statement follows since

$$p \in A \implies B_{\frac{1}{2}}^{\text{discr}}(p) = \{p\} \subseteq A.$$

11. Now let (X, d) be any metric space and let $R \in \mathbb{R}_{>0}$. Then

3.1 Open and closed sets

- (a) Each open ball $B_R(p)$ is actually an open subset. Indeed, for any $x_0 \in B_R(p)$, take $\epsilon > 0$ such that $\epsilon < R - d(p, x_0)$ (note that $R - d(p, x_0) > 0$ because $x_0 \in B_R(p)$). Then we have

$$B_\epsilon(x_0) = \{x \in X : d(x, x_0) < \epsilon\}.$$

Applying the triangle inequality we deduce

$$d(x, p) \leq d(x, x_0) + d(x_0, p) < \epsilon + d(p, x_0) < R,$$

when $x \in B_\epsilon(x_0)$. From this it follows that if $x \in B_\epsilon(x_0)$, then $x \in B_R(p)$, so the open ball $B_\epsilon(x_0)$ is a subset of $B_R(p)$. This concludes the proof that $B_R(p)$ is open.

- (b) Similarly we could prove that each closed ball $\overline{B}_R(x_0)$ is actually closed.
- (c) The subsets \emptyset and X are both open and closed in X . This is pretty much immediate from the definitions. The subset \emptyset is open since there are no elements in X , so the property of being open follows without having anything to check. The subset X is also open because $B_\epsilon(p)$ is a subset of X for any choice of p and ϵ .

The last example inspires us to observe that the collection of open subsets of a given metric space satisfies some formal properties that we are now going to list and prove.

Lemma 3.1.4 (Formal properties of open sets). — *Let (X, d_X) be a metric space.*

1. *The subsets \emptyset and X are open.*
2. *An arbitrary union of open sets is open.*
3. *A finite intersection of open sets is open.*

To see why we are only allowed to take an intersection of finitely many sets, as an example take $(X, d) = (\mathbb{R}, d_1)$ and the infinite collection of subsets $A_k = (-\frac{1}{k}, \frac{1}{k})$ for $k = 1, 2, 3, \dots$. Then A_k is open for every k , whilst

$$\bigcap_{k=1}^{\infty} A_k = \{0\},$$

which is not open in (\mathbb{R}, d_1) (because any ball of positive radius centred at 0 contains elements other than 0).

It may be worth at this point recalling the notation for taking intersections and unions of sets.

Remark 3.1.5. — If A, B are any two sets, then we define:

$$A \cap B = \{x \text{ such that } x \in A \text{ and } x \in B\},$$

$$A \cup B = \{x \text{ such that } x \in A \text{ or } x \in B\}.$$

If A_1, A_2, \dots, A_N are sets, then we define

$$\bigcap_{i=1}^N A_i = A_1 \cap A_2 \cap \dots \cap A_N = \{x : x \in A_1, x \in A_2, \dots, x \in A_N\},$$

$$\bigcup_{i=1}^N A_i = A_1 \cup A_2 \cup \dots \cup A_N = \{x : x \in A_1 \text{ or } x \in A_2 \dots \text{ or } x \in A_N\}.$$

More generally, if I is a set (to be regarded as a set of indices) and A_i is a set for all $i \in I$, then we define

$$\bigcap_{i \in I} A_i = \{x \text{ such that } x \in A_i \text{ for every } i \in I\},$$

$$\bigcup_{i \in I} A_i = \{x \text{ such that } x \in A_i \text{ for some } i \in I\}.$$

As an example, take $I = \mathbb{N} \setminus \{0\}$ and if $A_k = (-\frac{1}{k}, \frac{1}{k})$ for all $k \in \mathbb{N}$. Then

$$\bigcap_{i \in I} A_i = \{0\}$$

is the example that we have just discussed immediately before to show that the intersection of infinitely many open sets might fail to be open.

Proof of Lemma 3.1.4. 1. This was observed before..

2. Suppose the subsets A_i are open in (X, d) for all $i \in I$. For all $p \in \bigcup_{i \in I} A_i$, we must have that $p \in A_j$ for some $j \in I$. Because A_j is open, there exists an $\epsilon > 0$ such that

$$B_\epsilon(p) \subseteq A_j \subseteq \bigcup_{i \in I} A_i.$$

This proves then that $\bigcup_{i \in I} A_i$ is open in (X, d) .

3.1 Open and closed sets

3. Suppose that we have open sets A_1, A_2, \dots, A_N which are all open. Assume $p \in A_1 \cap \dots \cap A_N$. For all $j = 1, \dots, N$, because A_j is open in (X, d) , we can find an $\epsilon_j > 0$ and a ball $B_{\epsilon_j}(p) \subseteq A_j$. Now define $\epsilon = \min(\epsilon_1, \epsilon_2, \dots, \epsilon_N)$. Then $\epsilon > 0$ and we have the inclusion

$$B_\epsilon(p) = \bigcap_{i=1}^N B_{\epsilon_i}(p) \subseteq \bigcap_{i=1}^N A_i.$$

Because this argument is valid for all $p \in A_1 \cap \dots \cap A_N$, the latter is an open subset of (X, d) . □

Similar formal properties are enjoyed by the collection of closed sets, by simply taking the complement and observing that

$$\left(\bigcap_{i \in I} A_i \right)^c = \bigcup_{i \in I} A_i^c \quad \text{and} \quad \left(\bigcup_{i \in I} A_i \right)^c = \bigcap_{i \in I} A_i^c.$$

(The upper c on a subset S of X here is used as a shorthand notation for $S^c = X \setminus S$, the complement of S in X).

Corollary 3.1.6. — 1. The subsets \emptyset, X are closed in X .

2. The arbitrary intersection of closed sets is closed.
3. A finite union of closed sets is closed.

Example 3.1.7. — The set $A_k = [0, 1 - \frac{1}{k}]$ is closed for every $k \geq 1$. However, the infinite union $\bigcup_{k=1}^{\infty} A_k$ equals $[0, 1)$, which is *not* closed in (\mathbb{R}, d_1) .

So far we have explained how to see if a subset of a metric space is and if it is not open. It is often easier to check if a subset is closed, due to the following lemma that relates the property of being closed to the convergence of sequences.

Lemma 3.1.8. — Let (X, d) be a metric space, and let $A \subseteq X$. The following are equivalent:

1. The subset A is closed.
2. For every sequence (x_n) of A , if (x_n) converges to $\ell \in X$ then $\ell \in A$.

We may informally phrase this lemma by saying that “closed” for a subset means that it is closed for the operation of taking limits of sequences.

Proof. Suppose A isn't closed. Then $X \setminus A$ is not open, so there exists $\ell \in X \setminus A$ such that for every $\epsilon > 0$ we have $B_\epsilon(\ell) \not\subseteq X \setminus A$. Therefore set

$$\begin{aligned} \epsilon_1 = 1 & \quad B_1(\ell) \not\subseteq X \setminus A & \text{ then find an } x_1 \in B_1(\ell) : & \quad x_1 \in A \\ \epsilon_2 = \frac{1}{2} & \quad B_{1/2}(\ell) \not\subseteq X \setminus A & \text{ then find an } x_2 \in B_{1/2}(\ell) : & \quad x_2 \in A \\ & \quad \vdots & & \\ \epsilon_n = \frac{1}{n} & \quad B_{1/n}(\ell) \not\subseteq X \setminus A & \text{ then find an } x_n \in B_{1/n}(\ell) : & \quad x_n \in A \\ & \quad \vdots & & \end{aligned}$$

Then we have constructed $\ell \in X \setminus A$ and a sequence (x_n) that converges to ℓ , which contradicts the second statement.

Conversely, assume that the second statement is false. Then there exists a sequence (x_n) of A and $\ell \in X \setminus A$ such that (x_n) converges to ℓ . Then because (x_n) converges to ℓ for every $\epsilon > 0$, there exists n such that $x_n \in B_\epsilon(\ell)$ and $x_n \notin X \setminus A$, so $X \setminus A$ is not open, hence A is not closed. \square

We will now exploit Lemma 3.1.8 to give more examples of closed and of non-closed subsets.

Example 3.1.9. — 1. $[-1, 1)$ is not closed in (\mathbb{R}, d_1) . Let

$$(x_n) = \left(1 - \frac{1}{n+1}\right).$$

Then $(x_n) \in [-1, 1)$ but its limit $\lim_{n \rightarrow \infty} x_n = 1 \notin [-1, 1)$.

2. Let $(X, d) = (C[0, 1], d_{L^\infty})$ and

$$A = \{f \in X : \exists x' \text{ such that } f(x')^2 < 1\}.$$

We claim that A is not closed. Take $(f_n) = (1 - \frac{1}{n+1})$ for every $x \in [0, 1]$ and $g(x) = 1$. Then $f_n \xrightarrow{d_{L^\infty}} g$ and $f_n \in A \forall n \in \mathbb{N}$, but we have $g \notin A$.

3. Take $Q \subseteq (\mathbb{R}, d_1)$. Then Q is not closed. To see this, take any $x \in \mathbb{R} \setminus Q$. Let (x_n) be the decimal approximation of x to N places. For example, if $x = \pi$, then we have

$$\begin{aligned} x_1 &= 3.1 \\ x_2 &= 3.14 \\ x_3 &= 3.141 \\ &\vdots \end{aligned}$$

3.2 Introduction to topology

Then $x_n \in \mathbb{Q}$ for all $n \in \mathbb{N}$, but $x \notin \mathbb{Q}$.

4. Let $p < q$, then $\ell^p \subseteq \ell^q$ and ℓ^p is *not* closed in (ℓ^q, d_q) .

To see this, take any element $B \in \ell^q$ such that $B \notin \ell^p$. For example, take $(B_n) = \left(\frac{1}{n^{1/p}}\right)$. Then define the sequence of sequences

$$A_{k,n} := \begin{cases} \frac{1}{n^{1/p}} & n < k \\ 0 & n \geq k. \end{cases}$$

Then for every k we have $(A_{k,n})_{n \in \mathbb{N}} = A_k \in \ell^p$ (because only finitely many terms of the sequence are different from zero), and we also have

$$\lim_{k \rightarrow \infty} A_k = B$$

in the distance d_q (Exercise: prove this in detail using the definition of limit!). However, by definition $B \notin \ell^p$. This completes the proof that ℓ^p is not closed in (ℓ^q, d_q) for $p < q$.

3.2. Introduction to topology

Here we allow ourselves to make a small digression in a topic that has by now become central in modern mathematics and that you will encounter in your further studies. Inspired by the formal properties observed enjoyed by the collection of open sets (the properties that we proved in Lemma 3.1.4), we now define another notion of a “space”, more general than that of a metric space: the notion of a topological space.

Definition 3.2.1. — A *topology* \mathcal{U} on X is a collection of subsets U_i of X , called the *open subsets* of X , that satisfies the following properties.

(T₁) The subsets $\emptyset, X \in \mathcal{U}$.

(T₂) If $U_i \in \mathcal{U}$ for all $i \in I \Rightarrow \bigcup_{i \in I} U_i \in \mathcal{U}$.

(T₃) If $U_1, \dots, U_N \in \mathcal{U} \Rightarrow \bigcap_{i=1}^N U_i \in \mathcal{U}$.

A *topological space* is a pair (X, \mathcal{U}) where X is a set and \mathcal{U} is a topology on X .

One thing that might confuse at first is that the *elements* of a topology on a set X are *subsets* of X .

Remark 3.2.2. — Let X be a set. Let $\mathcal{P}(X)$ be the set of all subsets of X , in other words

$$\mathcal{P}(X) := \{A : A \subseteq X\}.$$

Then a topology is a subset \mathcal{U} of $\mathcal{P}(X)$ that satisfies the conditions T_1 , T_2 and T_3 .

Example 3.2.3. — For example, if X is a finite set, then the number of element $\#\mathcal{P}(X)$ equals $2^{\#X}$. In particular, the number of subsets of $\mathcal{P}(X)$ equals $2^{2^{\#X}}$. The number of subsets \mathcal{U} of $\mathcal{P}(X)$ such that $\emptyset, X \in \mathcal{U}$ equals $2^{2^{\#X}-2}$. This number gives a very coarse upper bound on the number of different topologies on a finite set X , for it ignores the two constraints T_2 and T_3 .

Example 3.2.4 (The Main Example of a Topological Space in MATH241). — Let (X, d_X) be any metric space. Then we set

$$\mathcal{U} = \{A \subseteq X \text{ such that } A \text{ is open in } (X, d_X)\}.$$

By Lemma 3.1.4, the pair (X, \mathcal{U}) is a topological space. We call \mathcal{U} the topology *generated/induced by the distance* d_X .

The previous example points to the fact that the notion of a topological space is more general than the notion of a metric space, the main notion of this module.

Remark 3.2.5. — Note that for any set X , the set $\mathcal{P}(X)$ is always a topology on X (the three axioms T_1, T_2, T_3 are trivially satisfied). As we have seen in part 10 of the list of examples above, if X is equipped with the *discrete* metric, all its subsets are open. We conclude that the topology $\mathcal{P}(X)$ is generated/induced by the discrete metric. (For this reason, the topology $\mathcal{P}(X)$ is sometimes referred to as the *discrete topology* on X .)

On the opposite end of the spectrum, another topology that is always available is the smallest possible topology, also known as the *trivial topology* (or indiscrete topology), which consists only of $\mathcal{U} = \{\emptyset, X\}$. (It is also straightforward to prove that this topology satisfies the three axioms).

Let's see some more examples of topologies on finite sets.

Example 3.2.6. — 1. Let $X = \{A\}$ be a set containing a unique element A . Then $\mathcal{U} = \{\emptyset, \{A\}\}$ is the only possible topology on X .

3.2 Introduction to topology

2. If $X = \{A, B\}$ is a set with 2 elements, then the possible topologies on X are

$$\mathcal{U}_1 = \{\emptyset, \{A, B\}\} \text{ - the trivial topology on } X,$$

$$\mathcal{U}_2 = \{\emptyset, \{A\}, \{A, B\}\},$$

$$\mathcal{U}_3 = \{\emptyset, \{B\}, \{A, B\}\},$$

$$\mathcal{U}_4 = \{\emptyset, \{A\}, \{B\}, \{A, B\}\} \text{ - the topology induced by the discrete metric on } X.$$

3. Let $X = \{A, B, C\}$. By the formula above the number of possible topologies on X is less than or equal to 64. We list a few possible collection of subsets:

$$\mathcal{U}_1 = \{\emptyset, X\} \text{ - the trivial topology}$$

$$\mathcal{U}_2 = \{\emptyset, \{A\}, \{B\}, \{C\}, \{A, B\}, \{A, C\}, \{B, C\}, X\} \text{ - the discrete topology}$$

$$\mathcal{U}_3 = \{\emptyset, X, \{A, B\}, \{A, C\}\}$$

Note that \mathcal{U}_3 is not actually a topology since

$$\{A, B\} \cap \{A, C\} = \{A\} \notin \mathcal{U}_3.$$

In fact, one can spend some time to verify (or check on Wikipedia) that there are 29 different topologies on X ; this is far fewer than the total number of collections of subsets of $\mathcal{P}(X)$ that contain \emptyset and X , which is $2^{(2^3-2)} = 2^6 = 64$.

Exercise 3.2.7. — Let X be a finite set. Prove that if a topology is generated by a metric, then it must be the discrete topology $\mathcal{P}(X)$ (see Remark 3.2.5). Does the metric need to necessarily be the discrete metric?

We will now see how the notions of convergence and continuity for metric spaces may be rephrased purely in topological terms. We start with convergence.

Theorem 3.2.8. — Let (X, d) be a metric space, (x_n) a sequence of X , and let $\ell \in X$. Then the following are equivalent:

1. (x_n) converges to ℓ .
2. for every open subset $U \subseteq X$ with $\ell \in U$, there exists an $N \in \mathbb{N}$ such that $x_n \in U$ for every $n > N$.

Proof. Suppose (x_n) converges to ℓ . For any open subset $U \subseteq X$ containing ℓ , since U is open we can find an $\epsilon > 0$ such that $B_\epsilon(\ell) \subseteq U$. Then since (x_n) converges to ℓ we find an N such that $x_n \in B_\epsilon(\ell) \subseteq U$ whenever $n > N$.

The converse is obvious — just take the ball $B_\epsilon(\ell)$ as the open subset. □

And now we see how continuity can be expressed only referring to the open sets, without the need for a distance.

Theorem 3.2.9. — *Let (X, d_X) and (Y, d_Y) be metric spaces, and let $f: X \rightarrow Y$. Then the following are equivalent:*

1. f is continuous,
2. for every open set $U \subseteq Y$, the inverse image $f^{-1}(U)$ is open in X .

Before we write the proof, let's recall how we expressed the notion of f being continuous using open balls and the direct image of a set. A function $f: X \rightarrow Y$ is continuous if for all $x_0 \in X$, for all $\epsilon > 0$ there exists $\delta > 0$ such that

$$f(B_\delta^{d_X}(x_0)) \subseteq B_\epsilon^{d_Y}(f(x_0)). \quad (3.1)$$

By applying f^{-1} to both sides of that inclusion, this is the same as requiring the inclusion

$$B_\delta^{d_X}(x_0) \subseteq f^{-1}(B_\epsilon^{d_Y}(f(x_0))). \quad (3.2)$$

We will use (3.2) instead of (3.1) in the proof below.

Proof. Let's start our proof. Assume that f satisfies the second condition. For $x_0 \in X$ we prove continuity of f at x_0 . For the $\epsilon > 0$ dictated by the definition of continuity of f at x_0 , take $U = B_\epsilon(f(x_0))$ for the open set. The second condition guarantees then that the inverse image $f^{-1}(B_\epsilon(f(x_0)))$ is open, which means that there exists $\delta > 0$ such that

$$B_\delta(x_0) \subseteq f^{-1}(B_\epsilon(f(x_0))).$$

This inclusion proves then that f is continuous at x_0 .

Now assume that f is continuous. Taking $U \subseteq Y$ open, we need to show that $f^{-1}(U)$ is open in X . For $x_0 \in f^{-1}(U)$, we aim at finding a ball centered at x_0 that is completely contained in the inverse image $f^{-1}(U)$.

3.3 Equivalent distances

1. Because U is open, we can find $\epsilon > 0$ such that $B_\epsilon(f(x_0)) \subseteq U$.
2. Because f is continuous at x_0 , for the $\epsilon > 0$ given above we can find a $\delta > 0$ such that $B_\delta(x_0) \subseteq f^{-1}(B_\epsilon(f(x_0)))$.

Therefore we have found $\delta > 0$ such that the ball $B_\delta(x_0)$ is contained in $f^{-1}(B_\epsilon(f(x_0))) \subseteq f^{-1}(U)$. This concludes our proof. \square

In the following remark we observe that the above characterisation of continuous functions could have been made with closed subsets instead of open subsets.

Remark 3.2.10. — For a function $f: X \rightarrow Y$ for every open subset $U \subseteq Y$, $f^{-1}(U)$ is open if and only if for every closed set $C \subseteq Y$, $f^{-1}(C)$ is closed.

This follows since $X \setminus f^{-1}(C) = f^{-1}(Y \setminus C)$, so $f^{-1}(Y \setminus C)$ is open if and only if $f^{-1}(C)$ is closed. Moreover, by definition C is closed if and only if $Y \setminus C$ is open.

Note that the second conditions of Theorems 3.2.8 and 3.2.9 only requires knowledge of the open subsets, and as such it can be used to extend the notions of convergence and that of continuity for *topological spaces*:

Definition 3.2.11. — A sequence (x_n) in a topological space *converges to the limit* ℓ if for every open subset $U \subseteq X$ with $\ell \in U$, there exists an $N \in \mathbb{N}$ such that $x_n \in U$ for every $n > N$.

A function f of topological spaces is *continuous* if the inverse image via f of every open set is open.

Note that, whilst' this formulation of continuity is very elegant, it is usually harder (and totally inconvenient for metric spaces) to use this in exercises to prove that a given function is continuous (or that a given sequences converges).

3.3. Equivalent distances

Because convergence and continuity for metric spaces only depend on the underlying notions of open sets, it makes sense to introduce the following notion of equivalence for distances on the same set, with the idea that two distances are equivalent if the notions of convergence and continuity defined using one are the same as those defined using the other.

Definition 3.3.1. — Let X be a set. Let d, d' be two different distances on X . Then we say that d and d' are *equivalent* whenever the open sets of (X, d) coincide with those of (X, d') . We write $d \sim d'$ to denote that two distances are equivalent.

(We could have said that that d is equivalent to d' when the topology on X induced by d is the same as the topology on X induced by d'). We shall see later in 3.3.7 that the different distances d_p on \mathbb{R}^n are all equivalent according to this definition.

First we shall see some consequences of the notion of equivalent distances. From Theorem 3.2.8 and Theorem 3.2.9, we immediately deduce the following corollaries:

Corollary 3.3.2. — Let (X, d) and (X, d') be metric spaces with $d \sim d'$, and let (x_n) be a sequence of X and $\ell \in X$. Then we have

$$x_n \xrightarrow{d} \ell \iff x_n \xrightarrow{d'} \ell.$$

Corollary 3.3.3. — Let $(X, d_X), (X, d'_X)$ and $(Y, d_Y), (Y, d'_Y)$ be metric spaces where $d_X \sim d'_X$ and $d_Y \sim d'_Y$. Then

$$\underbrace{f: (X, d_X) \rightarrow (Y, d_Y)}_{\text{is continuous}} \iff \underbrace{f: (X, d'_X) \rightarrow (Y, d'_Y)}_{\text{is continuous}}$$

Before we go any further, let's give a simple example of two metrics that are *not* equivalent.

Example 3.3.4. — Let $X = \mathbb{R}$. Then d_{discr} and d_1 are not equivalent. Indeed, the subset $[0, 1)$ is open with the discrete metric (all subsets are!) but it is not open with the standard metric d_1 .

We will now aim to prove Corollary 3.3.6, which gives an easy sufficient condition for two distances to be equivalent. That criterion will follow immediately from this Lemma.

Lemma 3.3.5. — Let $(X, d), (X, d')$ be metric spaces on the same underlying set. Suppose there exists $C > 0$ such that

$$d(x, y) \leq C \cdot d'(x, y)$$

for every x and y . Then $U \subseteq (X, d)$ is open implies that $U \subseteq (X, d')$ is also open.

3.3 Equivalent distances

Proof. For every $p \in U$, since U is open in (X, d) , we can find an $R > 0$ such that

$$B_R^d(p) = \{x : d(x, p) < R\} \subseteq U.$$

Because of the inequality $d(x, p) \leq C \cdot d'(x, p)$, we have that

$$d'(x, p) < \frac{R}{C} \implies d(x, p) < R,$$

hence the inclusion of balls

$$B_{R/C}^{d'}(p) \subseteq B_R^d(p),$$

so $B_{R/C}^{d'}(p) \subseteq U$ and therefore U is also open in (X, d') . \square

From the previous result, we immediately deduce the following.

Corollary 3.3.6. — *Suppose d, d' are distances on a set X . Then if there exist constants C and $C' > 0$ such that*

$$d(x, y) \leq C \cdot d'(x, y) \quad \text{and} \quad d'(x, y) \leq C' \cdot d(x, y)$$

for every $x, y \in X$, then d is equivalent to d' .

From this we can for example deduce that the different distances d_p that we have introduced in Chapter 1 are all equivalent on \mathbb{R}^n (in particular, they are all equivalent to the standard Euclidean distance d_2). We have observed at the end of Chapter 1 that in general (\mathbb{R}^n, d_p) and (\mathbb{R}^n, d_q) are *not* isometric. (And in fact one could even prove that they are never isometric unless of course when p equals q).

Corollary 3.3.7. — *On \mathbb{R}^n , the distances d_p and d_q are equivalent for all $p, q \geq 1$ including $p, q = \infty$.*

Proof. Suppose $p > q$. Then in Assignment 1 in B2.1 we have proved the inequalities:

$$d_\infty(x, y) \leq d_p(x, y) \leq d_q(x, y) \leq d_1(x, y) \tag{3.3}$$

and

$$d_1(x, y) \leq n \cdot d_\infty(x, y). \tag{3.4}$$

(See solutions to Assignment 1).

The result follows immediately from the above inequalities and Corollary 3.3.6 with $d = d_p$, $d' = d_q$, $C = 1$ and $C' = n$. \square

Corollary 3.3.7 is a very special feature that occurs for "finite dimensional spaces". We will now see how the other examples of metric spaces that we have introduced in Chapter 1 (spaces of sequences and spaces of continuous functions) behave very differently.

We first look at the case of the space of functions $C[0, 1]$, and compare the two distances d_{L^1} and d_{L^∞} .

Lemma 3.3.8. — *The inequality $d_{L^1}(f, g) \leq d_{L^\infty}(f, g)$ holds for all $f, g \in C[0, 1]$,*

Proof.

$$\begin{aligned} \int_0^1 |f(x) - g(x)| \, dx &\leq \int_0^1 \max |f(x) - g(x)| \, dx \\ &= (1 - 0) \cdot \max_{x \in [0,1]} |f(x) - g(x)| \\ &= d_{L^\infty}(f, g) \end{aligned}$$

□

This, by Lemma 3.3.5, says that a subset of $C[0, 1]$ that is open with the distance d_{L^1} is also open with the distance d_{L^∞} . The converse is not true, as we see in the next example.

Remark 3.3.9. — The distances d_{L^1} and d_{L^∞} are not equivalent. Set

$$f_n(x) = x^n, \quad f(x) = 0.$$

Then we have already seen that $f_n \xrightarrow{d_{L^1}} f$ whilst $f_n \not\xrightarrow{d_{L^\infty}} f$. So by Corollary 3.3.2, the distances d_{L^1} and d_{L^∞} are not equivalent.

To show that d_{L^1} and d_{L^∞} are not equivalent: we have exhibited a sequence that converges with the former distance, but not with the latter. The following exercise illustrates that the converse cannot happen.

Exercise 3.3.10. — Prove that if $f_n \xrightarrow{d_{L^\infty}} f$, then $f_n \xrightarrow{d_{L^1}} f$. Hint: use the definition of convergence and Lemma 3.3.8.

Exercise 3.3.11. — More generally, show that if d, d' are distances on X such that there exists $C > 0$ such that

$$d(x, y) \leq C \cdot d'(x, y)$$

for all $x, y \in X$, then if (x_n) is a convergent sequence in (X, d') , it must also converge in (X, d) (and to the same limit).

3.3 Equivalent distances

We now analyse the spaces of sequences, and see that if we endow the same space ℓ^p with different metrics d_q , the resulting metric spaces are *not* equivalent (note the contrast with Corollary 3.3.7).

Remark 3.3.12. — We have already discussed that, for fixed $p \geq 1$, the space of sequences ℓ^p can be endowed with distances d_q and $d_{q'}$ for all $p \leq q < q'$. These two distances are *not* equivalent, as we will see in the following two examples.

By taking the limit for $n \rightarrow \infty$ of both sides of Inequality (3.3) (both sides can be considered as sequences in n for n the dimension of \mathbb{R}^n), we deduce the inequality

$$d_{q'}(A, B) \leq d_q(A, B), \quad (3.5)$$

for all $A, B \in \ell^p$. A reverse inequality cannot be obtained by the same reasoning, because the right hand side of Inequality (3.4) tends to ∞ for $n \rightarrow \infty$!

Example 3.3.13. — Consider the space of sequences ℓ^1 — we will produce a sequence of elements of ℓ^1 that converges in d_∞ but not in d_1 . By Corollary 3.3.2, we deduce that d_1 and d_∞ are not equivalent on ℓ^1 .

Define $((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}}$ by

$$\begin{aligned} (A_{1,n})_{n \in \mathbb{N}} &= (1, 0, 0, 0, 0, \dots) \\ (A_{2,n})_{n \in \mathbb{N}} &= \left(\frac{1}{2}, \frac{1}{2}, 0, 0, 0, \dots \right) \\ (A_{3,n})_{n \in \mathbb{N}} &= \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, \dots \right) \\ &\vdots \\ (A_{k,n})_{n \in \mathbb{N}} &= \underbrace{\left(\frac{1}{k}, \dots, \frac{1}{k}, 0, \dots \right)}_{k \text{ times}} \end{aligned}$$

and $(B_n)_{n \in \mathbb{N}} = (0, 0, \dots)$. Then it is not difficult to see that

$$\begin{aligned} ((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}} &\xrightarrow{d_\infty} (B_n)_{n \in \mathbb{N}} \text{ for } k \rightarrow \infty, \quad \text{but} \\ ((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}} &\not\xrightarrow{d_1} (B_n)_{n \in \mathbb{N}} \text{ for } k \rightarrow \infty. \end{aligned}$$

(Indeed, the supremum of the difference of A_k and B equals $1/k$, but the sum of the absolute values of the difference of A_k and B equals 1 for all k).

Example 3.3.14. — Building on the previous example, we claim that on ℓ^p the distances d_q and $d_{q'}$ are *not* equivalent for all $p \leq q < q'$. (The previous example describes the case where $p = q = 1$ and $q' = \infty$.) Again we reason using Corollary 3.3.2.

On ℓ^p , by Inequality (3.5) we have that convergence with the distance d_q implies convergence with the distance $d_{q'}$. To see that the reverse implication does not hold, consider the sequences

$$\begin{aligned} ((A_{k,n})_{n \in \mathbb{N}})_{k \in \mathbb{N}} &= \left(\underbrace{\frac{1}{\sqrt[q]{k}}, \dots, \frac{1}{\sqrt[q]{k}}}_{k \text{ times}}, \dots, 0, \dots \right) \\ (B_n) &= (0, 0, 0, 0, \dots) \end{aligned}$$

We claim that (A_k) converges to B in the $d_{q'}$ -metric, but not in the d_q -metric. To see this, we compute

$$\begin{aligned} \lim_{k \rightarrow \infty} d_{q'}(A_k, B) &= \lim_{k \rightarrow \infty} \left(k \cdot \frac{1}{k^{q'/q}} \right)^{1/q'} \\ &= \lim_{k \rightarrow \infty} \left(\frac{1}{k^{\frac{q'}{q}-1}} \right)^{1/q'} = \begin{cases} 0 & \text{for } q' > q \\ 1 & \text{when } q' = q. \end{cases} \end{aligned}$$

Finally, in the following exercise we observe that any of the product distances of two metric spaces (which we defined in Chapter 1) are equivalent. The argument goes along the same lines of Corollary 3.3.7.

Exercise 3.3.15. — Let (X, d_X) and (Y, d_Y) be metric spaces. Prove the inequalities

$$D_\infty((x_1, y_1)(x_2, y_2)) \leq D_2((x_1, y_1)(x_2, y_2)) \leq D_1((x_1, y_1)(x_2, y_2))$$

and

$$D_1((x_1, y_1)(x_2, y_2)) \leq 2 \cdot D_\infty((x_1, y_1)(x_2, y_2)).$$

Use this to deduce that the three product distances D_1, D_2, D_∞ are equivalent on $X \times Y$.

3.4. Homeomorphisms

In the previous section we have discussed the topology of a metric space and we then defined a notion of when two distances on the same set ought

3.4 Homeomorphisms

to be considered equivalent from the point of view of convergence and continuity. What about the situation when we have not only two different distances, but also two different sets? Can we make a formal definition of a function that identifies two sets as well as the notions of open subsets therein? The answer is provided by the important notion of *homeomorphism* that we introduce in the next definition.

The key idea is to build on Theorem 3.2.9 and observe that one way to identify two sets and the underlying open sets is to have a bijection that is continuous and whose inverse is also continuous.

