

Minimal analysis II:

A concise introduction to real analysis of multiple variables

Tonatiuh Sánchez-Vizuet
The University of Arizona

Last update: August 7, 2025

Chapter 1

Preliminaries

1.1 Euclidean space

The *Euclidean space*, denoted by \mathbb{R}^n , is defined by the set:

$$\mathbb{R}^n = \{x = (x_1, x_2, \dots, x_n) : x_i \in \mathbb{R}\}.$$

That is, \mathbb{R}^n is the set of n -tuples of real numbers. The Euclidean space \mathbb{R}^n is a vector space with addition:

$$x + y = (x_1 + y_1, x_2 + y_2, \dots, x_n + y_n),$$

and scalar multiplication:

$$\alpha x = (\alpha x_1, \alpha x_2, \dots, \alpha x_n).$$

Remark 1.1. We should always *think of vectors as columns*, i.e. a vector $x \in \mathbb{R}^n$ should be pictured as a matrix with n rows and 1 column. However, for notational convenience, in this text we will often write a vector in terms of its components as a horizontal list (as we did above), this will make the notation a little simpler by sparing us from writing transpose symbols over and over. However you should always think of vectors as a column.

Additionally, the Euclidean space has *scalar product* given by:

$$x \cdot y = x_1 y_1 + x_2 y_2 + \dots + x_n y_n = \sum_{i=1}^n x_i y_i.$$

The scalar product product, also called the *dot product* or *inner product*, induces the *Euclidean norm*:

$$|x| = \sqrt{x \cdot x} = \sqrt{\sum_{i=1}^n (x_i)^2}.$$

We say that vectors $x, y \in \mathbb{R}^n$ are *orthogonal* if $x \cdot y = 0$. In this case, we say write that $x \perp y$ or $y \perp x$. Given two non-zero vectors x, y we can use the dot product to define the *projection* and the *rejection* of x onto y respectively by

$$\text{proj}_y x := \left(\frac{y \cdot x}{|y|^2} \right) y \quad \text{and} \quad \text{oproj}_y x \equiv (\text{proj}_y x)^\perp := x - (\text{proj}_y x).$$

Note that

$$\begin{aligned}\text{proj}_y x \cdot \text{oproj}_y x &= \left(\frac{y \cdot x}{|y|^2} \right) y \cdot \left(x - (\text{proj}_y x) \right) \\ &= \left(\frac{y \cdot x}{|y|^2} \right) y \cdot \left(x - \left(\frac{y \cdot x}{|y|^2} \right) y \right) \\ &= \left(\frac{x \cdot y}{|y|} \right)^2 - \left(\frac{x \cdot y}{|y|^2} \right)^2 |y|^2 = 0,\end{aligned}$$

and therefore $\text{proj}_y x$ is orthogonal to $\text{oproj}_y x$ which justifies the notation $\text{oproj}_y x \equiv (\text{proj}_y x)^\perp$. Hence, from the definition of the projection and the rejection we see that given any x, y we can decompose x into two orthogonal components as

$$x = \text{proj}_y x + (\text{proj}_y x)^\perp.$$

Vector addition, subtraction and projection all have geometric interpretations depicted in Figure 1.1.

The Euclidean norm and the dot product have the following properties

Proposition 1.1. *Let $x, y \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$.*

1. **Positive definiteness**

$$|x| \geq 0. \text{ Moreover, } |x| = 0 \text{ if and only if } x = 0.$$

2. **Absolute homogeneity**

$$|\alpha x| = |\alpha| |x|.$$

3. **Cauchy-Schwarz inequality**

$$|x \cdot y| \leq |x| |y|.$$

4. **Triangle inequality**

$$|x + y| \leq |x| + |y|.$$

Proof. The proofs of properties 1 and 2 are left as exercises. For property 3, we observe that if $x = 0$, then both sides of the inequality are zero. Suppose then that $x \neq 0$ and let w be the projection of x onto y :

$$w = \left(\frac{y \cdot x}{|x|^2} \right) x.$$

Then:

$$\begin{aligned}0 \leq |y - w|^2 &= (y - w) \cdot (y - w) \\ &= \left(y - \frac{y \cdot x}{|x|^2} x \right) \cdot \left(y - \frac{y \cdot x}{|x|^2} x \right), \\ &= |y|^2 - 2 \frac{(y \cdot x)^2}{|x|^2} + \frac{(y \cdot x)^2}{|x|^4} |x|^2 \\ &= |y|^2 - \frac{(y \cdot x)^2}{|x|^2}.\end{aligned}$$

From this, it follows that

$$|y|^2 \leq \frac{(y \cdot x)^2}{|x|^2}$$

and the inequality follows after taking the square root of this expression.




Figure 1.1: Geometric representation of vector addition, subtraction and projection. The vectors u and v are depicted in grey, while the subtraction, addition and projections are depicted in black.

For property 4, we see that

$$|x + y|^2 = (x + y) \cdot (x + y) = |x|^2 + 2x \cdot y + |y|^2 \leq |x|^2 + 2|x||y| + |y|^2,$$

where the last inequality follows from the Cauchy-Schwarz inequality. Therefore, we have:

$$|x + y|^2 \leq (|x| + |y|)^2.$$

Taking the square root above leads to the triangle inequality. \square

Remark 1.2. From the proof of Proposition 1.1 and Figure 1.1, we can observe that equality in point 3 holds if and only if one of the vectors x or y is a scalar multiple of the other one. Indeed, if y is a scalar multiple of x , then $(\text{proj}_y x)^\top = y - w = 0$ and $y = w$.

Similarly, equality holds in property 4 if and only if $x \cdot y = |x||y|$, that is, when one of the vectors x or y is a scalar multiple of the other and $x \cdot y > 0$. Geometrically, this means that y lies on the line generated by x and in the same direction.

1.2 Geometry of \mathbb{R}^n

The **canonical basis** of \mathbb{R}^n consists of the vectors e_1, e_2, \dots, e_n , where:

$$e_i = (0, 0, \dots, \overset{i\text{-th}}{1}, \dots, 0).$$

The scalar product can be used to define the **angle between two vectors** θ by the relation

$$\theta := \cos^{-1} \left(\frac{x \cdot y}{|x||y|} \right).$$

Note that, in view of the Cauchy-Schwarz inequality, the magnitude of the argument above remains bounded by 1, and therefore the angle is well defined. From our definition of the angle between two vectors it follows that orthogonal vectors with respect to the Euclidean inner product are perpendicular to each other.

A straight **line** in \mathbb{R}^n going through a point x^* and pointing in the direction of a given vector x_0 is the set of points of the form

$$L := \{x \in \mathbb{R}^n : x = x^* + \alpha x_0\},$$




Figure 1.2: Left: A hyperplane is completely determined by prescribing one point x^* in the plane and the normal vector n . Right: when the point x^* is taken to be the origin and the normal vector is parallel to the n -th canonical basis, the points above and below are referred to as upper \mathbb{R}_+^n and lower \mathbb{R}_-^n half space.

with $x^*, x_0 \in \mathbb{R}^n$ and $\alpha \in \mathbb{R}$. The **parametric equation of the line** passing through x and y is:

$$\gamma(t) = (1 - t)x + ty, \quad t \in \mathbb{R}.$$

Note that $\gamma(0) = x$ and $\gamma(1) = y$, therefore, restricting $t \in [0, 1]$ we obtain the parametrization of the line segment from x to y .

A **hyperplane** is a set of the form:

$$P := \{x \in \mathbb{R}^n : x \cdot x_0 = c_0\},$$

where $x_0 \in \mathbb{R}^n$, $x_0 \neq 0$, and $c_0 \in \mathbb{R}$ are both *fixed*. The hyperplane orthogonal to $n \in \mathbb{R}^n$ passing through $x^* \in \mathbb{R}^n$ is given by:

$$P := \{x : (x - x^*) \cdot n = 0\}.$$

In this context, the vector n is referred to as the **normal vector** to the plane. Alternatively, the **parametric equation of a hyperplane** can be obtained by observing that a hyperplane passing through the point x^* in \mathbb{R}^n is the set of all linear combinations of $n - 1$ linearly independent vectors v_1, \dots, v_{n-1}

$$\Gamma(\alpha_1, \dots, \alpha_{n-1}) = x^* + \sum_{i=1}^{n-1} \alpha_i v_i,$$

for $\alpha_1, \dots, \alpha_{n-1} \in \mathbb{R}$. A hyperplane P divides \mathbb{R}^n into two half-spaces:

$$\{x : x \cdot x_0 > c\} \quad \text{and} \quad \{x : x \cdot x_0 < c\}.$$

If $x_0 = e_n$ and $c = 0$, these are called the upper and lower half-spaces and denoted \mathbb{R}_+^n and \mathbb{R}_-^n , respectively, as depicted in Figure 1.2. If $x_0 \in \mathbb{R}^n$ and $r > 0$, the **sphere** of radius r around x_0 is given by the set:

$$S_r(x_0) = \{x : |x - x_0| = r\},$$

while the **open ball** (or, for simplicity, simply the ball) of radius r around x_0 is given by:

$$B_r(x_0) = \{x : |x - x_0| \leq r\}.$$

We say that $A \subset \mathbb{R}^n$ is **convex** set if, for all $x, y \in A$, the segment from x to y is completely contained in A . The set $A \subset \mathbb{R}^n$ is said to be a **star-shaped set** if there exists a point $O \in A$ such that, for $x \in A$, the




Figure 1.3: Left: In a convex set, all straight segments connecting interior points remain inside of the set. Center: Non-convex sets contain at least one pair of points such that the straight segment connecting them is not completely contained in the set. Right: Star-shaped domains may not be convex, but they contain at least one point that is connected to every other point through segments that remain inside of the set, such as the point O above. All convex-sets are star shaped.

segment from O to x is completely contained in A . Examples of these kinds of sets are depicted in Figure 1.3. We will study convex sets in more depth later on.

Let $I_i \subset \mathbb{R}$ be a bounded interval with left endpoint a_i and right endpoint b_i for every $1 \leq i \leq n$. A **rectangle** in \mathbb{R}^n is a set of the form:

$$R = I_1 \times I_2 \times \cdots \times I_n,$$

in words, a rectangle is the Cartesian product of n bounded intervals I_i in \mathbb{R} . If each I_i is an open interval, we say R is an open rectangle. If each I_i is closed, we say R is a closed rectangle. A rectangle in \mathbb{R}^n is sometimes known as a **hypercube**.

In \mathbb{R}^n , we will denote by S^{n-1} the **unit sphere** given by

$$S^{n-1} := \{x \in \mathbb{R}^n : |x| = 1\}.$$

Note that, even if it is embedded in the n -dimensional space \mathbb{R}^n , the sphere itself is intrinsically $n - 1$ dimensional. This justifies the superscript $n - 1$ in the notation.

1.3 Sequences in \mathbb{R}^n

A sequence $\{x_k\}_{k=1}^{\infty} \subset \mathbb{R}^n$ is a countable subset of \mathbb{R}^n that we tag using a natural number as index. Note that since for every natural number we assign a point $x_k \in \mathbb{R}^n$, a sequence can also be considered as a function $f : \mathbb{N} \rightarrow \mathbb{R}^n$ given by $f(k) = x_k$.

Remark 1.3. In \mathbb{R}^n , the indexing of elements of a sequence and the tagging of the components of each element of the sequence can be a little confusing. When dealing with sequences, we will use a subindex to denote the index of an element of a sequence, and a superindex to denote the component of an element of a sequence. Therefore in the expression

$$x_k = (x_k^1, x_k^2, \dots, x_k^i, \dots, x_k^n),$$

x_k^i denotes the i -th component of the n -th element of a sequence. Moreover, each of the coordinates of x_k defines a sequence $\{x_k^i\}_k$ in \mathbb{R} . To avoid confusions when raising a component to a power, we will use the notation

$$(x_k^i)^m$$

to mean *the i -th component, of the k -th term of the sequence raised to the m -th power.*

Definition 1.1. Let $A \subset \mathbb{R}$. We say that $\{x_k\}_{k=1}^{\infty}$ is a sequence in A if $x_k \in A$ for all k .

Definition 1.2. We say that $\{x_k\}_{k=1}^{\infty}$ is a **bounded sequence** if there exists $M > 0$ such that

$$|x_k| \leq M \quad \text{for all } k.$$

Equivalently, $\{x_k\}_{k=1}^{\infty}$ is bounded if there exists a rectangle R such that $x_k \in R$ for all k . Furthermore, $\{x_k\}_{k=1}^n$ is bounded in \mathbb{R}^n if and only if each every component-wise sequence $\{x_k^i\}_{k=1}^{\infty}$ is bounded in \mathbb{R} .

Definition 1.3 (convergence). We say that the sequence $\{x_k\}$ converges to $L \in \mathbb{R}^n$ if, for every $\epsilon > 0$, there exists N such that, if $k \geq N$,

$$|L - x_k| < \epsilon.$$

If the sequence $\{x_k\}$ converges to L , we call it the limit of $\{x_k\}$ and write:

$$L = \lim_{k \rightarrow \infty} x_k \quad \text{or} \quad x_k \rightarrow L.$$

Furthermore, the limit of a sequence is unique (Exercise (8)).

It is not very difficult to verify the following statements; each of them characterizes the convergence of a sequence.

1. The sequence $\{x_k\}$ converges to $L \in \mathbb{R}^n$ if, for every $\epsilon > 0$, there exists N such that, for $k \geq N$, $x_k \in B_{\epsilon}(L)$.
2. The sequence $\{x_k\}$ converges to $L \in \mathbb{R}^n$ if, for every open rectangle R containing L , there exists N such that, for $k \geq N$, $x_k \in R$.

However, in practice, the following proposition that connects the convergence of a sequence of vectors to convergence of the sequences of components is very useful.

Proposition 1.2. *The sequence $\{x_k\}_{k=1}^{\infty}$ converges to $L = (L^1, L^2, \dots, L^n)$ in \mathbb{R}^n if and only if each $\{x_k^i\}_{k=1}^{\infty}$ converges to L^i in \mathbb{R} , $i = 1, 2, \dots, n$.*

Proof. $\boxed{\Rightarrow}$ Suppose $x_k \rightarrow L$, and let $\epsilon > 0$. Let N be such that $k \geq N$ implies $|x_k - L| < \epsilon$. Then, for $k \geq N$,

$$|x_k^i - L^i| \leq \sqrt{(x_k^1 - L^1)^2 + \dots + (x_k^i - L^i)^2 + \dots + (x_k^n - L^n)^2} < \epsilon.$$

$\boxed{\Leftarrow}$ Now suppose each $x_k^i \rightarrow L^i$, and let $\epsilon > 0$. Take N_i such that, for $k \geq N_i$,

$$|x_k^i - L^i| < \epsilon/\sqrt{n}.$$

Let $N = \max_{1 \leq i \leq n} \{N_i\}$ and $L = (L^1, \dots, L^n)$. Then, if $k \geq N$,

$$|x_k - L| = \sqrt{\sum_{i=1}^n (x_k^i - L^i)^2} < \sqrt{\sum_{i=1}^n (\epsilon^2/n)} = \epsilon.$$

□

The following proposition classifies closed sets in terms of sequences.

Proposition 1.3. *A set $A \subset \mathbb{R}^n$ is closed if and only if, for every sequence $\{x_k\}_{k=1}^{\infty}$ of points in A that converges to some $L \in \mathbb{R}^n$, it holds that $L \in A$. In other words, a set is closed if and only if it contains all its limit points.*

Proof. \Rightarrow Suppose A is closed, and let (x_k) in A be a sequence that converges to L . Given $\epsilon > 0$, since $x_k \rightarrow L$, there exists K such that $x_K \in B_{\epsilon}(L)$. Since $x_K \in A$, we have shown that $B_{\epsilon}(L) \cap (A \setminus \{L\}) \neq \emptyset$. If $L \notin A$ it would follow that $B_{\epsilon}(L) \cap A^c \neq \emptyset$ and therefore L would be a limit point of A by the definition B.4, but since A was assumed to be closed, then L would have to belong to A , which is a contradiction. Therefore, we must have $L \in A$.

\Leftarrow Now suppose every sequence of points in A that converges has its limit in A . Let x be an accumulation point of A . For each $k \geq 1$, let $x_k \in A$ such that $|x_k - x| < 1/k$. Such x_k must exist because $B_{1/k}(x) \cap A \neq \emptyset$. Then x_k is a sequence in A , and $x_k \rightarrow x$, so by our hypothesis $x \in A$. □

It is easy to verify that convergent sequences are bounded (Exercise 9). However, while the converse is clearly false (imagine, for instance, the sequence $x_k = e_1$ if k is even and $x_k = -e_1$ if k is odd, which is clearly bounded and not convergent), the following theorem, known as the **Bolzano-Weierstrass theorem**, applies. A **subsequence** of $\{x_k\}_{k=1}^{\infty}$ is an infinite subset of elements of the original sequence that we tag by the index k_l as $\{x_{k_l}\}$. That is, a subsequence of $\{x_k\}$ is a sequence whose terms are taken from the terms of $\{x_k\}$, respecting their order.

Theorem 1.1 (Bolzano-Weierstrass). *Every bounded sequence has a convergent subsequence.*

In our proof, we will assume the theorem is valid in \mathbb{R} .

Proof. If $\{x_k\}_{k=1}^{\infty} \subset \mathbb{R}^n$ is bounded, each component-wise sequence $\{x_k^i\}_{k=1}^{\infty} \subset \mathbb{R}$ is bounded. Hence, by the Bolzano-Weierstrass theorem in \mathbb{R} , the sequence $\{x_k^1\}_{k=1}^{\infty}$ has a subsequence that converges to a real number L^1 . Call this subsequence $\{x_{k_{l_1}}^1\}$ and select the subsequence of vectors with the same index $\{x_{k_{l_1}}\}$.

This new sequence may not be convergent, but it has two important properties: 1) by construction, its first component is convergent, and 2) as a subset of the original sequence, it is also bounded and therefore the sequence of its second components $\{x_{k_{l_1}}^2\}$ is bounded. Once again, the Bolzano-Weierstrass theorem guarantees the existence of a subsequence $\{x_{k_{l_2}}^2\}$ that converges to some number L^2 . We take the subsequence of vectors with the same index $\{x_{k_{l_2}}\}$ and note that, by construction, the subsequence is bounded and its first two components converge.

Proceeding inductively in this fashion, after n steps we would have produced a subsequence $\{x_{k_{l_n}}\} \subset \{x_k\}$ such that its i -th component converges to the number L^i . Defining $L = (L^1, \dots, L^n) \in \mathbb{R}^n$ we see that the original sequence $\{x_k\}$ contains a subsequence $\{x_{k_{l_n}}\}$ such that

$$x_{k_{l_n}}^i \rightarrow L^i, \quad \forall i$$

and thus, by Proposition 1.2, it follows that the sequence of vectors $x_{k_{l_n}} \rightarrow L$. □

The Bolzano-Weierstrass theorem allows us to prove the following property of closed sets, which will be useful later on.

Theorem 1.2. Let $A \subset \mathbb{R}^n$ be a non-empty closed set and $x \in \mathbb{R}^n$. Then there exists a point $y \in A$ such that $|x - y|$ is minimized.

Proof. First of all, we observe that for any fixed $x \in \mathbb{R}^n$, the facts that 1) A is not empty, and 2) the Euclidean norm $|\cdot|$ is non negative, guarantee that the set of real numbers

$$\{d \in \mathbb{R} : d = |x - y| \text{ for } y \in A\}$$

is not empty and bounded from below by 0. It therefore it has an infimum that we shall denote

$$r_0 := \inf\{d \in \mathbb{R} : d = |x - y| \text{ for } y \in A\}.$$

Then, by the characterization of the infimum given in Theorem A.2, for all $k \geq 1$, there exists $y_k \in A$ such that:

$$r_0 \leq |x - y_k| < r_0 + \frac{1}{k}. \quad (1.1)$$

Some simple algebra shows that this implies that $|y_k| \leq r_0 + 1 + |x|$. Hence, the sequence $\{y_k\}$ is bounded, and by the Bolzano-Weierstrass theorem, it has a convergent subsequence. Let $\{y_{k_l}\} \subset A$ denote the subsequence and y be its limit. Since A is closed, Proposition 1.3 implies $y \in A$.

We shall now prove that y is indeed the point for which $|x - y|$ is minimized. To see this, given $\epsilon > 0$, take a term y_m in the convergent subsequence such that $|y_m - y| < \epsilon/2$ and $1/m < \epsilon/2$. Then:

$$r_0 \leq |x - y| \leq |x - y_m| + |y_m - y| \underset{\text{From (1.1)}}{\leq} r_0 + \frac{1}{m} + \frac{\epsilon}{2} < r_0 + \epsilon,$$

which implies that

$$0 \leq |x - y| - r_0 < \epsilon.$$

Since $\epsilon > 0$ is arbitrary, we use the ϵ -principle (A.1) to conclude that $r_0 = |x - y|$. □

Remark 1.4. Clearly, if $x \in A$, then

$$\inf\{d \in \mathbb{R} : d = |x - y| \text{ for } y \in A\} = 0.$$

Furthermore, since A is closed, if $x \notin A$, then x is not an accumulation point of A , and there exists $r > 0$ such that $B_r(x) \cap A = \emptyset$. Thus, $r_0 \geq r > 0$.

Definition 1.4. We say that $\{x_k\}_{k=1}^{\infty}$ is a **Cauchy sequence** if, for every $\epsilon > 0$, there exists N such that, if $k, l \geq N$, then:

$$|x_k - x_l| < \epsilon.$$

In other words, (x_k) is a Cauchy sequence if its terms get arbitrarily close to each other.

If a sequence converges, then it is a Cauchy sequence. To verify this, suppose $x_k \rightarrow L$. Then, given $\epsilon > 0$, there exists N such that, if $k \geq N$, $|x_k - L| < \epsilon/2$. Therefore, if $k, l \geq N$,

$$|x_k - x_l| \leq |x_k - L| + |L - x_l| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

Conversely, if $\{x_k\}$ is a Cauchy sequence, then it converges. This follows from the Bolzano-Weierstrass theorem and is left as an exercise for the reader (Exercises 22–24).

1.4 Compact sets

In this section, we study a very important concept in analysis and topology: compactness. We will explore compact sets and their relationship with sequences in \mathbb{R}^n .

Definition 1.5. Let $A \subset \mathbb{R}^n$ and let I be a collection of indices that may be uncountable. A **cover** for A is a collection $\{U_\alpha\}_{\alpha \in I}$ of open sets such that:

$$A \subset \bigcup_{\alpha \in I} U_\alpha.$$

If $\{U_\alpha\}_{\alpha \in I}$ is a cover for A , and I' is a (potentially uncountable) collection of indices, a **subcover** is a subset of $\{U_\alpha\}_{\alpha \in I}$, say $\{U_{\alpha_\beta}\}_{\alpha_\beta \in I'} \subset \{U_\alpha\}_{\alpha \in I}$, such that:

$$A \subset \bigcup_{\beta \in I'} U_{\alpha_\beta}.$$

We say that A is **compact** if every cover for A contains a finite subcover.

Remark 1.5. Note that the compactness of a set A **does not** depend on whether it is possible to cover the set with a single open set. Compactness hinges on whether from **any possible** family of opens sets containing A it is possible to select a finite number of them that still cover A .

Example 1.1 (A non-compact set). Consider the open set $A := (0, 1) \subset \mathbb{R}$. Even though it is possible to cover this set with a single open set (for instance with itself), it is not compact. To prove that, we must exhibit a family of open sets $\{U_\alpha\}$ such that A is contained in their union, but such that no finite subset of $\{U_\alpha\}$ can still cover A completely.

We then define the countable family of open intervals $U_n := (0, 1 - 1/n)$ for every $n \in \mathbb{N}$. Starting from the empty interval $U_1 = (0, 0)$, as n grows the length of the intervals U_n increases until, eventually, in the limit $n \rightarrow \infty$, the sets cover A . This follows from the Archimedian property of the natural numbers, since for every $x \in (0, 1)$ there exists some $M \in \mathbb{N}$ such that $x < 1 - 1/M$. Hence $(0, 1) \subset \bigcup_{n=1}^{\infty} U_n$.

Now consider, any finite subset of this collection (i.e. any finite subcover) and let I be the set containing all the indices of the sets U_n included in the subcollection. Since the number of sets is finite, there exists one finite N such that $N \geq i$ for all $i \in I$. The set corresponding to this index is of the form $(0, 1 - 1/N)$ and has the property that $U_i \subset U_N$ for all $i \in I$. Therefore $\bigcup_{i \in I} U_i = U_N$. However, since $1/(N + 1) < 1/N$ the point $x = 1 - 1/(N + 1) \in (0, 1)$ does not belong to U_N . Therefore, no finite collection of sets of the form $(0, 1 - 1/n)$ can cover the interval $(0, 1)$ and thus it is not compact.

Definition 1.6. A set $A \subset \mathbb{R}^n$ is said to be **sequentially compact** if for every sequence $\{x_n\}_{n=1}^{\infty} \subset A$ there exists a subsequence $\{x_{n_k}\}_{n_k=1}^{\infty} \subset \{x_n\}_{n=1}^{\infty}$ and a point $x \in A$ such that

$$x_{n_k} \rightarrow x.$$

In words, a set is sequentially compact if every sequence of points in the set contains a convergent subsequence whose limit also belongs to the set.

Definition 1.7. We say that A is a **bounded set** if it is contained in a ball $B_M(0)$, for some $M > 0$.

Equivalently, A is bounded if there exists a rectangle R such that $A \subset R$.