Definition 3.4.1. — Let $(X, d_X), (Y, d_Y)$ be metric spaces. We say that $f: X \rightarrow Y$ is a *homeomorphism* when

1. f is bijective and
2. both f and f^{-1} are continuous.

We say that (X, d_X) and (Y, d_Y) are *homeomorphic* when there exists a homeomorphism $f: X \rightarrow Y$.

The second condition means that the bijection f identifies the open sets of (X, d_X) with those of (Y, d_Y) by the mappings

$$U \mapsto f(U), \quad V \mapsto f^{-1}(V).$$

Note that the notion of homeomorphism makes perfect sense more generally for topological spaces.

We will drop the distance from this definition when the meaning is clear. In general we have the motto "Two spaces are homeomorphic when we can deform one to the other without tearing or breaking."

The first general example of a homeomorphism is that of an isometry.

Example 3.4.2. — Let $f: (X, d_X) \rightarrow (Y, d_Y)$ be an isometry. Then f is a homeomorphism. We know that f is a bijection $X \rightarrow Y$. Moreover, f is continuous because for all $\epsilon > 0$, taking $\delta = \epsilon$ we have $d_Y(f(x), f(x_0)) < \epsilon$ for $d_X(x, x_0) < \delta$ because $d_Y(f(x), f(x_0)) = d_X(x, x_0)$. For the same reason, f^{-1} is also continuous.

Remark 3.4.3. — Note that "being isometric" and "being homeomorphic" define two different equivalence relations on the collection of all metric spaces, and by Example 3.4.2, each equivalence class of the former is contained in an equivalence class of the latter.

Let's now see examples of homeomorphisms that are not isometries.

- Example 3.4.4.** — 1. The intervals $[0, 1]$ and $[0, 2]$ with the distance d_1 are homeomorphic. $f: X \rightarrow Y$ given by $f(x) = 2x$ is a homeomorphism with inverse $g(y) = \frac{1}{2}y$.
2. The interval $(-\pi/2, \pi/2)$ and \mathbb{R} with the distance d_1 are homeomorphic. A possible homeomorphism is given by $\tan: (-\pi/2, \pi/2) \rightarrow \mathbb{R}$ with continuous inverse given by \arctan .
3. The sets $X = \{f \in C[0, 1] : 0 \leq f(x) \leq 1\}$ and $Y = \{f \in C[0, 1] : 0 \leq f(x) \leq 2\}$ are homeomorphic with the distance d_{L^∞} . A possible homeomorphism is $\phi: X \rightarrow Y$ given by $\phi(f) = 2f$.
4. If we look at letters, O and D with the usual distance induced from \mathbb{R}^2 we could prove that they are homeomorphic while I and P are not! (We will later explain how this can be proved).
5. The interval $[0, 1]$ and \mathbb{R} (both with the distance d_1) are not homeomorphic. Also, the intervals $[0, 1]$ and $(0, 1)$ (still with the distance d_1) are not homeomorphic. We will see later why.

Example 3.4.5. — Consider $\text{Id}: (X, d) \rightarrow (X, d')$ where Id is the identity function i.e. defined by $\text{Id}(x) = x$. Then Id is a homeomorphism $\iff d$ and d' are equivalent.

Indeed, for all open sets $U \subseteq X$, $\text{Id}^{-1}(U) = U$ so we see that U is open in $(X, d) \iff U$ is open in (X, d') .

Establishing whether two spaces (topological or metric) are or not homeomorphic is, in general, a difficult problem. One of the main points of the mathematical field of *topology* is precisely to give methods to settle this. One way to establish that two topological spaces are not homeomorphic, is to try to find some property that is preserved under homeomorphisms. Such properties are called *topological properties*. In the following chapter we will study the notion of *compactness* and we will prove that it is a topological property.

4.

Completeness and compactness

Important notions to learn from this section:

1. The notion of a Cauchy sequence and that of a complete metric space.
2. Completeness of the metric space (\mathbb{R}^n, d_p) .
3. Theorem 4.1.6 relating completeness and closedness.
4. The notions of limit inferior and limit superior for a sequence.
5. The notion of a contraction, and the Contraction Mapping Theorem (CMT).
6. The notion of a compact metric space.
7. Compactness is preserved by continuous surjections. The Min/Max property of compact metric spaces. Compactness is a topological property.
8. Theorem 4.4.11 characterising the compact subsets of (\mathbb{R}^n, d_p) .
9. Compact implies complete.

In Year 1 (MATH101) you defined the set of real numbers using one of their axiomatic characterisations. The idea was to observe that \mathbb{R} is a *field* (it has two operations that satisfy certain properties), it is *ordered* (it has an ordering that respects the operations in the field), and it is *complete*. The last property was stated without proof. An important fact that was highlighted

but not proven is the fact that the set of real numbers is *unique* with respect to the above three properties. (More precisely, if F is any other complete ordered field, there exists a unique field isomorphism $\mathbb{R} \rightarrow F$ (defined in MATH247), and that field isomorphism preserves the order).

Among those listed above, the property that we here want to generalise is *completeness*. For the set of real numbers, that property was stated as follows.

Theorem 4.0.1 (Fact 1.3 in MATH101). — *Let $\emptyset \neq A \subseteq \mathbb{R}$. Then the set of upper bounds*

$$Ub(A) = \{x \in \mathbb{R} : x \geq a \text{ for all } a \in A\}$$

is either empty or it has a minimum called the supremum of A , denoted by $\sup(A)$. (If $Ub(A)$ is empty we also say that $\sup(A) = +\infty$).

Similarly, the set of lower bounds

$$Lb(A) := \{x \in \mathbb{R} : x \leq a \text{ for all } a \in A\}$$

is either empty or it has a maximum called the infimum of A , denoted by $\inf(A)$. (If $Lb(A)$ is empty we also say that $\inf(A) = -\infty$).

The property of the real numbers that is highlighted in the above result, and which was called completeness, makes essential use of the ordering of real numbers, and as such it is not viable for generalisation to arbitrary metric spaces. For this reason, in MATH241 we will use the word *complete* for a different notion that we will introduce via the notion of Cauchy convergence of sequence. When the metric space is \mathbb{R} with the distance d_1 , this new notion is equivalent to the property stated in Theorem 4.0.1.

After having introduced the notion of completeness, we discuss and prove the main result that makes essential use of that property, which is the celebrated *contraction mapping theorem* (CMT). The CMT will be used in later chapters to prove some other important results in real analysis.

The last part of this chapter is devoted to the notion of *compactness* for metric spaces. Compactness is formally a generalisation of finiteness for a metric space. It is important theoretically, because it allows to distinguish metric spaces that are not homeomorphic (when one is compact and the other isn't). It is also crucial in applications, because of the Min/Max property (which generalises the analogue property seen in Year 1).

4.1 Cauchy convergence and completeness

4.1. Cauchy convergence and completeness

The main idea is to define an alternative notion of convergence of sequences. One unsatisfactory issue in the notion of convergence of a sequence (x_n) in a metric space X is that one has to introduce an extra $\ell \in X$ to make sense of it. Can we make a notion of convergence that only refers to the elements of (x_n) without first having to guess a potential limit ℓ ? One possible way is given by the notion of Cauchy convergence.

Definition 4.1.1. — Let (X, d) be a metric space and let (x_n) be a sequence of X . We say that it is *Cauchy convergent* (or just *Cauchy*) if for every $\epsilon > 0$ there exists an $N \in \mathbb{N}$ such that $d(x_n, x_m) < \epsilon$ for every $n, m > N$.

The following shows that Cauchy convergence generalises the notion of convergence that we introduced in Chapter 2.

Lemma 4.1.2. — *If (x_n) is convergent, then (x_n) is Cauchy.*

Proof. Since (x_n) is convergent, there exists ℓ such that for every $\epsilon' > 0$ there exists an $N \in \mathbb{N}$ such that $d(x_n, \ell) < \epsilon'$ for all $n > N$. Take $\epsilon' = \epsilon/2$. Then for every $n, m > N$, we have

$$d(x_n, x_m) \leq d(x_n, \ell) + d(x_m, \ell) < \epsilon/2 + \epsilon/2 = \epsilon.$$

□

At this point, one may wonder if all Cauchy sequences converge. The answer, in general, is negative.

Example 4.1.3. — 1. Let $(X, d) = (\mathbb{Q}, d_1)$. Then $\pi = 3.1415\dots \notin \mathbb{Q}$. Let (x_n) be the decimal approximation of π to the n -th digit. Then (x_n) is Cauchy.

This follows since $|x_n - x_m| < \frac{2}{10^N}$ for $m, n > N$. Therefore, for every $\epsilon > 0$, we can find $N > 0$ such that $\frac{2}{10^N} < \epsilon$ which shows that $|x_n - x_m| < \epsilon$. This sequence however does *not* converge in (\mathbb{Q}, d_1) (because it converges to the irrational number π in (\mathbb{R}, d_1)).

2. Let $(X, d) = ((0, 1), d_1)$, $(x_n) = (\frac{1}{n+1})$. Then (x_n) is Cauchy but it does not converge (in $((0, 1), d_1)$). Let's prove this more conceptually than in the previous example.

First of all (x_n) converges to 0 in (\mathbb{R}, d_1) , of which $((0, 1), d_1)$ is a subspace metric. Therefore the sequence (x_n) is Cauchy, because it converges in (\mathbb{R}, d_1) . However (x_n) does not converge in $(0, 1)$ because it converges to 0 in (\mathbb{R}, d_1) , and because limits in metric spaces are unique.

We call *complete* the metric spaces for which the answer is always (= for all sequences) affirmative:

Definition 4.1.4. — A metric space (X, d) is *complete* if every Cauchy sequence converges.

Example 4.1.3 shows that the metric spaces (\mathbb{Q}, d_1) and $((0, 1), d_1)$ are not complete. The main example of a complete metric space is (\mathbb{R}, d_1) , a proof will soon be given in Theorem 4.2.1.8.

Here are some further important examples of complete metric spaces that we will consider in this module.

Example 4.1.5. — 1. For all n, p , the metric space (\mathbb{R}^n, d_p) is complete. (A proof is given in Corollary 4.2.2.3).

2. The space of continuous functions $(C[0, 1], d_{L^\infty})$ is complete. (A proof will be given in Chapter 5, but you may already assume this fact in this week's exercises).

3. The space of sequences (ℓ_p, d_p) is complete (we will not prove this).

4. Every discrete metric space (X, d_{discr}) is complete.

Proof. (That every discrete metric space is complete.) To see the last example is complete, let (x_n) be a Cauchy sequence in X . Take $\epsilon = 1/2$, so there exists an N such that $d_{\text{discr}}(x_n, x_m) < 1/2$ for every $n, m > N$, so (x_n) is constant after N , so (x_n) converges. \square

Assuming these results for now, here is a way to produce lots of other examples of complete metric spaces.

Theorem 4.1.6. — Let (X, d) be a metric space, and $A \subseteq X$ be a subset, and let d_A be the distance induced by d on the subset A .

1. If (A, d_A) is complete, then A is closed in (X, d_X) .
2. If (X, d) is complete and A is closed in (X, d) , then (A, d_A) is also complete.

4.1 Cauchy convergence and completeness

Note that this implies that if (X, d) is complete and $Y \subset X$, then (Y, d) is complete if and only if $Y \subset X$ is closed.

Proof. Both assertions follow as an application of Lemma 3.1.8.

For the first assertion, let (a_n) be a sequence of A which converges to $\ell \in X$. By Lemma 3.1.8, to prove that A is closed it is enough to prove that $\ell \in A$. Then since (a_n) converges in X , it is Cauchy in X , and therefore it is Cauchy in A too. Then since (A, d_A) is complete, (a_n) converges to some $\ell' \in A$. Finally, by uniqueness of the limit, we must have that $\ell = \ell' \in A$.

For the second assertion, suppose A is closed. Let (a_n) be a Cauchy sequence of A . Then (a_n) is Cauchy in X too. Since (X, d) is complete, (a_n) converges to some $\ell \in X$. Then since A is closed, by Lemma 3.1.8 the limit ℓ must belong to A . Therefore (a_n) is convergent in A and so A is complete. \square

We observe that two metric spaces may be homeomorphic, one of them be complete, but not the other.

Remark 4.1.7. — We have already observed that $\underbrace{((-\pi/2, \pi/2), d_1)}_{\text{Not complete}}$ is homeomorphic to $\underbrace{(\mathbb{R}, d_1)}_{\text{Complete}}$, from which we see that completeness really depends on the metric, not only on the underlying topology (the collection of open subsets).

It is natural at this point to ask whether it is possible for two equivalent metrics on the same set to have different Cauchy sequences. (Note that from the previous Chapter we know that this does not happen for convergent sequences!).

Exercise 4.1.8. — Give an example of a set X and two equivalent distances d, d' and a sequence (x_n) that is Cauchy for d but not for d' . (Hint: combine the example of Remark 4.1.7 with the 4th part of Exercise 1.6.11).

Exercise 4.1.9. — (To be compared with Lemma 3.3.5 and with Exercise 3.3.11).

Let X be a set and d, d' be two distances on X such that there exists $C > 0$ such that

$$d(x, y) \leq C \cdot d'(x, y)$$

for all $x, y \in X$. Prove that if (x_n) is a Cauchy sequence in (X, d') , then (x_n) is also Cauchy in (X, d) .

We conclude this section by giving an important example of a non complete metric space. (The details of the example should be considered as EXTRA material.)

Example 4.1.10. — The space $C[0, 1]$ is *not* complete when endowed with the distance d_{L^1} . Indeed, consider the sequence of functions

$$f_n(x) = \begin{cases} 0 & \text{for } 0 \leq x \leq \frac{n}{2(2+n)} \\ (n+2)x - \frac{n}{2} & \text{for } \frac{n}{2(2+n)} \leq x \leq \frac{1}{2} \\ 1 & \text{for } x \geq \frac{1}{2} \end{cases}$$

Hint: draw a picture of f_n . The value of each function f_n is 1 after $1/2$, and it is zero before $\frac{1}{2} - \frac{1}{n+1} = \frac{n}{2(n+2)}$, and in between it is the unique line that connects the points

$$\left(\frac{n}{2(n+2)}, 0\right) \text{ and } \left(\frac{1}{2}, 1\right).$$

In particular, each f_n is a continuous function $[0, 1] \rightarrow \mathbb{R}$. We will show that (f_n) is a Cauchy sequence, and that it does not converge in $C[0, 1]$. The idea is that if (f_n) converged in d_{L^1} , its limit should be the *discontinuous* function

$$g(x) = \begin{cases} 0 & \text{when } 0 \leq x < \frac{1}{2} \\ 1 & \text{when } \frac{1}{2} \leq x \leq 1. \end{cases}$$

To prove that (f_n) is Cauchy we observe that the sequence

$$I(n) = \int_0^1 |f_n(x)| \, dx = \frac{1}{2} + \int_{\frac{n}{2(n+2)}}^{\frac{1}{2}} f_n(x) \, dx = \frac{1}{2} + \frac{1}{2(n+2)} = \frac{n+3}{2n+4}$$

of integrals is Cauchy in (\mathbb{R}, d_1) . This implies that (f_n) is Cauchy because $f_n(x) \leq f_m(x)$ for all $x \in [0, 1]$.

Finally, we prove that the sequence (f_n) does not converge in $(C[0, 1], d_{L^1})$. If it did, let g be that limit. We will find a contradiction by proving that g must equal 0 on the interval $[0, 1/2]$ and that it must equal 1 on the interval $[1/2, 1]$.

Indeed, since g is the limit of (f_n) , we have

$$d_{L^1}(f_n, g) = \int_0^{\frac{1}{2}} |g(x) - f_n(x)| \, dx + \int_{\frac{1}{2}}^1 |g(x) - f_n(x)| \, dx \rightarrow 0$$

4.2 Completeness of \mathbb{R}^N

Because both summands are ≥ 0 , each of them converges to zero.

For the second summand, this implies

$$\int_{\frac{1}{2}}^1 |g(x) - 1| dx < \epsilon$$

for all $\epsilon > 0$, which implies

$$\int_{\frac{1}{2}}^1 |g(x) - 1| dx = 0,$$

which implies that g equals 1 on the interval $[\frac{1}{2}, 1]$ because of Lemma 1.4.5 and the assumption that g is continuous.

A similar result is achieved for the first summand. Observe that on $C[0, 1/2]$ (the set of continuous functions from $[0, 1/2]$ to \mathbb{R}) we also have a distance d_{L^1} defined as for $C[0, 1]$. Then by restricting f_n and g to the interval $[0, 1/2]$ we obtain elements of $C[0, 1/2]$, and by uniqueness of the limit in $(C[0, 1/2], d_{L^1})$, and the fact that

$$\int_0^{\frac{1}{2}} |g(x) - f_n(x)| dx \rightarrow 0$$

we deduce that g equals 0 on the interval $[0, 1/2]$, a contradiction.

4.2. Completeness of \mathbb{R}^N

This section is entirely devoted to the proof that (\mathbb{R}^N, d_p) is a complete metric space.

4.2.1. The case $N = 1$ — The one-dimensional case essentially follows from material seen in Year 1, which we now review in detail. (We will rely on 4.0.1, which we will assume without proof).

Remark 4.2.1.1. — Suppose (x_n) is an *increasing* sequence of real numbers (i.e. $x_n \leq x_{n+1}$ for all n). Then

1. (x_n) not bounded above $\Rightarrow \lim_{n \rightarrow \infty} x_n = +\infty$
2. (x_n) bounded above $\Rightarrow \exists \ell \in \mathbb{R} : \lim_{n \rightarrow \infty} x_n = \ell$

Similarly, if (y_n) is a *decreasing* sequence of real numbers, then if it is not bounded below, it has limit $-\infty$, and if it is bounded below, it has a finite limit.

Proof. Let $A = \{x_n, n \in \mathbb{N}\}$ and assume that A is bounded above (the proof of the unbounded case is similar and left to the reader). Define $\ell := \sup(A)$. We want to show that $\ell = \lim_{n \rightarrow \infty} x_n$.

First of all, note that $x_n \leq \ell$ for all n since the supremum is an upper bound. Since further to this the supremum is the *minimum* among all upper bounds, for every $\epsilon > 0$, there exists an $N \in \mathbb{N}$ such that $x_N > \ell - \epsilon$. Then, since (x_n) is increasing, for every $n > N$ we have that

$$\ell - \epsilon < x_N \leq x_n \leq \ell$$

so in particular $\lim_{n \rightarrow \infty} x_n = \ell$. □

While all *monotonic* real sequences (i.e. increasing or decreasing) always admit a limit (which may be $\pm\infty$), the situation for an arbitrary real (A_n) is more complicated, as this is very often not the case. This motivates the following definitions:

Definition 4.2.1.2. — We define the *limit superior* and the *limit inferior* of a sequence of real numbers (A_n) respectively to be

$$\limsup_{n \rightarrow \infty} (A_n) := \begin{cases} \lim_{m \rightarrow \infty} \left(\sup_{n \geq m} A_n \right) & \text{when } (A_n) \text{ is bounded above, and} \\ +\infty & \text{otherwise;} \end{cases}$$

$$\liminf_{n \rightarrow \infty} (A_n) := \begin{cases} \lim_{m \rightarrow \infty} \left(\inf_{n \geq m} A_n \right) & \text{when } (A_n) \text{ is bounded below, and} \\ -\infty & \text{otherwise.} \end{cases}$$

Remark 4.2.1.3. — By Remark 4.2.1.1, these limits always exist (although they might possibly equal $\pm\infty$).

Indeed, let (A_n) be a real sequence and define

$$B_m = \sup_{n \geq m} A_n,$$

then since

$$B_m = \sup_{n \geq m} A_n \geq \sup_{n \geq m+1} A_n = B_{m+1},$$

we see that (B_m) is a decreasing sequence and hence it has a limit. (Note that if the sequence (A_n) is not bounded above, then the sequence (B_m) is a sequence of elements in $\mathbb{R} \cup \{+\infty\}$).

(A similar discussion holds for the $\lim \inf$).

4.2 Completeness of \mathbb{R}^N

Example 4.2.1.4. — Let $(A_n) = ((-1)^n \frac{n+2}{n+1})$. The first terms of this sequence are:

$$\left(2, -\frac{3}{2}, \frac{4}{3}, -\frac{5}{4}, \frac{6}{5}, -\frac{7}{6}, \dots\right)$$

Let $(B_m) = \sup_{n \geq m} (A_n)$, so that

$$(B_m) = \left(2, \frac{4}{3}, \frac{4}{3}, \frac{6}{5}, \frac{6}{5}, \dots\right)$$

and so we see that $\lim_{m \rightarrow \infty} B_m = 1$. Similarly, if we set $C_m = \inf_{n \geq m} (A_n)$, we see that $\lim_{m \rightarrow \infty} C_m = -1$. We can see this more directly if we notice that for all $k \in \mathbb{N}$ we have $A_{2k+1} = -\frac{2k+3}{2k+2} \rightarrow -1$, whilst $A_{2k} = \frac{2k+2}{2k+1} \rightarrow +1$.

Remark 4.2.1.5. — Since the infimum of a subset of \mathbb{R} is always less than or equal to the supremum of that set, we have on taking limits that

$$\limsup_{n \rightarrow \infty} (A_n) \geq \liminf_{n \rightarrow \infty} (A_n)$$

Lemma 4.2.1.6. — Let (A_n) be a sequence of real numbers, and let $\ell \in \mathbb{R}$. Then the following are equivalent:

1. (A_n) converges to ℓ .
2. $\liminf_{n \rightarrow \infty} (A_n) = \limsup_{n \rightarrow \infty} (A_n) = \ell$

Proof. Suppose that (A_n) converges to ℓ . Then for every $\epsilon > 0$ we find an N such that $\ell - \epsilon < A_n < \ell + \epsilon$ for every $n > N$. We deduce then

$$\Rightarrow \ell - \epsilon \leq \sup_{n \geq N} A_n \leq \ell + \epsilon.$$

This means that

$$\limsup (A_n) = \ell.$$

A similar argument shows that $\liminf (A_n) = \ell$.

Conversely, suppose that $\liminf_{n \rightarrow \infty} (A_n) = \limsup_{n \rightarrow \infty} (A_n) = \ell$. Then for every $\epsilon > 0$ we find N_1, N_2 so that

$$\ell - \epsilon < \sup_{n \geq N_1} A_n < \ell + \epsilon$$

$$\ell - \epsilon < \inf_{n \geq N_2} A_n < \ell + \epsilon$$

From this, for $N \geq \max(N_1, N_2)$ we obtain the chain of inequalities

$$\ell - \epsilon < \inf_{n \geq N} A_n \leq \sup_{n \geq N} A_n < \ell + \epsilon,$$

in particular

$$\ell - \epsilon < A_n < \ell + \epsilon$$

for all $n > N$, hence the real sequence (A_n) converges to ℓ as claimed. \square

Exercise 4.2.1.7. — Show the inequalities

$$\limsup_{n \rightarrow \infty} (A_n + B_n) \leq \limsup_{n \rightarrow \infty} A_n + \limsup_{n \rightarrow \infty} B_n.$$

$$\liminf_{n \rightarrow \infty} (A_n + B_n) \geq \liminf_{n \rightarrow \infty} A_n + \liminf_{n \rightarrow \infty} B_n.$$

(Once one knows that there are such inequalities, to remember which way they go one can use the example $(A_n) = ((-1)^n)$ and $(B_n) = ((-1)^{n+1})$.)

Hint: use the inequality $\max(a_1 + b_1, a_2 + b_2) \leq \max(a_1, a_2) + \max(b_1, b_2)$. Taking limits gives the same inequality for sup, and taking limits again gives the inequality for lim sup. (The inequality for the limit inferior follows similarly).

We deduce completeness of (\mathbb{R}, d_1) by exploiting the properties of lim inf and lim sup (which is assuming Theorem 4.0.1 behind the scenes).

Theorem 4.2.1.8. — The metric space (\mathbb{R}, d_1) is complete.

Proof. Let (a_n) be a Cauchy sequence of (\mathbb{R}, d_1) . Let

$$\underline{A} = \liminf_{n \rightarrow \infty} (a_n), \quad \overline{A} = \limsup_{n \rightarrow \infty} (a_n).$$

By Remark 4.2.1.5 we have that $\underline{A} \leq \overline{A}$. To prove that they are equal, we aim to prove the inequality $\overline{A} \leq \underline{A} + \epsilon$ for all $\epsilon > 0$.

Since (a_n) is Cauchy, for any $\epsilon > 0$, we can find an N such that

$$-\epsilon/2 < a_n - a_m < \epsilon/2$$

for every $n, m > N$. On taking lim sup we deduce the inequalities

$$-\epsilon/2 \leq \overline{A} - a_m \leq \epsilon/2 \tag{4.1}$$

for all $m > N$. Now by the definition of lim inf, there is some $k > N$ such that

$$a_k \leq \underline{A} + \epsilon/2$$

4.2 Completeness of \mathbb{R}^N

By combining this with the second inequality of (4.1), we find

$$\bar{A} \leq a_k + \epsilon/2 \leq \underline{A} + \epsilon,$$

which concludes our proof. \square

4.2.2. The case $N > 1$ — Let us now move to the proof that (\mathbb{R}^N, d_p) is complete. (Here we use N for the exponent \mathbb{R}^N to avoid confusions with the index n of a sequence $(x_n)_{n \in \mathbb{N}}$!) The idea of the proof is to reduce to the case when $N = 1$ by proving that a sequence converges (resp. it is Cauchy) in \mathbb{R}^N if and only if all its coordinates converge (resp. are Cauchy sequences) in \mathbb{R} .

Lemma 4.2.2.1. — A sequence (x_n) of (\mathbb{R}^N, d_p) converges to $\ell \in \mathbb{R}^N$ if and only if for all $j = 1, \dots, N$, the sequence of j -th coordinates $(x_{j,n})$ of (x_n) converges to the j -th coordinate ℓ_j of ℓ in (\mathbb{R}, d_1) .

Proof. By Corollaries 3.3.2 and 3.3.7 we may assume $p = \infty$.

\implies For all $\epsilon > 0$, the definition of convergence of $(x_n) \rightarrow \ell$ gives M such that

$$\max_{j=1, \dots, N} |x_{j,n} - \ell_j| < \epsilon$$

for all $n > M$. This implies that each coordinate

$$|x_{j,n} - \ell_j| < \epsilon$$

for all $n > M$.

\impliedby The fact that each coordinate $(x_{j,n})$ converges to ℓ_j gives, for all $\epsilon > 0$ a M_j such that

$$|x_{j,n} - \ell_j| < \epsilon$$

for all $n > M_j$. Taking M to be the maximum of M_1, \dots, M_N we have

$$\max_{j=1, \dots, N} |x_{j,n} - \ell_j| < \epsilon$$

for all $n > M$.

\square

A similar argument holds for Cauchy sequences.

Lemma 4.2.2.2. — A sequence (x_n) of (\mathbb{R}^N, d_p) is Cauchy if and only if for all $j = 1, \dots, N$, the sequence of j -th coordinates $(x_{j,n})$ is Cauchy in (\mathbb{R}, d_1) .

For the proof of this Lemma we cannot argue exactly as in the proof of the previous one. Indeed, it is in general false that equivalent metrics always preserve the Cauchy property of a sequence (this was already observed in Exercise 4.1.8).

Proof. From the inequalities $d_\infty(x, y) \leq d_p(x, y)$ and $d_p(x, y) \leq N \cdot d_\infty(x, y)$ combined with Exercise 4.1.9 we deduce that a sequence is Cauchy with the distance d_p if and only if it is Cauchy with the distance d_∞ .

Therefore, we can assume $p = \infty$. The proof is then very similar to the proof of 4.2.2.1, and we leave it as an exercise. □

Corollary 4.2.2.3. — The metric space (\mathbb{R}^N, d_p) is complete for all N and $p \geq 1$ (including $p = \infty$).

Proof. Assume (x_n) is Cauchy. Then each coordinate sequence $(x_{j,n})$ is Cauchy for $j = 1, \dots, N$ by Lemma 4.2.2.2. By Theorem 4.2.1.8 each such sequence admits a limit l_j in (\mathbb{R}, d_1) . We conclude that (x_n) converges to $l = (l_1, \dots, l_N)$ by applying Lemma 4.2.2.1. □

4.3. The Contraction Mapping Theorem (CMT)

The main reason why completeness is important in this module is that it is the essential ingredient in the existence part of the proof of the Contraction Mapping Theorem, the main result on metric spaces that we discuss in MATH241.

Definition 4.3.1. — Let X be a set, $f: X \rightarrow X$ a function and let $p \in X$. We say that p is a fixed point of f if $f(p) = p$.

An interesting, general problem in mathematics is to decide if a given f has a unique fixed point.

Example 4.3.2. — 1. Let $X = \mathbb{R}$. Assume f is continuous, and satisfies $f(a) < a, f(b) > b$ for some $a, b \in \mathbb{R}$. Then $f: \mathbb{R} \rightarrow \mathbb{R}$ has a fixed point in $[a, b]$ by the Intermediate Value Theorem (IVT) from Year 1.

4.3 The Contraction Mapping Theorem (CMT)

2. Suppose $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is linear, so that $f(x) = Ax$ for some $n \times n$ matrix A . Then $x = 0$ is a fixed point. Suppose we have another fixed point $x \neq 0$.

$$\begin{aligned} Ax = x &\iff Ax - x = 0 \\ &\iff (A - \text{Id})x = 0 \end{aligned}$$

Therefore 0 is the unique fixed point $\iff (A - \text{Id})$ is invertible \iff the matrix A does not have 1 as an eigenvalue.

Also, the set of fixed points has the structure of a vector space, indeed they form the eigenspaces of the eigenvalue 1.

Definition 4.3.3. — Let (X, d) be a metric space. Then $f: X \rightarrow X$ is a *contraction* when there exists $0 \leq L < 1$ such that $d(f(x), f(y)) \leq L \cdot d(x, y)$ for every $x, y \in X$.

Theorem 4.3.4 (Contraction Mapping Theorem). — *Suppose (X, d) is a complete metric space. If $f: (X, d) \rightarrow (X, d)$ is a contraction, then f has a unique fixed point.*

Remark 4.3.5. — The definition of a contraction cannot be relaxed to

$$d(f(x), f(y)) < d(x, y)$$

for all $x, y \in X$, or the existence of a fixed point in Theorem 4.3.4 would in general be lost.

To illustrate this, consider $f: [1, \infty) \rightarrow [1, \infty)$ given by $f(x) = x + \frac{1}{x}$. Then

$$\begin{aligned} |f(x) - f(y)| &= \left| x + \frac{1}{x} - y - \frac{1}{y} \right| = \left| \frac{x^2y + y - xy^2 - x}{xy} \right| \\ &= \left| \frac{(x - y)(xy - 1)}{xy} \right| \\ &= |x - y| \left| \frac{xy - 1}{xy} \right|. \end{aligned}$$

Then clearly $\left| \frac{xy-1}{xy} \right| < 1$ whenever $xy > 1$, so f satisfies our proposed condition. However f is not a contraction as it clearly has no fixed points, despite $[1, \infty)$ being complete with the distance d_1 .

As an example, we see how one can verify that f is a contraction on the intervals in the real line.

Lemma 4.3.6 (Criterion for a real function of one real variable to be a contraction.). — Let $f: [a, b] \rightarrow [a, b]$ be differentiable with $|f'(x)| \leq L < 1$ for every $x \in [a, b]$. Then f is a contraction when $[a, b]$ is endowed with the distance d_1 .