Theorem 1.3. *In the Euclidean space \mathbb{R}^n the following three statements are equivalent:*

1. *The set A is sequentially compact.*
2. *The set A is compact.*
3. *The set A is closed and bounded.*

Remark 1.6. *The statements 1 and 2 are equivalent to each other in the more general setting of a metric space. The fact that \mathbb{R}^n , the statements 2 and 3 are equivalent to each other is known as the **Heine-Borel theorem**.*

Proof. We will proceed by proving the implications $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 1$. The first implication will be the hardest one, while the third one will seem to be easy thanks to the power of the Bolzano-Weierstrass theorem.

1 \Rightarrow 2 We must show that if A is sequentially compact, then any open cover of A can be reduced to a finite subcover. We will then start with a potentially uncountable cover, first reduce it to a countable cover, and finally argue by contradiction to show that it is possible to obtain a finite subcover.

Step 1. Let $\{U_\alpha\}$ be a potentially uncountable cover for A and pick $x \in A$. Since $\{U_\alpha\}$ is a cover of A , there exists at least one $U_\beta \in \{U_\alpha\}$ such that $x \in U_\beta$. Recalling that the boundary ∂U_β is a closed set (see Proposition B.4 in the Appendix), and that $x \notin \partial U_\beta$, then Proposition 1.2 and Remark 1.4 guarantee that the minimum distance from x to ∂U_β is strictly positive. Let us denote this distance by $d > 0$ and note that, by construction, $B_d(x) \subset U_\beta$.

Since the set

$$\mathbb{Q}^n := \{x = (x_1, \dots, x_n) : x_i \in \mathbb{Q} \text{ for all } 1 \leq i \leq n\}$$

of vectors with rational components is dense in \mathbb{R}^n (Exercise 7), we can then pick $p \in \mathbb{Q}$ and $q \in \mathbb{Q}^n$ such that

$$|x - q| < p < d/2.$$

Clearly, $x \in B_p(q)$ and, for any other $z \in B_p(q)$

$$|x - z| \leq |x - q| + |q - z| < p + p < d/2 + d/2 = d,$$

so that $B_p(q) \subset B_d(x) \subset U_\beta$.

We have then proven that for every $x \in A$ there exists a ball with rational radius centered at a rational point that contains x and is contained in an element of the cover $\{U_\alpha\}$. **Note that there can be at most countably many of such balls** (as the rational numbers are countable). Hence, if for every $x \in A$ we select the open set $\{U_\beta\}$ containing the “rational ball” appearing in the argument above, we will end up with a countable subset of the original open cover that still contains every element of A . Let us then denote this countable subcover by $\{U_n\}_{n=1}^\infty$.

Step 2. Having obtained a countable subcover, we will now proceed by contradiction to show that there exists a finite subcover. Assume that not finite subset of the open cover $\{U_n\}_{n=1}^\infty$ can cover A . Pick an arbitrary point $x_1 \in A$ and note that, since $A \subset \{U_n\}_{n=1}^\infty$ there exists some \tilde{U}_1 such that $x_1 \in \tilde{U}_1$, however, since there is no finite subcover, there must exist some $x_2 \in A$ such that $x_2 \notin \tilde{U}_1$. On the other hand, given that $A \subset \{U_n\}_{n=1}^\infty$ there exists some \tilde{U}_2 such that $x_2 \in \tilde{U}_2$. Once again, due to the fact that there is no finite subcover, we can find some $x_3 \in A$ such that $x_3 \notin \tilde{U}_1 \cup \tilde{U}_2$ and some \tilde{U}_3 such that $x_3 \in \tilde{U}_3$. Proceeding inductively in this way we obtain a sequence of points $\{x_k\} \subset A$ and a family of open sets $\{\tilde{U}_k\} \subset \{U_n\}$ such that

$$x_k \notin \bigcup_{i=1}^{k+1} \tilde{U}_i.$$

By construction, the open set \tilde{U}_k can contain at most k terms in the sequence.

Now let us look at the sequence of points $\{x_k\}_{k=1}^{\infty} \subset A$ that we have obtained. Since we assumed A to be sequentially compact, there exists a subsequence $\{x_{k_\ell}\}_{\ell=1}^{\infty} \subset \{x_k\}_{k=1}^{\infty}$ and a point $x \in A$ such that

$$x_{k_\ell} \rightarrow x.$$

Due to the fact that x is a limit point of the sequence, there must exist some \tilde{U}_m in the family such that $x \in \tilde{U}_m$ otherwise there would exist a positive number r such that for every $n \in \mathbb{N}$

$$\min\{|x - z| : z \in \overline{U}_n\} > r > 0. \quad (\text{the overline represents the closure of the set})$$

Since $x_{k_\ell} \in U_{k_\ell}$ this would imply that

$$|x - x_{k_\ell}| > r > 0 \quad \text{for all } k_\ell \in \mathbb{N},$$

in contradiction with the fact that $x_{k_\ell} \rightarrow x$. Therefore $x \in \tilde{U}_m$ for some m . Once again due to the fact that x is the limit of the sequence, the open set \tilde{U}_m must then contain infinitely many of the terms in the sequence $\{x_{k_\ell}\}_{\ell=1}^{\infty}$ but every set in the family $\{\tilde{U}_k\}$ could contain only finitely many terms in the sequence. We have arrived at a contradiction. Therefore, there must be a finite subcover for A , which ensures that it is compact.

2 \Rightarrow 3 We will argue by contraposition and prove that, if A is either not closed or unbounded then A can not be compact. In both cases we will build a cover for A that has no finite subcovers.

Part 1. Assume that A is not closed. Therefore there exists some $x \notin A$ that is an accumulation point of A . Then, for every $k \in \mathbb{N}$ there exists $x_k \in A \cap \overline{B_{1/k}(x)}$. Consider the sets:

$$U_k = \mathbb{R}^n \setminus \overline{B_{1/k}(x)}.$$

Each U_k is open because $\overline{B_{1/k}(x)}$ is closed and, since $x_k \rightarrow x$:

$$\bigcup_k U_k = \mathbb{R}^n \setminus \{x\}.$$

Since $x \notin A$, the collection $\{U_k : k \geq 1\}$ is a cover for A . However, for every finite m there we have that $1/m > 1/(m+1)$ and therefore $|x_{m+1} - x| < 1/(m+1) < 1/m$ and therefore $x_{m+1} \notin U_m$ which implies that no finite collection of U_k 's is a cover for A .

Part 2. Assume that A is not bounded and consider the collection $\{B_k(0) : k \in \mathbb{N}\}$. This collection is a cover for A because:

$$\bigcup_k B_k(0) = \mathbb{R}^n.$$

However, since A is unbounded, for every $k \in \mathbb{N}$, there exists $x_k \in A$ such that $x_k \notin B_k(0)$. Thus, no finite subset of this family can cover A and A is not compact.

3 \Rightarrow 1 Assume that A is bounded and let $\{x_k\}_{k=1}^{\infty}$ be a sequence of point in A . Since A is bounded, the sequence $\{x_k\}_{k=1}^{\infty}$ has a subsequence that converges, by the Bolzano-Weierstrass theorem. Since A is closed, the limit of this subsequence is in A . □

In plain English, the word “compact” suggests that something is not too large and also it is “dense” or “does not have holes”. As we shall see now, this notion carries over the the mathematical definition of compactness: a compact set must be bounded and must contain all its limit points (so there is no sequence of points in the set converging towards a missing point, i.e. “a hole”, in the set).

1.5 Exercises

1. Prove the first two parts of Proposition 1.1.

2. Prove the reverse triangle inequality: if $x, y \in \mathbb{R}^n$,

$$||x| - |y|| \leq |x - y|.$$

3. Prove the parallelogram identity: if $x, y \in \mathbb{R}^n$,

$$|x|^2 + |y|^2 = \frac{1}{2} (|x + y|^2 + |x - y|^2).$$

Explain how this identity relates to a parallelogram (Hint: use the geometric interpretation of vector addition and subtraction).

4. Let V be a subspace of \mathbb{R}^n and $x \in \mathbb{R}^n$. If $y_1, y_2 \in V$ are such that:

$$x - y_1 \perp z \quad \text{and} \quad x - y_2 \perp z \quad \text{for all } z \in V,$$

show that $y_1 = y_2$. (Hint: Compute $|y_1 - y_2|$.)

5. Show that, if $x_1, x_2 \in \mathbb{R}^n$, the set:

$$\{x \in \mathbb{R}^n : |x - x_1| = |x - x_2|\}$$

is a hyperplane.

6. Show that the intersection of two rectangles in \mathbb{R}^n is empty or is another rectangle.

7. Prove that the set of n -dimensional vectors with rational components

$$\mathbb{Q}^n := \{x = (x_1, \dots, x_n) : x_i \in \mathbb{Q} \text{ for all } 1 \leq i \leq n\}$$

is dense in \mathbb{R}^n . Namely, that for every $x \in \mathbb{R}^n$ and $\epsilon > 0$ there exists $q \in \mathbb{Q}^n$ such that $|x - q| < \epsilon$. You can use the fact that \mathbb{Q} is dense in \mathbb{R} .

8. Let (x_k) be a sequence in \mathbb{R}^n such that $x_k \rightarrow L$ and $x_k \rightarrow M$. Show that $L = M$.

9. Show that, if $\{x_k\}_{k=1}^{\infty}$ converges, then it is bounded.

10. Show that a sequence (x_k) is Cauchy in \mathbb{R}^n if and only if each sequence (x_k^i) is Cauchy in \mathbb{R} .

11. If (x_k) is a Cauchy sequence, then it is bounded.

12. Let (x_k) be a Cauchy sequence such that a subsequence converges, say $x_{k_l} \rightarrow L$. Show that $x_k \rightarrow L$.

13. Conclude, from the previous problems, that every Cauchy sequence in \mathbb{R}^n converges. (Use the Bolzano-Weierstrass theorem.)

14. Show that every infinite and bounded set in \mathbb{R}^n has an accumulation point.

15. Consider, in \mathbb{R} , the cover:

$$\left\{ \left(\frac{1}{2^n}, \frac{3}{2^n} \right) : n = 1, 2, \dots \right\}$$

of the set:

$$\left\{ 1, \frac{1}{2}, \frac{1}{3}, \dots \right\}.$$

Show that this cover has no finite subcovers.

16. Consider, in \mathbb{R}^n , the cover $\{A_n\}_n$,

$$A_n = \left\{ x \in \mathbb{R}^n : \frac{1}{2^n} < |x| < \frac{3}{2^n} \right\},$$

for the punctured ball $B_1^*(x) = \{x : 0 < |x| \leq 1\}$. Show that this cover has no finite subcovers.

17. Let $A_1 \supset A_2 \supset \cdots$ be compact non-empty subsets of \mathbb{R}^n . Show that:

$$\bigcap_i A_i \neq \emptyset.$$

18. Show that the previous statement is false if the A_i are only closed.

19. Show that, if $x \in \mathbb{R}^n$ and $E \subset \mathbb{R}^m$ is compact, then $\{x\} \times E$ is compact in \mathbb{R}^{n+m} .

20. Show that if $E \subset \mathbb{R}^n$ and $F \subset \mathbb{R}^m$ are compact, then $E \times F$ is compact in \mathbb{R}^{n+m} . (Hint: Use the previous problem.)

21. Let $E \subset F \subset \mathbb{R}^n$. Prove that if E is closed and F is compact, then E is compact.

Chapter 2

Functions of Several Variables

2.1 Basic definitions

We will consider functions $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$. Since each coordinate of $x = (x_1, \dots, x_n) \in A$ can be viewed as an independent variable of f , such functions are commonly referred to as functions of several variables. In fact, we write, for $x \in A$,

$$f(x) = f(x_1, x_2, \dots, x_n).$$

Given that $f(x) \in \mathbb{R}^m$, expressing the vector $f(x)$ in terms of its coordinates we get

$$f(x) = (f_1(x), f_2(x), \dots, f_m(x)),$$

where every $f^i(x)$ is referred to as a **component function**.

Recall that the **image** of f is the set

$$f(A) = \{f(x) \in \mathbb{R}^m : x \in A\},$$

and the **preimage** or **inverse image** of $B \subset \mathbb{R}^m$ under f is the set

$$f^{-1}(B) = \{x \in A : f(x) \in B\}.$$

Let's take $A \subset \mathbb{R}^n$ and $B \subset \mathbb{R}^m$. If $f : A \rightarrow B$ and $g : B \rightarrow \mathbb{R}^p$, then the composition $g \circ f : A \rightarrow \mathbb{R}^p$ is given by

$$(g \circ f)(x) = g(f(x)).$$

A function $f : A \rightarrow \mathbb{R}^m$ is said to be **injective** if $x \neq y$ implies that $f(x) \neq f(y)$, or equivalently, $f(x) = f(y)$ only if $x = y$.

Proposition 2.1. *If $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is injective, then there exists a function $f^{-1} : f(A) \subset \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that $f^{-1}(y) = x$ if and only if $f(x) = y$.*

Proof.

\Rightarrow Assume that f is injective and there exists a function $f^{-1} : f(A) \subset \mathbb{R}^m \rightarrow \mathbb{R}^n$ such that $f^{-1}(y) = x$. Then, since f is injective, we have that $y = f(f^{-1}(y)) = f(x)$, as desired.

\Leftarrow Assume that f is injective, and pick $y \in f(A)$ such that $y = f(x)$. This implies that there is at least one $x \in \mathbb{R}^n$ such that $f(x) = y$. However, the injectivity of f implies that this x is unique. Therefore for every $y \in f(A)$, we can define the inverse function univocally by the relation $f^{-1}(y) = x$, as desired. \square

When the inverse of a function exists, we say that f is **bijective** or **invertible**. Moreover, in this case, $f^{-1} \circ f : A \rightarrow A$ is the identity function on A and $f \circ f^{-1} : f(A) \rightarrow f(A)$ is the identity on $f(A)$.

The **projections** $\pi_i : \mathbb{R}^n \rightarrow \mathbb{R}$ are given by $\pi_i(x) = x_i$. Note that, for $f : A \rightarrow \mathbb{R}^m$,

$$f^i = \pi_i \circ f$$

for each $i = 1, \dots, m$.

If x_0 is an accumulation point of A , we say that the function $f : A \rightarrow \mathbb{R}^m$ has a **limit** at x_0 if there exists $L \in \mathbb{R}^m$ such that, for every $\epsilon > 0$, there exists $\delta > 0$ such that, if $0 < |x - x_0| < \delta$, then $|f(x) - L| < \epsilon$.

If this vector L exists, it is unique (Exercise 1) and is called the limit of f at x_0 . We write:

$$\lim_{x \rightarrow x_0} f(x) = L.$$

In the definition of the limit, we note that x_0 does not necessarily belong to A . If $x_0 \notin A$, then f is not defined at x_0 . In fact, even when f is defined at x_0 , it is possible that $f(x_0) \neq L$.

The relationship between the limit of a function and the limit of a sequence is given by the following proposition.

Proposition 2.2. *Let $f : A \rightarrow \mathbb{R}^m$ and x_0 be an accumulation point of A . Then:*

$$\lim_{x \rightarrow x_0} f(x) = L$$

if and only if, for every sequence (x_k) in A that converges to x_0 with $x_k \neq x_0$ for all k , the sequence $(f(x_k))$ in \mathbb{R}^m converges to L .

We leave the proof of this equivalence as an exercise (Exercise 2).

2.2 Continuity

Let $f : A \rightarrow \mathbb{R}^m$. We say that f is **continuous** at $x_0 \in A$ if, for every $\epsilon > 0$, there exists $\delta > 0$ such that, if

$$|x - x_0| < \delta \quad \text{then} \quad |f(x) - f(x_0)| < \epsilon.$$

The definition of continuity is local, meaning that it must be verified at each point. If a function $f : A \rightarrow \mathbb{R}^m$ is continuous at every point in its domain, we simply say that it is continuous. Similarly, if $B \subset A$, we say that f is continuous on B if it is continuous at each point $x \in B$.

The proof of the following proposition is not too complicated and will be left as an exercise (Exercise 3).

Proposition 2.3. *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuous if and only if each of its components $f_i(x)$ is continuous.*

Example 2.1 (Projections). The projections $\pi_i : \mathbb{R}^n \rightarrow \mathbb{R}$ are continuous: given $x \in \mathbb{R}^n$ and $\epsilon > 0$, if $\delta = \epsilon$ and $|x - x_0| < \delta$, then

$$|\pi_i(x) - \pi_i(x_0)| = |x^i - x_0^i| \leq \left(\sum_{i=1}^n (x^i - x_0^i)^2 \right)^{1/2} = |x - x_0| < \epsilon.$$

The proof of the following proposition is essentially identical to the equivalent statement for functions of a single variable. Alternatively, the proof can be done component-wise using Proposition 2.3.

Proposition 2.4. Let $f, g : A \rightarrow \mathbb{R}^m$ be continuous at $x_0 \in A$. Then:

1. $f + g$ is continuous at x_0 .
2. $\lambda f + \mu g$ is continuous at x_0 , for $\lambda, \mu \in \mathbb{R}$.
3. fg is continuous at x_0 .
4. If $m = 1$ and $g(x_0) \neq 0$, then f/g is defined in an open neighborhood around x_0 and is continuous at x_0 .

The composition of two continuous functions is continuous.

Proposition 2.5. Let $A \subset \mathbb{R}^n$, $B \subset \mathbb{R}^m$ and consider $f : A \rightarrow B$ and $g : B \rightarrow \mathbb{R}^p$. If f is continuous at $x_0 \in A$ and g is continuous at $f(x_0)$. Then $g \circ f : A \rightarrow \mathbb{R}^p$ is continuous at x_0 .

Proof. Given $\epsilon > 0$, let $\eta > 0$ be such that, if $|y - f(x_0)| < \eta$, then $|g(y) - g(f(x_0))| < \epsilon$. Such η exists because g is continuous at $f(x_0)$. Now, by the continuity of f at x_0 , there exists $\delta > 0$ such that, if $|x - x_0| < \delta$, then $|f(x) - f(x_0)| < \eta$.

Thus, if $|x - x_0| < \delta$, then $|f(x) - f(x_0)| < \eta$ and

$$|g \circ f(x) - g \circ f(x_0)| = |g(f(x)) - g(f(x_0))| < \epsilon.$$

Therefore, $g \circ f$ is continuous at x_0 . \square

The relationship between continuity and sequences is very important and is given by the following proposition, analogous to Proposition 2.2. Its proof follows an analogous argument and will also be left as an exercise (Exercise 8)

Proposition 2.6. Let $f : A \rightarrow \mathbb{R}^m$ and $x_0 \in A$. Then f is continuous at x_0 if and only if, for every sequence $\{x_k\}_{k=1}^{\infty} \subset A$ that converges to x_0 , we have that the sequence $f(x_k) \rightarrow f(x_0)$.

A function satisfying the second property in the proposition above is said to be **sequentially continuous**.

The Proposition 2.6 above, together with Proposition 2.2, implies that a function f continuous at an accumulation point x_0 must have the limit $f(x_0)$ at x_0 (Exercise 11).

The previous results refer to the continuity of a function locally (at a point). The following proposition, however, analyzes the global continuity of a function, that is, on its entire domain. To avoid some technical difficulties, **in the proposition below we will consider that A is the largest possible set where the function f can be defined**

Proposition 2.7. $f : A \rightarrow \mathbb{R}^m$ is continuous if and only if, for every open set $V \subset \mathbb{R}^m$, the inverse image $f^{-1}(V)$ is open.

Proof.

\Rightarrow Suppose f is continuous and let $V \subset \mathbb{R}^m$ be open.

If $f^{-1}(V) = \emptyset$, since \emptyset is open, we are done. Let us then consider that $f^{-1}(V) \neq \emptyset$, thus there exists some $x \in f^{-1}(V)$. Since V is open, there exists $\epsilon > 0$ such that

$$B_\epsilon(f(x)) \subset V.$$

Since f is continuous at x , there exists $\delta_x > 0$ such that, if $|x - y| < \delta_x$, then $|f(x) - f(y)| < \epsilon$, i.e., $f(B_{\delta_x}(x)) \subset B_\epsilon(f(x)) \subset V$. Hence, since $B_{\delta_x}(x) \subset A$ we have

$$f(B_{\delta_x}(x) \cap A) \subset B_\epsilon(f(x)) \subset V,$$

so that $B_{\delta_x}(x) \cap A \subset f^{-1}(V)$. Since this argument is valid for all points $x \in f^{-1}(V)$ it follows that $f^{-1}(V)$ is open.

◁ Now suppose that for every open set $V \subset \mathbb{R}^m$, the inverse image $f^{-1}(V)$ is open in \mathbb{R}^n .

Let $x \in A$ and $\epsilon > 0$, and consider the open set $V := B_\epsilon(f(x))$. Then $f^{-1}(V)$ is open. Since $x \in f^{-1}(V)$, there exists $\delta > 0$ such that $B_\delta(x) \subset f^{-1}(V)$. This implies that

$$B_\delta(x) \cap A \subset f^{-1}(V),$$

that is,

$$f(B_\delta(x) \cap A) \subset B_\epsilon(f(x)).$$

This means that, if $|x - y| < \delta$, then $|f(x) - f(y)| < \epsilon$, and therefore f is continuous at x . \square

The previous proposition provides us with a very useful criterion for understanding the topological properties of continuous functions.

Proposition 2.8. *If A is compact and $f : A \rightarrow \mathbb{R}^m$ is continuous, then $f(A)$ is compact.*

Proof. Let $\{V_\alpha\}$ be a cover for $f(A)$, so that

$$f(A) \subset \cup_\alpha V_\alpha.$$

By Proposition 2.7, since f is continuous for each α , the inverse image $f^{-1}(V_\alpha)$ is continuous. Therefore

$$A \subset f^{-1}(\cup_\alpha V_\alpha) = \cup_\alpha f^{-1}(V_\alpha).$$

The equality above can be easily proven from the definition of the inverse image (Exercise 5).

Then $\{f^{-1}(V_\alpha)\}$ is a cover for A , and since A is compact, it has a finite subcover, say $\{f^{-1}(V_i)\}_{i=1}^N$. Since

$$A \subset \cup_{i=1}^N f^{-1}(V_i),$$

then

$$f(A) \subset f(\cup_{i=1}^N f^{-1}(V_i)) = \cup_{i=1}^N V_i,$$

where the equality follows easily from the definition. We have then proved that $\{V_i\}_{i=1}^N$ is a finite subcover for $f(A)$ and $f(A)$ is compact. \square

We say that $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a **bounded function** if there exists $M > 0$ such that

$$|f(x)| \leq M|x| \quad \text{for all } x \in A$$

A function $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ for which there exists $M > 0$ such that

$$|f(x) - f(y)| \leq M|x - y|$$

is called a **Lipschitz function**, and the infimum of such M is called the Lipschitz constant of f . Depictions of Lipschitz and non-Lipschitz functions are shown in Figure 2.1.

Boundedness of a functions pertains how the function affects the size of its argument *compared to itself* (i.e. regardless of what the function may do to a neighbor): bounded functions *do not stretch* their arguments too much. On the other hand, the property of being Lipschitz does involve the neighbors of a point: a Lipschitz function does not take neighboring points too far apart. This sounds too similar to continuity...

Theorem 2.1. *If a function $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is Lipschitz on A , then it is continuous on A .*




Figure 2.1: Left: The function $f(x, y) = \sqrt{x^2 + y^2}$ is Lipschitz since the slope of any secant line connecting two points on its graph is uniformly bounded. Right: On the other hand, the function $f(x, y) = \sqrt{|x| + |y|}$ is not Lipschitz, since the slope of the secant lines connecting two points in the graph become unboundedly steep as the points approach the origin, forming a needle-like shape called a **cusp**.

Proof. Let $x \in \mathbb{R}^n$. Given $\epsilon > 0$, we take $\delta = \epsilon/M$, where M is the Lipschitz constant of f . Then, if $|x - y| < \delta$,

$$|f(x) - f(y)| \leq M|x - y| < M \frac{\epsilon}{M} = \epsilon.$$

□

For this reason, the property of being Lipschitz is sometimes referred to as **Lipschitz continuity**. In fact, this property is even a little stronger than continuity, as we shall see in the next section.

2.3 Linear functions

We say that $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a **linear function** if, for $x, y \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}$,

$$f(x + y) = f(x) + f(y) \quad \text{and} \quad f(\lambda x) = \lambda f(x).$$

In this section, we will review the basic properties of linear functions.

First, we must observe that if f is linear then necessarily $f(0) = 0$, since

$$f(0) = f(0x) = 0f(x) = 0.$$

Also, for any linear combination,

$$f(\alpha_1 x_1 + \alpha_2 x_2 + \cdots + \alpha_k x_k) = \alpha_1 f(x_1) + \alpha_2 f(x_2) + \cdots + \alpha_k f(x_k). \quad (2.1)$$

From equation 2.1, we can conclude that, for $x = (x_1, \dots, x_n) \in \mathbb{R}^n$,

$$f(x) = x_1 f(e_1) + x_2 f(e_2) + \cdots + x_n f(e_n),$$

so the values $f(e_1), f(e_2), \dots, f(e_n)$ define the function f on all of \mathbb{R}^n . In general, if u_1, u_2, \dots, u_n form a basis of \mathbb{R}^n , then the vectors $f(u_1), \dots, f(u_n)$ define f in \mathbb{R}^n .