With the same hypotheses as above, if there is a point $x_0 \in [a, b]$ such that $|f'(x_0)| > 1$ then we can show that f is not a contraction by taking two points sufficiently close to x_0 .

Proof. By the Mean Value Theorem, for every $x_1, x_2 \in [a, b]$ there exists $\xi \in (x_1, x_2)$ such that

$$|f(x_2) - f(x_1)| = |f'(\xi)| \cdot |x_2 - x_1| \leq L \cdot |x_2 - x_1|.$$

□

To prove the Contraction Mapping Theorem, we will first prove that a contraction is a continuous function.

Lemma 4.3.7. — Suppose $f: (X, d) \rightarrow (X, d)$ is a contraction. Then f is continuous.

Proof. For every $\epsilon > 0$, if $L = 0$ in the above definition, this implies that f is constant, hence continuous. If, on the other hand, we have $0 < L < 1$, we take $\delta = \frac{\epsilon}{L}$. Then

$$d(f(x), f(y)) \leq L \cdot d(x, y) < L \cdot \frac{\epsilon}{L} = \epsilon$$

whenever $d(x, y) < \delta$.

□

Note that the proof above never uses that $0 \leq L < 1$.

Proof of Theorem 4.3.4. Let (X, d) be a complete metric space, and suppose $f: X \rightarrow X$ is a contraction.

1. *Uniqueness.* Suppose that x, y are fixed by f . Then

$$d(x, y) = d(f(x), f(y)) \leq L \cdot d(x, y)$$

where $L \in [0, 1)$. This implies $d(x, y) = 0 \iff x = y$.

4.3 The Contraction Mapping Theorem (CMT)

2. For every $x \in X$, we define a sequence

$$(x_n) = (x, f(x), f(f(x)), \dots, f^n(x), \dots)$$

whose first element is x , and whose next element is iteratively obtained by applying f to the previous element.

Then (x_n) is Cauchy. To show this, we start by computing the distance of two consecutive elements of the sequence:

$$\begin{aligned} d(x_{n+1}, x_n) &= d(f(x_n), f(x_{n-1})) \leq L \cdot d(x_n, x_{n-1}) \\ &\leq L^n \cdot d(x_1, x_0). \end{aligned}$$

Now assume that $n > m$. Then

$$\begin{aligned} d(x_n, x_m) &\leq d(x_n, x_{n-1}) + d(x_{n-1}, x_{n-2}) + \dots + d(x_{m+1}, x_m) \\ &= \sum_{i=0}^{n-m-1} d(x_{m+i+1}, x_{m+i}) \\ &\leq d(x_1, x_0) \cdot \sum_{i=0}^{n-m-1} L^{m+i} \\ &= d(x_1, x_0) \cdot L^m \cdot \frac{1 - L^{n-m}}{1 - L} \leq d(x_1, x_0) \cdot \frac{L^m}{1 - L} \end{aligned}$$

so for every $\epsilon > 0$, we take an N such that

$$\frac{d(x_1, x_0)}{1 - L} \cdot L^N < \epsilon,$$

which we can do since $0 \leq L < 1$. By the previous calculation, we have

$$d(x_n, x_m) < \epsilon$$

for all $n, m > N$, which means that (x_n) is Cauchy.

3. If p is the limit of the above sequence (x_n) (whose existence is guaranteed because (X, d) is complete and (x_n) is Cauchy), then p is a fixed point of f . This follows because

$$p = \lim_{n \rightarrow \infty} (x_n) = \lim_{n \rightarrow \infty} (x_{n+1}) = \lim_{n \rightarrow \infty} f(x_n) = f\left(\lim_{n \rightarrow \infty} x_n\right) = f(p),$$

where the penultimate equality follows because f is continuous.

□

Note that the proof of uniqueness doesn't need the metric space to be complete. In other words, if f is a contraction, we have proved that f has *at most* one fixed point (it may have none when the metric space is not complete).

Example 4.3.8. — 1. Let $f: [0, 3/2] \rightarrow [0, 3/2]$ be given by $f(x) = \frac{x^2+3}{4}$. f is a contraction: $f'(x) = \frac{1}{2}x$, so $|f'(x)| \leq 3/4$. $[0, 3/2]$ is complete with the d_1 distance. Therefore f has a unique fixed point.

In this example we can even calculate that fixed point explicitly, and independently of the CMT: $f(x) = x \Rightarrow x^2 - 4x + 3 = 0$ which has the unique solution $x = 1$ in the interval $[0, 3/2]$.

2. Let $\phi: (C[0, 1], d_{L^\infty}) \rightarrow (C[0, 1], d_{L^\infty})$ be given by

$$\phi(f)(x) = x + \frac{1}{5} \left(f(x) + f\left(\frac{e^x - 1}{e - 1}\right) \right).$$

We ask again the question of whether ϕ admits a unique fixed point. In this case we will be able to answer the question by applying the CMT, and I wouldn't know how to find the fixed point explicitly as in the previous example.

The metric space $(C[0, 1], d_{L^\infty})$ is complete, so by the CMT, to show that ϕ has a fixed point, it is enough to show that it is a contraction. Note that the image of the function $x \mapsto \frac{e^x - 1}{e - 1}$ for $x \in [0, 1]$ is indeed $[0, 1]$. Then

$$\begin{aligned} d_{L^\infty}(\phi(f), \phi(g)) &= \frac{1}{5} \max \left| f(x) + f\left(\frac{e^x - 1}{e - 1}\right) - g(x) - g\left(\frac{e^x - 1}{e - 1}\right) \right| \\ &\leq \frac{1}{5} \max |f(x) - g(x)| + \frac{1}{5} \max \left| f\left(\frac{e^x - 1}{e - 1}\right) - g\left(\frac{e^x - 1}{e - 1}\right) \right| \\ &= \frac{1}{5} d_{L^\infty}(f, g) + \frac{1}{5} d_{L^\infty}(f, g) \\ &= \frac{2}{5} d_{L^\infty}(f, g) \end{aligned}$$

which shows that ϕ is indeed a contraction as claimed.

We conclude with one remark, where we observe that the fixed point of Theorem 4.3.4 can be effectively approximated.

Remark 4.3.9. — In the proof of Theorem 4.3.4, we showed the inequality

$$d(x_n, x_m) \leq d(x_1, x_0) \cdot \frac{L^m}{1 - L}, \tag{4.2}$$

4.4 Compactness

for $n > m$, where x_k was defined by $f^k(x)$ for all $k \in \mathbb{N}$ and x was an arbitrary point of X .

From (4.2) and taking the limit for $n \rightarrow \infty$, we deduce the following estimate

$$d(\ell, x_m) \leq d(x_1, x_0) \cdot \frac{L^m}{1-L}, \quad (4.3)$$

(to take the limit, here we used the fact that for all $y \in X$, the function $d_y: X \rightarrow \mathbb{R}$ defined by $d_y(p) = d(y, p)$ is continuous from (X, d) to (\mathbb{R}, d_1)).

We may use (4.3) to approximate the fixed point ℓ of Theorem 4.3.4, in the following sense.

For any $\epsilon > 0$, we may construct a value y whose distance from the fixed point ℓ of f is smaller than ϵ , by picking an arbitrary point $x \in X$, and then taking m large enough so that

$$d(f(x), x) \cdot \frac{L^m}{1-L} < \epsilon.$$

(It is clear that this can be done, for we can make L^m arbitrarily small, since $0 \leq L < 1$). Then by (4.3), the value $y = f^m(x)$ has the property that the distance $d(y, \ell)$ is smaller than ϵ .

4.4. Compactness

In this section we introduce the notion of compactness for metric spaces. This is a *topological property* for metric spaces, i.e. if two metric spaces are homeomorphic, one is compact if and only if the other is. Compactness is key in applications because of its min/max property: a continuous function from a *compact* metric space to \mathbb{R} always admits a maximum and a minimum. (This generalises the Min/Max Theorem for functions $f: [a, b] \rightarrow \mathbb{R}$ from Year 1).

Note that the standard notion of compactness introduced in many textbooks looks different (and more abstract) than the one we introduced here, and it makes use of open covers. The notion we introduced here makes use of subsequences, and it is sometimes referred to as *sequential compactness*. (However, one could prove that the two notions of compactness are equivalent for metric spaces).

Definition 4.4.1. — Let $(x_n)_{n \in \mathbb{N}}$ be a sequence in X . Let $(n_k)_{k \in \mathbb{N}}$ be a strictly increasing sequence of natural numbers. We say that $(x_{n_k})_{k \in \mathbb{N}}$ is a *subsequence* of $(x_n)_{n \in \mathbb{N}}$.

For example, a subsequence of

$$(0, 1, 4, 9, 16, 25, \dots, n^2, \dots) \tag{4.4}$$

is

$$(0, 4, 16, 36, \dots, (2k)^2, \dots)$$

(where $n_k = 2k$) and also

$$(4, 9, 16, 25, \dots, (k+2)^2, \dots)$$

(where $n_k = k+2$), but not

$$(1, 0, 9, 4, 25, 16, \dots).$$

Also, a subsequence must be a sequence, so the vector

$$(0, 1, 4, 9) \in \mathbb{R}^4$$

is also not a subsequence of the sequence defined in (4.4).

Note that if a sequence (x_n) of a metric space (X, d) converges to some limit ℓ , then all of its subsequences also converge to the same ℓ .

Definition 4.4.2. — We say that a metric space (X, d) is *compact* if for every sequence of X there exists a subsequence that converges in (X, d) .

Informally, we say that a metric space is compact if every sequence admits a convergent subsequence.

Remark 4.4.3. — In many textbooks this property is referred to as *sequential compactness*, and it could be defined more generally for topological spaces. In fact, for metric spaces, we have that (X, d) is sequentially compact \iff it is *compact* (in a sense that uses open covers, which we won't discuss here). Since in this module we only discuss sequential compactness, we will simply call it compactness throughout.

Example 4.4.4. — Let (X, d) be a metric space and assume that X has finitely many elements. Then X is compact. If this wasn't the case, then each of the elements of X would be visited only finitely many times by the sequence. This would imply that the set of indices of the sequence is finite, but that set of indices is the set of natural numbers, which isn't finite!

4.4 Compactness

Proposition 4.4.5. — *If $X \subseteq (\mathbb{R}, d_1)$ is such that (X, d_1) is compact, then X has a maximum and a minimum.*

Proof. We first prove that X is bounded. Assume it is not, then for every $R > 0$ there exists $x \in X$ such that $|x| > R$. Take $R = 1$. Then we can find an element $x_1 \in X$ such that $|x_1| > 1$ by the above. Then take $R = |x_1| + 1$, and find $x_2 \in X$ such that $|x_2| > |x_1| + 1$. Continuing this way, we construct a sequence (x_n) in X such that $|x_{n+1}| > |x_n| + 1$.

The sequence (x_n) has no subsequence that is Cauchy. Indeed, the distance between two consecutive elements of the sequence (x_n) is by construction always > 1 . Because a convergent sequence is Cauchy, we deduce that (x_n) has no convergent subsequence, contradicting the hypothesis that (X, d_1) is compact.

Because X is bounded, the supremum $\sup(X)$ is not $+\infty$. Then take a sequence (x_n) defined by the property that $\sup(X) - \frac{1}{n} \leq x_n \leq \sup(X)$. Then (x_n) converges to $\sup(X)$ in (\mathbb{R}, d_1) . Since (X, d_1) is compact, some subsequence of (x_n) must converge to $\sup(X)$ in (X, d_1) . By the uniqueness of the limit, we deduce that $\sup(X) \in X$. (A similar argument works for the infimum/minimum.) \square

We will deduce the two most important results on compactness from the following Key Theorem.

Theorem 4.4.6 (Key Result). — *If $f: (X, d_X) \rightarrow (Y, d_Y)$ is a continuous function, if (X, d_X) is compact, then so is $(f(X), d_Y)$.*

Proof. Take an arbitrary sequence (y_n) of $f(X)$, so that there is a sequence (x_n) of X such that $f(x_n) = (y_n)$ for every n . Since X is compact, we can find a subsequence (x_{n_k}) which converges to, say, ℓ . Then, since f is continuous, $f(x_{n_k}) \rightarrow f(\ell)$ implies that $y_{n_k} \rightarrow f(\ell)$.

We have therefore found a convergent subsequence of (y_n) , which proves that $f(X)$ is compact. \square

Corollary 4.4.7. — *(Compactness is a topological property) If (X, d_X) is homeomorphic to (Y, d_Y) , then*

$$(X, d_X) \text{ is compact} \iff (Y, d_Y) \text{ is compact.}$$

The analogue statement where the word “compact” is replaced with the word “complete” is false, as illustrated in Remark 4.1.7.

Proof. Suppose $f: X \rightarrow Y$ is a homeomorphism. Then f is both continuous and surjective. Since X is compact, so is $f(X) = Y$. Then use the same argument for f^{-1} . □

By combining Proposition 4.4.5 with 4.4.6 we deduce the following important result.

Corollary 4.4.8. — *Suppose (X, d_X) is compact, and that $f: X \rightarrow \mathbb{R}$ is continuous (with the usual distance d_1 on \mathbb{R}). Then f has a minimum and a maximum.*

Proof. The subset $f(X) \subseteq \mathbb{R}$ is compact, so by Proposition 4.4.5 it has a minimum \underline{y} and a maximum \bar{y} . Since f is surjective onto $f(X)$, we can pick elements \underline{x} and \bar{x} which map to \underline{y} and \bar{y} respectively. □

4.4.9. Compactness for subsets of \mathbb{R}^N — Here we discuss a characterisation of the compact subsets of \mathbb{R}^N . Let's start by addressing the question of what metric subspaces of (\mathbb{R}, d_1) are compact. This will follow from some Year 1 material. We start with a preliminary (general) lemma.

Lemma 4.4.10. — *Let (X, d) be a metric space. Suppose that $C \subseteq X$ and (C, d) is compact. Then C is a closed subset of (X, d) .*

Proof. If C were not closed, we could find a sequence (x_n) that converges to $\ell \in X \setminus C$. Then, because of the uniqueness of the limit, we have disproved that every subsequence of (x_n) converges in C . This means that C is not compact. □

Theorem 4.4.11 (Bolzano-Weierstrass Theorem). — *Let $X \subseteq (\mathbb{R}, d_1)$. Then the following are equivalent:*

1. (X, d_1) is compact.
2. X is closed and bounded.

Proof. In Year 1 (MATH101) you have seen a result with the same name, which stated that every bounded real sequence admits a convergent subsequence (Corollary 3.21 of MATH101 notes).

Assume X is closed and bounded in (\mathbb{R}, d_1) . Let (x_n) be a sequence in X . Then that sequence is bounded, so by the Year 1 Bolzano-Weierstrass theorem, (x_n) admits a convergent subsequence. Because X is closed, the

4.4 Compactness

limit of that subsequence must be an element of X . This proves that (X, d_1) is compact.

The other implication follows from Proposition 4.4.5 and Lemma 4.4.10. \square

Theorem 4.4.12 (Multidimensional Bolzano-Weierstrass). — *The compact subsets of (\mathbb{R}^N, d_p) are the closed and bounded subsets.*

We have already seen the case where $N = 1$, though we currently have not yet defined what bounded means in the multidimensional context. Here it comes:

Definition 4.4.13. — We say a subset $X \subseteq \mathbb{R}^N$ is bounded if there exists an $R > 0$ such that $X \subseteq B_R(0)$.

Proof of Theorem 4.4.12. Assume $X \subseteq (\mathbb{R}^N, d_p)$ is closed and bounded. Take a sequence $(x_k)_{k \in \mathbb{N}}$ of X . We need to find a subsequence $(x_{k_m})_{m \in \mathbb{N}}$ that converges in (\mathbb{R}^N, d_p) . Since X is closed, this sequence will then also converge in X .

Begin by taking the bounded sequence $(x_{1,k})$ in (\mathbb{R}, d_1) of the first coordinates of the sequence (x_k) . Then by applying the Year 1 version of Bolzano-Weierstrass (Corollary 3.21 of MATH101 notes) there is a convergent subsequence $(x_{1,k_{m_1}})$. Now do the same again, restricting to the first two coordinates to get a sequence $(x_{2,k_{m_1}})$. Again by Year 1 Bolzano-Weierstrass this admits a convergent subsequence $(x_{2,k_{m_2}})$. Repeating this procedure N times produces a subsequence of (x_n) whose coordinates *all* converge as sequences in (\mathbb{R}, d_1) , and we conclude that this subsequence then converges in (\mathbb{R}^N, d_p) by applying Lemma 4.2.2.1. Furthermore, that subsequence converges in (X, d_p) because X is closed in (\mathbb{R}^N, d_p) .

If X is not closed we know by 4.4.10 that it cannot be compact. So assume X is not bounded. The argument will be very similar to the one we gave for the $N = 1$ case.

Because $X \not\subseteq B_1(0)$ there exists $x_1 \notin B_1(0)$ such that $x_1 \in X$.

Now take $B_{R_2}(0)$ for $R_2 = d_p(x_1, 0) + 1$. Then $X \not\subseteq B_{R_2}(0)$. Then we find $x_2 \notin B_{R_2}(0)$ so that $x_2 \in X$. Take $R_3 = d_p(x_2, 0) + 1$ and continue this process. Doing this, we find a sequence (x_n) whose distance of any two consecutive elements is larger than 1. That sequence cannot admit any Cauchy subsequence, hence it cannot admit any convergent subsequence. \square

Here is an example of compact/noncompact spaces.

Example 4.4.14. — By the multidimensional Bolzano-Weierstrass theorem, closed balls in (\mathbb{R}^N, d_p) are compact, and open balls are not.

The above characterisation of compact subsets does not generalise to other metric spaces, as we show in the following example.

Remark 4.4.15. — The closed ball $\overline{B}_1((0, 0, \dots)) \subseteq (\ell^p, d_p)$ of radius 1 centered at the origin is *not* compact. Take the sequence $(x_k)_{k \in \mathbb{N}}$ in ℓ^p defined by

$$\begin{aligned} x_1 &= (x_{n,1})_{n \in \mathbb{N}} = (1, 0, 0, \dots) \\ x_2 &= (x_{n,2})_{n \in \mathbb{N}} = (0, 1, 0, \dots) \\ &\vdots \\ x_k &= (x_{n,k})_{n \in \mathbb{N}} = (0, \dots, 0, 1, 0, \dots) \\ &\dots \end{aligned}$$

where we take the sequences to be zero everywhere except for a 1 in the k -th entry. Then $d_\infty(x_{k_1}, x_{k_2}) = 1$ whenever $k_1 \neq k_2$, and

$$d_p(x_{k_1}, x_{k_2}) = \sqrt[p]{2}.$$

Therefore any subsequence $(x_{k_m})_{m \in \mathbb{N}}$ satisfies the property that for all $N > 0$, there are $n_1 \neq n_2$ such that the distance of two elements of the subsequence equals $\sqrt[p]{2}$. We deduce that no subsequence is Cauchy, hence no subsequence converges.

Note that we have never formally defined what it means for a subset of a spaces of sequences ℓ^p for some p to be *bounded*. In analogy with the definition of a bounded subset of \mathbb{R}^n we could define a subset of ℓ^p to be bounded when it is contained in some open ball centred at the zero function. The previous example shows that, with this definition of boundedness, there are closed and bounded subsets of (ℓ^p, d_p) that fail to be compact. We will see in Chapter 5 that the same occurs for $(C[0, 1], d_{\infty})$.

4.4.16. Compactness and completeness — We conclude the chapter by briefly discussing the relation between compactness and completeness.

Lemma 4.4.17. — *Let (X, d) be a metric space. Suppose (x_n) is a Cauchy sequence and a subsequence (x_{n_k}) converges to ℓ . Then (x_n) converges to ℓ .*

4.4 Compactness

Proof. Since (x_n) is Cauchy, for every $\epsilon > 0$ we can find an N_1 such that $d(x_n, x_m) < \epsilon/2$ for every $m, n > N_1$. Also, since (x_{n_k}) converges to ℓ , for every $\epsilon > 0$ we can find a K such that $d(x_{n_k}, \ell) < \epsilon/2$ for every $k > K$. Take a $k_1 \geq K$ such that $n_{k_1} > N_1$ (this can be done because, by definition of a subsequence, the sequence n_k is strictly increasing). Then take $N = n_{k_1}$. It follows that

$$d(x_n, \ell) \leq d(x_n, x_{n_{k_1}}) + d(x_{n_{k_1}}, \ell) < \epsilon/2 + \epsilon/2 = \epsilon$$

for every $n > N$. □

Here comes the relation between compact and complete.

Corollary 4.4.18. — *A compact metric space is complete.*

The converse implication is clearly false. For example, we have seen that (\mathbb{R}, d_1) is complete, but we also know it isn't compact (for example by the Bolzano-Weierstrass theorem).

Proof. A Cauchy sequence in a compact metric space has a convergent subsequence by definition of compactness. By Lemma 4.4.17, the original Cauchy sequence must also converge. □

5.

Spaces of continuous functions

Important notions to learn from this section:

1. The notion of uniform and pointwise convergence.
2. The space of bounded real functions is complete with d_{L^∞} .
3. The subspace of continuous real functions on a compact metric space is complete.
4. Uniform convergence respects continuity and integrals/derivatives.
5. The radius of convergence of a power series and the fundamental Theorem 5.2.3.
6. The Peano-Picard theorem on local existence/uniqueness for Cauchy problems, and the strategy of its proof with the CMT.

In this Chapter we explore more details of the theory developed in the previous chapters for the particular case when X is the space of continuous functions endowed with its maximum (or equivalently supremum) distance. The main example to keep in mind is that of $X = C[0, 1]$.

The key notion here is that of *uniform convergence*, which we will make sense of for sets $B(Y)$ of bounded functions $Y \rightarrow \mathbb{R}$, for Y an arbitrary set. We will prove that the metric space of bounded functions (with the metric of the supremum) is complete. We then deduce that if (Y, d) is a compact

metric space, then the space $C(Y)$ of continuous functions $Y \rightarrow \mathbb{R}$ with the metric of the supremum is also complete. We prove this by showing that $C(Y) \subseteq B(Y)$ is a closed subset, namely by proving that the uniform limit of continuous functions must be continuous.

Uniform convergence has other benefits (as opposed to pointwise convergence, a more rudimental notion that we also introduce for comparison): namely it behaves well with respect to integration and derivation. (In order to develop this theory more generally, one would have to have the notion of a measure on the set Y . Therefore we will limit ourselves to $Y = [a, b]$. The interested student will find more details on the general case in MATH365 - Measure Theory). We will use this theory to prove the existence of the (uniform) limit of power series.

After we have scratched the surface of the theory of spaces of continuous functions, we discuss an application of the contraction mapping theorem and of the fact that $(C[a, b], d_{L^\infty})$ is complete, which is the local existence and uniqueness theorem for Cauchy problems. A Cauchy problem is an ordinary differential equation in 1 (real) variable, paired with an initial value (the initial value assignment is there to give the problem a chance of having a finite number of solutions, and possibly exactly 1).

5.1. Uniform convergence

We start by defining the notions of uniform and pointwise convergence for sequences of functions to the real numbers

Definition 5.1.1. — Let X be a set and let $f: X \rightarrow \mathbb{R}$ be a real-valued function, and $(f_n: X \rightarrow \mathbb{R})$ be a sequence of real-valued functions.

1. We say that $(f_n)_{n \in \mathbb{N}}$ *converges pointwise* to f if $\lim_{n \rightarrow \infty} f_n(x) = f(x)$ for every point $x \in X$. In other words, (f_n) is pointwise convergent to f if for every $x \in X$, for every $\epsilon > 0$ there exists $N_x \in \mathbb{N}$ such that

$$|f_n(x) - f(x)| < \epsilon$$

for every $n > N_x$. (With the notation, we are emphasising that the number N_x may depend on $x \in X$!)

2. We say that $(f_n)_{n \in \mathbb{N}}$ *converges uniformly* to f if for every $\epsilon > 0$ there

5.1 Uniform convergence

exists $N \in \mathbb{N}$ such that

$$\sup_{x \in X} |f_n(x) - f(x)| < \epsilon$$

for every $n > N$.

The difference between the two notions is that, in the former, we are allowed to choose N depending on the point x , whereas in the latter the number N is chosen independently (or uniformly) for all x .

Remark 5.1.2. — If a sequence (f_n) converges uniformly to f , in particular it converges pointwise to f . The converse is not true as we show in the following example.

Example 5.1.3. — Let (X, d_X) be the metric space $([0, 1], d_1)$ and let us consider the sequence of functions (f_n) defined by

$$\begin{aligned} f_n: [0, 1] &\rightarrow \mathbb{R} \\ x &\mapsto x^n \end{aligned}$$

Let $f: [0, 1] \rightarrow \mathbb{R}$ be given by

$$f(x) = \begin{cases} 0 & \text{for all } 0 \leq x < 1 \\ 1 & \text{when } x = 1. \end{cases}$$

Then, the sequence (f_n) converges pointwise to f . Indeed, fix $x \in [0, 1]$. If $x \in [0, 1)$ we have that $f_n(x) = x^n$ which converges to $f(x) = 0$. If $x = 1$ instead, $f_n(x) = 1$ for every n , hence $f_n(x)$ converges to $f(x) = 1$.

However, for all $n \in \mathbb{N}$, we have

$$\sup_{x \in [0, 1]} |f_n(x) - f(x)| = 1$$

(exercise) so (f_n) does not converge uniformly to f .

A quicker argument to reach the same conclusion is given in Remark 5.1.15.

Exercise 5.1.4. — Let $(f_n: X \rightarrow \mathbb{R})$ be a sequence of real-valued functions and $f: X \rightarrow \mathbb{R}$ a real-valued function. Determine if (f_n) converges uniformly or pointwise to f in the following cases.

1. Let $X = [0, 1]$, let $f_n(x) = \frac{nx}{n+1}$ and $f(x) = x$ for every $x \in X$.

2. Let $X = [-1, 1]$, let $f_n(x) = x^{2n}$ for every $x \in X$ and let

$$f(x) = \begin{cases} 1 & \text{for } x = \pm 1 \\ 0 & \text{for } |x| < 1 \end{cases}$$

Exercise 5.1.5. — Let $X = [0, 1]$ and $f_n(x) = \sum_{k=0}^n x^k$, for every $x \in X$. Does the sequence (f_n) converge pointwise/uniformly to some limit function $f: X \rightarrow \mathbb{R}$?

One important theoretical difference between uniform and pointwise convergence is that the former preserves continuity (as we shall soon see in this section), whereas the latter in general does not.

In MATH241 we have learned that often, when the word “converge” is used, there is often a metric space lurking behind the scenes, and that word means that there is some sequence that converge in that metric space. The notion of pointwise convergence does not resonate with this idea — we will see in a Section B problem that pointwise convergence is a notion of convergence for sequence of functions that does *not* come from convergence in any metric space.

Uniform convergence is more in line with the general theory of metric spaces. It can be rephrased as convergence in the following metric space.

Definition 5.1.6. — We define the set of *bounded real-valued functions* on a set X

$$B(X) := \text{bound}(X, \mathbb{R}) := \{f: X \rightarrow \mathbb{R} : f \text{ bounded}\}$$

(Where *bounded* for a function $f: X \rightarrow \mathbb{R}$ means that there exists $R > 0$ such that $|f(x)| < R$ for all $x \in X$.) For every two functions $f, g \in B(X)$, we define

$$d_{L^\infty}(f, g) = \sup_{x \in X} |f(x) - g(x)|$$

Arguing as in Chapter 1, it is not difficult to see that d_{L^∞} defines a distance on the set $B(X)$.

It immediately follows from the definitions that a sequence (f_n) in $B(X)$ converges uniformly to $f \in B(X)$ if and only if f_n converges in $B(X)$ to f , with respect to the distance d_{L^∞} that we have just introduced. We really need the functions to be bounded, for otherwise the supremum might be infinity!

5.1 Uniform convergence

Example 5.1.7. — Let $X = [0, 1]$ and let $f: X \rightarrow \mathbb{R}$ be given by

$$x \mapsto \begin{cases} \frac{1}{x} & \text{for all } x \in (0, 1], \\ 0 & \text{for } x = 0 \end{cases}$$

and $g: X \rightarrow \mathbb{R}$ be the constant function at 0. Then, we have

$$\sup_{x \in (0, 1)} |f(x) - g(x)| = \infty.$$

When X itself comes with some distance d , we can also define the following spaces of functions.

Definition 5.1.8. — Let (X, d) be a metric space. Define the set of *continuous functions*

$$C(X) := \{f: X \rightarrow \mathbb{R} : f \text{ continuous}\}$$

and the subset of *bounded and continuous functions*

$$BC(X) := \{f: X \rightarrow \mathbb{R} : f \text{ bounded and continuous}\}.$$

Remark 5.1.9. — For d a fixed distance on X , the space of functions $BC(X)$ is a subset of $B(X)$, and we can endow it with the subspace distance d_{L^∞} as seen in Chapter 1.

In general, the function d_{L^∞} does *not* define a distance on the set $C(X)$ because $\sup_{x \in X} |f(x)|$ might equal infinity. However, when (X, d) is *compact*, we have seen that a continuous function $f: X \rightarrow \mathbb{R}$ from a compact set always has maximum and minimum. We deduce that if (X, d) is compact (for example, when $X = [0, 1]$), we have the equality $C(X) = BC(X)$.

The central result of this section is that the metric space $(B(X), d_{L^\infty})$ is always complete. This follows, with some work, from the fact that (\mathbb{R}, d_1) is complete.

Theorem 5.1.10. — *The metric space of bounded functions on X :*

$$(B(X), d_{L^\infty})$$

is complete.

Proof. Let $(f_n)_{n \in \mathbb{N}}$ be a Cauchy sequence in $B(X)$: for every $\epsilon > 0$ there exists $N' \in \mathbb{N}$ such that

$$d_{L^\infty}(f_n, f_m) < \epsilon$$

for every $m, n > N'$.

Fix a point $x \in X$ and consider the sequence $(f_n(x))_{n \in \mathbb{N}}$ in (\mathbb{R}, d_1) . Because the sequence (f_n) is Cauchy, so is the sequence $f_n(x)$, and because (\mathbb{R}, d_1) is complete we deduce that the sequence $(f_n(x))$ converges. Therefore, we can define a function $f: X \rightarrow \mathbb{R}$ by the formula

$$f(x) = \lim_{n \rightarrow \infty} f_n(x).$$

To conclude the proof, we need to show that f is bounded, and that it is the uniform limit of the sequence (f_n) .

The function f is bounded by the following argument. Because (f_n) is Cauchy, for $\epsilon = 1/2$ we can find M such that

$$|f_n(x) - f_m(x)| < \frac{1}{2}$$

for all $x \in X$ and all $n, m \geq M$. In particular, we deduce by taking the limit $n \rightarrow \infty$ that

$$|f(x) - f_M(x)| \leq \frac{1}{2}$$

for all x . We deduce that, for all $x \in X$, we have

$$|f(x)| \leq |f(x) - f_M(x)| + |f_M(x)| \leq \frac{1}{2} + d_{L^\infty}(f_M, 0)$$

(for $0 \in C[0, 1]$ the zero function), which implies that f is bounded.