Recall that $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ can be written in terms of its component functions as

$$f(x) = (f_1(x), \dots, f_m(x)).$$

Let us denote the action of the i -th component function on the j -th basis vector e_j as

$$a_{ij} := f_i(e_j), \quad (2.2)$$

and let's write the i -th component of the vector $f(x)$ as

$$f_i(x) = a_{i1}x_1 + a_{i2}x_2 + \cdots + a_{in}x_n. \quad (2.3)$$

If we gather all the constants appearing above in one vector we see that the expression above can be expressed as the dot product

$$f_i(x) = a_i \cdot x \quad \text{where} \quad a_i := (a_{i,1}, \dots, a_{i,n}). \quad (2.4)$$

Moreover, recalling the rules for matrix-vector multiplication, we also see that the expression (2.3) shows that the full vector $f(x)$ can be expressed in the form

$$\begin{pmatrix} f_1(x) \\ f_2(x) \\ \vdots \\ f_m(x) \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}.$$

So that **every linear function f is determined by the matrix multiplication**

$$f(x) = \mathbf{A}x$$

where the ij -th entry of the matrix \mathbf{A} is given by the action of the i -th component function f_i on the j -th basis vector e_j as defined in (3.3).

Consider a fixed vector $b \in \mathbb{R}^m$ and a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$. An **affine function** is a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ of the form

$$f(x) = \mathbf{A}x + b.$$

Theorem 2.2. *If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is linear, then*

1. f is bounded and,
2. f is Lipschitz.

Proof. We start by computing the norm

$$\begin{aligned} |f(x)| &= \sqrt{(f_1(x))^2 + \cdots + (f_m(x))^2} \\ &= \sqrt{(a_1 \cdot x)^2 + \cdots + (a_m \cdot x)^2} && \text{(By (2.4))} \\ &\leq \sqrt{(|a_1||x|)^2 + \cdots + (|a_m||x|)^2} && \text{(By Cauchy-Schwarz)} \\ &= |x| \sqrt{|a_1|^2 + \cdots + |a_m|^2}. \end{aligned}$$

Therefore, since the term multiplying $|x|$ in the last line above is a constant, we can define

$$M := \sqrt{|a_1|^2 + \cdots + |a_m|^2}$$

and boundedness has been proven.

To show that the function is Lipschitz, we now consider $x, y \in \mathbb{R}^n$ and compute

$$\begin{aligned} |f(x) - f(y)| &= |f(x - y)| && \text{(By linearity)} \\ &\leq M|x - y| && \text{(since the function is bounded).} \end{aligned}$$

□

The previous result, together with Theorem 2.1 imply the following:

Corollary 2.3. *If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is linear then f is continuous.*

2.4 Uniform continuity

We say that a function $f : A \rightarrow \mathbb{R}^m$ is **uniformly continuous** if, for every $\epsilon > 0$, there exists $\delta > 0$ such that $|x - y| < \delta$ implies

$$|f(x) - f(y)| < \epsilon, \quad \text{for all } x, y \in A.$$

The difference between a continuous function and a uniformly continuous function is that, in the latter case, for each $\epsilon > 0$, the number $\delta > 0$ in the definition of continuity is independent of the point where we want to verify continuity.

From the final comments of the previous section, we see that linear functions are uniformly continuous. In general, Lipschitz functions are uniformly continuous. The proof of this statements is left as an exercise (Exercise 12).

Example 2.2. Consider the function $f : \mathbb{R}^2 \setminus \{(0, 0)\} \rightarrow \mathbb{R}$ given by $f(x) = 1/|x|$. This function is continuous but, it is not uniformly continuous: for any $\delta > 0$, if $\delta \geq 1$, we take $x_1 = (1, 0)$ and $x_2 = (1/2, 0)$, then $|x_1 - x_2| < \delta$ and $|f(x_1) - f(x_2)| = 1$; whereas, if $\delta < 1$, taking $x_1 = (\delta, 0)$ and $x_2 = \delta/2$, we obtain $|x_1 - x_2| < \delta$ and

$$|f(x_1) - f(x_2)| = 1/\delta > 1.$$

In the previous example, f is not bounded in any neighborhood of 0. However, as we will show briefly, this can not happen if the function is uniformly continuous. In fact, something even stronger holds: a uniformly continuous function has a limit at every accumulation point of its domain. To prove this result we will make use of the following lemma, whose proof will be left as an exercise (Exercise 10).

Lemma 2.1. *Let $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ be uniformly continuous, and $\{x_k\} \subset A$ be a Cauchy sequence. Then the sequence of images $\{f(x_k)\}$ is Cauchy.*

Theorem 2.4. *Let $f : A \rightarrow \mathbb{R}^m$ be uniformly continuous and x_0 be an accumulation point of A . Then f has a limit at x_0 .*

Proof. Let $\epsilon > 0$ be given and x_0 be an accumulation point of A . We will consider two cases:

Case 1: $x_0 \in A$.

Since f is uniformly continuous, there is a fixed $\delta > 0$ such that, for all $x \in A$ such that $|x - x_0| < \delta$ it follows that $|f(x) - f(x_0)| < \epsilon$. This proves, by the definition of the limit of a function, that

$$\lim_{x \rightarrow x_0} f(x) = f(x_0).$$




Figure 2.2: Left: When the accumulation point $x_0 \in A$, it is possible to evaluate $f(x_0)$ and use the δ given by the uniform continuity of f to ensure that all the points in the ball $B_\delta(x_0)$ will be mapped by f to the interior of the ball $B_\epsilon(L)$, which proves that $L = \lim_{x \rightarrow x_0} f(x) = f(x_0)$. Right: When $x_0 \notin A$ it is not possible to evaluate f at x_0 . However, the fact that x_0 is an accumulation point of A allows us to find a point x_k arbitrarily close, where it is possible to evaluate $f(x_k)$. Shifting the ball from the previous argument to this new center, we can then use the δ given from the uniform continuity to ensure that points within a distance δ of x_0 will be mapped to the ball $B_\epsilon(L)$.

Case 1: $x_0 \notin A$.

The previous case argument is remarkably simple: uniform continuity tells us that it is enough to build a ball of radius δ around x_0 , and all the images of points in this ball will be within ϵ distance to $f(x_0)$. However, if $x_0 \notin A$, then $f(x_0)$ is not defined, and the previous argument does not hold. However, the geometric idea remains: since x_0 is an accumulation point of A there is a point in A arbitrarily close to x_0 where we can evaluate the function, and shifting the ball of radius δ to be centered around this point will also work just as before. We now formalize this idea.

Since x_0 is an accumulation point of A , we can pick a sequence $\{x_k\} \subset A$ such that $x_k \rightarrow x_0$. Since f is uniformly continuous, Lemma 2.1 guarantees that the sequence $\{f(x_k)\} \subset \mathbb{R}^m$ is also Cauchy. Given that the Euclidean space is complete, we then know that there exists some $L \in \mathbb{R}^m$ such that $f(x_k) \rightarrow L$. We will show that this L is indeed the limit for the function as $x \rightarrow x_0$ by showing that, for some point in the sequence $\{x_k\}$ that is "close enough" to x_0 , there is a ball $B_\delta(x_k)$ that gets mapped into the ball $B_\epsilon(L) \subset \mathbb{R}^m$.

Since f is uniformly continuous, there exists $\delta > 0$ such that if $|x - y| < \delta$ then $|f(x) - f(y)| < \epsilon/2$. We then go back to the sequence $\{x_k\}$ and fix K such that, for $k \geq K$,

$$|x_k - x_0| < \delta/2 \quad \text{and} \quad |f(x_k) - L| < \epsilon/2.$$

Then, if $|x - x_0| < \delta/2$,

$$|x - x_k| \leq |x - x_0| + |x_0 - x_k| < \delta/2 + \delta/2 = \delta,$$

and thus $|f(x) - f(x_k)| < \epsilon/2$. Therefore, if $0 < |x - x_0| < \delta/2$, we have

$$|f(x) - L| < |f(x) - f(x_k)| + |f(x_k) - L| < \epsilon/2 + \epsilon/2 = \epsilon,$$

which proves that

$$\lim_{x \rightarrow x_0} f(x) = L.$$

□

Theorem 2.4 above provides a necessary condition on a set A to guarantee that a continuous function defined on A will be uniformly continuous: A must be closed. In fact, if x_0 is an accumulation point of A that does not belong to A , then the function

$$x \mapsto \frac{1}{|x - x_0|}$$

is a continuous function that cannot be uniformly continuous (Exercise 15). If, in addition, the set is bounded (i.e., A is compact), then every continuous function will be uniformly continuous, as we will now show.

Theorem 2.5. *If $f : A \rightarrow \mathbb{R}^m$ is continuous and A is compact, then f is uniformly continuous.*

Proof. Let $\epsilon > 0$ be given. For each $x \in A$, let $\delta_x > 0$ be such that $|y - x| < \delta_x$ implies $|f(y) - f(x)| < \epsilon/2$. The collection of open balls

$$\{B_{\delta_x/2}(x) : x \in A\}$$

is a cover of A . Since A is compact, there exist x_1, \dots, x_k such that

$$A \subset B_{\delta_{x_1}/2}(x_1) \cup \dots \cup B_{\delta_{x_k}/2}(x_k).$$

We will show that $\delta = \frac{1}{2} \min\{\delta_{x_1}, \dots, \delta_{x_k}\}$ satisfies the definition of uniform continuity. Let $x, y \in A$ such that $|x - y| < \delta$. If i is such that $x \in B_{\delta_{x_i}/2}(x_i)$, then $|f(x) - f(x_i)| < \epsilon/2$. Now,

$$|y - x_i| \leq |y - x| + |x - x_i| < \delta + \frac{\delta_{x_i}}{2} \leq \frac{\delta_{x_i}}{2} + \frac{\delta_{x_i}}{2} = \delta_{x_i},$$

and thus $|f(y) - f(x_i)| < \epsilon/2$. Then,

$$|f(x) - f(y)| \leq |f(x) - f(x_i)| + |f(x_i) - f(y)| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

□

In particular, Theorem 2.18 guarantees that a continuous function on a closed rectangle is uniformly continuous. This result will be very important later when we study the integral of functions defined on rectangles in \mathbb{R}^n .

The converse of Theorem 2.18 is false (Exercise 16).

2.5 Exercises

1. Show that if $f : A \rightarrow \mathbb{R}^m$ has limits L and M at x_0 , then $L = M$.
2. Prove Proposition 2.2.
3. Show that the function $f : A \rightarrow \mathbb{R}^m$ is continuous at $x \in A$ if and only if each of its components $f^i : A \rightarrow \mathbb{R}$ is continuous at x .
4. Consider the function in \mathbb{R}^2 defined by

$$f(x, y) = \begin{cases} \frac{xy}{x^2+y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0). \end{cases}$$

- (a) Show that each of the functions

$$x \mapsto f(x, y_0) \quad \text{and} \quad y \mapsto f(x_0, y)$$

is continuous in \mathbb{R} for any fixed $x_0, y_0 \in \mathbb{R}$.

- (b) Show that, despite the previous part, the function f is not continuous at $(0, 0)$.

5. Consider a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and a family of sets $\{V_\alpha\} \subset \mathbb{R}^m$. Use the definition of the inverse image of a set to prove that
- (a) $\cup_\alpha f^{-1}(V_\alpha) = f^{-1}(\cup_\alpha V_\alpha)$.
- (b) $\cap_\alpha f^{-1}(V_\alpha) = f^{-1}(\cap_\alpha V_\alpha)$.
6. Let $f : A \rightarrow \mathbb{R}^m$ be continuous at $x_0 \in A$ such that $f(x_0) \neq 0$. Then there exists $\alpha > 0$ and an open set $U \subset \mathbb{R}^n$ such that $x_0 \in U$ and $|f(x)| > \alpha$ for all $x \in U \cap A$.

7. Prove Proposition 2.4 (you will need the previous exercise to prove property 4).

8. Prove Proposition 2.6.

9. Let $f : A \rightarrow \mathbb{R}^m$ be continuous. Show that the function $|f| : A \rightarrow \mathbb{R}$ given by $|f|(x) = |f(x)|$ is continuous.

10. Let $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ be uniformly continuous, and $\{x_k\} \subset A$ be a Cauchy sequence. Prove that the sequence of images $\{f(x_k)\}$ is Cauchy.

11. Let $f : A \rightarrow \mathbb{R}^m$ and x_0 be an accumulation point of A . Show that f is continuous at x_0 if and only if

$$\lim_{x \rightarrow x_0} f(x) = f(x_0).$$

12. Show that a Lipschitz function is uniformly continuous.

13. Let $E \subset \mathbb{R}^n$ be compact and $f : E \rightarrow \mathbb{R}^m$ be continuous. Show that there exist x_1, \dots, x_n and y_1, \dots, y_n in E such that

$$f^i(x_i) = \max\{f^i(x) : x \in E\} \quad \text{and} \quad f^i(y_i) = \min\{f^i(x) : x \in E\}.$$

That is, each of the components of f attains its maximum and minimum in E .

14. Similarly to the previous problem, show that if $E \subset \mathbb{R}^n$ is compact and $f : E \rightarrow \mathbb{R}^m$ is continuous, then there exist $x', x'' \in E$ such that

$$|f(x')| = \max\{|f(x)| : x \in E\} \quad \text{and} \quad |f(x'')| = \min\{|f(x)| : x \in E\}.$$

15. Let $A \subset \mathbb{R}^n$ and $x_0 \notin A$ be an accumulation point of A . Show that the function $f : A \rightarrow \mathbb{R}$ given by

$$f(x) = \frac{1}{|x - x_0|}$$

is continuous but not uniformly continuous.

16. Give an example of an unbounded set $A \subset \mathbb{R}^n$ such that every continuous function on A is uniformly continuous.

Chapter 3

Differentiability

3.1 Introduction

In this section, we will define the derivative of a function at a point. Let us recall that, for a *real-valued function of a real argument* $f : \mathbb{R} \rightarrow \mathbb{R}$, the derivative of f at a point $x_0 \in \mathbb{R}$ is given, if it exists, by the limit

$$f'(x_0) = \lim_{x \rightarrow x_0} \frac{f(x) - f(x_0)}{x - x_0}.$$

If we attempt to mimic the definition above for *vector-valued functions of a vector variable* $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we will face several difficulties. First of all, the objects appearing in the Newton quotient are vectors in potentially different spaces:

$$f(x) - f(x_0) \in \mathbb{R}^m \quad \text{while} \quad x - x_0 \in \mathbb{R}^n.$$

Therefore, arithmetic operations between them may not make any sense. Moreover, even if $m = n > 1$, the division of two vectors is not defined. A natural attempt for solving this issue would be to propose using the norm $|x - x_0|$ in the denominator. However, the original definition is sensitive to the sign of the difference $x - x_0$ and replacing this by the norm would interfere with this. In particular, using the norm $|x - x_0|$ would result in an object that is *different* from the derivative if $m = n = 1$, as clearly in general

$$\lim_{x \rightarrow x_0} \frac{f(x) - f(x_0)}{x - x_0} \neq \lim_{x \rightarrow x_0} \frac{f(x) - f(x_0)}{|x - x_0|}.$$

Another reasonable-but-unsuccessful attempt would be to notice that, in the alternative the alternative (but equivalent) definition of the derivative of $f : \mathbb{R} \rightarrow \mathbb{R}$

$$f'(x_0) = \lim_{h \rightarrow 0} \frac{f(x_0 + h) - f(x_0)}{h},$$

the argument $x_0 + h$ can be interpreted as a perturbation of size h around the point x_0 . Thus the derivative measures the reaction of a function to infinitesimally small perturbations of its argument. With this in mind, we could take a point $x_0 \in \mathbb{R}^n$, a vector $x \in \mathbb{R}^n$ and a real number $h \in \mathbb{R}$, and consider small perturbations of the form $x_0 + hx$. This would lead to a definition involving the term

$$\lim_{h \rightarrow 0} \frac{f(x_0 + hx) - f(x_0)}{h}. \tag{3.1}$$

The term above is a *much* better approach: it avoids the issue of “division” by vectors and the absolute value problem; it even recovers the “correct” definition of derivative if $m = n = 1$. It, however, presents the

limitation when $n > 1$, of measuring only the change in the *particular direction* of x . In fact, whenever the limit (3.1) exists, it is known as the **directional derivative** of f at x_0 in the direction of x (we will come back to this concept a little later). However, we would still like to obtain an object that gives us information about how the function f reacts to perturbations in **any** direction. To obtain that, we will have to resort to a approximation argument.

Let us first introduce some **asymptotic notation**. Consider a function $f : \mathbb{R} \rightarrow \mathbb{R}$. We will say that

1. f is **big O** of x , and write $f(x) = O(x)$ if

$$\lim_{x \rightarrow 0} \frac{f(x)}{x} = C \neq 0$$

for some non zero constant $C \in \mathbb{R}$.

2. f is **little o** of x , and write $f(x) = o(x)$ if

$$\lim_{x \rightarrow 0} \frac{f(x)}{x} = 0.$$

We remark that the equalities of the form $f(x) = O(x)$ or $f(x) = o(x)$ **are not statements about the equality between two functions**. Instead, they are statements about the behavior of the function f as its arguments approach zero.

The following arithmetic properties are easy to verify (Exercise 1):

Proposition 3.1. Consider distinct functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$. The following hold

1. If $f(x) = O(x)$ and $g(x) = O(x)$, then $f(x) \pm g(x) = O(x)$.
2. If $f(x) = O(x)$ and $g(x) = o(x)$, then $f(x) \pm g(x) = O(x)$.
3. $f(x) = o(x)$ and $g(x) = o(x)$, then $f(x) \pm g(x) = o(x)$.
4. $f(x) = O(x)$ and $g(x) = O(x)$, then $f(x)g(x) = o(x)$.
5. $f(x) = O(x)$ and $g(x) = o(x)$, then $f(x)g(x) = o(x)$.

3.2 The derivative

Lets consider real-valued functions functions of a single variable. If $f'(x_0)$ is the derivative of f at x_0 , then

$$f(x) \approx f(x_0) + f'(x_0)(x - x_0)$$

is a linear approximation of f near x_0 ; in fact, it is a very good approximation. This follows from the definition of the derivative as a limit: given any $\epsilon > 0$, there exists $\delta > 0$ such that, if

$$0 < |x - x_0| < \delta \quad \text{then} \quad \left| \frac{f(x) - f(x_0)}{x - x_0} - f'(x_0) \right| < \epsilon,$$

which implies that for any $\epsilon > 0$

$$0 \leq \lim_{x \rightarrow x_0} \frac{|f(x) - f(x_0) - f'(x_0)(x - x_0)|}{|x - x_0|} < \epsilon.$$

By the ϵ -principle, this implies that the limit is zero. This fact can be expressed in asymptotic notation compactly as

$$f(x) = f(x_0) + f'(x_0)(x - x_0) + o(|x - x_0|).$$

This observation motivates us to define the differentiability of a function in \mathbb{R}^n as follows.

Definition 3.1. Let U be an open set in \mathbb{R}^n and $f : U \rightarrow \mathbb{R}^m$. We say that f is differentiable at $x_0 \in U$ if there exists a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

$$f(x) = f(x_0) + T(x - x_0) + o(|x - x_0|). \quad (3.2)$$

Remark 3.1. We insist that (3.2) should be understood as a limiting statement as $x \rightarrow x_0$. The value of the function f can be approximated by the affine function

$$f(x_0) + T(x - x_0)$$

at the expense of a small error $o(|x - x_0|)$ that vanishes faster than the distance between x and x_0 . The definition 3.1 can be expressed in several equivalent ways:

1. The function f is differentiable at x_0 if there exists a linear transformation T such that, for every $\epsilon > 0$, there exists $\delta > 0$ such that if $|x - x_0| < \delta$, $x \in U$, then

$$|f(x) - f(x_0) - T(x - x_0)| < \epsilon|x - x_0|.$$

2. The function f is differentiable at x_0 if and only if there exists a linear transformation T such that

$$\lim_{x \rightarrow x_0} \frac{|f(x) - f(x_0) - T(x - x_0)|}{|x - x_0|} = 0,$$

3. The function f is differentiable at x_0 if and only if there exists a linear transformation T such that

$$\lim_{h \rightarrow 0} \frac{|f(x + h) - f(x) - Th|}{|h|} = 0. \quad (3.3)$$

We shall now prove that, whenever this transformation exists, it is unique. We will first prove an auxiliary result that establishes the intuitive fact that if two linear mappings coincide on a neighborhood, they must be equal everywhere.

Lemma 3.1. Let $T, S : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be linear transformations such that

$$T(x) = S(x) \quad \text{for all} \quad |x| = 1.$$

Then $T(x) = S(x)$ for all $x \in \mathbb{R}^n$.

Proof. If the argument $x = 0$, the linearity of the mappings ensures that $T(x) = 0 = S(x)$. Thus, we will assume that $x \neq 0$ but is otherwise arbitrary and observe that

$$T(x) = T\left(|x| \frac{x}{|x|}\right) = |x|T\left(\frac{x}{|x|}\right) \stackrel{\text{Since } \left|\frac{x}{|x|}\right|=1}{=} |x|S\left(\frac{x}{|x|}\right) = S\left(|x| \frac{x}{|x|}\right) = S(x).$$

□

We can now prove the following

Theorem 3.1. *If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at x_0 , then the linear function T in the equation (3.2) is unique.*

Proof. Suppose that $T, S : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are linear functions both satisfying (3.2), namely

$$f(x_0) + S(x - x_0) + o(|x - x_0|) = f(x) = f(x_0) + T(x - x_0) + o(|x - x_0|).$$

This implies that

$$(T - S)(x - x_0) = o(|x - x_0|).$$

Clearly, if $x = x_0$ it follows that $T(x - x_0) = 0 = S(x - x_0)$, so we assume that $x \neq x_0$, define

$$t := |x - x_0| \quad y := t \frac{x - x_0}{|x - x_0|},$$

and express the equality above as

$$(T - S)(ty) = o(t).$$

The expression above implies that

$$0 = \lim_{t \rightarrow 0} \frac{(T - S)(ty)}{t} = \lim_{t \rightarrow 0} \frac{t(T - S)(y)}{t} = \lim_{t \rightarrow 0} (T - S)(y).$$

Since the rightmost expression does not depend on t this implies that

$$(T - S)(y) = 0.$$

However, y is an arbitrary unit vector, therefore we can use Lemma 3.1 to assert that $T(x) = S(x)$ for all $x \in \mathbb{R}^n$. \square

The linear function T in (3.2) is called the **derivative** of f at x_0 and is denoted by $Df(x_0)$. Using this notation we can rewrite equation (3.2) as

$$f(x) = f(x_0) + \underbrace{Df(x_0)}_{:=T}(x - x_0) + o(|x - x_0|).$$

Note that *the entire expression* $Df(x_0)$ denotes a linear function that then takes a vector $x \in \mathbb{R}^n$ as an argument and produces the vector $Df(x_0)(x) \in \mathbb{R}^m$ as a result.

As we established in equation (2.3), any linear transformation f can be represented as a matrix whose entries are the result of applying the component functions of f to the basis vectors of \mathbb{R}^n . The particular values appearing in the matrix will depend on the basis that is used. The matrix induced by the transformation $Df(x_0)$ in the standard basis of \mathbb{R}^n is called the **Jacobian**, and is denoted by $Jf(x_0)$. When the matrix representation of $Df(x_0)$ is not written using the canonical basis, we will denote the matrix by $f'(x_0)$.

We will now introduce the derivatives of a few functions.

Proposition 3.2.

1. If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is constant, then

$$Df(x_0) = 0$$

for each $x_0 \in \mathbb{R}^n$. Note that **this is not the number zero**, but the linear transformation that assigns the vector $0 \in \mathbb{R}^m$ to every vector $x \in \mathbb{R}^n$. This transformation is represented, in any basis, by a matrix with n rows and m columns all of whose entries are zero.

2. If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is linear, then

$$Df(x_0) = f$$

for each $x_0 \in \mathbb{R}^n$. Therefore, if F is the matrix representing f in the canonical basis, $Jf(x_0) = F$.

3. If we define the function $\text{sum} : \mathbb{R}^n \rightarrow \mathbb{R}$ by $\text{sum}(x) := \sum_{i=1}^n x_i$, then

$$D\text{sum}(x_0) = \text{sum}$$

for each $x_0 \in \mathbb{R}^n$. Therefore, $Jf(x_0) = (1, \dots, 1)$. This is a matrix with 1 row and n columns, i.e. a row vector.

4. For $n > 1$, we define the function $\text{mult} : \mathbb{R}^n \rightarrow \mathbb{R}$ by $\text{mult}(x) := \prod_{i=1}^n x_i = x_1 \cdot x_2 \cdot \dots \cdot x_n$, then

$$D\text{mult}(x_0)(x) = \sum_{i=1}^n \left(\prod_{j \neq i} x_0^j \right) x^i$$

for each $x_0 \in \mathbb{R}^n$. Therefore, $Jf(x_0) = \left(\prod_{j \neq 1} x_0^j, \dots, \prod_{j \neq i} x_0^j, \dots, \prod_{j \neq n} x_0^j \right)$. This is a matrix with 1 row and n columns, i.e. a row vector.

Proof. Our proofs will not be constructive, at this stage we will simply show that the proposed transformation satisfy the property (3.3) from the definition, and therefore from the uniqueness guaranteed by Lemma 3.1, they are the derivative of each of the functions.

1. If f is constant,

$$\frac{|f(x_0 + h) - f(x_0) - 0|}{|h|} = 0$$

for all $x_0, h \in \mathbb{R}^n, h \neq 0$. And therefore (3.3) is verified with $Df(x_0) = 0$.