Now we prove that f is actually the uniform limit of the sequence (f_n) . Recall that, because the sequence (f_n) is Cauchy, for all $\epsilon > 0$ there exists $N' \in \mathbb{N}$ such that

$$|f_n(x) - f_m(x)| < \epsilon$$

for all $x \in X$ and for all $n, m > N'$. Fix x and n and take the limit for $m \rightarrow \infty$, then we get

$$\begin{aligned} \lim_{m \rightarrow \infty} |f_n(x) - f_m(x)| &= \left| f_n(x) - \lim_{m \rightarrow \infty} f_m(x) \right| \\ &= |f_n(x) - f(x)| \end{aligned}$$

Therefore, for every $\epsilon > 0$ we can take $N = N'$, and then deduce that

$$|f_n(x) - f(x)| \leq \epsilon$$

for every $n > N$ and for every $x \in X$. This shows that (f_n) converges uniformly to f . □

5.1 Uniform convergence

When the set X is endowed with a distance d , we can make sense of $BC(X)$, the set of functions $f: X \rightarrow \mathbb{R}$ that are bounded and continuous. It turns out that the subset $BC(X)$ is closed into $B(X)$, as follows from the next result.

Theorem 5.1.11. — *Let $(f_n: (X, d) \rightarrow \mathbb{R})_{n \in \mathbb{N}}$ be a sequence of continuous functions and let $f: X \rightarrow \mathbb{R}$. If (f_n) converges uniformly to f , then f is continuous.*

Proof. Let $x_0 \in X$. For every $\epsilon > 0$, since (f_n) converges uniformly to f we can find $N \in \mathbb{N}$ such that

$$|f_N(x) - f(x)| < \frac{\epsilon}{3}$$

for every $x \in X$. Moreover, since f_N is continuous at x_0 , we can find $\delta > 0$ such that

$$|f_N(x) - f_N(x_0)| < \frac{\epsilon}{3}$$

for every x with $d(x, x_0) < \delta$. For such a δ , we have:

$$\begin{aligned} |f(x) - f(x_0)| &\leq |f(x) - f_N(x)| + |f_N(x) - f_N(x_0)| + |f_N(x_0) - f(x_0)| \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \\ &= \epsilon. \end{aligned}$$

For every x such that $d(x, x_0) < \delta$. This proves that f is continuous at $x_0 \in X$. □

By applying the sequential characterisation of closed subsets, we immediately deduce the following.

Corollary 5.1.12. — *Let (X, d) be a metric space. Then $BC(X)$ is a closed subset of $(B(X), d_{L^\infty})$.*

Proof. Let (f_n) be a sequence in $BC(X)$ and assume that f_n converges with the distance d_{L^∞} to $f \in B(X)$. Then, f_n converges uniformly to f and by Theorem 5.1.11 we deduce that $f \in BC(X)$. □

Remember that from Chapter 4 we know that a subset of a complete metric space is complete if and only if the corresponding subset is closed. We deduce that:

Corollary 5.1.13. — *The metric space $(BC(X), d_{L^\infty})$ is complete.*

We conclude our reasoning by observing that if (X, d) is a *compact* metric space, then by the Min/Max Theorem we deduce that $BC(X) = C(X)$ and therefore

Corollary 5.1.14. — *If (X, d) is a compact metric space, then the metric space $(C(X), d_{L^\infty})$ is complete.*

From this we in particular deduce that the metric space $(C[a, b], d_{L^\infty})$ of continuous functions $f: [a, b] \rightarrow \mathbb{R}$ (introduced in Chapter 1 for $a = 0$ and $b = 1$) is a complete metric space, as was claimed in Chapter 4.

Remark 5.1.15. — Let $f_n: [0, 1] \rightarrow \mathbb{R}$ be the sequence defined as

$$f_n(x) = x^n.$$

Then, there is no subsequence $(f_{n_k})_{k \in \mathbb{N}}$ of (f_n) which is a Cauchy sequence in $(C[0, 1], d_{L^\infty})$, which implies that the unit (closed) ball $\overline{B}_1^{d_{L^\infty}}(0)$ centered at the zero function 0 in $C[0, 1]$ is not compact.

Indeed, if there existed a subsequence $(f_{n_k})_{k \in \mathbb{N}}$ which is Cauchy, then since $(C[0, 1], d_{L^\infty})$ is complete, there would be a function $f \in C[0, 1]$ continuous such that (f_{n_k}) converges uniformly to f . But (f_n) converges pointwise to the function with value 1 on $x = 1$ and 0 everywhere else, which is not continuous and gives a contradiction.

Remark 5.1.16. — More generally, using the same reasoning of the previous remark, we can show that any sequence $(f_n)_{n \in \mathbb{N}}$ of continuous functions on $[0, 1]$ that converges pointwise to a discontinuous function, does not have any convergent subsequence in d_{L^∞} .

5.1.17. Results on uniform convergence — Here we go back to the theory of continuous functions on closed and bounded intervals $[a, b]$ of the real line and explore other benefits of the notions of uniform convergence. (What we need of $[a, b]$ that is not available on an arbitrary set is the notion of a *measure*, which allows us to discuss integrals of functions. The notion of a measure and of integrals of functions will be generalised to an arbitrary set X in MATH365).

We have seen that uniform convergence preserves continuity, and we will see that it also preserves integrals and (with some care) also derivatives.

5.1 Uniform convergence

Theorem 5.1.18. — Let $f_n: [a, b] \rightarrow \mathbb{R}$ define a sequence (f_n) of continuous function and assume that (f_n) converges uniformly to a function $f: [a, b] \rightarrow \mathbb{R}$. Then, for all points $x_1, x_2 \in [a, b]$ the following equality holds:

$$\lim_{n \rightarrow \infty} \int_{x_1}^{x_2} f_n(x) \, dx = \int_{x_1}^{x_2} f(x) \, dx$$

Here is an example to show that pointwise convergence would not allow us to obtain the same result.

Example 5.1.19. — Let $[a, b] = [0, 1]$ and consider the sequence of functions:

$$f_n(x) = \begin{cases} 0 & \text{when } x \geq \frac{1}{n} \text{ or } x = 0 \\ n - n^2x & \text{when } 0 < x < \frac{1}{n} \end{cases}$$

notice that f_n is not continuous. Then, we have $\int_0^1 f_n(x) \, dx = \frac{1}{2}$. However, f_n converges pointwise to the function f that is constantly equal to 0. In particular, we have that the limit

$$\lim_{n \rightarrow \infty} \int_0^1 f_n(x) \, dx = \frac{1}{2} \neq 0 = \int_0^1 0 \, dx.$$

This example shows that in the case of pointwise convergence we can not swap integration and taking the limit, while in the case of uniform convergence we can.

Proof of Theorem 5.1.18. By Lemma 2.2.7 all we need to prove is that for every two points x_1 and x_2 in \mathbb{R} the integral function

$$I: (C[x_1, x_2], d_{L^\infty}) \rightarrow (\mathbb{R}, d_1)$$

is continuous.

Fix $\epsilon > 0$ and take $\delta = \frac{\epsilon}{|x_1 - x_2|}$, then

$$\begin{aligned} \int_{x_1}^{x_2} |f(x) - g(x)| \, dx &\leq \int_{x_1}^{x_2} \max |f(x) - g(x)| \, dx \\ &< \delta |x_1 - x_2| = \epsilon \end{aligned}$$

when $d_{L^\infty}(f, g) < \delta$. Then,

$$\lim_{n \rightarrow \infty} I(f_n) = I\left(\lim_{n \rightarrow \infty} f_n\right) = I(f).$$

□

In order to obtain the result that uniform convergence preserves the derivatives, we are going to need the following crucial result from Year 1.

Theorem 5.1.20 (Fundamental Theorem of Calculus — Theorem 6.9 and discussion thereafter in the notes of MATH101.). — *Let $\phi: [a, b] \rightarrow \mathbb{R}$ be a continuous function.*

1. *Let us consider the function $F: [a, b] \rightarrow \mathbb{R}$ given by*

$$F(x) = \int_a^x \phi(t) \, dt$$

Then F is differentiable and $F' = \phi$.

2. *Assume $G: [a, b] \rightarrow \mathbb{R}$ is a differentiable function, such that $G' = \phi$. Then, the following equality holds*

$$\int_a^b \phi(t) \, dt = G(b) - G(a).$$

Here comes the last result of this section.

Theorem 5.1.21. — *Let $f_n: [a, b] \rightarrow \mathbb{R}$ define a sequence of differentiable functions and assume that (f_n) converges uniformly to a function $f: [a, b] \rightarrow \mathbb{R}$ and that (f'_n) is continuous and that it converges uniformly to a function $g: [a, b] \rightarrow \mathbb{R}$ (that a posteriori is also continuous). Then, f is differentiable and $g = f'$.*

Proof of Theorem 5.1.21. We have that

$$\begin{aligned} f(x) &= \lim_{n \rightarrow \infty} f_n(x) \\ &= \lim_{n \rightarrow \infty} \left(f_n(a) + \int_a^x f'_n(t) \, dt \right) \\ &= f(a) + \int_a^x g(t) \, dt \end{aligned}$$

Where the second equality follows from the second part of Theorem 5.1.20, and the third follows from Theorem 5.1.18.

Moreover, applying the first part of Theorem 5.1.20 and thanks to the continuity of g , we obtain that $f'(x) = g(x)$ for all x . □

The following exercise is to observe that in the previous theorem, the uniform convergence of the derivatives is not automatically guaranteed from the uniform convergence of the original sequence of functions.

5.2 Power series

Exercise 5.1.22. — Consider the sequence of functions $f_n: [a, b] \rightarrow \mathbb{R}$ defined by

$$f_n(x) = \frac{\sin(nx)}{\sqrt{n}}.$$

Does (f_n) converge uniformly? If so, what is the derivative of the uniform limit? Does (f'_n) converge uniformly?

5.2. Power series

The goal of this section is to give a precise meaning to the infinite sum

$$\sum_{n=0}^{\infty} a_n x^n \tag{5.1}$$

as a function for assigned coefficients

$$a_0, a_1, \dots, a_n, \dots \in \mathbb{R},$$

and then to study its properties (continuity, its integral and its derivative). (Formula (5.1) describes a power series centered at 0. We will only discuss this case. The case when the power series is centered at an arbitrary $x_0 \in \mathbb{R}$ is obtained by replacing the monomial x^n with $(x - x_0)^n$).

In order to do so, we will apply the results of the previous section, on uniform convergence, to the sequence $S_N: \mathbb{R} \rightarrow \mathbb{R}$ of functions defined by

$$S_N(x) = \sum_{n=0}^N a_n x^n,$$

and called the partial (=up to N) sums of the initial power series. (Each of them is a degree N polynomial function of 1 real variable).

Definition 5.2.1. — Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of real numbers. The *radius of convergence* of the corresponding power series $\sum_{n=0}^{\infty} a_n x^n$ is the nonnegative real number

$$R = \frac{1}{\limsup_{n \in \mathbb{N}} \sqrt[n]{|a_n|}}$$

where by convention we set

$$R = \begin{cases} \infty & \text{if } \limsup_{n \in \mathbb{N}} \sqrt[n]{|a_n|} = 0 \\ 0 & \text{if } \limsup_{n \in \mathbb{N}} \sqrt[n]{|a_n|} = \infty \end{cases}$$

Definition 5.2.2. — Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of real numbers. For every $x \in \mathbb{R}$ and for every $N \in \mathbb{N}$ we define the N -th *partial sum* of the power series $\sum_{n=0}^{\infty} a_n x^n$ defined by (a_n) at x as the real number:

$$S_N(x) = \sum_{n=0}^N a_n x^n$$

Varying x this gives a real-valued function $S_N: \mathbb{R} \rightarrow \mathbb{R}$.

Theorem 5.2.3. — Assume that the radius of convergence R of the power series $\sum_{n=0}^{\infty} a_n x^n$ is greater than 0. Then, for all $0 < \delta < R$, the sequence of functions $(S_N)_{N \in \mathbb{N}}$ is a Cauchy sequence in the metric space $(C[-\delta, \delta], d_{L^\infty})$.

From this we immediately deduce:

Corollary 5.2.4. — Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of real numbers and assume that the radius of convergence R of the corresponding power series $\sum_{n=0}^{\infty} a_n x^n$ is larger than 0. Then the power series defines a continuous function

$$S: (-R, R) \rightarrow \mathbb{R}.$$

Proof. The first part follows from $(C[-\delta, \delta], d_{L^\infty})$ being complete and the fact that the uniform limit of the continuous functions S_N is continuous.

The fact that S is defined on $(-R, R)$ follows because the latter is the union of all the intervals $[-\delta, \delta]$ for $\delta < R$. □

This allows us to make sense of the expression (5.1) as a function.

Definition 5.2.5. — Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of real numbers with radius of convergence $R > 0$. Then, we define the *sum of the power series* $\sum_{n=0}^{\infty} a_n x^n$, as the uniform limit S of the sequence $(S_N)_{N \in \mathbb{N}}$. For every point $x \in (-R, R)$ we write

$$\sum_{n=0}^{\infty} a_n x^n =: S(x)$$

for the value of S at x .

The fact that a power series is defined as the uniform limit of the partial sums, allows to differentiate and integrate them in a simple way.

Corollary 5.2.6. — Let $\sum_{n=0}^{\infty} a_n x^n$ be a power series with radius of convergence $R > 0$ and let $S: (-R, R) \rightarrow \mathbb{R}$ be the sum of the corresponding power series. Then:

5.2 Power series

1. The function S is differentiable and $((n+1)a_{n+1})$ defines a power series with radius of convergence R such that

$$S'(x) = \sum_{n=0}^{\infty} (n+1)a_{n+1}x^n$$

2. If F is a function $F: (-R, R) \rightarrow \mathbb{R}$ such that $F'(x) = S(x)$, then

$$F(x) = \sum_{n=0}^{\infty} \frac{a_n}{n+1} x^{n+1} + c$$

for some constant $c \in \mathbb{R}$. (In particular, F is the sum of a power series with radius of convergence R .)

Proof. 1. For every $N \in \mathbb{N}$, let S_N be the N -th partial sum of the power series, which is a uniformly convergent function in $C[-\delta, \delta]$ for every $\delta < R$. We have that the derivative S'_N is defined by

$$S'_N(x) = \sum_{n=1}^N n a_n x^{n-1}$$

and its radius of convergence is given by

$$\frac{1}{\limsup \sqrt[n]{n|a_n|}} = \frac{1}{\limsup \sqrt[n]{|a_n|}}$$

since the limit of $\sqrt[n]{n}$ for $n \rightarrow \infty$ equals 1. Therefore, the sequence $(S'_N)_{N \in \mathbb{N}}$ is also uniformly convergent on the interval $[-\delta, \delta]$. Applying Theorem 5.1.21 we conclude that

$$S' = \lim_{N \rightarrow \infty} S'_N$$

in the distance d_{L^∞} , which proves our result by the definition of the sum of the power series $\sum_{n=0}^{\infty} (n+1)a_{n+1}x^n$.

2. This can be proven similarly, this time applying Theorem 5.1.18. □

Example 5.2.7. — Let us compute the sum:

$$\sum_{n=1}^{\infty} n x^{n-1} = \sum_{n=1}^{\infty} \frac{d}{dx}(x^n) = \frac{d}{dx} \sum_{n=1}^{\infty} x^n = \frac{d}{dx} \left(\frac{1}{1-x} \right) = \frac{1}{(1-x)^2}$$

which has radius of convergence equal to 1.

5.2.8. — In order to prove the main result Theorem 5.2.3, we first take a brief detour to review some convergence criteria for series of nonnegative real numbers from Year 1.

Proposition 5.2.9 (Ratio test). — Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of nonnegative real numbers.

1. If $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} < 1$, then the series $\sum_{n=0}^{\infty} a_n$ converges.
2. If $\lim_{n \rightarrow \infty} \frac{a_{n+1}}{a_n} > 1$, then the series $\sum_{n=0}^{\infty} a_n$ diverges.

Proposition 5.2.10 (Root test). — Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of nonnegative real numbers.

1. If $\lim_{n \rightarrow \infty} \sqrt[n]{a_n} < 1$, then the series $\sum_{n=0}^{\infty} a_n$ converges.
2. If $\lim_{n \rightarrow \infty} \sqrt[n]{a_n} > 1$, then the series $\sum_{n=0}^{\infty} a_n$ diverges.

Notice that if the limit involved in the ratio test (respectively the root test) is 1 or if it does not exist, then the test is inconclusive.

Example 5.2.11. — Let $0 < t < 1$ be a fixed real number and consider the series

$$1 + 2t + t^2 + 2t^3 + t^4 + 2t^5 + \dots$$

Then the series converges to $\frac{1}{1-t^2} + \frac{2t}{1-t^2} < \infty$ (as the sum of two geometric series).

One can check that the ratio test is inconclusive (the ratio of two consecutive elements oscillates between $t/2$ and $2t$), while the root test gives

$$\lim_{n \rightarrow \infty} \sqrt[n]{2t^n} = \lim_{n \rightarrow \infty} \sqrt[n]{t^n} = t < 1$$

which confirms that the series converges.

Here is a proposition that motivates our reasoning and shows that the root test is more powerful than the ratio test. We won't prove this proposition as we won't need it, but the interested reader is encouraged to attempt it as an exercise.

Proposition 5.2.12. — Let (a_n) be a sequence of nonnegative real numbers. Then, the following inequalities hold

$$\liminf \frac{a_{n+1}}{a_n} \leq \liminf \sqrt[n]{a_n} \leq \limsup \sqrt[n]{a_n} \leq \limsup \frac{a_{n+1}}{a_n}.$$

5.2 Power series

Remark 5.2.13. — In particular, the proposition implies that the root test is more effective than the ratio test. For instance, the inequalities of Proposition 5.2.12 applied to the example discussed in Example 5.2.11 become

$$\frac{x}{2} \leq x \leq x \leq 2$$

The following result gives the only test that we will use to prove our convergence result for power series. The result is an improvement of the ratio and of the root test that we described above.

Proposition 5.2.14. — (*Enhanced root test*) Let $(a_n)_{n \in \mathbb{N}}$ be a sequence of non-negative real numbers.

1. If $\limsup_{n \rightarrow \infty} \sqrt[n]{a_n} < 1$, then the series $\sum_{n=0}^{\infty} a_n$ converges.
2. If $\limsup_{n \rightarrow \infty} \sqrt[n]{a_n} > 1$, then the series $\sum_{n=0}^{\infty} a_n$ diverges.

Example 5.2.15. — If we consider the series $\sum_{n=1}^{\infty} \frac{1}{n}$, the series diverges while the (enhanced) root test is inconclusive. On the other hand, the series $\sum_{n=1}^{\infty} \frac{1}{n^2}$ converges while, again, the (enhanced) root test is inconclusive.

Proof of Proposition 5.2.14. Assume that $\limsup \sqrt[n]{a_n} > 1$ and take a subsequence $(a_{n_k})_{k \in \mathbb{N}}$ such that $\sqrt[n_k]{a_{n_k}} \geq 1$ for every $k \in \mathbb{N}$. Then

$$\sum_{n=0}^{\infty} a_n \geq \sum_{k=0}^{\infty} a_{n_k} \geq \sum_{k=0}^{\infty} 1 = \infty$$

On the other hand, assume that $\limsup \sqrt[n]{a_n} < L$ for some $L < 1$. We now show that the sequence (b_m) defined by $b_m = \sum_{n=0}^m a_n$ is a Cauchy sequence in (\mathbb{R}, d_1) . Since (\mathbb{R}, d_1) is complete, this implies that b_m converges.

Because $\limsup \sqrt[n]{a_n} < L$, we can find M such that $\sqrt[n]{a_n} \leq L$ for every $n \geq M$. For every $\epsilon > 0$ take $N \geq M$ such that $\frac{L^N}{1-L} < \epsilon$. Assume that $m_1 < m_2$, then we have

$$|b_{m_2} - b_{m_1}| = \left| \sum_{n=m_1+1}^{m_2} a_n \right| \leq \sum_{n=m_1+1}^{m_2} |a_n| \leq \sum_{n=N}^{\infty} L^n = \frac{L^N}{1-L} < \epsilon$$

For all $N < m_1 < m_2$. This shows that the sequence (b_m) is Cauchy, and therefore it concludes our proof. \square

We are now ready to prove Theorem 5.2.3.

Proof of Theorem 5.2.3. We start by noting that, if we fix indices m_1, m_2 such that $m_1 < m_2$, we have

$$\begin{aligned} |S_{m_2}(x) - S_{m_1}(x)| &\leq \sum_{n=m_1+1}^{m_2} |a_n| |x|^n \\ &\leq \sum_{n=m_1+1}^{m_2} |a_n| \delta^n \end{aligned}$$

Following the proof of Proposition 5.2.14 (part 1), we fix L with $\delta < L < R$. Then, for every $\epsilon > 0$, by the definition of lim sup there exists N such that

$$\sqrt[n]{|a_n|} \leq \frac{1}{L}$$

for all $n > N$, and such that

$$\frac{\left(\frac{\delta}{L}\right)^N}{1 - \frac{\delta}{L}} < \epsilon.$$

Then, for all m_1, m_2 such that $N < m_1 < m_2$ we have:

$$\begin{aligned} |S_{m_2}(x) - S_{m_1}(x)| &\leq \sum_{n=N}^{\infty} |a_n| \delta^n \\ &\leq \sum_{n=N}^{\infty} \left(\frac{\delta}{L}\right)^n \\ &= \frac{\left(\frac{\delta}{L}\right)^N}{1 - \frac{\delta}{L}} \\ &< \epsilon \end{aligned}$$

This concludes the proof that the sequence (S_N) is Cauchy in $C[-\delta, \delta]$. \square

Remark 5.2.16. — One may wonder if the radius of convergence R that we defined above is optimal. Or, more precisely, if it is possible that the sequence $(S_N(x))$ converges for some x with $|x| > R$ (notation as above). This is not possible. Indeed, for such x , by the definition of R we have

$$\limsup_{n \rightarrow \infty} \sqrt[n]{|a_n x^n|} > 1$$

so the sequence $a_n x^n$ does not tend to zero, therefore the series

$$\sum_{n=0}^{\infty} a_n x^n$$

does not converge by the basic criterion of convergence for series.

5.3 Differential equations

Exercise 5.2.17. — Determine the radius of convergence of the power series $\sum_{n=0}^{\infty} n^{2-n} x^n$.

Exercise 5.2.18. — Let $x \neq 0$ and consider the series

$$S(x) = \sum_{n=0}^{\infty} \left(-n \frac{x^2 - 1}{x} \right)^n.$$

Find an appropriate change of variable to write s as the sum of a power series, and determine its radius of convergence. On what subset of \mathbb{R} is the limit function well-defined?

Exercise 5.2.19. — Determine the radius of convergence of the power series $\sum_{n=0}^{\infty} a_n x^n$, where

$$a_n := \begin{cases} 2^n & n \text{ even} \\ 3^n & n \text{ odd.} \end{cases}$$

Then determine the full interval of convergence of the power series (i.e. study if the power series converges at the points $x = \pm R$).

5.3. Differential equations

The goal of this section is to show how to use the contraction mapping theorem as a tool for finding existence and uniqueness results for local solutions to some particular examples of differential equations.

We will show that a first order ordinary differential equation paired with an initial value (which we will call a "Cauchy problem") is equivalent to an integral equation (called Volterra integral equation). Then we will see that the integral equation can be formulated as a fixed point problem for some function between spaces of continuous functions.

We start by providing some examples of differential equations, and by highlighting that typically a differential equation does not have a unique solution.

Example 5.3.1. — Let us consider the following two differential equations

$$y'(x) = x^2 \tag{5.2}$$

$$y'(x) = y(x). \tag{5.3}$$

We can solve each equation in an elementary way. For the first one, we are looking for a function whose derivative is $f(x) = x^2$. For the second case,

we are looking for functions with the property that their derivatives equals the function itself.

A solution of (5.2) is given by

$$y(x) = \frac{x^3}{3} + c$$

for $c \in \mathbb{R}$ an arbitrary real number. A solution of (5.3) is given by

$$y(x) = ce^x,$$

for c an arbitrary real number.

Note that in neither example is the solution unique.

A Cauchy problem consists of an (ordinary) differential equation in 1 variable paired up with the constraint that the unknown function must equal a certain value when calculated at a given point.

Example 5.3.2. — Let's start from the second example above, and let's pair it up with the constraint that a solution y must equal y_0 at the point x_0 (for some choice of real numbers x_0 and y_0):

$$\begin{cases} y'(x) = y(x) \\ y(x_0) = y_0, \end{cases}$$

then we have a *unique* solution given by

$$y(x) = y_0 e^{x-x_0}.$$

Consider then the following system of equations

$$\begin{cases} y'(x) = f(x) \\ y(x_0) = y_0, \end{cases}$$

where the first line is Equation (5.2). This system has a *unique* solution given by

$$y(x) = y_0 + \int_{x_0}^x f(t) dt,$$

The examples above show that it is unreasonable to expect a differential equation to have a unique solution, but one has better chances when the differential equation is paired up with some additional constraint. This idea is enshrined in the following notion.

5.3 Differential equations

Definition 5.3.3. — A *Cauchy Problem* (CP) is the data of a point $(x_0, y_0) \in \mathbb{R}^2$, of a pair $a, b \in \mathbb{R}_{\geq 0} \cup \{+\infty\}$, and of a continuous function


$$f: [x_0 - a, x_0 + a] \times [y_0 - b, y_0 + b] \rightarrow \mathbb{R}. \quad (5.4)$$

A *local solution* to that Cauchy problem is a differentiable function

$$y: [x_0 - \alpha, x_0 + \alpha] \rightarrow [y_0 - b, y_0 + b]$$

for some $0 < \alpha \leq a$ such that

$$\begin{cases} y'(x) = f(x, y(x)) & \forall x \in [x_0 - \alpha, x_0 + \alpha] \\ y(x_0) = y_0. \end{cases}$$



A *local solution* y is said to be a (*global*) *solution* if one can take $\alpha = a$.

Our aim in this section is to prove that, under certain hypotheses, a Cauchy problem admits a unique local solution. This will be proven by applying the Contraction Mapping Theorem from the previous chapter. Here is the result that we will prove.

Theorem 5.3.4 (Peano–Picard). — *Let (x_0, y_0, a, b, f) be the defining data of a Cauchy problem. If $\frac{\partial f}{\partial y}$ exists and it is continuous on some open subset of \mathbb{R}^2 containing (x_0, y_0) , then the Cauchy problem*

$$\begin{cases} y'(x) = f(x, y(x)), \\ y(x_0) = y_0 \end{cases}$$

admits a unique local solution.

(This result follows immediately from a more precise and detailed version below, called Theorem 5.3.7).

The idea is to reformulate the Cauchy problem so that solutions correspond to fixed points of a certain contraction from a space of continuous functions to itself. This is achieved essentially by means of the Fundamental Theorem of Calculus, Theorem 5.1.20. The reformulation of a Cauchy problem is in terms of an integral equation, known under the name of Volterra Integral Equation (VIE), the integral equation

$$y(x) = y_0 + \int_{x_0}^x f(t, y(t)) dt$$

with unknown y . The data to define the VIE problem are exactly the same used to define a CP, but the solutions are defined differently (although we will eventually prove that the solutions for CP and VIE defined by the same data coincide).

Definition 5.3.5. — Assume that (x_0, y_0, a, b, f) are the data that define a Cauchy problem (see 5.3.3).

A *local solution* to the Volterra Integral Equation that they define is a continuous function

$$y: [x_0 - \alpha, x_0 + \alpha] \rightarrow [y_0 - b, y_0 + b]$$

for some $\alpha \in (0, a]$, such that

$$y(x) = y_0 + \int_{x_0}^x f(t, y(t)) dt \tag{5.5}$$

for every $x \in [x_0 - \alpha, x_0 + \alpha]$.

(And the local solution is said to be *global* if α can be chosen to equal a).

We now apply the fundamental theorem of calculus to prove that local solutions to the CP correspond bijectively to local solutions to the VIE.

Lemma 5.3.6. — *Let (x_0, y_0, a, b, f) be the defining data of a CP, and let*

$$y: [x_0 - \alpha, x_0 + \alpha] \rightarrow [y_0 - b, y_0 + b]$$

be a function, for some $\alpha \in (0, a]$. Then, y is a local solution to the CP if and only if y is a local solution to the VIE defined by the same data.

Proof. Assume that y is a solution to the Cauchy problem. Then we have

$$\begin{aligned} y_0 + \int_{x_0}^x f(t, y(t)) dt &= y_0 + \int_{x_0}^x y'(t) dt \\ &= y_0 + [y(t)]_{x_0}^x \\ &= y_0 + y(x) - y(x_0) \\ &= y(x), \end{aligned}$$

5.3 Differential equations

where in the first and the last equality we used the fact that y is a local solution for the Cauchy problem. In the second equality we applied the second part of the Fundamental Theorem of Calculus 5.1.20.

Conversely, assume that y satisfies the Volterra Integral Equation, then by the first part of the Fundamental Theorem of Calculus the function

$$y(x) = y_0 + \int_{x_0}^x f(t, y(t)) dt$$

is differentiable, and its derivative $y'(x)$ is given by $f(x, y(x))$. Because $\int_{x_0}^{x_0} = 0$, we deduce that $y(x_0) = y_0$. \square

Observe that a solution to the VIE (5.5) is naturally the fixed point of the function that maps y to the function defined by

$$y_0 + \int_{x_0}^x f(t, y(t)) dt.$$

Let us define carefully the metric space (of functions) and the function that this is the fixed point of.

For a given Cauchy problem, define the set

$$X = \{y: [x_0 - \alpha, x_0 + \alpha] \rightarrow [y_0 - b, y_0 + b], y \text{ is continuous}\} \subseteq C[x_0 - \alpha, x_0 + \alpha]$$

for some $\alpha \in (0, a]$. Then, we define a function

$$F: X \rightarrow X$$

sending a $y \mapsto F(y)$, where the latter is the defined as

$$F(y)(x) = y_0 + \int_{x_0}^x f(t, y(t)) dt.$$

Then, finding a local solution to the Volterra integral equation is the same as finding a fixed point of $F: X \rightarrow X$.

Theorem 5.3.7. — *Let (x_0, y_0, a, b, f) be the defining data of a Cauchy problem and assume $a, b < \infty$. Then, if $f: \mathbb{R} \rightarrow \mathbb{R}$ is differentiable with respect to its second variable and $\frac{\partial f}{\partial y}$ is continuous on the rectangle*

$$R := [x_0 - a, x_0 + a] \times [y_0 - b, y_0 + b],$$

there exists a unique local solution to the Cauchy problem.

More precisely, if we set:

$$M = \max_{(x,y) \in \mathbb{R}} |f(x,y)|$$

$$L = \max_{(x,y) \in \mathbb{R}} \left| \frac{\partial f}{\partial y}(x,y) \right|$$

then for any α with

$$\alpha < \min \left(a, \frac{b}{M}, \frac{1}{L} \right)$$

There exists a unique local solution

$$y: [x_0 - \alpha, x_0 + \alpha] \rightarrow [y_0 - b, y_0 + b]$$

to the Cauchy problem.

Proof. We prove that the space (X, d_{L^∞}) is complete, that the function

$$F: X \rightarrow X$$

is well defined and that F is a contraction. Then, by the Contraction Mapping Theorem we can conclude that a unique solution to the Volterra integral equation exists, which is equivalent to say that the Cauchy problem has a unique solution.

We have that X equals the closed ball centered at the constant function y_0 and of radius b , hence X is closed in $(C[x_0 - \alpha, x_0 + \alpha], d_{L^\infty})$ (for we have seen that all closed balls in a metric space are closed subsets). We conclude that X is complete with the metric d_{L^∞} from the fact that the metric space $(C[x_0 - \alpha, x_0 + \alpha], d_{L^\infty})$ is itself complete.