2. If f is linear, then $f(x_0 + h) = f(x_0) + f(h)$, so that

$$\frac{|f(x_0 + h) - f(x_0) - f(h)|}{|h|} = 0$$

for all $x_0, h \in \mathbb{R}^n, h \neq 0$.

3. We start by noting that the function sum is linear. Indeed, for $x, y \in \mathbb{R}^n$ and $\alpha, \beta \in \mathbb{R}$ we have

$$\begin{aligned} \text{sum}(\alpha x + \beta y) &= \text{sum}((\alpha x_1, \dots, \alpha x_n) + (\beta y_1, \dots, \beta y_n)) \\ &= \text{sum}(\alpha x_1 + \beta y_1, \dots, \alpha x_n + \beta y_n) \\ &= \sum_{i=1}^n (\alpha x_i + \beta y_i) \\ &= \alpha \sum_{i=1}^n x_i + \beta \sum_{i=1}^n y_i \\ &= \alpha \text{sum}(x) + \beta \text{sum}(y). \end{aligned}$$

We can therefore apply the previous point to conclude that $D\text{sum} = \text{sum}$. Therefore, when applied to a vector x we have that

$$D\text{sum}(x) = \text{sum}(x) = x_1 + \dots + x_n = (1, \dots, 1) \cdot (x_1, \dots, x_n) = (1, \dots, 1)(x_1, \dots, x_n)^\top,$$

where the superscript “ \top ” denotes matrix transposition (recall that we consider vectors to be columns). From the uniqueness of the derivative we then conclude that

$$J_{\text{sum}} = (1, \dots, 1).$$

4. We first observe that

$$\text{mult}(x_0 + h) = \begin{cases} \text{mult}(x_0) + \text{mult}(h) + \sum_{i=1}^n \left(\prod_{j \neq i} x_0^j \right) h^i & \text{if } n = 2 \\ \text{mult}(x_0) + \text{mult}(h) + \sum_{i=1}^n \left(\prod_{j \neq i} x_0^j \right) h^i + \sum_{i=1}^n \left(\prod_{j \neq i} h^j \right) x_0^i & \text{if } n > 2 \end{cases}.$$

Therefore we have

$$\begin{aligned} \frac{|\text{mult}(x_0 + h) - \text{mult}(x_0) - \sum_{i=1}^n \left(\prod_{j \neq i} x_0^j \right) h^i|}{|h|} &\leq \begin{cases} \frac{|\prod_{i=1}^n h^i|}{|h|} & \text{if } n = 2, \\ \frac{|\prod_{i=1}^n h^i|}{|h|} + \frac{|\sum_{i=1}^n \left(\prod_{j \neq i} h^j \right) x_0^i|}{|h|} & \text{if } n > 2. \end{cases} \end{aligned} \quad (3.4)$$

Analyzing separately the last two terms in the inequality above we have that

$$\frac{|\prod_{i=1}^n h^i|}{|h|} = \underbrace{\frac{(\prod_{i=1}^n |h^i|^2)^{1/2}}{|h|}}_{\text{By the geometric-arithmetic mean inequality}} \leq \left(\frac{\sum_{i=1}^n (h^i)^2}{n|h|} \right)^{n/2} \leq \left(\frac{1}{n} \right)^{n/2} |h|^{n/2},$$

and also

$$\frac{|\sum_{i=1}^n \left(\prod_{j \neq i} h^j \right) x_0^i|}{|h|} \leq \underbrace{\frac{\sum_{i=1}^n \left(\prod_{j \neq i} |h^j| \right) |x_0^i|}{|h|}}_{\text{Since } |z^i| \leq |z| \forall z \in \mathbb{R}^n} \leq \frac{\sum_{i=1}^n |h|^{n-1} |x_0^i|}{|h|} \leq n |h|^{n-2} |x_0|.$$

And therefore the right hand side of the inequality (3.17) vanishes as $h \rightarrow 0$ which, by the uniqueness of the derivative proves the result.

Finally, applying D_{mult} to a vector $x \in \mathbb{R}^n$ we observe that

$$\begin{aligned} D_{\text{mult}}(x_0)(x) &= \sum_{i=1}^n \left(\prod_{j \neq i} x_0^j \right) x^i \\ &= \left(\prod_{i=2}^n x_0^i, \dots, \prod_{j \neq i} x_0^j, \dots, \prod_{i=1}^{n-1} x_0^i \right) \cdot (x^1, \dots, x^n) \\ &= \left(\prod_{i=2}^n x_0^i, \dots, \prod_{j \neq i} x_0^j, \dots, \prod_{i=1}^{n-1} x_0^i \right) (x^1, \dots, x^n)^\top \end{aligned}$$

proving that

$$J_{\text{mult}}(x_0) = \left(\prod_{i=2}^n x_0^i, \dots, \prod_{j \neq i} x_0^j, \dots, \prod_{i=1}^{n-1} x_0^i \right).$$

□

Just as in the single-variable case, differentiability is a stronger condition than continuity, as established by the following proposition, whose proof will be left as an Exercise (3.3).

Proposition 3.3. *If $f : U \rightarrow \mathbb{R}^m$ is differentiable at $x_0 \in U$, then it is continuous at x_0 .*

In Proposition 2.3 we established that a function f is continuous if and only if each of its components f^i is continuous. We have an analogous result for differentiability.

Proposition 3.4. *Let $f : U \rightarrow \mathbb{R}^m$ and $x_0 \in U$. Then f is differentiable at x_0 if and only if each $f^i : U \rightarrow \mathbb{R}$, $i = 1, \dots, m$, is differentiable at x_0 . In that case,*

$$Df(x_0)(x) = (Df^1(x_0)(x), \dots, Df^m(x_0)(x)).$$

That is, the components of the derivative of f are precisely the derivatives of the components of f .

Proof. First, suppose that f is differentiable at x_0 and, for each i , let $T^i = \pi_i \circ Df(x_0)$. Then,

$$\frac{|f^i(x_0 + h) - f^i(x_0) - T^i h|}{|h|} \leq \frac{|f(x_0 + h) - f(x_0) - Df(x_0)(h)|}{|h|} \rightarrow 0$$

as $h \rightarrow 0$, so that each f^i is differentiable at x_0 and

$$Df^i(x_0) = \pi_i \circ Df(x_0).$$

Now suppose that each f^i is differentiable at x_0 and let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be the transformation

$$Tx = (Df^1(x_0)(x), \dots, Df^m(x_0)(x)).$$

Given $\epsilon > 0$, we choose $\delta > 0$ such that, if $0 < |h| < \delta$, each

$$\frac{|f^i(x_0 + h) - f^i(x_0) - Df^i(x_0)(h)|}{|h|} < \frac{\epsilon}{\sqrt{m}}.$$

Then, if $0 < |h| < \delta$,

$$\frac{|f(x_0 + h) - f(x_0) - Th|}{|h|} = \frac{1}{|h|} \sqrt{\sum_{i=1}^m (f^i(x_0 + h) - f^i(x_0) - Df^i(x_0)(h))^2} < \epsilon.$$

Thus, f is differentiable at x_0 and $Df(x_0) = T$.

□

Theorem 3.2 (Chain Rule). *Let $U \subset \mathbb{R}^n$ and $V \subset \mathbb{R}^m$ be open sets, $f : U \rightarrow V$ be differentiable at $x_0 \in U$, and $g : V \rightarrow \mathbb{R}^p$ be differentiable at $f(x_0) \in V$. Then, $g \circ f : U \rightarrow \mathbb{R}^p$ is differentiable at x_0 , and*

$$D(g \circ f)(x_0) = Dg(f(x_0)) \circ Df(x_0).$$

That is, the derivative of the composition of two functions is given by the composition of their derivatives.

Proof. Define $y_0 := f(x_0)$, $T := Df(x_0)$, and $S := Dg(y_0)$. We will show that $D(g \circ f)(x_0) = S \circ T$ by verifying the approximation property (3.2). Namely, that

$$g \circ f(x) - g(y_0) - S \circ T(x - x_0) = o(|x - x_0|).$$

Define the auxiliary functions

$$\phi(x) := f(x) - f(x_0) - T(x - x_0), \quad (3.5a)$$

$$\psi(y) := g(y) - g(y_0) - S(y - y_0), \quad (3.5b)$$

$$\rho(x) := g \circ f(x) - g(y_0) - S \circ T(x - x_0). \quad (3.5c)$$

Since f and g are differentiable at x_0 and y_0 we have that

$$\phi(x) = o(|x - x_0|), \quad \text{and} \quad \psi(y) = o(|y - y_0|), \quad (3.6)$$

and we want to prove that

$$\rho(x) = o(|x - x_0|).$$

First, observe that

$$\begin{aligned} \rho(x) &= g(f(x)) - g(y_0) - S(T(x - x_0)) \\ &= g(f(x)) - g(y_0) - S(f(x) - y_0) + S(\phi(x)) \\ &= \psi(f(x)) + S\phi(x). \end{aligned}$$

We will consider each of the two terms above separately. It will be useful to recall that every linear function is bounded (Theorem 2.2), therefore there exist positive constants M_S, M_T such that

$$|Sy| \leq M_S|y| \quad \forall y \in \mathbb{R}^m \quad \text{and} \quad |Tx| \leq M_T|x| \quad \forall x \in \mathbb{R}^n.$$

From the first equation in (3.6) we see that, given $\epsilon > 0$ there exists $\delta_f > 0$ such that if $|x - x_0| < \delta_f$ we have $|\phi(x)| \leq \epsilon|x - x_0|/M_S$. From this, it follows that

$$|S\phi(x)| \leq M_S|\phi(x)| \leq M_S\epsilon|x - x_0|/M_S = \epsilon|x - x_0|,$$

which implies that $S\phi(x) = o(|x - x_0|)$.

For the other term, we observe that the second equality in (3.6) implies that there exists $\delta_g > 0$ such that, as long as $|y - y_0| < \delta_g$, we have that

$$|\psi(y)| < \frac{\epsilon}{\max\{\epsilon, M_T\}}|y - y_0|.$$

The need for the somewhat strange denominator $\max\{\epsilon, M_T\}$ will become apparent soon. The important

part is that we can bound $|\psi(y)|$ by a multiple of $\epsilon|y - y_0|$. Hence letting $|x - x_0| < \min\{\delta_f, \delta_g\}$

$$\begin{aligned}
|\psi(f(x))| &\leq \frac{\epsilon}{\max\{\epsilon, M_T\}} |y - y_0| \\
&= \frac{\epsilon}{\max\{\epsilon, M_T\}} |f(x) - f(x_0)| \\
&= \frac{\epsilon}{\max\{\epsilon, M_T\}} |\phi(x) + T(x - x_0)| && \text{(From (3.5a))} \\
&\leq \frac{\epsilon}{\max\{\epsilon, M_T\}} (|\phi(x)| + |T(x - x_0)|) \\
&\leq \frac{\epsilon}{\max\{\epsilon, M_T\}} (\epsilon|x - x_0| + M_T|x - x_0|) && \text{(From (3.6))} \\
&\leq \frac{\epsilon}{\max\{\epsilon, M_T\}} \max\{\epsilon, M_T\} |x - x_0| \\
&= \epsilon|x - x_0|
\end{aligned}$$

and therefore $\psi(f(x)) = o(|x - x_0|)$. In view of Proposition 3.1, the previous discussion proves that

$$\rho(x) = \psi(f(x)) + S\phi(x) = o(|x - x_0|),$$

and therefore $g \circ f : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is differentiable and by the uniqueness of the derivative it follows that

$$D(f \circ g)(x_0) = S \circ T(x_0) = Dg(f(x_0)) \circ Df(x_0).$$

□

3.3 Directional and partial derivatives

So far, we have only defined the notion of differentiability, introduced the derivative as the unique linear transformation that satisfies the approximation property (3.2), proved some properties about this transformation, but we have not defined how to actually compute the derivative. In this section, we will finally study how to compute derivatives by exploiting the differentiability of a function through its component-wise differentiability. We begin by considering real-valued functions.

Definition 3.2. Let $U \subset \mathbb{R}^n$ be an open set, $f : U \rightarrow \mathbb{R}^m$, $x_0 \in U$, and $u \in \mathbb{R}^n$ a unit vector. If the limit

$$\lim_{t \rightarrow 0} \frac{f(x_0 + tu) - f(x_0)}{t}, \quad t \in \mathbb{R}, \quad (3.7)$$

exists, we call it the **directional derivative** of f in the direction u , denoted by $D_u f(x_0)$.

The derivative of a function, in fact, contains the information of all its directional derivatives. To extract it, we must evaluate the derivative by providing the desired direction as an argument, as we now prove.