In the remainder of this proof, we will need to consider inequalities among integrals. It will be convenient to adopt the convention

$$\int_a^b g(t) dt = - \int_b^a g(t) dt.$$

Note that we will need to compute integrals from a to b where $a > b$, so that (for example), the integral of a continuous positive function will be negative.

To show that $F: X \rightarrow X$ is well-defined, we need to prove that $F(y)$ is continuous and that $y_0 - b \leq F(y)(x) \leq y_0 + b$ for every $x \in [x_0 - \alpha, x_0 + \alpha]$.

The fact that $F(y)$ is continuous follows from the fact that it is differentiable. Indeed, by definition, and assuming that y is continuous, $F(y)$

5.3 Differential equations

is the integral from x_0 to x of a continuous function, and therefore by the fundamental theorem of calculus $F(y)$ is differentiable with respect to the variable x .

In order to show that

$$|F(y)(x) - y_0| \leq b$$

for every $x \in [x_0 - \alpha, x_0 + \alpha]$, it is enough to pick $\alpha \leq \frac{b}{M}$ as in the statement of the theorem, and then consider the following inequalities

$$\begin{aligned} |F(y)(x) - y_0| &= \left| \int_{x_0}^x f(t, y(t)) dt \right| \\ &\leq M |x - x_0| \\ &\leq M\alpha \\ &\leq b. \end{aligned}$$

To conclude the proof, we show that F is a contraction. In particular, we will use the hypothesis of continuity of the partial derivatives.

Fix $t_0 \in [x_0 - \alpha, x_0 + \alpha]$ and apply the Mean Value Theorem (from Year 1) to the differentiable function of 1 variable

$$f(t_0, -): [y_0 - b, y_0 + b] \rightarrow \mathbb{R}.$$

Hence for given values $b_1, b_2 \in [y_0 - b, y_0 + b]$ there exists $\xi \in (b_1, b_2)$ such that

$$\frac{f(t_0, b_2) - f(t_0, b_1)}{b_2 - b_1} = \frac{\partial f}{\partial y}(t_0, \xi).$$

So, in particular

$$|f(t, b_2) - f(t, b_1)| \leq L |b_2 - b_1|.$$

Now we get

$$\begin{aligned} d_{L^\infty}(F(y_1), F(y_2)) &= \max_{x_0 - \alpha \leq x \leq x_0 + \alpha} \left| \int_{x_0}^x f(t, y_1(t)) - f(t, y_2(t)) dt \right| \\ &\leq L |x - x_0| \cdot d_{L^\infty}(y_1, y_2) \\ &\leq \alpha \cdot L \cdot d_{L^\infty}(y_1, y_2), \end{aligned}$$

which concludes the proof that F is a contraction, because by our assumption $\alpha \cdot L < 1$. □

Since the result uses the Contraction Mapping Theorem to prove the existence and uniqueness of a local solution, the proof of the CMT also provides an effective way of approximating that local solution:

Remark 5.3.8. — With the notation as above, if we take any $u \in X$ and apply F , we get a sequence

$$u, F(u), F(F(u)), \dots, F^n(u), \dots$$

which is a Cauchy sequence that converges to the solution to the Cauchy problem in the metric d_{L^∞} .

If we want to approximate the unique local solution y on the interval $[x_0 - \alpha, x_0 + \alpha]$, we can use that F is a contraction of constant $\alpha \cdot L$ combined with Remark 4.3.9 to conclude that, for all differentiable $u: [x_0 - \alpha, x_0 + \alpha] \rightarrow [y_0 - b, y_0 + b]$, the inequality

$$d_{L^\infty}(y, F^m(u)) \leq d_{L^\infty}(u, F(u)) \cdot \frac{(\alpha \cdot L)^m}{1 - \alpha \cdot L}$$

holds. Because $\alpha \cdot L$ is smaller than 1, the right hand side of the above inequality can be made arbitrarily small by taking m sufficiently large,

We would like now to give some counterexamples when the hypothesis of the theorem fails. Note that if we keep continuity of the main defining datum of a Cauchy problem, then the following general result (which we are not going to prove) guarantees the existence of local solutions.

Remark 5.3.9 (Peano's existence Theorem). — If f is a continuous function there exists a (non necessarily unique) local solution to the corresponding Cauchy problem. This result will not be discussed in this course.

For a counterexample we will therefore focus on the failure of uniqueness of local solutions.

The following counterexample to uniqueness makes use of a method known as "solving differential equation by separation of variables" that we have not discussed/explained in this module. The example has been written in a way that it should be understandable even if you have not been exposed to that method. Learning how to apply the method of separation of variables to solve differential equations goes beyond the scopes of this module.

Example 5.3.10. — Let us consider the Cauchy problem

$$\begin{cases} y'(x) = 3y^{\frac{2}{3}}(x) \\ y(x_0) = y_0 \end{cases}$$

5.3 Differential equations

Here the defining function f is $f(x, y) = 3y^{\frac{2}{3}}$. It is independent of x and it is continuous on \mathbb{R}^2 , and its partial derivative $\frac{\partial f}{\partial y}(x, y) = 2y^{-\frac{1}{3}}$ is defined everywhere except at $y = 0$.

Given a point (x_0, y_0) in \mathbb{R}^2 , if $y_0 \neq 0$ we can find a rectangle R around (x_0, y_0) such that the Cauchy problem satisfies the hypothesis of Theorem 5.3.7, while if $y_0 = 0$ this is not possible.

To find explicit solutions to the Cauchy problem we first divide the differential equation by $3y^{\frac{2}{3}}(x)$ and take the integral from x_0 to x , thus obtaining:

$$\int_{x_0}^x \frac{y'(t)}{3y^{\frac{2}{3}}(t)} dt = \int_{x_0}^x 1 dt = x - x_0.$$

(NOTE: in order for this to make sense, we need to assume that $y(t)$ is different from zero for t in the interval $[x_0, x]$, for otherwise the denominator vanishes. This is OK when $y(x_0) \neq 0$ and for x sufficiently close to x_0 , because of the continuity assumption for y).

To compute the first integral, we perform the change of variables:

$$y(t) = z, \text{ with } y'(t)dt = dz,$$

and apply the integration by substitution theorem, to obtain:

$$\begin{aligned} \int_{x_0}^x \frac{y'(t)}{3y^{\frac{2}{3}}(t)} dt &= \int_{y(x_0)}^{y(x)} \frac{z^{-\frac{2}{3}}}{3} dz \\ &= \left[z^{\frac{1}{3}} \right]_{y_0}^{y(x)} \\ &= (y(x))^{\frac{1}{3}} - y_0^{\frac{1}{3}}. \end{aligned}$$

Therefore, we get that solutions of the Cauchy problem have the form:

$$y(x) = \left(x - x_0 + y_0^{\frac{1}{3}} \right)^3$$

This argument proves that this is the unique local solution of the given Cauchy problem when $y(x_0) = y_0 \neq 0$.

We now analyse the case when $y_0 = 0$. One can check that the previous solution is still valid (by verifying that it satisfies the CP), but now the constant function $y(x) = 0$ is also a solution! We can blend these two families of solutions together, and deduce that for every $c, d > 0$ with $-c < x_0 < d$

the function

$$y_{c,d}(x) := \begin{cases} 0 & -c \leq x \leq d \\ (x+c)^3 & x \leq -c \\ (x-d)^3 & x \geq d \end{cases}$$

is a solution to the original CP! (In fact, to check this assertion it is just enough to verify that $y_{c,d}$ satisfies the differential equation and the initial value condition, there is no need to understand how the solution was found).

Therefore $y_{c,d}$ gives a family (parametrised by positive real numbers c and d) of solutions and in particular, for every $\epsilon > 0$ there are infinitely many solutions on the interval $[x_0 - \epsilon, x_0 + \epsilon]$ for the given Cauchy problem when $y_0 = 0$.

We conclude this section by proposing some exercises to consolidate the understanding of the theory that we have discussed.

Exercise 5.3.11. — Consider the Cauchy problem

$$\begin{cases} y'(x) = \cos(y(x)^2(x-1)) \\ y(0) = 1. \end{cases}$$

1. Does it satisfy the hypotheses of Theorem 5.3.7?
2. Describe the equivalent Volterra integral equation.

Exercise 5.3.12. — Consider the Volterra integral equation given by

$$y(x) = 1 + \int_1^x (t \cdot y(t) + \sqrt{y(t)}) dt$$

1. Describe the equivalent Cauchy problem.
2. Does it satisfy the hypotheses of Theorem 5.3.7?

Exercise 5.3.13. — Let

$$X = \left\{ y \in C \left[-\frac{1}{12}, \frac{1}{12} \right] : y(x) \in [0, 2], \quad \forall x \in \left[-\frac{1}{12}, \frac{1}{12} \right] \right\}$$

1. Show that (X, d_{L^∞}) is a complete metric space.

5.3 Differential equations

2. Show that the function

$$F: C\left[-\frac{1}{12}, \frac{1}{12}\right] \rightarrow C\left[-\frac{1}{12}, \frac{1}{12}\right]$$

Sending y to the function $F(y)$ defined by the formula

$$F(y)(x) = 1 + \int_0^x (y^2(t) + y(t) + t) dt$$

is well-defined.

3. Prove that $F(y)(x) \in [0, 2]$ for all $y \in X$ and $x \in [-\frac{1}{12}, \frac{1}{12}]$.
4. Prove that F defines a contraction on X .
5. Explain why the above implies the existence of a solution for the associated Cauchy problem.

6.

Introduction to multi-variable calculus

Important notions to learn from this section:

1. Linear functions $\mathbb{R}^n \rightarrow \mathbb{R}^m$ and how to represent them with matrices.
2. Partial and directional derivatives. The Jacobian matrix.
3. Differentiable functions and the notion of differential.
4. All partial derivatives exist and are continuous \implies differentiable \implies continuous.
5. The operator norm and the Chain Rule.
6. The mean value theorem (MVT) for vector-valued functions of several variables.

This chapter will be devoted to the study of the first basic elements of multi-variable calculus. The main theme is to understand the behaviour of non-linear functions between higher dimensional real vector spaces \mathbb{R}^n and \mathbb{R}^m in terms of linear functions, which are very well understood and for which we have all the results from linear algebra at our disposal. The idea is that some of the behaviour of nonlinear functions, at least locally, is captured by their linear approximation. The need to find a linear function that best approximates a nonlinear one leads us to introduce the concept of *differential*. The philosophy of linearisation will reach its apex in the next chapter.

From now on we abandon the generality of arbitrary metric spaces, and work with the usual Euclidean spaces \mathbb{R}^n with the distance d_2 . (One needs the structure of a vector space to make sense of a linear function, and we won't deal with infinite dimensional vector spaces, so we may as well work with coordinate and work directly with \mathbb{R}^n . A similar theory could be developed from a more general standpoint in the infinite dimensional case, by considering normed spaces that are complete, also known as Banach spaces).

For $v \in \mathbb{R}^n$, it will be convenient to denote by $\|v\| := d_2(v, 0)$ (called: the *norm* of the vector v), the distance of the vector v from the origin of \mathbb{R}^n .

6.1. A quick recap of linear algebra

We first present a recollection of basic results from linear algebra.

Definition 6.1.1. — Let $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a function, we say that f is *linear* if the following properties hold:

1. For every pair of vectors (v, w) in \mathbb{R}^n :

$$f(v + w) = f(v) + f(w).$$

2. For every vector $v \in \mathbb{R}^n$ and every constant $\lambda \in \mathbb{R}$

$$f(\lambda v) = \lambda f(v).$$

Remark 6.1.2. — There is a natural bijection between linear functions

$$f: \mathbb{R}^n \rightarrow \mathbb{R}^m$$

and $m \times n$ matrices with coefficients in \mathbb{R} . Given an $m \times n$ matrix

$$A = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn} \end{pmatrix}$$

we define a linear function $L_A: \mathbb{R}^n \rightarrow \mathbb{R}^m$ via the rule:

$$L_A(x_1, \dots, x_n) = \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn} \end{pmatrix} \cdot \begin{pmatrix} x_1 \\ x_2 \\ \cdots \\ x_n \end{pmatrix}$$

6.1 A quick recap of linear algebra

where the multiplication is matrix multiplication, and all vectors are written as column vectors.

Conversely, if $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear function, there exists a unique matrix A such that $L_A = f$. Let $e_j = (0, \dots, 1, \dots, 0)$ be the vector in \mathbb{R}^n with 1 in the j -th coordinate and 0 everywhere else, then the coefficient a_{ij} of A is given by the i -th coordinate of the vector $f(e_j)$ in \mathbb{R}^m .

Example 6.1.3. — Let $f: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be the linear function given by the assignment

$$f(x_1, x_2, x_3) = (2x_1 - x_2, x_3)$$

Then, applying f to e_j for $j = 1, 2, 3$ we get:

$$f(1, 0, 0) = (2, 0) \quad f(0, 1, 0) = (-1, 0) \quad f(0, 0, 1) = (0, 1).$$

Therefore, if we define A to be the matrix

$$A = \begin{pmatrix} 2 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Then, we see that $L_A = f$.

6.1.4. — Let $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $g: \mathbb{R}^m \rightarrow \mathbb{R}^l$ be linear functions and let A and B be the matrices such that $f = L_A$ and $g = L_B$. Then, the composition $g \circ f$ is a linear function as well and we have that

$$g \circ f = L_{B \cdot A}$$

where \cdot denotes the matrix multiplication. Notice that A is an $m \times n$ matrix and B is an $l \times m$ matrix, so that $B \cdot A$ is an $l \times n$ matrix.

6.1.5. — Recall that, given a linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ associated to a matrix A , the function f is a bijection if and only if A is *invertible* and the matrix A is invertible if and only if $m = n$ and the determinant of A is non-zero (and this can only happen when $m = n$).

We now recall the definition of the determinant of a matrix and how to compute it. We denote by $M_{n,m}(\mathbb{R})$ the set of $n \times m$ matrices with coefficients in \mathbb{R} and simply by $M_n(\mathbb{R})$ the set of $n \times n$ matrices.

Definition 6.1.6. — For every n , we define the *determinant function*

$$\det: M_n(\mathbb{R}) \rightarrow \mathbb{R}$$

recursively, as follows. If $n = 1$ and a is a real number, we define $\det(a) = a$. Let $n > 1$ and let A be an $n \times n$ matrix, and let us fix $i \in \{1, \dots, n\}$. Then the determinant of A is the real number:

$$\det(A) = \sum_{j=1}^n (-1)^{i+j} a_{ij} \det(A_{ij})$$

Where A_{ij} is the $(n - 1) \times (n - 1)$ matrix obtained by removing the i -th row and the j -th column from A . One can prove that this definition does not depend on the choice of i .

Example 6.1.7. — Let A be the four by four matrix with coefficients in \mathbb{R} given by:

$$A = \begin{pmatrix} 1 & 0 & 0 & 2 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 2 & 7 \\ 1 & 2 & 0 & 1 \end{pmatrix}$$

Then, choosing the second row for the first step and the first row for the second step, we compute the determinant of A as:

$$\begin{aligned} \det(A) &= (-1)^{(2+4)} \det \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 2 \\ 1 & 2 & 0 \end{pmatrix} \\ &= (-1)^{1+1} \det \begin{pmatrix} 0 & 2 \\ 2 & 0 \end{pmatrix} \\ &= -4 \end{aligned}$$

6.1.8. — We now list some properties of the determinant of a matrix. Let n be a positive integer and let us consider the set $M_n(\mathbb{R})$ of $n \times n$ matrices with coefficients in \mathbb{R} . First, we fix some notation, we denote by Id the *identity* matrix, the diagonal matrix given by:

$$\text{Id} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix}$$

Moreover, if v_1, \dots, v_n are column vectors in \mathbb{R}^n , we denote by (v_1, \dots, v_n) the $n \times n$ matrix with i -th column given by the entries of v_i

6.1 A quick recap of linear algebra

1. The determinant $\det(\text{Id})$ of the identity matrix is equal to 1.
2. The determinant is a multi-linear function. Let v_1, \dots, v_n and w be column vectors in \mathbb{R}^n and let $\lambda \in \mathbb{R}$, then
$$\det(v_1, \dots, v_i + w, \dots, v_n) = \det(v_1, \dots, v_i, \dots, v_n) + \det(v_1, \dots, w, \dots, v_n)$$
$$\det(v_1, \dots, \lambda v_i, \dots, v_n) = \lambda \det(v_1, \dots, v_n)$$
for all $i \in \{1, \dots, n\}$.
3. The determinant is an alternating function. In other words, if we consider a set v_1, \dots, v_n of vectors in \mathbb{R}^n , the equality:

$$\det(v_1, \dots, v_j, \dots, v_i, \dots, v_n) = -\det(v_1, \dots, v_i, \dots, v_j, \dots, v_n)$$

holds, for every $j < i$. Where the second matrix is obtained by switching the i -th and the j -th column.


In fact, the three properties above are *defining properties* of the determinant. In other words, the determinant is the unique function

$$M_n(\mathbb{R}) \rightarrow \mathbb{R}$$

that satisfies properties 1,2 and 3 above.

6.1.9. — We now give a geometric interpretation of the determinant. For simplicity, we fix $n = 3$ and we denote by e_i the column vector with 1 in the i -th place and 0 everywhere else, for $i = 1, 2, 3$.


The parallelepiped in \mathbb{R}^3 spanned by the three vectors $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$ and $e_3 = (0, 0, 1)$ is a cube (see figure, image courtesy of Wikipedia) with edges of length 1, hence whose volume equals 1.



Notice that the matrix (e_1, e_2, e_3) is the identity matrix and $\det(\text{Id}) = 1$. More generally, if A is an invertible matrix in $M_3(\mathbb{R})$ and we put

$$v_i = Ae_i \text{ for } i = 1, 2, 3,$$

then, by the multilinear property of the determinant, the absolute value $|\det(A)|$ of the determinant of A is equal to the volume of the parallelepiped spanned by v_1, v_2 and v_3 in \mathbb{R}^3 , see figure (image courtesy of Wikipedia).



Moreover, the sign of the determinant is positive when A “preserves the orientation” and negative when A “changes the orientation”, by the alternating property of the determinant.

In particular we have that A is not invertible, hence $\det(A) = 0$, if and only if the columns of A are linearly dependent, which can be interpreted geometrically by saying that the parallelepiped spanned by v_1, v_2 and v_3 is *degenerate*. In other words, the determinant of A is zero if and only if the volume of the corresponding solid figure is 0.

What we have written here for the case $n = 3$ remains valid for the case of arbitrary n , where the 3-dimensional parallelepiped is replaced by an arbitrary n -dimensional parallelotope.

Remark 6.1.10. — We point out that, although the determinant is a multilinear function, it is not linear. In other words, if A and B are n by n matrices and λ is a real number, we have in general that

$$\det(\lambda A) \neq \lambda \det(A)$$

$$\det(A + B) \neq \det(A) + \det(B).$$

Even more, the following relation occurs:

$$\det(\lambda A) = \lambda^n \det(A)$$

which implies that the determinant is a linear function if and only if $n = 1$.

However, one can show that the determinant respects products in the following sense:

$$\det(A \cdot B) = \det(A) \det(B).$$

where \cdot denotes the product of matrices.

6.2 Differentials

Example 6.1.11. — Let Id be the identity matrix in $M_n(\mathbb{R})$ for $n > 1$. Then,

$$\det(\text{Id} + \text{Id}) = \det(2 \cdot \text{Id}) = 2^n \det(\text{Id}) = 2^n.$$

On the other hand

$$\det(\text{Id}) + \det(\text{Id}) = 2.$$

We conclude with something that is not strictly speaking linear algebra, but that still fits in this section as it describes the behaviour of linear functions

Lemma 6.1.12. — Let $L: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear function. Then L is continuous.

Here the source and target are endowed with the standard Euclidean distance d_2 .

Proof. Since the distance d_∞ is equivalent to the Euclidean distance d_2 , we can use the distance d_∞ to prove the continuity of L .

Let A be the matrix corresponding to L , then the i -th component of $L(x_1, \dots, x_n)$ is given by $\sum_j a_{ij}x_j$. Let us define

$$M := \max_{1 \leq i \leq m} \left| \sum_j a_{ij} \right|.$$

We aim to prove that L is continuous at every point x in \mathbb{R}^n . For every $\epsilon > 0$, we let $\delta = \frac{\epsilon}{M}$. Then for all $y \in \mathbb{R}^n$ satisfying the inequality

$$\max_{1 \leq j \leq n} |y_j - x_j| < \delta,$$

we have

$$\begin{aligned} \max_{1 \leq i \leq m} \left| \sum_j a_{ij} (y_j - x_j) \right| &\leq M \max_{1 \leq j \leq n} |y_j - x_j| \\ &\leq M\delta = \epsilon \end{aligned}$$

This proves continuity at x , and concludes our proof. \square

6.2. Differentials

We now come back to the study of more general, possibly non-linear, functions $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$. From now on, unless otherwise stated, we will consider all spaces \mathbb{R}^k equipped with the Euclidean metric d_2 .

The generalisation of the theory of derivatives from $m = 1$ to arbitrary $m > 1$ is pretty straightforward. Not so the case of generalising $n = 1$ to arbitrary $n > 1$. We will see that a function from a source of dimension > 1 may admit all directional derivatives, yet be discontinuous! We will then introduce a better notion to linearly approximate a function, called the *differential* (or total derivative).

A function $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is uniquely determined by its components

$$f_1, \dots, f_m: \mathbb{R}^n \rightarrow \mathbb{R}$$

so that, for every vector $x = (x_1, \dots, x_n)$ in \mathbb{R}^n , the value of f at x is given by the vector $(f_1(x), \dots, f_m(x))$ in \mathbb{R}^m .

Example 6.2.1. — Let $f: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be given by

$$f(x_1, x_2, x_3) = (x_1^2, \sin(x_2x_3), e^{x_1}),$$

then the components of f are:

$$\begin{aligned} f_1(x_1, x_2, x_3) &= x_1^2 \\ f_2(x_1, x_2, x_3) &= \sin(x_2x_3) \\ f_3(x_1, x_2, x_3) &= e^{x_1}. \end{aligned}$$

Here is a first result that indicates that “separating the variables in the target space” is harmless.

Proposition 6.2.2. — Let $U \subset \mathbb{R}^n$ be a subset and $f = (f_1, \dots, f_m): U \rightarrow \mathbb{R}^m$ be a function and let $x \in U$. Then, f is continuous at x if and only if each f_i is continuous at x , for $i = 1, \dots, m$.

An $\epsilon - \delta$ proof would be possible, but unnecessarily long at this stage. We can give a quicker proof by relying on the characterisation of continuity by means of sequences (Lemma 2.2.7) and on the fact that we have already shown in Chapter 4 that a sequence converges in \mathbb{R}^m if and only if all of its coordinates converge.

Proof. By Lemma 2.2.7 the function f is continuous at x if and only if for all sequences $(y_k)_{k \in \mathbb{N}}$ such that $y_k \rightarrow x$, we have $f(y_k) \rightarrow f(x)$. By Lemma 4.2.2.1 the sequence $(f(y_k))_{k \in \mathbb{N}}$ converges to $f(x)$ if and only if for all $1 \leq i \leq m$ the sequence $(f_i(y_k))_{k \in \mathbb{N}}$ converges to $f_i(x)$. By applying again Lemma 2.2.7, the last condition is equivalent to the fact that each function f_i is continuous at x . □

6.2 Differentials

While the previous example shows that continuity can be checked coordinate-wise on the target, an analogue statement fails if one tries to separate the coordinates of a function on the source, as the following example illustrates.

Example 6.2.3. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined as

$$f(x_1, x_2) = \begin{cases} \frac{x_1 x_2}{x_1^2 + x_2^2} & (x_1, x_2) \neq (0, 0) \\ 0 & (x_1, x_2) = (0, 0) \end{cases}$$

If we restrict f to the $\{x_2 = 0\}$ -axis we obtain the function of 1 real variable

$$\begin{aligned} f(\cdot, 0): \mathbb{R} &\rightarrow \mathbb{R} \\ t &\mapsto f(t, 0) \end{aligned}$$

which is constantly zero. Similarly, f restricted to the $\{x_1 = 0\}$ -axis is the constant function with value 0, hence the restriction of f to both axis is continuous at $(0, 0)$. However, f is not continuous at $(0, 0)$.

Indeed, take the sequence $y_n = (\frac{1}{n}, \frac{1}{n})$ which converges to the point $(0, 0)$ in \mathbb{R}^2 . Then, $f(y_n) = \frac{1}{2}$ for every n . Hence:

$$\lim_{n \rightarrow \infty} f(y_n) = \frac{1}{2}.$$

On the other hand,

$$f\left(\lim_{n \rightarrow \infty} y_n\right) = f(0, 0) = 0.$$

which proves that f is not continuous at $(0, 0)$.

The point of the previous example is that the origin can be approached from several directions, not only moving along the fundamental axis.

Let us now review the notion of differentiability, starting from the case of functions of 1 real variable.

Definition 6.2.4. — Let $I \subset \mathbb{R}$ be an open interval and $x \in I$. We say that a function $f: I \rightarrow \mathbb{R}$ is *differentiable* at x if the limit:

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x)}{h}$$

exists. If f is differentiable at x we call the above limit the *derivative* of f at x and, we denote it by $f'(x)$.

Here we review some observations that you should already be familiar with from Year 1.

Exercise 6.2.5. — You should be familiar from Year 1 with the fact that functions may be continuous and not differentiable, and with the fact that the derivative of a differentiable function may be discontinuous.

1. Show that the function $x \mapsto |x|$ is not differentiable at $x = 0$.
2. Show that the function

$$x \mapsto \begin{cases} x^2 \sin(1/x) & \text{when } x \neq 0 \\ 0 & \text{when } x = 0 \end{cases}$$

is differentiable, but its derivative is discontinuous at $x = 0$.

(See later in this Chapter for a solution).

Directional derivatives — Here is a possible generalisation of the notion of the derivative for functions $\mathbb{R}^n \rightarrow \mathbb{R}^m$. We will allow for a more general set-up, where the function is defined on some open subset U of \mathbb{R}^n .

Definition 6.2.6. — Let $x = (x_1, \dots, x_n) \in U \subseteq \mathbb{R}^n$ with U an open subset, and let $f: U \rightarrow \mathbb{R}^m$ be a function. We define the *first order partial derivative* of f at x with respect to x_j as the limit

$$\frac{\partial f}{\partial x_j}(x) = \lim_{t \rightarrow 0} \frac{f(x_1, \dots, x_j + t, \dots, x_n) - f(x)}{t}$$

(when that limit exists).

The notion of partial derivatives should already be familiar from Calculus II.

Example 6.2.7. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be defined by

$$f(x_1, x_2) = (x_1^2, x_1 e^{x_2}, x_1 x_2)$$

then we have

$$\begin{aligned} \frac{\partial f}{\partial x_1}(x_1, x_2) &= (2x_1, e^{x_2}, x_2) \\ \frac{\partial f}{\partial x_2}(x_1, x_2) &= (0, x_1 e^{x_2}, x_1). \end{aligned}$$

6.2 Differentials

Example 6.2.8. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be the function defined by

$$f(x_1, x_2) = x_1^2 - 3e^{x_1 x_2}$$

then we have

$$\begin{aligned}\frac{\partial f}{\partial x_1}(x_1, x_2) &= 2x_1 - 3x_2 e^{x_1 x_2} \\ \frac{\partial f}{\partial x_2}(x_1, x_2) &= -3x_1 e^{x_1 x_2}.\end{aligned}$$

Example 6.2.9. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ be the function defined by

$$f(x_1, x_2) = (x_1^2, x_1^3, \cos(x_1 - x_2))$$

then we have

$$\begin{aligned}\frac{\partial f}{\partial x_1}(x_1, x_2) &= (2x_1, 3x_1^2, -\sin(x_1 - x_2)) \\ \frac{\partial f}{\partial x_2}(x_1, x_2) &= (0, 0, \sin(x_1 - x_2)).\end{aligned}$$

The partial derivatives of a function can be conveniently collected in a matrix, called the Jacobian matrix.

Definition 6.2.10. — With the same notation as in Definition 6.2.6, assume that the first order partial derivative of f at x with respect to x_j exists, for every $j \in \{1, \dots, n\}$.

Then, we define the *Jacobian matrix* of f at x as the $m \times n$ matrix given by:

$$Jf(x) = \begin{pmatrix} \frac{\partial f_1}{\partial x_1}(x) & \cdots & \frac{\partial f_1}{\partial x_n}(x) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1}(x) & \cdots & \frac{\partial f_m}{\partial x_n}(x) \end{pmatrix}.$$

Note that, as discussed in Chapter 6.1, the Jacobian matrix defines a linear function $L_{Jf(x)}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ by matrix multiplication: $v \mapsto Jf(x) \cdot v$.

Example 6.2.11. — Let us consider the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ defined in Example 6.2.7. We have shown that for every vector $x \in \mathbb{R}^2$, the first order partial derivatives of f at x exist. Moreover, the Jacobian matrix of f at x is given by

$$Jf(x) = \begin{pmatrix} 2x_1 & 0 \\ e^{x_2} & x_1 e^{x_2} \\ x_2 & x_1 \end{pmatrix}.$$

In the following, we will need the notion and notation for "the length" or "the norm" of a vector.

Remark 6.2.12. — Recall that the Euclidean metric d_2 gives a way to measure lengths of vectors $v \in \mathbb{R}^n$. The length of v is denoted by $\|v\|$, which we read as "the norm of v " defined by

$$\|v\| = d_2(v, 0) = \sqrt{v_1^2 + \dots + v_n^2}.$$

For later use, we note that the function

$$\|\cdot\| : \mathbb{R}^n \rightarrow \mathbb{R}$$

(calculating the norm of vectors) is continuous, which we leave as an easy exercise to the reader. (It also follows from the more general (and difficult) result that an arbitrary distance $d: X \rightarrow X \rightarrow \mathbb{R}_{\geq 0}$ is continuous – see the Enrichment problem of the third homework sheet).

The partial derivatives encode the information of the "rate of growth" of a function along the fundamental directions. That notion can be generalised to describe the rate of growth along any direction:

Definition 6.2.13. — Let $x \in U \subseteq \mathbb{R}^n$ with U open in \mathbb{R}^n , and $f: U \rightarrow \mathbb{R}^m$ be a function. Let $v \in \mathbb{R}^n$ be a vector of length 1 (i.e. such that $\|v\| = 1$).

We define the *directional derivative* of f at x with direction v as the limit:

$$D_v f(x) = \lim_{t \rightarrow 0} \frac{f(x + tv) - f(x)}{t}$$

when that limit exists.

Remark 6.2.14. — Assume the same hypotheses of Definition 6.2.13. Let $1 \leq i \leq n$ and the fundamental i -th direction, corresponding to the vector e_i (consisting of all zeroes except for a 1 in position i). Calculating the directional derivative of f along the direction e_i , we obtain

$$D_{e_i} f(x) = \lim_{t \rightarrow 0} \frac{f(x + te_i) - f(x)}{t} = \frac{\partial f}{\partial x_i}(x).$$

In other words, the directional derivative along the fundamental direction e_i is the same thing as the partial derivative with respect to the i -th variable.

In the following, we show an example of a function that is not continuous at the origin, yet it admits *all* the directional derivatives at that point.

6.2 Differentials

Example 6.2.15. — The idea to construct this example is to produce a function whose discontinuity at a point x cannot be detected by varying x along a straight line, i.e. by some perturbation of type $x + tv$ for some vector v and arbitrary small $t \in \mathbb{R}$.

Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be the function defined as

$$f(x_1, x_2) = \begin{cases} 1 & x_1, x_2 > 0 \text{ and } x_2 < x_1^2 \\ 0 & \text{otherwise.} \end{cases}$$

We claim that, for every direction v , the directional derivative along v of the function f at the origin $(0, 0)$ is equal to 0. Indeed, because $D_{-v}(f)(x) = -D_v(f)(x)$, we may assume without loss of generality that $v_1 \geq 0$. Then if $v_2 \leq 0$, then $f(tv) = 0$ for all t . If $v_2 > 0$, we have $f(tv) = 0$ for all t when $v_1 = 0$ and for all $t \leq \frac{v_2}{v_1^2}$ when $v_1 > 0$. Therefore, for every direction $v = (v_1, v_2) \in \mathbb{R}^2$, we have

$$\lim_{t \rightarrow 0} \frac{f(tv) - f(0, 0)}{t} = \lim_{t \rightarrow 0} \frac{f(tv)}{t} = \lim_{t \rightarrow 0} \frac{0}{t} = 0,$$

and this proves our claim.

However, f is not continuous at $(0, 0)$. Indeed let us consider the sequence

$$y_n = \left(\frac{1}{n}, \frac{1}{n^3} \right),$$

then $f(y_n) = 1$ for every n , so that:

$$\lim_{n \rightarrow \infty} f(y_n) = 1.$$

However, y_n converges to the origin, $(0, 0)$. Therefore

$$f\left(\lim_{n \rightarrow \infty} y_n\right) = 0.$$

6.2.16. Differentiability. — The notion of directional derivatives at a point is enough to provide information on the rate of growth of along the lines passing through that point. Yet, as we have seen in the previous section, it does not quite capture the whole infinitesimal behaviour of that function. For example, a function might have all directional derivatives at a point and yet fail to be continuous at that point.

We will start again from the notion of differentiability for functions of 1 real variable, and we will then generalise that notion into a notion that better captures the infinitesimal behaviour of the function.

Let $f: I \rightarrow \mathbb{R}$ be a function on some open interval $I \subset \mathbb{R}$. Then f is differentiable at a point $x_0 \in \mathbb{R}$ according to Definition 6.2.4 if and only if there exists a real number $\ell \in \mathbb{R}$ such that

$$\lim_{h \rightarrow 0} \frac{f(x_0 + h) - f(x_0) - \ell \cdot h}{h} = 0$$

In this equivalent formulation, the number ℓ is the derivative $f'(x_0)$ of f at x_0 .

The linear function

$$L_\ell: \mathbb{R} \rightarrow \mathbb{R}$$

$$h \mapsto \ell \cdot h$$

(multiplication by ℓ) is what we are going to call the *differential* of f at x_0 .

This interpretation can naturally be generalised to higher dimensions in the following manner.

Definition 6.2.17. — Let $U \subseteq \mathbb{R}^n$ be an open subset, let $x_0 \in U$ and let $f: U \rightarrow \mathbb{R}^m$ be a function. We say that f is *differentiable* at x_0 if there exists a linear function $L: \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

$$\lim_{h \rightarrow 0} \frac{\|f(x_0 + h) - f(x_0) - L(h)\|}{\|h\|} = 0.$$

The linear function L is called the *differential* of f at x_0 and it is denoted by Df_{x_0} or by $Df(x_0)$.

The above definition requires the existence of some linear function approximating the function f at the point x_0 . The first order of business is to observe that if such a function $L: \mathbb{R}^n \rightarrow \mathbb{R}^m$ exists, then it is unique.

Remark 6.2.18. — Let L_1 and L_2 be two linear functions $\mathbb{R}^n \rightarrow \mathbb{R}^m$ as in Definition 6.2.17. Then $L_1 = L_2$.

Indeed, by the triangle inequality we have

$$\|L_1(h) - L_2(h)\| \leq \|f(x_0 + h) - f(x_0) - L_1(h)\| + \|f(x_0 + h) - f(x_0) - L_2(h)\|.$$

Combining this with the fact that both L_1 and L_2 satisfy the definition of "being a differential for f at x_0 ", we deduce

$$\lim_{h \rightarrow 0} \frac{\|(L_1 - L_2)(h)\|}{\|h\|} = 0.$$

For all $v \in \mathbb{R}^n$ and $t \in \mathbb{R}$, we then deduce

$$0 = \lim_{t \rightarrow 0} \frac{\|(L_1 - L_2)(tv)\|}{\|tv\|} = \lim_{t \rightarrow 0} \frac{\|(L_1 - L_2)(v)\|}{\|v\|}$$

6.2 Differentials

(by linearity of $L_1 - L_2$ and because $\|tw\| = t\|w\|$). The right hand side of the last equality is independent of t , so we deduce $\|(L_1 - L_2)(v)\| = 0$, hence that $(L_1 - L_2)(v) = 0$.

Since this is valid for all v , we deduce that the linear function $L_1 - L_2$ equals zero, hence that $L_1 = L_2$.

The differential admits the following geometric interpretation, which we are only mentioning in passing, and without proof.

Remark 6.2.19. — Let $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a function, differentiable at a point $x_0 \in \mathbb{R}^n$, then the set of points given by

$$\{(x, f(x_0) + Df_{x_0}(x - x_0)) : x \in \mathbb{R}^n\}$$

can be interpreted geometrically as the hyperplane tangent to the graph of f at the point $(x_0, f(x_0)) \in \mathbb{R}^n \times \mathbb{R}^m$.

Most importantly for our discussion, is the fact that the differential contains the information of all directional derivatives:

Proposition 6.2.20. — Let $f: U \rightarrow \mathbb{R}^m$ be a function defined on an open subset $U \subseteq \mathbb{R}^n$, differentiable at x_0 , and let $v \in \mathbb{R}^n$ be a vector with $\|v\| = 1$ and $x_0 \in U$. Then, the differential of f at x_0 calculated at v coincides with the directional derivative of f at x_0 along v . In other words

$$Df_{x_0}(v) = D_v f(x_0).$$

Proof. We have that

$$\lim_{h \rightarrow 0} \frac{\|f(x_0 + h) - f(x_0) - Df_{x_0}(h)\|}{\|h\|} = 0$$

with $h \in \mathbb{R}^n$. In particular, we may write $h = tv$ for $v \in \mathbb{R}^n$ with $\|v\| = 1$ and $t \in \mathbb{R}$ and deduce

$$\lim_{t \rightarrow 0} \frac{\|f(x_0 + tv) - f(x_0) - Df_{x_0}(tv)\|}{\|tv\|} = 0$$

By continuity of the function $\|\cdot\|$ (see Remark 6.2.12), the left hand side is equal to

$$\left\| \lim_{t \rightarrow 0} \frac{(f(x_0 + tv) - f(x_0) - Df_{x_0}(tv))}{|t|} \right\|$$

which implies that

$$\lim_{t \rightarrow 0} \frac{(f(x_0 + tv) - f(x_0) - Df_{x_0}(tv))}{|t|} = 0.$$

because the only vector of norm 0 is the origin. The latter is equivalent to

$$\lim_{t \rightarrow 0} \frac{f(x_0 + tv) - f(x_0)}{t} = \lim_{t \rightarrow 0} \frac{Df_{x_0}(tv)}{t}.$$

The left hand side is equal to $D_v f(x_0)$ by definition of directional derivative, while the right hand side is equal to $Df_{x_0}(v)$ by linearity of the differential. □

We have defined the differential as a linear function, and linear functions are described by matrices (see Remark 6.1.2). We are now ready to reveal (and then prove) that the matrix that represents the differential is the Jacobian matrix!

Corollary 6.2.21. — *Let x_0, U, f be as above. Then, for every h in \mathbb{R}^n we have that*

$$Df_{x_0}(h) = Jf(x_0) \cdot h.$$

(In other words, the Jacobian matrix represents the differential in the standard basis of \mathbb{R}^n and \mathbb{R}^m).

Proof. If we evaluate Df_{x_0} along the fundamental direction $e_i = (0, \dots, 1, \dots, 0)$ (the 1 is in i -th position) we obtain the i -th column of the matrix corresponding to Df_{x_0} . Now, since e_i is a vector of length 1, we have that

$$Df_{x_0}(e_i) = D_{e_i} f(x_0) = \begin{pmatrix} \frac{\partial f_1}{\partial x_i}(x_0) \\ \dots \\ \frac{\partial f_m}{\partial x_i}(x_0) \end{pmatrix}$$

which is the i -th column of the Jacobian matrix, which precisely means that the Jacobian matrix represents the differential according to Remark 6.1.2. □

We are now ready to prove that differentiability implies continuity.

Proposition 6.2.22. — *Let x_0, U, f be as above. Because f is differentiable at x_0 , then f is continuous at x_0 .*

6.2 Differentials

Proof. To show that f is continuous at x_0 , it is enough to prove that

$$\lim_{h \rightarrow 0} \|f(x_0 + h) - f(x_0)\| = 0$$

We have that

$$\begin{aligned} \|f(x_0 + h) - f(x_0)\| &= \|f(x_0 + h) - f(x_0) - Df_{x_0}(h) + Df_{x_0}(h)\| \\ &\leq \|f(x_0 + h) - f(x_0) - Df_{x_0}(h)\| + \|Df_{x_0}(h)\| \end{aligned}$$

where the last step is the triangle inequality. Now, since f is differentiable at x_0 , the limit

$$\lim_{h \rightarrow 0} \frac{\|f(x_0 + h) - f(x_0) - Df_{x_0}(h)\|}{\|h\|}$$

vanishes, which implies that

$$\lim_{h \rightarrow 0} \|f(x_0 + h) - f(x_0) - Df_{x_0}(h)\| = 0$$

On the other hand, Df_{x_0} is continuous by Lemma 6.1.12. Therefore, if we fix $\epsilon > 0$, we can find a $\delta > 0$ such that

$$\|f(x_0 + h) - f(x_0) - Df_{x_0}(h)\| + \|Df_{x_0}(h)\| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

for all $0 < \|h\| < \delta$, which concludes our proof. \square

We know already from Year 1 that the converse implication does not hold in general: a function might be continuous and fail to be differentiable.

Example 6.2.23. — The function f defined by $f(x) = |x|$ is continuous but not differentiable at 0.

Example 6.2.24. — The function $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ of Example 6.2.15 has all directional derivatives, but it is not even continuous at $(0, 0)$. In particular, f is not differentiable.

The fact that a function admits all directional derivatives and is continuous, is not enough to guarantee that it is differentiable. Here is an example to illustrate this.

Example 6.2.25. — Here is an example of a continuous function, which has all directional derivatives, but is not differentiable at the origin.

Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined as

$$f(x_1, x_2) = \begin{cases} \frac{x_1 x_2^2}{x_1^2 + x_2^2} & (x_1, x_2) \neq (0, 0) \\ 0 & (x_1, x_2) = (0, 0) \end{cases}$$

Then, f is continuous at $(0,0)$. Indeed, first of all we observe that the following inequality holds

$$\left| \frac{x_1 x_2^2}{x_1^2 + x_2^2} \right| \leq \sqrt{x_1^2 + x_2^2}$$

since the absolute values of x_1 and x_2 are bounded above by $\sqrt{x_1^2 + x_2^2}$. In particular, for every $\epsilon > 0$, choosing $\delta = \epsilon$ and x_1, x_2 such that $\sqrt{x_1^2 + x_2^2} < \delta$, we get that

$$|f(x_1, x_2)| \leq \sqrt{x_1^2 + x_2^2} < \delta = \epsilon$$

by the above argument.

For the directional derivative at $(0,0)$, fix a vector $v = (v_1, v_2)$ of length 1 in \mathbb{R}^2 , then we have

$$\lim_{t \rightarrow 0} \frac{f(t(v_1, v_2)) - f(0,0)}{t} = \lim_{t \rightarrow 0} \frac{t^3 v_1 v_2^2}{t^3} = v_1 v_2^2$$

On the other hand, to show that f is not differentiable at $(0,0)$ we first compute the Jacobian matrix of f at $(0,0)$

$$Jf(0,0) = \left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2} \right) = (0,0)$$

Then, if f was differentiable at $(0,0)$ we would have that $Df_{(0,0)}(h) = 0$ for any h . This would imply:

$$\lim_{h \rightarrow 0} \frac{|f(h) - f(0,0) - Df_{(0,0)}(h)|}{\|h\|} = \lim_{h \rightarrow 0} \frac{|f(h)|}{\|h\|} = 0$$

However, if we take the sequence $(\frac{1}{n}, \frac{1}{n})$ and substitute it into the limit, we see that

$$\lim_{n \rightarrow \infty} \frac{|f(1/n, 1/n)|}{\|(1/n, 1/n)\|} = \lim_{n \rightarrow \infty} \frac{\frac{1}{n^3}}{\left(\frac{2}{n^2}\right)^{\frac{3}{2}}} = \frac{1}{\sqrt{8}}$$

is non-zero, hence f is not differentiable at $(0,0)$.

In line with the fact that the generalisation from $M = 1$ to arbitrary $M > 1$ isn't hard, let's now try to "separate the variables in the target". We start by reviewing the notation of the Jacobian in the case when the target is the real line.

6.2 Differentials

Remark 6.2.26. — If $M = 1$ and f is differentiable at $x_0 \in \mathbb{R}^n$, then the Jacobian matrix of f at x_0 reduces to a row vector

$$Jf(x_0) = \left(\frac{\partial f}{\partial x_1}(x_0), \dots, \frac{\partial f}{\partial x_n}(x_0) \right)$$

which is usually called the *gradient* of f at x_0 and denoted by $\nabla f(x_0)$. Moreover, the normalized vector

$$\frac{\nabla f(x_0)}{\|\nabla f(x_0)\|}$$

can be interpreted as the "direction of maximal growth of f " (we will not prove/discuss this fact here because we won't need it. You should be familiar with this result from Year 1, specifically from Calculus II).

We are now ready to discuss the fact that a function is differentiable if and only if so are all its components.

Proposition 6.2.27. — Let $f: U \rightarrow \mathbb{R}^m$ be a function defined on an open subset $U \subseteq \mathbb{R}^n$, and let $x_0 \in U$. The function f is differentiable at x_0 if and only if its components f_1, \dots, f_m are all differentiable at x_0 . Furthermore, the differential satisfies

$$Df_{x_0} = (Df_{1\ x_0}, \dots, Df_{m\ x_0})$$

(Or, in matrix terms, the Jacobian matrix Jf_{x_0} has first row $\nabla f_1(x_0)$, second row $\nabla f_2(x_0)$, ..., m -th row $\nabla f_m(x_0)$).

Proof. (Sketch)

We need to prove that

$$\lim_{h \rightarrow 0} \frac{\|f(x_0 + h) - f(x_0) - Df_{x_0}(h)\|}{\|h\|} = 0$$

if and only if

$$\lim_{h \rightarrow 0} \frac{|f_j(x_0 + h) - f_j(x_0) - Df_{j\ x_0}(h)|}{\|h\|} = 0$$

for all $j = 1, \dots, m$.

By now we know a strategy to simplify this problem: change the metric from d_2 to (for example) d_∞ . The details are then left to the reader as an exercise. \square

We conclude by giving a convenient criterion to establish if a function is differentiable.

Theorem 6.2.28. — Let $f: U \rightarrow \mathbb{R}^m$ be a function defined on an open subset $U \subseteq \mathbb{R}^n$ and let $x_0 \in U$. If f has all partial derivatives on U , and they are all continuous at x_0 , then f is differentiable at x_0 .

For brevity, we will say that a function that satisfies the hypothesis of the theorem above is continuously differentiable at x_0 :

Definition 6.2.29. — Let $f: U \rightarrow \mathbb{R}^m$ be a function defined on an open subset $U \subseteq \mathbb{R}^n$ and let $x_0 \in U$. We say that f is *continuously differentiable* at x_0 if there exists an open subset $V \subset U$ containing x_0 such that f admits all partial derivatives on V , and all partial derivatives are continuous at x_0 .

We will motivate this definition in Section 6.3.1 (see in particular Theorem 6.3.1.2).

Proof. (Sketch.)

By the previous proposition, and because being continuous and admitting partial derivatives can equivalently be checked on all components (in the target space \mathbb{R}^m), we may reduce the problem to the case $m = 1$.

The case $n = 1$ is trivial (the partial derivative equals the derivative which then gives the differential). We will give the proof in the case $n = 2$, the case of larger n is not substantially harder, it only makes the notation more complicated. Also, to simplify the notation we may assume (after possibly translating the function in the source and in the target space) that the point where we prove differentiability is $x_0 = (0, 0)$ and that $f(0, 0) = 0$. Furthermore, the property of differentiability of f will not change if we subtract from f a linear function. By subtracting the linear function defined by the Jacobian matrix of f at the origin, we may also assume that both partial derivatives of f at $(0, 0)$ are zero.

After all this preparation, in order to prove that f is differentiable at $(0, 0)$ and with differential equal to zero, we need to prove that

$$\lim_{h \rightarrow 0} \frac{|f(h)|}{\|h\|} = 0$$

or equivalently that for all $\epsilon > 0$ there exists $\delta > 0$ such that

$$|f(h)| \leq \epsilon \sqrt{h_1^2 + h_2^2}$$

for all (h_1, h_2) such that $0 < \sqrt{h_1^2 + h_2^2} < \delta$.

6.2 Differentials

The strategy is that we are going to separate the increments of h_1 from the increments of h_2 , and use the triangle inequality:

$$|f(\mathbf{h})| \leq |f(\mathbf{h}) - f(h_1, 0)| + |f(h_1, 0) - f(0, 0)|. \quad (6.1)$$

Now we apply the Mean Value Theorem from Year 1 to deduce

$$f(h_1, h_2) - f(h_1, 0) = \frac{\partial f}{\partial x_2}(h_1, c_2) \cdot h_2, \quad f(h_1, 0) - f(0, 0) = \frac{\partial f}{\partial x_1}(c_1, 0) \cdot h_1.$$

for some $|c_1| \leq |h_1|$ and $|c_2| \leq |h_2|$.

By applying the continuity of the partial derivatives, we find $\delta > 0$ such that

$$\left| \frac{\partial f}{\partial x_1}(v_1, v_2) \right| \leq \epsilon/2, \quad \left| \frac{\partial f}{\partial x_2}(v_1, v_2) \right| \leq \epsilon/2$$

for all $\sqrt{v_1^2 + v_2^2} < \delta$ (remember that we have assumed that the partial derivatives are both zero at the origin!)

Applying all these considerations to Equation (6.1) we obtain

$$|f(\mathbf{h})| \leq |h_2| \cdot \epsilon/2 + |h_1| \cdot \epsilon/2 \leq \epsilon \sqrt{h_1^2 + h_2^2}$$

for all (h_1, h_2) such that $0 < \sqrt{h_1^2 + h_2^2} < \delta$, because $\sqrt{h_1^2 + c_2^2} \leq \sqrt{h_1^2 + h_2^2}$, and because $\sqrt{0^2 + c_1^2} \leq \sqrt{h_1^2 + h_2^2}$. \square

The converse implication of the above result does not hold: a function might be differentiable, yet have discontinuous partial derivatives. This is already true for functions of 1 real variable, and it is an example that should already be familiar from Year 1.

Example 6.2.30. — Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be the function defined as

$$f(x) = \begin{cases} x^2 \sin\left(\frac{1}{x}\right) & x \neq 0 \\ 0 & x = 0 \end{cases}$$

Then, f is differentiable at 0 and $Df_0(\mathbf{h}) = 0$ for every $h \in \mathbb{R}$. Indeed

$$\begin{aligned} \lim_{h \rightarrow 0} \frac{|f(h) - f(0) - 0 \cdot h|}{|h|} &= \lim_{h \rightarrow 0} \left| \frac{h^2 \sin\left(\frac{1}{h}\right)}{h} \right| \\ &= \lim_{h \rightarrow 0} \left| h \sin\left(\frac{1}{h}\right) \right| \\ &= 0. \end{aligned}$$

Where the last equality follows since \sin is a bounded function. However, if we compute the derivative of $f(x)$ for $x \neq 0$, we obtain

$$f'(x) = 2x \sin\left(\frac{1}{x}\right) - \cos\left(\frac{1}{x}\right)$$

which does not have a limit for $x \rightarrow 0$. Therefore, the derivative f' is not continuous at 0.

6.3. Properties of differentials and of differentiable functions

In this section we discuss some fundamental properties of differentiable functions.

In the first subsection we discuss the notion of operator norm, and the continuity of the differential.

In the second subsection is the main result, the Chain Rule, which allows to calculate the differential of the composite of two differentiable functions in terms of the differentials of the two functions.

In the last subsection we generalise the Year 1 Mean Value Theorem to the case of functions of several variables.

6.3.1. The operator norm and continuity of the differential — Why did we call “continuously differentiable” a function with continuous partial derivatives? Shouldn’t “continuously differentiable” refer to a function whose differential is continuous when we vary the point where the differential is calculated? It turns out that these two notions are equivalent, as we shall prove in this section.

First we need to make some sense of the second definition. We start by thinking again about the differential of a function $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$. For fixed $P \in \mathbb{R}^n$, we defined Df_P as a linear function $\mathbb{R}^n \rightarrow \mathbb{R}^m$, and we observed that this function is represented by the Jacobian matrix Jf_P . This produces a function

$$Jf: \mathbb{R}^n \rightarrow \text{Mat}(m, n)$$

that maps every point P to the Jacobian matrix $Jf(P)$. On the source space we have a natural Euclidean distance, and on the target space we have the “operator distance” defined as follows.

6.3 Properties of differentials and of differentiable functions

Definition 6.3.1.1 (Operator norm). — Let A be a matrix in $M_{m,n}(\mathbb{R})$ (matrices with real entries and m rows and n columns). The *operator norm* of A is defined as

$$\|A\|_{\text{op}} = \max_{w \neq 0} \frac{\|Aw\|}{\|w\|} = \max_{\|v\|=1} \|Av\|.$$

(The second equality is obtained because multiplication times A is linear. The fact that the maximum exists follows from the fact that the function $v \mapsto \|Av\|$ is continuous, and from the fact that the set of points having distance 1 from the origin in \mathbb{R}^n is compact).

This definition gives a distance on the set $M_{m,n}(\mathbb{R})$ of $m \times n$ matrices by setting

$$d_{\text{op}}(A, B) = \|A - B\|_{\text{op}}$$

(We leave it to the reader to check that this satisfies the axioms of a distance on the set $M_{m,n}(\mathbb{R})$).

With the operator distance on the set of matrices, it now makes sense to ask whether the function Jf is continuous.

Theorem 6.3.1.2. — Let $U \subseteq \mathbb{R}^n$ be open, let $x_0 \in U$ and let $f: U \rightarrow \mathbb{R}^m$ be a function that has all partial derivatives $\frac{\partial f}{\partial x_j}$ on U for $j = 1, \dots, n$.

Then f is continuously differentiable at x_0 (Definition 6.2.29) if and only if $Jf: U \rightarrow \text{Mat}(m, n)$ is continuous at x_0 . (Here U is given the standard Euclidean distance, and $\text{Mat}(m, n)$ is given the operator distance defined above).

As a warm-up, and to fix notation, let's verify that the above result is valid in the case of *linear* functions.

Example 6.3.1.3. — (Linear functions are continuously differentiable). Let $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ be *linear*. Then:

1. the partial derivatives of f are constant functions, hence they are continuous;
2. the function $P \mapsto Jf_P$ is constant, hence continuous.

The remainder of this section is devoted to the proof of this theorem, and it can be regarded as extra material (which can safely be skipped on a first reading).

We will first prove that the operator distance on $\text{Mat}(m, n)$ is equivalent to other distances on the same set. We can regard the elements of that set as vectors in $\mathbb{R}^{n \times m}$ and endow the latter with any of the distances d_p (which we already know are equivalent distances). It turns out that the operator distance is equivalent to these distances:

Lemma 6.3.1.4. — The following inequalities hold:

$$\max_{i,j} |a_{i,j}| \leq \|A\|_{\text{op}} \leq \sqrt{m \cdot n} \max_{i,j} |a_{i,j}|.$$

Proof. The first inequality is easiest to prove. Let p, q be such that

$$|a_{p,q}| := \max_{i,j} |a_{i,j}|.$$

By multiplying A times the element $e_q = (0, 0, \dots, 1, \dots, 0)$ of the canonical basis, we obtain:

$$\|A \cdot e_q\| = \sqrt{\sum_{i=1}^m a_{i,q}^2} \geq |a_{p,q}|$$

for all $q = 1, \dots, n$. Applying the definition of operator norm (and the fact that $\|e_q\| = 1$), we deduce

$$\|A\|_{\text{op}} \geq \|A \cdot e_q\| \geq |a_{p,q}| = \max_{i,j} |a_{i,j}|.$$

The second inequality is more complicated to prove. We rely on the following inequality (left as an exercise.):

$$\sum_{i=1}^k |x_i| \leq \sqrt{k} \sqrt{\sum_{i=1}^k x_i^2}. \quad (6.2)$$

(Hint: expand the inequality

$$\sum_{1 \leq i < j \leq k} (|x_i| - |x_j|)^2 \geq 0.)$$

Now take a unit vector v such that $\|A\|_{\text{op}} = \|A \cdot v\|$, and let for convenience $L := \max_{i,j} |a_{i,j}|$. We have then:

$$\|A \cdot v\| = \left\| \left(\sum_{j=1}^n a_{1,j} v_j, \dots, \sum_{j=1}^n a_{m,j} v_j \right) \right\| \leq L \cdot \left\| \left(\sum_{j=1}^n |v_j|, \dots, \sum_{j=1}^n |v_j| \right) \right\| =$$

6.3 Properties of differentials and of differentiable functions

$$= L \cdot \sqrt{m} \cdot \sum_{j=1}^n |v_j| \leq \sqrt{m \cdot n} \cdot L \cdot \|v\| = \sqrt{m \cdot n} \cdot L = \sqrt{m \cdot n} \max_{i,j} |\alpha_{i,j}|,$$

where we used Equation (6.2) for the last inequality. This concludes the proof. \square

Using the inequalities that we proved in the previous Lemma, the proof of Theorem 6.3.1.2 is now straightforward.

Proof. (Sketch). First of all, by using that the distances d_2 and d_∞ are equivalent on \mathbb{R}^m , we may change to that distance. We then apply the two inequalities of the lemma.

From the inequality

$$\|Jf(x) - Jf(x_0)\|_{\text{op}} \leq \sqrt{m \cdot n} \max_{i,j} \left\| \frac{\partial f_i}{\partial x_j}(x) - \frac{\partial f_i}{\partial x_j}(x_0) \right\|$$

we deduce that if f is continuously differentiable at x_0 , then the function Jf is continuous.

The converse is obtained from the inequalities

$$\left| \frac{\partial f_i}{\partial x_j}(x) - \frac{\partial f_i}{\partial x_j}(x_0) \right| \leq \max_{i,j} \left\| \frac{\partial f_i}{\partial x_j}(x) - \frac{\partial f_i}{\partial x_j}(x_0) \right\| \leq \|Jf(x) - Jf(x_0)\|_{\text{op}}.$$

\square

6.3.2. The Chain Rule — The Chain Rule allows to express the differential of the composite of two functions in terms of the differentials of each of the functions. Its proof relies on the notion of operator norm that we introduced in the previous subsection.

Theorem 6.3.2.1 (The Chain Rule). — Let $U \subseteq \mathbb{R}^n$ and $V \subseteq \mathbb{R}^m$ be open subsets and let x_0 be a point in U . Let $f: U \rightarrow V$ and $g: V \rightarrow \mathbb{R}^k$ be functions. Then, if f is differentiable at x_0 and g is differentiable at $f(x_0)$, the composite function $g \circ f$ is differentiable at x_0 and

$$D(g \circ f)(x_0) = Dg(f(x_0)) \circ Df(x_0).$$

Before proving the theorem, we state some useful corollaries.

Corollary 6.3.2.2. — Let $f: U \rightarrow \mathbb{R}^m$ and $g: V \rightarrow \mathbb{R}^k$ be as in the hypothesis of Theorem 6.3.2.1, then the following equality of matrices holds

$$J(g \circ f)(x_0) = Jg(f(x_0)) \cdot Jf(x_0)$$

(the right hand side is the product of the two matrices).

In other words, the following equalities holds

$$\frac{\partial(g \circ f)_i}{\partial x_j}(x_0) = \sum_{l=1}^m \frac{\partial g_i}{\partial x_l}(f(x_0)) \cdot \frac{\partial f_l}{\partial x_j}(x_0)$$

for all $1 \leq i \leq m$ and $1 \leq j \leq n$.

From the last equation, we deduce the following corollary.

Corollary 6.3.2.3. — Under the hypothesis of Theorem 6.3.2.1, if f has continuous partial derivatives at x_0 and g has continuous partial derivatives at $f(x_0)$, the composite function $g \circ f$ has continuous partial derivatives at x_0 .

The proof of the Chain Rule is quite involved (and it can safely be skipped on a first reading).

Proof of Theorem 6.3.2.1. Let $L = Jf(x_0)$ and $M = Jg(f(x_0))$. We want to show that, for every $\epsilon > 0$, there exists $\delta > 0$ such that, for every h with $0 < \|h\| < \delta$, the following inequality holds

$$\|g(f(x_0 + h)) - g(f(x_0)) - M \cdot L \cdot h\| \leq \epsilon \|h\|. \quad (6.3)$$

By hypothesis, we know that

$$\lim_{h \rightarrow 0} \frac{\|f(x_0 + h) - f(x_0) - L \cdot h\|}{\|h\|} = 0, \quad (6.4)$$

$$\lim_{k \rightarrow 0} \frac{\|g(f(x_0) + k) - g(f(x_0)) - M \cdot k\|}{\|k\|} = 0. \quad (6.5)$$

Let us fix some notation and denote by q and r the functions:

$$q(h) = f(x_0 + h) - f(x_0) - L \cdot h$$

$$r(k) = g(f(x_0) + k) - g(f(x_0)) - M \cdot k.$$

Using the triangular inequality, the Left Hand Side of (6.3) is less than or equal to the sum:

$$\|r(f(x_0 + h) - f(x_0))\| + \|M \cdot q(h)\|.$$

Now, given $\epsilon' > 0$, by (6.4) there exists a $\delta' > 0$ such that, taking

$$k := f(x_0 + h) - f(x_0)$$

6.3 Properties of differentials and of differentiable functions

such that $\|k\| < \delta'$, the following inequality holds

$$\|r(f(x_0 + h) - f(x_0))\| \leq \epsilon' \|k\|.$$

By applying (6.5) with $\epsilon = \delta'$, we find $\delta'' > 0$ such that the inequality $\|k\| \leq \delta'$ holds indeed for all $\|h\| < \delta''$. By applying again the triangle inequality and the definition of operator norm, we obtain

$$\begin{aligned} \|r(f(x_0 + h) - f(x_0))\| &\leq \epsilon' \|f(x_0 + h) - f(x_0)\| \\ &\leq \epsilon' \|f(x_0 + h) - f(x_0) - L \cdot h\| + \|L \cdot h\| \\ &\leq \epsilon' \|h\| (1 + \|L\|_{\text{op}}) \end{aligned}$$

for all h such that $\|h\| < \delta''$.

Similarly, for a small enough $\delta''' > 0$ and every h with $0 < \|h\| \leq \delta'''$, the following inequality holds

$$\begin{aligned} \|M \cdot q(h)\| &\leq \|M\|_{\text{op}} \cdot \|f(x_0 + h) - f(x_0) - L \cdot h\| \\ &\leq \epsilon' \|h\| \cdot \|M\|_{\text{op}}. \end{aligned}$$

Putting everything together, we have that

$$\text{Left Hand Side of (6.3)} \leq \epsilon' \|h\| (1 + \|L\|_{\text{op}} + \|M\|_{\text{op}}).$$

For small enough values of $\|h\|$. Therefore, given $\epsilon > 0$, if we take ϵ' to be

$$\epsilon' = \frac{\epsilon}{1 + \|L\|_{\text{op}} + \|M\|_{\text{op}}}.$$

Taking δ to be the minimum of δ'' and δ''' , we deduce that

$$\|g(f(x_0 + h)) - g(f(x_0)) - M \cdot L \cdot h\| < \epsilon \|h\|$$

for all $0 < \|h\| < \delta$, which concludes our proof. \square

From the Chain Rule, we immediately deduce some other elementary properties of the differential.

Corollary 6.3.2.4. — Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function at $x_0 \in \mathbb{R}^n$ and assume that $f(x_0) \neq 0$. Then, the function

$$\frac{1}{f}: U \rightarrow \mathbb{R}$$

is differentiable at x_0 , where U is the open subset of \mathbb{R}^n where f does not vanish. Moreover, the following equality holds

$$D\left(\frac{1}{f}\right)(x_0) = -\frac{Df(x_0)}{f^2(x_0)}.$$

Proof. (Sketch). Apply the Chain Rule to $g \circ f$, where the function g is defined by $g(y) = \frac{1}{y}$. □

The proofs of the next properties are left as exercises.

Exercise 6.3.2.5. — Let $f, g: U \rightarrow \mathbb{R}^m$ be functions differentiable at $x_0 \in U$ and let $\lambda, \mu \in \mathbb{R}$, then the linear combination

$$\lambda f + \mu g$$

of f and g is differentiable at x_0 . Moreover, one has

$$D(\lambda f + \mu g)(x_0) = \lambda Df(x_0) + \mu Dg(x_0)$$

(Hint: compose with the function $(y, z) \mapsto \lambda y + \mu z$ and apply the Chain Rule.)

Exercise 6.3.2.6. — (Leibniz rule or product rule).

Let $f, g: U \rightarrow \mathbb{R}$ be functions differentiable at x_0 and let

$$\begin{aligned} f \cdot g: \mathbb{R}^n &\rightarrow \mathbb{R} \\ x &\mapsto f(x)g(x) \end{aligned}$$

be the function given by the point-wise product of f and g in \mathbb{R} . Then, $f \cdot g$ is differentiable at x_0 and

$$D(f \cdot g)(x_0) = g(x_0)Df(x_0) + f(x_0)Dg(x_0)$$

Hint: Apply the chain rule to the composite of

$$x \mapsto (f(x), g(x))$$

and

$$(y, z) \mapsto y \cdot z.$$

We conclude with a few examples of how the Chain Rule and the above properties can be used to calculate differentials.

Example 6.3.2.7. — Let $h: U \rightarrow \mathbb{R}^2$ be given by

$$h(x_1, x_2) = \left(e^{x_1 \sin(x_2)}, \log(x_1 x_2) \right)$$

for $U := \{(x_1, x_2) : x_1 \cdot x_2 > 0\} \subseteq \mathbb{R}^2$.

6.3 Properties of differentials and of differentiable functions

We compute the differential of h using the Chain Rule. Indeed, h can be written as the composition of g and f , where

$$\begin{aligned} f: U &\rightarrow \mathbb{R}^2 \\ (x_1, x_2) &\mapsto (x_1 \sin(x_2), x_1 x_2) \\ g: V &\rightarrow \mathbb{R}^2 \\ (y_1, y_2) &\mapsto (e^{y_1}, \log(y_2)) \end{aligned}$$

for $V = \{(y_1, y_2) : y_2 > 0\} \subseteq \mathbb{R}^2$.

Then we compute the Jacobian matrices of f and g

$$\begin{aligned} Jf(x_1, x_2) &= \begin{pmatrix} \sin x_2 & x_1 \cos x_2 \\ x_2 & x_1 \end{pmatrix} \\ Jg(y_1, y_2) &= \begin{pmatrix} e^{y_1} & 0 \\ 0 & \frac{1}{y_2} \end{pmatrix}. \end{aligned}$$

Then we can compute the Jacobian of h as

$$\begin{aligned} Jh(x_1, x_2) &= Jg(f(x_1, x_2)) \cdot Jf(x_1, x_2) \\ &= \begin{pmatrix} e^{x_1 \sin(x_2)} \sin(x_2) & e^{x_1 \sin(x_2)} x_1 \cos(x_2) \\ \frac{1}{x_1} & \frac{1}{x_2} \end{pmatrix} \end{aligned}$$

The linear function associated to this matrix gives the differential of the function h at each point (x_1, x_2) .

Example 6.3.2.8. — Let us consider the function:

$$\begin{aligned} f: \mathbb{R}^2 &\rightarrow \mathbb{R}^2 \\ (x_1, x_2) &\mapsto (x_1 x_2, x_1 + x_2^3) \end{aligned}$$

and compute the differential of f at the point $P = (1, -1)$ applied to the vector $h = (2, 3)$. We first compute the Jacobian matrix:

$$Jf(x_1, x_2) = \left(\frac{\partial f_i}{\partial x_j}(x_1, x_2) \right) = \begin{pmatrix} x_2 & x_1 \\ 1 & 3x_2^2 \end{pmatrix}$$

which, calculated at the point P , gives the matrix:

$$Jf(P) = \begin{pmatrix} -1 & 1 \\ 1 & 3 \end{pmatrix}.$$

Therefore, we obtain

$$f(P) \cdot h = \begin{pmatrix} 1 \\ 11 \end{pmatrix}.$$

Exercise 6.3.2.9. — Let us consider the function:

$$f: \mathbb{R}^2 \rightarrow \mathbb{R}$$

$$(x_1, x_2) \mapsto x_2 \int_1^{x_1^2} e^{-t^2} dt$$

and compute the differential of f at the point $P = (1, e)$ applied to the vector $(1, 2)$.

Notice that, given a function of the form

$$g: \mathbb{R} \rightarrow \mathbb{R}$$

$$x \mapsto \int_c^x h(t) dt,$$

the derivative of g is given by $\frac{dg}{dx} = h(x)$ by the fundamental theorem of calculus.

Therefore, we have that

$$\frac{\partial f}{\partial x_1}(x_1, x_2) = 2x_1 x_2 e^{-(x_1^2)^2},$$

$$\frac{\partial f}{\partial x_2}(x_1, x_2) = \int_1^{x_1^2} e^{-t^2} dt.$$

Hence, the Jacobian matrix of f at P is given by

$$Jf(P) = \left(2e \cdot e^{-1}, \int_1^1 e^{-t^2} dt \right) = (2, 0).$$

In particular, the value of the differential of f at P applied to h is

$$Jf(P) \cdot h = (2, 0) \cdot \begin{pmatrix} 1 \\ 2 \end{pmatrix} = 2.$$

Exercise 6.3.2.10. — Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined by

$$f(x_1, x_2) = \begin{cases} x_1 x_2 \sin\left(\frac{1}{x_1 x_2}\right) & x_1, x_2 \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Determine all points of \mathbb{R}^2

1. where f is continuous,
2. where f is differentiable,
3. where both partial derivatives of f exist and are continuous.

6.3 Properties of differentials and of differentiable functions

6.3.3. The Mean Value Theorem — In this subsection we recall the Mean Value Theorem for functions $f: \mathbb{R} \rightarrow \mathbb{R}$ and then generalise it to the case of functions of several variables. Contrarily to what we have seen so far, increasing the dimension of the source space does not make a big difference, and the biggest difficulties arise when trying instead to increase the dimension of the target.

The Mean Value Theorem is the result that is used to prove that a differentiable function with zero derivative on an interval is constant, and we will see the analogue of that result for multivariable functions in Corollary 6.3.3.8.

Theorem 6.3.3.1 (Year 1 Mean Value Theorem). — Let $f: [a, b] \rightarrow \mathbb{R}$ be a continuous function that is differentiable on (a, b) . Then, there exists a constant $c \in (a, b)$ such that

$$f(b) - f(a) = f'(c)(b - a).$$

We aim to generalise this result to higher dimensions.

Definition 6.3.3.2. — Let u and v be points in \mathbb{R}^n . We define the *closed segment* from u to v as the set of points:

$$[u, v] = \{(1 - t)u + tv : t \in [0, 1]\}.$$

Similarly, we define the *open segment* from u to v as:

$$(u, v) = \{(1 - t)u + tv : t \in (0, 1)\}.$$

The following is the generalisation of Theorem 6.3.3.1 to a function of multiple variables.

Corollary 6.3.3.3. — Let $U \subseteq \mathbb{R}^n$ be an open subset and let $f: U \rightarrow \mathbb{R}$ be a differentiable function. Let $u, v \in U$ be such that the segment $[u, v]$ is contained in U . Then, there exists a constant $c \in (u, v)$ such that

$$f(u) - f(v) = Jf(c) \cdot (u - v).$$

Proof. We define a function:

$$\begin{aligned} \phi: [0, 1] &\rightarrow \mathbb{R} \\ t &\mapsto f((1 - t)u + tv). \end{aligned}$$

Then, ϕ is differentiable on $[0, 1]$. Moreover, by Theorem 6.3.3.1 there exists $t_0 \in (0, 1)$ such that:

$$f(v) - f(u) = \phi(1) - \phi(0) = \phi'(t_0) = Jf((1 - t_0)u + t_0v) \cdot (v - u),$$

where the last equality follows from the Chain Rule. Taking c equal to $(1 - t_0)u + t_0v$ completes the proof. \square

The analogue of the above result does not hold when we replace the target space to \mathbb{R}^m with $m > 1$, as we illustrate in the next Example.

Example 6.3.3.4. — Consider the function

$$f: [0, 2\pi] \rightarrow \mathbb{R}^2$$

$$x \mapsto \begin{pmatrix} \cos(x) \\ \sin(x) \end{pmatrix}$$

Then, we have that

$$Jf(x) = \begin{pmatrix} -\sin(x) \\ \cos(x) \end{pmatrix}.$$

Since there is no value c such that $\sin(c) = \cos(c) = 0$, the Jacobian matrix is never zero. In particular

$$f(2\pi) - f(0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \neq Jf(c) \cdot 2\pi$$

for every $c \in [0, 2\pi]$.

The previous Example shows the most obvious generalisation of the Mean Value Theorem does not hold for functions $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ with $m > 1$. However, we have the following replacement:

Lemma 6.3.3.5. — Let $U \subseteq \mathbb{R}^n$ be an open subset and let $f: U \rightarrow \mathbb{R}^m$ be a differentiable function. Given $u, v \in U$ such that $[u, v]$ is contained in U , the following inequality holds

$$\|f(v) - f(u)\| \leq \sup_{c \in (u, v)} \|Jf(c)\|_{\text{op}} \cdot \|v - u\|.$$

Here the supremum $\sup_{c \in (u, v)} \|Jf(c)\|_{\text{op}}$ may equal $+\infty$. As a consequence of Theorem 6.3.1.2, we have seen that if f is additionally *continuously differentiable* on U , then $\|Jf\|$ is a continuous function on U , so the supremum $\sup_{c \in [u, v]} \|Jf(c)\|_{\text{op}}$ over the *closed* segment $[u, v]$ is in fact a maximum, because $[u, v]$ is compact (it is closed and bounded in \mathbb{R}^n).

6.3 Properties of differentials and of differentiable functions

Proof. By possibly translating the function f in the source, we may assume that $u = 0$ is the origin in \mathbb{R}^n . By translating the function in the target, we may assume that $f(u)$ is the origin in \mathbb{R}^m .

Defining $\phi(t) := \|f(tv)\|$, we have that ϕ is continuous on $[0, 1]$ and differentiable on $(0, 1)$ (note that the norm is continuous but *not* differentiable at the origin). By writing out

$$\phi(t) = \sqrt{f_1(tv)^2 + \dots + f_m(tv)^2}$$

and then taking the derivative and applying to it the Chain Rule and the Cauchy-Schwarz inequality, we deduce

$$|\phi'(t)| = \frac{|\sum_{i=1}^m f_i(tv) \cdot (Jf_i(tv) \cdot v)|}{\|f(tv)\|} \leq \frac{\|Jf(tv) \cdot v\| \cdot \|f(tv)\|}{\|f(tv)\|} = \|Jf(tv) \cdot v\|$$

for all $t \in (0, 1)$.

By combining the former inequality with the Year 1 Mean Value Theorem and the definition of operator norm, we deduce

$$\|f(v) - f(u)\| = \|f(v)\| = |\phi(1) - \phi(0)| \leq \sup_{c \in (u,v)} \|Jf(c)\|_{\text{op}} \cdot \|v\|,$$

which concludes our proof. \square

We are now in a position to prove an important consequence of the Mean Value Theorem. This states that if the open set U “consists of only one piece”, if the differential of f vanishes then f is constant on U . Without the extra hypothesis that U consists only of one piece, the function f might not be constant for it equals one constant on one piece of U and a different constant on another piece of U . Think of $U = (0, 1) \cup (2, 3)$ and f defined by

$$f(x) = \begin{cases} 0 & \text{for } x \in (0, 1) \\ 1 & \text{for } x \in (2, 3). \end{cases}$$

This function is not constant, but its derivative is constantly equal to zero.

We start by defining this notion of “consisting of only one piece”.

Definition 6.3.3.6. — Let $U \subseteq \mathbb{R}^n$ be a subset. We say that U is *path-connected* if for every two points P and Q in U there exists a *finite sequence*

$$P = x_0, x_1, \dots, x_k = Q$$

of points of U such that the segment $[x_i, x_{i+1}]$ is contained in U for $i = 0, \dots, k-1$.

Remark 6.3.3.7. — There are more general notions of path-connected and of connected that are valid for arbitrary metric (or topological) spaces (as opposed to just \mathbb{R}^n).

If $U \subseteq \mathbb{R}^n$ is path-connected according to our definition above, then one can prove that the only subsets that are simultaneously open and closed in U (with the Euclidean distance d_2) are the empty set and U itself – the latter is the usual notion of a “connected” metric or topological space.

Moreover, one could also show the converse implication if $U \subset \mathbb{R}^n$ is *open*, namely if such U is connected, then it is also path-connected.

Corollary 6.3.3.8. — Let $U \subseteq \mathbb{R}^n$ be open and path-connected. Let $f: U \rightarrow \mathbb{R}^m$ be a differentiable function and assume that $Df(x) = 0$ for every point $x \in U$. Then, f is a constant function.

Proof. It suffices to prove that for every two points P and Q in U the norm $\|f(P) - f(Q)\|$ is 0. Since U is path-connected we can find points

$$P = x_0, \dots, x_k = Q$$

such that $[x_i, x_{i+1}] \subset U$ for all $i = 0, \dots, k-1$. Then we have

$$\begin{aligned} \|f(P) - f(Q)\| &\leq \|f(x_0) - f(x_1)\| + \dots + \|f(x_{k-1}) - f(x_k)\| \\ &\leq 0 \cdot \|x_0 - x_1\| + \dots + 0 \cdot \|x_{k-1} - x_k\| \end{aligned}$$

where the first inequality is the triangle inequality, and the second inequality follows from Lemma 6.3.3.5 and the fact that the Jacobian of f is 0 everywhere. □

7.

The inverse function theorem and the implicit function theorem

Important notions to learn from this section:

1. The inverse function theorem (INFT).
2. What it means that an equation $f(x, y) = 0$ implicitly defines y as a function of x locally at a point (x_0, y_0) .
3. The implicit function theorem (IMFT).
4. (EXTRA) How the IMFT can be used to find maxima/minima of functions subject to some constraints (also known as the method of Lagrange multipliers).

This chapter will be devoted to the inverse and the implicit function theorems. In line with the spirit of the previous chapter, these two results are *linearisation* results. They say that, under suitable hypotheses, certain results that we are familiar with and that are valid for *linear* functions remain valid *locally* for nonlinear functions that can be approximated by means of linear ones.

Both results are quite delicate to prove, although it is not so hard to see that one Theorem holds if and only if the other does. The proof of the existence of the inverse (or of the implicit) function is obtained by applying the Contraction Mapping Theorem to a certain complete space of continuous functions. The proof uses the results that we have discussed in the previous Chapter: the operator norm, the Chain Rule, and the Mean Value Theorem.

An important application of the implicit function theorem is the Lagrange's multipliers method to find the (local) extremal points of a scalar function that is subject to some constraint. This will be discussed as extra material in the final section of this Chapter.

7.1. The inverse function theorem

Let $f: U \rightarrow \mathbb{R}^n$ be a continuous function, where U is an open subset of \mathbb{R}^n . The idea is that, if f can be approximated around a point $x_0 \in U$ by an *invertible* linear function, then f itself is locally invertible near x_0 . Before stating the main theorem of the section, we recall the following definition from the previous Chapter.

Definition 7.1.1. — Let $U \subseteq \mathbb{R}^n$ be an open subset, $f: U \rightarrow \mathbb{R}^m$ be a function and $x_0 \in U$. We say that f is *continuously differentiable* at x_0 if f has all first order partial derivatives on U and they are continuous at x_0 .

In the previous chapter we proved that if f is continuously differentiable, then it is differentiable.

We now state one of the two main results of this chapter.

Theorem 7.1.2 (INverse Function Theorem=INFT). — *Let $U \subseteq \mathbb{R}^n$ be an open subset and let $f: U \rightarrow \mathbb{R}^n$ be continuously differentiable. Let $x_0 \in U$ and assume $Df(x_0): \mathbb{R}^n \rightarrow \mathbb{R}^n$ is invertible. Then, there exist open subsets $V \subset U$ with $x_0 \in V$ and $W \subset \mathbb{R}^n$ with $f(x_0) \in W$ such that the restriction $f|_V: V \rightarrow W$ is invertible and its inverse is a continuously differentiable function on W .*

Observe that it would not make a difference to modify the statement of the theorem by replacing the target space \mathbb{R}^n with \mathbb{R}^m for arbitrary m (so m is possibly different from n). Indeed, there are no invertible linear functions $\mathbb{R}^n \rightarrow \mathbb{R}^m$ unless n equals m .

The proof of the INFT will be given later.

Remark 7.1.3. — Under the hypothesis of the INFT, call $g: W \rightarrow V$ the inverse of the restriction $f|_V: V \rightarrow W = f(V)$.

We claim that the differential of g at $f(x_0)$ equals the inverse of the differential of f at x_0 . (This statement is sometimes also included as a part of the INFT). This follows immediately from the Chain Rule.

7.1 The inverse function theorem

Indeed, by differentiating both sides of the equality

$$g \circ f|_V = \text{Id}_V: V \rightarrow V$$

we obtain

$$D(g \circ f)_{x_0} = D(\text{Id})_{x_0} = \text{Id}: \mathbb{R}^n \rightarrow \mathbb{R}^n$$

and applying the Chain Rule we deduce:

$$D(g \circ f)_{x_0} = Dg_{f(x_0)} \circ Df_{x_0} = \text{Id}: \mathbb{R}^n \rightarrow \mathbb{R}^n,$$

which implies that $Dg_{f(x_0)} = Df_{x_0}^{-1}$.

In light of Theorem 7.1.2, it is convenient to introduce the following definition.

Definition 7.1.4. — Let $U \subseteq \mathbb{R}^n$ be an open subset, $f: U \rightarrow \mathbb{R}^m$ be a function and $x_0 \in U$. A *local inverse* of f near x_0 is a triple (V, W, g) where $V \subseteq U$ and $W \subseteq \mathbb{R}^m$ are open subsets such that $x_0 \in V$ and $f(x_0) \in W$ and $g: W \rightarrow V$ is a function such that $g \circ f|_V = \text{id}_V$ and $f|_V \circ g = \text{id}_W$.

The chain Rule can be used to deduce what exactly goes wrong when the differential fails to be invertible.

Corollary 7.1.5. — Let $U \subseteq \mathbb{R}^n$ be an open set and $f: U \rightarrow \mathbb{R}^m$ be a function and $x_0 \in U$. If f is differentiable at x_0 and Df_{x_0} is not invertible, then if a local inverse of f exists, that local inverse is not differentiable at $f(x_0)$.

Note that this result follows directly from the Chain Rule, similarly to Remark 7.1.3, we are not using the INFT in the proof!

Proof. Assuming a local inverse g of f exists, we have

$$g \circ f|_V = \text{Id}_V, \quad f|_V \circ g = \text{Id}_W$$

for open sets V and W containing x_0 and $f(x_0)$ respectively. Taking differentials in the previous equality, we deduce:

$$D(g \circ f)_{x_0} = D(\text{Id})_{x_0} = \text{Id}, \quad D(f \circ g)_{f(x_0)} = D(\text{Id})_{f(x_0)} = \text{Id}$$

On the other hand, if the local inverse g was also differentiable, then we could apply the Chain Rule and deduce:

$$D(g \circ f)_{x_0} = Dg_{f(x_0)} \circ Df_{x_0}, \quad D(f \circ g)_{f(x_0)} = Df_{x_0} \circ Dg_{f(x_0)}.$$

In particular $Dg_{f(x_0)}$ and Df_{x_0} are inverses to each others, therefore Df_{x_0} is invertible, contradicting the hypothesis. \square

The inverse function theorem and the implicit function theorem

Remark 7.1.6. — Another way to express the result of Corollary 7.1.5 is to say that if $f: V \rightarrow W$ is a differentiable bijection with differentiable inverse between open subsets $V \subseteq \mathbb{R}^n$ and $W \subseteq \mathbb{R}^m$, then the differential must be invertible. In particular, because there is no linear bijection $\mathbb{R}^n \rightarrow \mathbb{R}^m$ if $n \neq m$, this result implies that differentiable bijections with differentiable inverse between open subsets of finite dimensional Euclidean vector spaces only exist when the dimension of the two vector spaces is the same!

Recall that a continuous bijection with continuous inverse is called a *homeomorphism*. One basic result of algebraic topology says that the same is true of homeomorphisms, *i.e.* a continuous bijection $V \rightarrow W$ with continuous inverse may only exist when $n = m$. This result is known under the name of “Theorem of invariance of the domain”.

We give now two examples of what may happen when the Jacobian matrix fails to be invertible.

Example 7.1.7. — Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be defined as $x \mapsto x^2$ and let $x_0 = 1$. Then, the local inverse of f at x_0 is the function $y \mapsto \sqrt{y}$. On the other hand, of $x_0 = -1$, a local inverse to f is given by the function $y \mapsto -\sqrt{y}$.

Finally, at the point $x_0 = 0$ the function f has no local inverse, because the restriction of f to every open interval containing 0 fails to be injective.

Example 7.1.8. — Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be defined as $x \mapsto x^3$. Then, f has a global inverse given by the function $g: \mathbb{R} \rightarrow \mathbb{R}$ mapping y to $\sqrt[3]{y}$. One can immediately check that this function fails to be differentiable at $f(0) = 0$. (And indeed, we have $f'(0) = 0$).

Let's now work an example where we combine the results obtained in Theorem 7.1.2 and in Corollary 7.1.5 to solve a standard exercise.

Example 7.1.9. — Consider the function $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined by

$$f(x_1, x_2) = (x_1^2 + x_2^2, 3x_2).$$

Does f admit a local differentiable inverse at $P = (0, 0)$? Does it admit a local differentiable inverse at $Q = (1, 1)$?

The Jacobian matrix of f is

$$\begin{pmatrix} 2x_1 & 2x_2 \\ 0 & 3 \end{pmatrix}.$$

7.1 The inverse function theorem

This matrix is not invertible when calculated at $P = (0,0)$, so by Corollary 7.1.5 (a straightforward consequence of the Chain Rule), the function f does not admit a local differentiable inverse at P . This matrix is invertible at $Q = (1,1)$, so by Theorem 7.1.2 the function f admits a local differentiable inverse at Q .

Once this is established, one may want to address the more difficult question of whether f is or it is not invertible near P . This is a question that the theory we have developed does not help to answer, and we have to resort to brute force. In this particular example, this is still doable and easy. For given $(y_1, y_2) \in \mathbb{R}^2$ we want to solve

$$\begin{cases} y_2 = 3x_2 \\ y_1 = x_1^2 + x_2^2 \end{cases}$$

in the unknown $(x_1, x_2) \in \mathbb{R}^2$. From the first equation we deduce $x_2 = y_2/3$. The second equation becomes

$$x_1^2 = y_1 - \frac{y_2^2}{9}.$$

This situation is similar to that of Example 7.1.7. For any given $k \neq 0$ there is no open set containing $(0,0)$ such that the equation $x_1^2 = k$ has exactly 1 solution. This proves that the function f is not locally invertible near P .

Example 7.1.10. — (A function that is differentiable but not *continuously differentiable*, and that has invertible derivative at a given point, but that is not locally invertible at that point).

Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be defined by

$$f(x) = \begin{cases} x^2 \sin(1/x) + x/2 & \text{for all } x \neq 0 \\ 0 & \text{when } x = 0. \end{cases}$$

This function is a perturbation (by adding the linear function $x \mapsto x/2$) of the prototypical example (seen in the previous chapter) of a function that is differentiable at the origin, but whose derivative is not continuous at the origin.

This function is differentiable at 0, and its derivative equals $1/2$, so the differential at the origin is $h \mapsto h/2$ and therefore it is an invertible function.

To show that the function f is not invertible on any open set containing zero, we compute its derivative for all $x \neq 0$:

$$f'(x) = 2x \sin(1/x) - \cos(1/x) + 1/2.$$

The inverse function theorem and the implicit function theorem

The derivative is continuous at all points, except at $x = 0$. Consider the sequences (x_n) and (y_n) defined by

$$x_n = \frac{1}{(2n+1)\pi} \quad y_n = \frac{1}{2n\pi}.$$

Then we have both x_n and $y_n \rightarrow 0$ for $n \rightarrow \infty$, and the derivatives

$$f'(x_n) = 3/2, \quad f'(y_n) = -1/2.$$

Because f' is continuous for $x \neq 0$, for all n the function f' is positive on small open intervals containing x_n and it is negative on small open intervals containing y_n . On the former intervals the function f is increasing, and on the latter intervals the function f is decreasing, thus not giving f any chance of being invertible on any open set that contains the origin.

From the theoretical standpoint, the inverse function theorem helps us establish that a given function is a homeomorphism, as we will now explain.

Definition 7.1.11. — Let $f: (X, d_X) \rightarrow (Y, d_Y)$ be a function of metric spaces. We say that f is a *local homeomorphism* if, for every $x_0 \in X$ there exist $U \subseteq X$ and $V \subseteq Y$ open subsets such that $x_0 \in U$ and the restriction $f|_U: U \rightarrow V$ is a homeomorphism.

Exercise 7.1.12. — Let $f: (X, d_X) \rightarrow (Y, d_Y)$ be a function. Show that f is a homeomorphism if and only if f is a local homeomorphism and f is bijective.

We conclude this section with two more corollaries of Theorem 7.1.2.

Corollary 7.1.13. — Let $f: U \rightarrow \mathbb{R}^n$ be continuously differentiable at every point of $U \subseteq \mathbb{R}^n$ open, and assume that Df_{x_0} is invertible for every $x_0 \in U$. Then f is a local homeomorphism. (And in particular $f(U)$ is open in \mathbb{R}^n).

Proof. The statement follows immediately from the Inverse Function Theorem 7.1.2. □

The following result, which is an immediate consequence of the previous result, gives an easy criterion to establish that certain functions are homeomorphisms.

7.1 The inverse function theorem

Corollary 7.1.14. — Let $f: U \rightarrow \mathbb{R}^n$ be injective and continuously differentiable at every point of $U \subseteq \mathbb{R}^n$ open, and assume that $Df(x)$ is invertible for every $x \in U$. Then $f(U)$ is open in \mathbb{R}^n and $f: U \rightarrow f(U)$ is a homeomorphism.

Remark 7.1.15. — For $U, V \subseteq \mathbb{R}^n$ one could define $f: U \rightarrow V$ to be a *local diffeomorphism* if: for every $x_0 \in U$ there exist open subsets $U' \subseteq U$ and $V' \subseteq V$ such that f restricts to a function $f|_{U'}: U' \rightarrow V'$ which is bijective, continuously differentiable and such that the inverse of $f|_{U'}$ is continuously differentiable. A diffeomorphism is a bijective local diffeomorphism.

Clearly, every diffeomorphism is in particular a homeomorphism, and every local diffeomorphism is a local homeomorphism.

The above corollary could be made stronger by saying that, under the same hypotheses, we can conclude that f is a local diffeomorphism.

We conclude this section with some exercises on the INFT.

Exercise 7.1.16. — In each of the following cases you are given a function $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and a point $P \in \mathbb{R}^n$. Decide if f admits a local differentiable inverse at P .

1. $f(x_1, x_2) = (x_1^2 - 2x_2, 2x_1^3)$, $P = (0, 2)$.
2. f as above and $P = (1, 1)$.
3. $f(x_1, x_2) = (x_1x_2, x_1^2e^{x_2}, x_2e^{x_1})$, $P = (0, 2)$.

When the answer to the previous question is negative, can you decide whether f admits a non-differentiable local inverse near the point P ?

Exercise 7.1.17. — Compute the differential of the local inverse of $f: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ at $f(P)$, where $f(x_1, x_2) = (x_1^2 - 2x_2, 2x_1^3)$ and $P = (1, 1)$.

Exercise 7.1.18. — Can you think of an example of a function $f: \mathbb{R} \rightarrow \mathbb{R}$ that is not even continuous at some point x_0 , but that is invertible? (Hint: it's easy!)

Exercise 7.1.19. — Prove the INFT when the dimension n of the source and target equals 1. Hint: prove that the function is either increasing or decreasing on some open interval containing x_0 . Then use the intermediate value theorem to prove that the image of an open interval is an open interval.

7.2. The implicit function theorem

The main idea of the implicit function theorem is to give a local, parametric description of a subset of \mathbb{R}^{n+m} that is described by m "locally independent" equations. The crucial point is to understand what we mean by "local independence" for a system of equations.

We first review the analogue theory in the linear case. There are two equivalent ways of defining a linear subspace L of \mathbb{R}^n : a parametric definition and an equational definition. We illustrate these two different ways in the following examples.

Example 7.2.1. — Let

$$L = \{(x, y, z) \in \mathbb{R}^3 : x + y + z = 0\} \subset \mathbb{R}^3,$$

then this definition of L is given in the form of "points satisfying a certain equation". However, we could equivalently define L by:

$$L = \{(x, y, -x - y) : (x, y) \in \mathbb{R}^2\} \subset \mathbb{R}^3.$$

We call such description *parametric*, with parameters $(x, y) \in \mathbb{R}^2$.

Example 7.2.2. — Similarly, the linear space defined in "equational form" by:

$$L = \{(x, y, z) : x + y + z = 0, \quad x + y + 2z = 0\} \subset \mathbb{R}^3$$

can be written in parametric form as

$$L = \{(x, -x, 0) : x \in \mathbb{R}\} \subset \mathbb{R}^3.$$

In this case the parameter is $x \in \mathbb{R}$.

Let us give a precise definition of what is meant by "parametric description" in the case of a linear subspace.

Definition 7.2.3. — Let $f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a linear function and consider the linear space:

$$L = \{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m : f(x, y) = 0\} \subseteq \mathbb{R}^{n+m}.$$

We say that the equation $f(x, y) = 0$ *implicitly defines y as a function of x* if there exists a function $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that:

$$L = \{(x, \phi(x)) : x \in \mathbb{R}^n\} \subset \mathbb{R}^{n+m}.$$

7.2 The implicit function theorem

We say that ϕ is the function that implicitly defines y in terms of x for the equation $f(x, y) = 0$.

(We say that the first presentation of L is in *equational form*, and that the second one is in *parametric form*).

Note that the implicit function ϕ of the above definition is always a linear function.

For a system of *linear* equations we know how to precisely define the notion of "independence" from linear algebra. The following linear algebra result how to give a parametric description of a linear subspace, in the case when the linear equations are independent.

Proposition 7.2.4. — *Let $f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a linear function, let A be the corresponding matrix $m \times (n + m)$ and let us write A as*

$$A = (A_x, A_y)$$

where A_x is a $m \times n$ matrix and A_y is a square $m \times m$ matrix.

If A_y is invertible, then the equation $f(x, y) = 0$ implicitly defines y as a function of x . (And moreover the implicit function is linear).

(If f is linear, the notion of "independence" for the system of equations $f(x, y) = 0$, mentioned in the opening of this chapter, is encoded in the hypothesis that the $m \times m$ matrix A_y is *invertible*).

Proof. We have that $f(x, y) = 0 \in \mathbb{R}^m$ if and only if $A_x \cdot x + A_y \cdot y = 0$. Since A_y is invertible this can be written equivalently as:

$$y = -A_y^{-1} \cdot A_x \cdot x. \tag{7.1}$$

Defining $\phi(x) := -A_y^{-1} \cdot A_x \cdot x$ we have that ϕ is linear (the composite of linear functions), and we have the equality

$$\{(x, y) \in \mathbb{R}^n \times \mathbb{R}^m : f(x, y) = 0\} = \{(x, \phi(x)) : x \in \mathbb{R}^n\}.$$

The latter is saying that ϕ is the implicit function defining y as a function of x as defined in Definition 7.2.3. \square

The main scope of this section is to *locally* generalise the result of Proposition 7.2.4 to the case of non-linear functions that are continuously differentiable. We start by precisely defining the notion of a "local parametrisation", in analogy with Definition 7.2.3.

The inverse function theorem and the implicit function theorem

Definition 7.2.5. — Let $U \subseteq \mathbb{R}^n \times \mathbb{R}^m$ be open and let $(x_0, y_0) \in U$. Let $f: U \rightarrow \mathbb{R}^m$ be a function such that $f(x_0, y_0) = 0$. Then, we say that the equation $f(x, y) = 0$ *implicitly defines* y as a function of x , locally at (x_0, y_0) , if there exist open sets $V \subset \mathbb{R}^n$ and $W \subset \mathbb{R}^m$ such that $x_0 \in V$ and $y_0 \in W$, with $V \times W \subseteq U$ and a function $\phi: V \rightarrow W$, such that:

$$\{(x, y) \in V \times W : f(x, y) = 0\} = \{(x, \phi(x)) : x \in V\}.$$

If this is the case, $\phi: V \rightarrow W$ is called the *local implicit function* (expressing y as a function of x) at (x_0, y_0) .

The following result is the second main theorem of this chapter, and it gives sufficient conditions for a function f so that the corresponding equation $f(x, y) = 0$ implicitly defines y as a function x locally at a given point.

Theorem 7.2.6 (Implicit Function Theorem = IMFT). — Let $U \subseteq \mathbb{R}^n \times \mathbb{R}^m$ be an open subset and let $(x_0, y_0) \in U$. Let $f: U \rightarrow \mathbb{R}^m$ be a function such that $f(x_0, y_0) = 0$. Assume that f is continuously differentiable on U and that, writing the Jacobian matrix as

$$Jf(x_0, y_0) = (J_x f(x_0, y_0), J_y f(x_0, y_0)),$$

the $m \times m$ matrix $J_y f(x_0, y_0)$ is invertible.

Then, the equation $f(x, y) = 0$ implicitly defines y as a function of x locally at (x_0, y_0) (as in Definition 7.2.5). Moreover, the local implicit function is continuously differentiable.

(The notion of “local independence” for the system of equations $f(x, y) = 0$, mentioned in the opening of this chapter, is encoded in the hypothesis that the $m \times m$ matrix $J_y f(x_0, y_0)$ is invertible).

The following is the standard textbook example to illustrate the implicit function theorem.

Example 7.2.7. — Let $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be the function defined by

$$f(x, y) = x^2 + y^2 - 1.$$

Then, the equation $f(x, y) = 0$ describes a unit circle in the plane, and it implicitly defines y as a function of x locally at $P_+ := (0, 1)$ and $P_- := (0, -1)$. On the one hand, we have that $\frac{\partial f}{\partial y}(x_0, y_0) = 2y_0$ which is non-zero at the points P_{\pm} . Therefore, we can apply the implicit function theorem.

7.2 The implicit function theorem

On the other hand, in this explicit simple situation we can give a formula for the implicit function, which is given by $\phi_{\pm}(x) = \pm\sqrt{1-x^2}$ at P_{\pm} .

Moreover, we claim that the equation $f(x, y) = 0$ does not implicitly define y as a function of x locally at the point $Q = (1, 0)$ (and at $Q' = (-1, 0)$). Indeed, for all $x > 0$ we have that the equation $x^2 + y^2 = 1$ has 2 solutions in y if $x < 1$ and it has 0 solutions in y for $x > 1$ (in order to define a function of y in x , it should have precisely 1 solution in y).

Here is another way to argue that $f(x, y) = 0$ does not implicitly define y as a *differentiable* function of x near Q . Suppose that one such function existed, and call it Φ . Then we have $f(x, \Phi(x)) = 0$. Taking the derivative with respect to x with the Chain Rule, we obtain

$$\frac{\partial f}{\partial x}(Q) + \frac{\partial f}{\partial y}(Q) \cdot \Phi'(x) = 0. \quad (7.2)$$

We can explicitly calculate $\frac{\partial f}{\partial x}(Q) = 2 \neq 0$, and also $\frac{\partial f}{\partial y}(Q) = 0$, and these two together are in contradiction with (7.2). This kind of reasoning will prove itself useful soon (in 7.2.9 and 7.2.10).

A first point that we would like to address is that the implicit function theorem is only a method to deduce the existence of an implicit function: it is possible that such function exists even when the hypotheses of the theorem are not satisfied.

Example 7.2.8. — The hypotheses of the implicit function theorem are not strictly necessary to deduce that the equation $f(x, y) = 0$ implicitly defines y as a function of x .

If $g: \mathbb{R} \rightarrow \mathbb{R}$ is any function (possibly discontinuous, etc.), then $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x, y) = y - g(x)$ will satisfy the fact that the equation $f(x, y) = 0$ defines y as a function of x : the function g itself!

Next, we are going to address the following two questions: (a) how can we compute the differential of the implicit function? and (b) assuming that the Jacobian $J_y f(x_0, y_0)$ is not invertible (but with the other hypothesis of the IMFT holding), what can we say about the non existence of a differentiable implicit function for the equation $f(x, y) = 0$ locally at (x_0, y_0) ? We start by addressing the first question.

Remark 7.2.9. — Let $U \subset \mathbb{R}^n \times \mathbb{R}^m$ be an open subset containing (x_0, y_0) and assume that $f: U \rightarrow \mathbb{R}^m$ is differentiable, that $f(x_0, y_0) = 0$, and that

The inverse function theorem and the implicit function theorem

the equation $f(x, y) = 0$ implicitly defines y as a function of x locally at (x_0, y_0) . Assume furthermore that the implicit function ϕ is differentiable at x_0 . Then, we claim that the Jacobian matrix of ϕ at x_0 equals:

$$J\phi(x_0) = -J_y f(x_0, y_0)^{-1} \cdot J_x f(x_0, y_0) \quad (7.3)$$

(compare this with the linear case, Equation (7.1)).

Indeed, for every point $x \in V$, we have $f(x, \phi(x)) = 0$ and we can apply the Chain Rule to the composite function $G(x) = f(F(x))$ with $F(x) = (x, \phi(x))$, obtaining:

$$\begin{aligned} 0 = JG(x_0) &= (J_x f(x_0, y_0), J_y f(x_0, y_0)) \cdot \begin{pmatrix} \text{Id} \\ J\phi(x_0) \end{pmatrix} \\ &= J_x f(x_0, y_0) + J_y f(x_0, y_0) \cdot J\phi(x_0). \end{aligned}$$

From this equation, we deduce (7.3).

These considerations allow us to derive a result that says, in some cases, what happens when the main hypothesis of the implicit function theorem (the invertibility of the Jacobian matrix $J_y f(x_0, y_0)$) fails. Recall from linear algebra that, given a matrix A together with its associated linear operator $L_A: \mathbb{R}^n \rightarrow \mathbb{R}^m$, we define the *rank* of A as the non-negative integer:

$$\text{rank}(A) = \dim(\text{Im}(L_A)).$$

Corollary 7.2.10. — *Let $U \subseteq \mathbb{R}^n \times \mathbb{R}^m$ be open and $(x_0, y_0) \in U$. Assume that $f: U \rightarrow \mathbb{R}^m$ is a function differentiable at (x_0, y_0) , and that the equation $f(x, y) = 0$ defines y as a differentiable function of x , locally at (x_0, y_0) . Then, if*

$$\text{rank}(J_x f(x_0, y_0)) = m,$$

the matrix $J_y f(x_0, y_0)$ is invertible.

Comment on this last result: the matrix $J_x f(x_0, y_0)$ is an $m \times n$ matrix, and as such its rank is smaller than or equal to the minimum of m and n . The previous result can be interpreted as saying that if the rank of $J_x f(x_0, y_0)$ equals m , and if the matrix $J_y f(x_0, y_0)$ is not invertible, then the equation $f(x, y) = 0$ does not define y as a differentiable function of x locally near (x_0, y_0) .

7.2 The implicit function theorem

Example 7.2.11. — (Of a function such that $J_y f(x_0, y_0)$ is not invertible, yet the equation $f(x, y) = 0$ defines y as a differentiable function of x).

Consider $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x, y) = y^2$. Then $f(0, 0) = 0$ and the partial derivative of f with respect to y is zero at $(0, 0)$. However, the equation $f(x, y) = 0$ defines y as a function of x : the constant zero function.

(Note that this does not contradict Corollary 7.2.10, for the derivative of f with respect to x vanishes at $(0, 0)$, so the rank of $J_x f(0, 0)$ equals $0 \neq 1$).

We conclude this section with a list of exercises on the IMFT.

Exercise 7.2.12. — Let $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$f(x, y) = x^3 y^2 + xy + 3x - 5.$$

Find the slope of the tangent line of the curve $\{(x, y) : f(x, y) = 0\}$ at the point $(1, 1)$. (Hint: parametrise the curve using the IMFT).

Exercise 7.2.13. — Let $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be defined as $f(x, y) = (y - 3)^2$. Does the equation $f(x, y) = 0$ define y as a function of x locally at the point $(0, 3)$?

Exercise 7.2.14. — Let $g: \mathbb{R} \rightarrow \mathbb{R}$ be a function. Define a function $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ by $f(x, y) = y - 2g(x) + 1$.

1. Does the equation $f(x, y) = 0$ necessarily satisfy the hypotheses of Theorem 7.2.6?
2. Does the equation $f(x, y) = 0$ implicitly define y as a function of x ?

Exercise 7.2.15. — Let $f(x, y, z) = x^4 + 4y^2 + 2xz^2 - 3ye^z - 1$. Does the equation $f(x, y, z) = 0$ implicitly define z as a function of (x, y) locally at $P = (0, 1, 0)$ or at $Q = (1, 0, 0)$? If so, is the implicit function differentiable?

Exercise 7.2.16. — Let $f = (f_1, f_2): \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be defined by

$$f_1(x, y, z) = x^2 - y^2 - z^2 - 2y, \quad f_2(x, y, z) = (x + y + z)^2.$$

Does the equation $f(x, y, z) = 0$ implicitly define the variables y, z as functions of x locally at the points $P_1 = (-1, 0, 1)$, $P_2 = (0, 0, 0)$ and $P_3 = (1, -2, 1)$? Hint: Applying directly Theorem 7.2.6 won't help!

7.3. Proof of the IMFT and of the INFT

This section is entirely devoted to proving the inverse function theorem and the implicit function theorem.

We start by showing that the INFT and the IMFT are equivalent, so it suffices proving one of them. Then we prove the IMFT.

7.3.1. IMFT and INFT are equivalent — Here we show how to prove the INFT assuming the IMFT and, conversely, how to prove the IMFT assuming the INFT. (To complete our proof, only the former is needed, so feel free to skip the second point below).

- We begin by $\text{IMFT} \implies \text{INFT}$. Let $F: U \rightarrow \mathbb{R}^n$ be a continuously differentiable function, with $U \subseteq \mathbb{R}^n$ open, and assume that $JF(x_0)$ is invertible. Define the function $f: U \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ by

$$f(x, y) = F(x) - y.$$

Note that the equation $f(x, y) = 0$ locally defines x as a function of y at the point $(x_0, F(x_0))$ if and only if F is locally invertible at the point x_0 .

Since F is continuously differentiable on U , the function f is also continuously differentiable on $U \times \mathbb{R}^n$. Moreover, the matrix

$$J_x f(x_0, F(x_0)) = JF(x_0)$$

is invertible. Therefore, if we apply Theorem 7.2.6 to the equation $f(x, y) = 0$ we find a local implicit function $G: V \rightarrow W$ such that the following equality holds:

$$\{(x, F(x)) : x \in W\} = \{(x, y) \in W \times V : f(x, y) = 0\} = \{(G(y), y) : y \in V\}.$$

Therefore, in $W \times V$ we have that $F(G(y)) = y$ and that $G(F(x)) = x$. In other words the restriction of F to W is invertible, with inverse given by G . Moreover, by Theorem 7.2.6 the inverse G is continuously differentiable.

- Conversely, for $\text{INFT} \implies \text{IMFT}$, we are given an open subset $U \subseteq \mathbb{R}^n \times \mathbb{R}^m$, a point $(x_0, y_0) \in U$ and a continuously differentiable function $f: U \rightarrow \mathbb{R}^m$ such that $f(x_0, y_0) = 0$ and we are assuming $J_y f(x_0, y_0)$

7.3 Proof of the IMFT and of the INFT

is invertible. We define the function:

$$F: U \rightarrow \mathbb{R}^n \times \mathbb{R}^m$$

$$(x, y) \mapsto (x, f(x, y)).$$

In particular, $F(x_0, y_0) = (x_0, 0)$ and the Jacobian matrix of F at (x, y) is given by

$$JF(x, y) = \begin{pmatrix} \text{Id} & 0 \\ J_x f(x, y) & J_y f(x, y) \end{pmatrix}.$$

This implies that the matrix $JF(x_0, y_0)$ is invertible. Therefore, we may apply Theorem 7.1.2 to the function F , to obtain a local inverse $G: V \rightarrow W$ of F at the point (x_0, y_0) . The function G consists of $n + m$ coordinates:

$$G(x, y) = (G_1(x, y), \dots, G_n(x, y), G_{n+1}(x, y), \dots, G_{n+m}(x, y)).$$

For convenience, we set $G'_1(x, y)$ to be the function defined by the first n components of G and $G'_2(x, y)$ to be the function defined by the last m components of G .

Let $\phi(x) = G'_2(x, 0)$ and set

$$V' := W \cap (\mathbb{R}^n \times \{0\}) \subseteq \mathbb{R}^n \times \{0\} = \mathbb{R}^n,$$

$$W' := V \cap (\{x_0\} \times \mathbb{R}^m) \subseteq \{x_0\} \times \mathbb{R}^m = \mathbb{R}^m.$$

We claim that $\phi: V' \rightarrow W'$ is the local implicit function for the equation $f(x, y) = 0$. In order to prove this, we need to show the following equality of sets:

$$\{(x, y) \in V' \times W' : f(x, y) = 0\} = \{(x, \phi(x)) : x \in V'\}.$$

Let (x, y) be a point of the Left Hand Side, then $y = G'_2(x, f(x, y))$ because $G \circ F$ is the identity on V . Moreover, since $f(x, y) = 0$ we have that:

$$y = G'_2(x, 0) = \phi(x).$$

Conversely let $(x, \phi(x))$ be a point of the Right Hand side. From the fact that $F(G(x, 0)) = (x, 0)$ we deduce the equalities $G'_1(x, 0) = x$ and $f(G'_1(x, 0), G'_2(x, 0)) = 0$ for all x . Using the definition of ϕ , we obtain

$$f(x, \phi(x)) = f(x, G'_2(x, 0)) = 0.$$

To conclude, notice that ϕ is continuously differentiable, since G'_2 is continuously differentiable by Theorem 7.1.2.

The inverse function theorem and the implicit function theorem

7.3.2. Proof of the IMFT, or Theorem 7.2.6 — This section is entirely devoted to the proof of the Implicit Function Theorem. The proof is an application of the Contraction Mapping Theorem, and it uses the Mean Value Theorem and the continuity of the differential, which we discussed in previous sections.

The details of the proof are quite involved, but the idea is similar to the proof of the existence-uniqueness theorem for Cauchy problems.

Assuming the hypotheses of Theorem 7.2.6, for all positive real numbers α and β we define the sets:

$$V_\alpha := \{x \in \mathbb{R}^n : |x_i - (x_0)_i| \leq \alpha, \quad \forall i = 1, \dots, n\} \subseteq \mathbb{R}^n,$$

$$W_\beta := \{y \in \mathbb{R}^m : |y_i - (y_0)_i| \leq \beta, \quad \forall i = 1, \dots, m\} \subseteq \mathbb{R}^m$$

(respectively the balls of radius α and β centered at x_0 and at y_0 in \mathbb{R}^n and \mathbb{R}^m with the distance d_∞). Moreover, we define the set of functions

$$\text{cont}(V_\alpha, W_\beta) = \{g: V_\alpha \rightarrow W_\beta : g \text{ is continuous}\}.$$

For every two functions g and h in $\text{cont}(V_\alpha, W_\beta)$, we define:

$$d_{L^\infty}(f, g) := \max_{x \in V_\alpha} \|f(x) - g(x)\|.$$

(The maximum exists because $x \mapsto \|f(x) - g(x)\|$ is continuous and V_α is compact).

Notice that d_{L^∞} defines a metric on $\text{cont}(V_\alpha, W_\beta)$ and that the metric space $(\text{cont}(V_\alpha, W_\beta), d_{L^\infty})$ is complete (Exercise 7.3.4).

Recall that $J_y f(x_0, y_0)$ is invertible by hypothesis. We define a map

$$\Omega: \text{cont}(V_\alpha, W_\beta) \rightarrow \text{cont}(V_\alpha, W_\beta)$$

mapping $\psi: V_\alpha \rightarrow W_\beta$ to the function:

$$\begin{aligned} \Omega(\psi): V_\alpha &\rightarrow W_\beta \\ x &\mapsto \psi(x) - J_y f(x_0, y_0)^{-1} \cdot f(x, \psi(x)). \end{aligned}$$

The key to the proof of the existence of a continuous implicit function is the observation that a function ϕ is a fixed point of Ω if and only if $\phi: V_\alpha \rightarrow W_\beta$ is a local implicit function for the equation $f(x, y) = 0$.

Our task is therefore to find α and β such that the function Ω satisfies the following properties:

1. The function Ω is well-defined.

7.3 Proof of the IMFT and of the INFT

2. The function Ω is a contraction.

Then, applying the Contraction Mapping Theorem, we will find open sets

$$V := \{x \in \mathbb{R}^n : |x_i - (x_0)_i| < \alpha, \quad \forall i = 1, \dots, n\} \subset \mathbb{R}^n,$$

$$W := \{y \in \mathbb{R}^m : |y_i - (y_0)_i| < \beta, \quad \forall i = 1, \dots, m\} \subset \mathbb{R}^m,$$

and a unique fixed point $\phi: V \rightarrow W$ that is continuous and that is a local implicit function (of y as a function of x) for the equation $f(x, y) = 0$ at (x_0, y_0) . (The fact that ϕ is continuously differentiable is left as an exercise).

We now prove that we can find $\alpha > 0$ and $\beta > 0$ in such a way that Ω satisfies Properties 1 and 2 above.

From now on, to simplify the notation, we will make the extra assumption that the Jacobian matrix

$$J_y f(x_0, y_0) = \text{Id}$$

equals the identity matrix. (This can be achieved by modifying the original function

$$(x, y) \mapsto f(x, y)$$

by

$$(x, y) \mapsto J_y f(x_0, y_0)^{-1} \cdot f(x, y).$$

In particular, we have that

$$\|J_y f(x_0, y_0)\|_{\text{op}} = \|J_y f(x_0, y_0)^{-1}\|_{\text{op}} = 1$$

and that

$$\begin{aligned} \Omega(\psi): V_\alpha &\rightarrow W_\beta \\ x &\mapsto \psi(x) - f(x, \psi(x)). \end{aligned}$$

Because f is continuously differentiable, we have in particular that $J_y f$ is defined and continuous on \mathcal{U} . By continuity at (x_0, y_0) , we may find $\delta > 0$ and $\beta > 0$ such that:

$$\|J_y f(x_0, y_0) - J_y f(x, y)\|_{\text{op}} \leq \frac{1}{2\sqrt{m}} \quad (7.4)$$

for all $(x, y) \in V_\delta \times W_\beta$.

Then, since $f(x_0, y_0) = 0$ and because the function $(x, y) \mapsto \|f(x, y)\|$ is continuous at (x_0, y_0) , we can find $0 < \alpha \leq \delta$ such that:

$$\|f(x, y_0)\| \leq \frac{\beta}{2}. \quad (7.5)$$

The inverse function theorem and the implicit function theorem

for all $x \in V_\alpha$. We claim that these α and β work for proving Properties 1. and 2. above.

1. For every continuous function $\psi: V_\alpha \rightarrow W_\beta$, the function $\Omega(\psi)$ is continuous because it is the composite of continuous functions. Moreover, we will prove the inequality

$$\|\Omega(\psi)(x) - y_0\| \leq \beta \quad (7.6)$$

for all $x \in V_\alpha$, which implies that $\Omega(\psi)(x) \in W_\beta$ because $d_\infty \leq d_2$. Therefore if we can prove (7.6), we deduce that $\Omega: \text{cont}(V_\alpha, W_\beta) \rightarrow \text{cont}(V_\alpha, W_\beta)$ is well-defined.

In order to prove (7.6), we observe that

$$\|\Omega(\psi)(x_0) - y_0\| = \|\psi(x) - f(x, \psi(x)) - y_0\|, \quad (7.7)$$

which is then equal to

$$\|\psi(x) - f(x, \psi(x)) - (y_0 - f(x, y_0)) - f(x, y_0)\|. \quad (7.8)$$

Fix x and define $G_x: W_\beta \rightarrow \mathbb{R}^m$ by $G_x(y) = y - f(x, y)$. Then the Jacobian matrix of G_x equals

$$JG_x(y) = \text{Id} - J_y f(x, y).$$

Therefore $JG_{x_0}(y_0) = \text{Id} - \text{Id} = 0$ and we have

$$\begin{aligned} \|JG_x(y)\|_{\text{op}} &= \|\text{Id} - J_y f(x, y)\|_{\text{op}} \\ &\leq \frac{1}{2\sqrt{m}}, \end{aligned} \quad (7.9)$$

for all $(x, y) \in V_\alpha \times W_\beta$ because of (7.4).

We conclude the proof of (7.6) by the following argument:

$$\begin{aligned} (7.8) &\leq \|\psi(x) - f(x, \psi(x)) - (y_0 - f(x, y_0))\| + \|f(x, y_0)\| \\ &= \|G_x(\psi(x)) - G_x(y_0)\| + \|f(x, y_0)\| \\ &\leq \frac{1}{2\sqrt{m}} \|\psi(x) - y_0\| + \|f(x, y_0)\| \\ &\leq \frac{\beta}{2} + \frac{\beta}{2} = \beta \end{aligned}$$

7.4 Lagrange multipliers (EXTRA)

for all $(x, y) \in V_\alpha \times W_\beta$. For the first summand of the second inequality we used the Mean Value Theorem combined with Inequality (7.9). For the last inequality we used the inequality $d_2 \leq \sqrt{m} \cdot d_\infty$ (exercise!) for distances in \mathbb{R}^m (for the first summand) and Inequality (7.5) (for the second summand). This concludes the proof of Part 1. (the fact that Ω is well defined).

2. Now we show Part 2, namely that Ω is a contraction. By the Mean Value Theorem combined with Inequality (7.9), for every ψ_1 and ψ_2 in $\text{cont}(V_\alpha, W_\beta)$ we have:

$$\|\Omega(\psi_1)(x) - \Omega(\psi_2)(x)\| \leq \frac{1}{2\sqrt{m}} \|\psi_1(x) - \psi_2(x)\|.$$

Since $\frac{1}{2\sqrt{m}} < 1$, we conclude that Ω is a contraction.

This concludes the proof of the existence of a continuous local implicit function. The fact that the function is also continuously differentiable is left as an exercise.

Exercise 7.3.3. — Prove that ϕ is also continuously differentiable.

Exercise 7.3.4. — Let V_α and W_β be as in 7.3.2. Prove that

$$(\text{cont}(V_\alpha, W_\beta), d_{L^\infty})$$

is a metric space, and then that it is complete.

Hint: Follow what we have done in the first part of Chapter 5. To show that the metric space is complete, prove that $\text{cont}(V_\alpha, \mathbb{R}^n)$ is complete and then that $\text{cont}(V_\alpha, W_\beta)$ is a closed subset of $\text{cont}(V_\alpha, \mathbb{R}^n)$.

7.4. Lagrange multipliers (EXTRA)

An important application of the implicit function theorem is a result that goes under the name of Lagrange's multipliers method.

The problem addressed by that method is the following. We are given an open subset $U \subseteq \mathbb{R}^{n+m}$, a number of *constraint functions* $g_1, \dots, g_m: U \rightarrow$

The inverse function theorem and the implicit function theorem

\mathbb{R} , and a *cost function* $f: U \rightarrow \mathbb{R}$ that we seek to minimise (or maximise) subject to the constraints

$$\begin{cases} g_1(x_1, \dots, x_{n+m}) = 0, \\ g_2(x_1, \dots, x_{n+m}) = 0, \\ \dots \\ g_m(x_1, \dots, x_{n+m}) = 0. \end{cases} \quad (7.10)$$

In other words, we want to find the minimum of f on the subset X of U that is defined by the system of equations (7.10).

An example of this that was discussed extensively in Year 1 is the case when $m = 0$ (i.e. when there are no constraints). When looking for local minima (and maxima) of a differentiable function f defined on some open subset U of \mathbb{R}^n , we can just consider the points of U where the differential of f (or equivalently its gradient, since f is scalar-valued) vanishes. This works because, if P is a local extremal point (a local minimum or maximum), then the differential Df_P is zero (it is the zero function). Let's start by reviewing this last assertion.

Remark 7.4.1. — If a function $f: U \rightarrow \mathbb{R}$ is differentiable and $x_0 \in U$ is a local extremal point of f , then $Df_{x_0} = 0$. (This is also known as Fermat's theorem).

Indeed, if x_0 is (for example) a local minimum of f , then $f(x) \geq f(x_0)$ for all $x \in V$ for some V open containing x_0 . Taking the limits

$$\lim_{t \rightarrow 0^+} \frac{f(x_0 + tv) - f(x_0)}{t} \geq 0 \quad \lim_{t \rightarrow 0^-} \frac{f(x_0 + tv) - f(x_0)}{t} \leq 0$$

we deduce that $Df_v(x_0) = 0$ for all directions v , hence the differential $Df(x_0)$ must equal zero (= the zero linear function).

The result of Lagrange's multipliers is obtained by combining the above idea with the Implicit Function Theorem.

Theorem 7.4.2. — (*Lagrange's multipliers*) In the above setup, assume furthermore that the function $g = (g_1, \dots, g_m): U \rightarrow \mathbb{R}^m$ is continuously differentiable, that $f: U \rightarrow \mathbb{R}^m$ is differentiable, and that the matrix $J_y g(x_0, y_0)$ is invertible for all $(x_0, y_0) \in U$, and let

$$X = \{(x, y) \in U : g(x, y) = 0\} \subseteq U \subseteq \mathbb{R}^{n+m}.$$

7.4 Lagrange multipliers (EXTRA)

If the restriction $f|_X$ of f to X has a local extremal point at some $P \in X$, then there exists $(\lambda_1, \dots, \lambda_m) \in \mathbb{R}^m$ such that

$$Jf(P) + \sum_{k=1}^m \lambda_k \cdot Jg_k(P) = 0. \quad (7.11)$$

(This is an equality of row vectors in \mathbb{R}^{n+m}).

It is perhaps worth noting that if a point is a local maximum (or minimum) for f on U , then it is also a local maximum (or minimum) for f on X , but the converse is not true (can you think of a simple example?).

Note that the Jacobian matrices $Jf = \nabla f$ and $Jg_k = \nabla g_k$ are $1 \times m$ matrices, hence row vectors (because the functions f, g_1, \dots, g_m are scalar-valued). One could prove that the vectors $\nabla g_1, \dots, \nabla g_m$ span the normal directions to X in \mathbb{R}^{n+m} (to prove this assertion we should also define the notion of "normal directions").

In the proof we use our usual convention that $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$.

Proof. (Sketch). Apply the IMFT and find a local implicit function ϕ that parameterises X near P , so there are V and W open subsets of \mathbb{R}^n and \mathbb{R}^m such that $P \in V \times W \subseteq U$, and an implicit function $\phi: V \rightarrow W$ such that

$$X \cap (V \times W) = \{(x, \phi(x)) : x \in V\}.$$

Now if f has a local extremal point at $P = (x_0, y_0) \in X$, then so does $(f \circ \phi): V \rightarrow \mathbb{R}$. The latter is a function from an open subset of \mathbb{R}^n (without constraints), and we can apply the usual criterion on the vanishing of the differential.

The Chain Rule, combined with the formula for the differential of the implicit function, gives:

$$0 = J(f \circ (\text{Id}, \phi))(x_0) = Jf(P) \cdot \begin{pmatrix} \text{Id} \\ J\phi(x_0) \end{pmatrix} = J_x f(P) + J_y f(P) \cdot (-J_y g(P)^{-1} \cdot J_x g(P)).$$

By defining

$$(\lambda_1, \dots, \lambda_m) := J_y f(P) \cdot J_y g(P)^{-1}$$

we obtain

$$J_x f(P) + \sum_{k=1}^m \lambda_k \cdot J_x g_k(P) = 0,$$

which is Equation (7.11), but only for the first n components of the vectors.

The inverse function theorem and the implicit function theorem

For the remaining m components, we have

$$J_y f(P) = J_y f(P) \cdot (J_y g(P))^{-1} \cdot J_y g(P) = (\lambda_1, \dots, \lambda_m) \cdot J_y g(p),$$

where the last equality follows from the definition of $(\lambda_1, \dots, \lambda_m)$. This last equality completes the proof of (7.11) for the remaining m components. \square

Remark 7.4.3. — The Lagrange multipliers method helps finding local extremal points (i.e. local minima and maxima). A further analysis is required to determine whether those local extremal points are actually global minima or maxima (just as in the case of the same question for the case of a differentiable function $f: U \rightarrow \mathbb{R}$ on an open subset $U \subseteq \mathbb{R}^n$).