Theorem 3.3. *If $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is differentiable at $x_0 \in \mathbb{R}^n$, then its directional derivatives $D_u f(x_0)$ are defined for every direction u and*

$$D_u f(x_0) = Df(x_0)(u).$$

Proof. Since f is given to be differentiable at x_0 , the linear mapping $Df(x_0) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is well defined. All we need to do is to verify that the limit (3.7) exists for any u and that it is equal to $Df(x_0)(u)$. To that avail,

we compute

$$\begin{aligned}
 D_u f(x_0) &= \lim_{t \rightarrow 0} \frac{f(x_0 + tu) - f(x_0)}{t} \\
 &= \lim_{t \rightarrow 0} \frac{Df(x_0)(tu) + o(|x - x_0|)}{t} && \text{(By (3.2))} \\
 &= \lim_{t \rightarrow 0} \left(\frac{tDf(x_0)(u)}{t} + \frac{o(|x - x_0|)}{t} \right) && \text{(The derivative is linear)} \\
 &= Df(x_0)(u) + \lim_{t \rightarrow 0} \frac{o(|x - x_0|)}{t} \\
 &= Df(x_0)(u) && \text{(Definition of } o(|x - x_0|)\text{)}.
 \end{aligned}$$

This proves the result. □

Definition 3.3. Lets consider for now *real-valued functions*, i.e. $f : \mathbb{R}^n \rightarrow \mathbb{R}$. In this case, if in the limit (3.7) appearing in the definition directional derivatives, we let u be one of the canonical basis vectors e_i , then $D_{e_i} f(x_0)$ is called the i -th **partial derivative** of f at x_0 . Partial derivatives are denoted in many different ways, the most common being

$$D_i f(x_0) \equiv f_{x_i}(x_0) \equiv \frac{\partial f}{\partial x_i}(x_0) \equiv \partial_{x_i} f(x_0).$$

In these notes, we will prefer to use the latter of the above, $\partial_{x_i} f(x_0)$ to refer to the partial derivative with respect to the x_i variable.

Substituting $u = e_i$ in t(3.7) we see that i -th partial derivative of f is given by

$$\partial_{x_i} f(x_0) = \lim_{t \rightarrow 0} \frac{f(x_0^1, \dots, x_0^i + t, \dots, x_0^n) - f(x_0^1, \dots, x_0^i, \dots, x_0^n)}{t}.$$

This observation provides an easy way to compute partial derivatives:

The partial derivative $\partial_{x_i} f$ is the “usual” derivative of f viewed as a function of the single variable x_i , and considering the rest of the variables as constants.

Since the partial derivatives of a function are essentially derivatives of single-variable functions, it is no surprise that the partial derivatives have very similar properties to those of the derivative of functions from \mathbb{R} to \mathbb{R} .

Definition 3.4. Let A be a set in \mathbb{R}^n and $x_0 \in A$. We say that $f : A \rightarrow \mathbb{R}$ has a **minimum** at x_0 if

$$f(x_0) \leq f(x) \quad \text{for all } x \in A.$$

We say that f has a **local minimum** at x_0 if there exists a ball $B_\epsilon(x_0) \subset U$ such that

$$f(x_0) \leq f(x) \quad \text{for all } x \in B_\epsilon(x_0) \cap A.$$

Similarly, we say that f has a **maximum** at x_0 if

$$f(x_0) \geq f(x) \quad \text{for all } x \in A,$$

and we say that f has **local maximum** at x_0 if there exists a ball $B_\epsilon(x_0) \subset U$ such that

$$f(x_0) \geq f(x) \quad \text{for all } x \in B_\epsilon(x_0) \cap A.$$

Proposition 3.5. *If $f : U \rightarrow \mathbb{R}$ has a local minimum or maximum at x_0 and its partial derivatives exist, then $\partial_{x_i} f(x_0) = 0$, for $i = 1, \dots, n$.*

The following is a weak version of the **mean value theorem**.

Proposition 3.6. *If $U \subset \mathbb{R}^n$ is open, $f : U \rightarrow \mathbb{R}$ has partial derivatives at each $x \in U$, $x_0 \in U$, and $t \in \mathbb{R}$ is such that*

$$(x_0^1, \dots, x_0^i + s, \dots, x_0^n) \in U$$

for all $s \in [0, t]$ (or $s \in [t, 0]$, if $t < 0$), then there exists c between x_0^i and $x_0^i + t$ such that

$$f(x_0^1, \dots, x_0^i + t, \dots, x_0^n) - f(x_0^1, \dots, x_0^i, \dots, x_0^n) = t \partial_{x_i} f(x_0^1, \dots, c, \dots, x_0^n).$$

The proofs of these propositions follow directly from their single-variable versions, and we leave them as exercises (Exercises 9 and 10).

We will make use of the fact that the partial derivatives behave just like the “regular” derivative (one that we know how to compute) to finally compute the *total* derivative of a function at a point x_0 .

Theorem 3.4. *Let $U \subset \mathbb{R}^n$ be open and $f : U \rightarrow \mathbb{R}^m$ be differentiable at $x_0 \in U$. Then each $\partial_{x_i} f(x_0)$ exists, and the ij -th entry of the **Jacobian** of f at x_0 (i.e. the matrix representation of $Df(x_0)$ in the canonical basis) is given by*

$$(Jf(x_0))_{i,j} = \partial_{x_j} f^i(x_0).$$

Proof. The fact that every partial derivative exists is a simple consequence of taking $u = e_i$ in Theorem 3.3. In order to compute the entries of the Jacobian matrix, we recall that in the discussion leading to Equation eqref we established that, given a basis $\{u_1, \dots, u_n\}$ for \mathbb{R}^n , the ij -th component of the matrix representation of a linear mapping $T = (T_1, \dots, T_m) : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is given by the action of the i -th component function of T on the j -th basis vector u_j .

Moreover, in Proposition 3.4 we proved that its i -th component function is simply the derivative of the i -th component of f , we have that

$$[Df(x_0)]_i = Df^i(x_0).$$

In view of these two observations, and since $Df(x_0)$ is linear, it follows that

$$(Jf(x_0))_{i,j} = [Df(x_0)]_i(e_j) = Df^i(x_0)(e_j) \underbrace{=}_{\text{definition}} \partial_{x_j} f^i(x_0).$$

□

The Jacobian of a real-valued function is particularly important and we will devote a few lines to its properties.

Definition 3.5. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable at x_0 . The Jacobian of f is known as the **gradient** of f and is denoted by

$$\nabla f(x_0) := Jf(x_0) \underbrace{=}_{\text{By Theorem 3.4}} (\partial_{x_1} f(x_0), \dots, \partial_{x_n} f(x_0)).$$

Remark 3.2. *Note that the representation of the gradient as “the vector of partial derivatives” appearing in the definition above, holds only when using the Cartesian canonical basis. In some applications is convenient to use a different coordinate system (spherical, cylindrical, etc.) and in those cases the gradient is not represented by the vector of partial derivatives.*

The symbol ∇ appearing in the definition of $\nabla f(x_0)$, is known as **nabla**. In Cartesian coordinates it is possible (and convenient) to understand it as the “vector of partial differential operators”

$$\nabla := (\partial_{x_1}, \dots, \partial_{x_n}).$$

However, just as in the remark above, this representation is only correct in Cartesian coordinates.

As we will show in the next two results, the gradient encodes important information about the optimal direction of growth/decrease of a scalar valued function.

Theorem 3.5. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable at x_0 . Then, there exists a direction $\tilde{u} \in S^{n-1}$ such that*

$$Df(x_0)(-\tilde{u}) \leq Df(x_0)(u) \leq Df(x_0)(\tilde{u}) \quad \text{for all } u \in S^{n-1}. \quad (3.8)$$

Moreover, if $Df(x_0) \neq 0$, then \tilde{u} is unique. In other words, if a real-valued function is differentiable and its derivative is not the zero mapping, then there exists a unique direction \tilde{u} of maximal growth at x_0 . Moreover, the anti parallel direction $-\tilde{u}$ is the unique direction of maximal descent.

Proof. Existence of the optimal direction follows easily from two observations: 1) since $Df(x_0)$ is linear, it must then be continuous (Corollary 2.3), and 2) the unit sphere S^{n-1} is compact (as it is closed and bounded). Therefore, by the extreme value theorem, there exists $\tilde{u} \in S^{n-1}$ such that

$$Df(x_0)(u) \leq Df(x_0)(\tilde{u}) \quad \text{for all } u \in S^{n-1}.$$

We now observe that $u \in S^{n-1}$ implies that $-u \in S^{n-1}$ and therefore we have

$$Df(x_0)(-u) \leq Df(x_0)(\tilde{u}).$$

However, since $Df(x_0)(-u) = -Df(x_0)(u)$, this implies that

$$Df(x_0)(u) \geq -Df(x_0)(\tilde{u}) = Df(x_0)(-\tilde{u}).$$

Putting these inequalities together we obtain (3.8).

Clearly, if $Df(x_0) = 0$ then any direction is both a maximum and a minimum. We will now assume that $Df(x_0) \neq 0$ and that there exist two directions $\tilde{u}_1, \tilde{u}_2 \in S^{n-1}$ satisfying (3.8). We first show that they are not antipodes, as if $\tilde{u}_1 = -\tilde{u}_2$ it would follow that

$$0 = Df(x_0)(\tilde{u}_1 + \tilde{u}_2) = Df(x_0)(\tilde{u}_1) + Df(x_0)(\tilde{u}_2) = 2Df(x_0)(\tilde{u}_1),$$

which would imply that $Df(x_0)(\tilde{u}_1) = 0$, but since \tilde{u}_1 is the direction of maximum growth, this contradicts the assumption that $Df(x_0) \neq 0$.

We can therefore assume that $|\tilde{u}_1 + \tilde{u}_2| \neq 0$ and define

$$v := \frac{\tilde{u}_1 + \tilde{u}_2}{|\tilde{u}_1 + \tilde{u}_2|} \in S^{n-1}.$$

We start by remarking that, since $\tilde{u}_1, \tilde{u}_2 \in S^{n-1}$ and these vectors are not multiples of each other we have that

$$1 = |v| = \frac{|\tilde{u}_1 + \tilde{u}_2|}{|\tilde{u}_1 + \tilde{u}_2|} < \frac{1 + 1}{|\tilde{u}_1 + \tilde{u}_2|} = \frac{2}{|\tilde{u}_1 + \tilde{u}_2|}, \quad (3.9)$$

where, to obtain the strict inequality, we used the fact that the equality case in the triangle inequality holds only if the two vectors are multiples of each other.




Figure 3.1: The function from Example 3.1 takes the constant value 0 along the x and y axes, hence its partial derivatives at $(0, 0)$ exist and equal 0. However the function is discontinuous at $(0, 0)$, and therefore is not differentiable.

We then compute

$$Df(x_0)(v) = \frac{1}{|\tilde{u}_1 + \tilde{u}_2|} (Df(x_0)(\tilde{u}_1) + Df(x_0)(\tilde{u}_2)) = \frac{2}{|\tilde{u}_1 + \tilde{u}_2|} Df(x_0)(\tilde{u}_1) \underset{\text{By (3.9)}}{>} Df(x_0)(\tilde{u}_1)$$

but this contradicts the maximality of \tilde{u}_1 . Therefore, the maximizer \tilde{u} must be unique. \square

Theorem 3.6. If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is differentiable at x_0 , the direction of maximal growth guaranteed by Theorem 3.5 is given by the normalized gradient

$$\frac{\nabla f(x_0)}{|\nabla f(x_0)|} = \frac{(\partial_{x_1} f(x_0), \dots, \partial_{x_n} f(x_0))}{((\partial_{x_1} f(x_0))^2 + \dots + (\partial_{x_n} f(x_0))^2)^{1/2}}.$$

Proof. Using the canonical basis we see that, for any $u \in S^{n-1}$ it follows that

$$Df(x_0)(u) = \nabla f(x_0) \cdot u = |\nabla f(x_0)| \cdot |u| \cos \theta,$$

where θ is the angle between u and the gradient $|\nabla f(x_0)|$. Since the magnitudes of $\nabla f(x_0)$ and all $u \in S^{n-1}$ are constant, it follows that the maximum value of $Df(x_0)(u)$ on the sphere will happen when $\theta = 0$; i.e. when $u = \lambda \nabla f(x_0)$ for $\lambda > 0$. Hence, the direction of maximum growth in the unit sphere must be given by

$$\tilde{u} = \frac{\nabla f(x_0)}{|\nabla f(x_0)|} = \frac{(\partial_{x_1} f(x_0), \dots, \partial_{x_n} f(x_0))}{((\partial_{x_1} f(x_0))^2 + \dots + (\partial_{x_n} f(x_0))^2)^{1/2}}.$$

\square

Example 3.1. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ (depicted in Figure 3.1) be given by

$$f(x, y) = \begin{cases} 1, & \text{if } 0 < y < x^2, \\ 0, & \text{otherwise.} \end{cases}$$

Then $f(x, 0) = f(0, y) = 0$ for all x, y , so that $\partial_x f(0, 0) = \partial_y f(0, 0) = 0$. However, f is not even continuous at $(0, 0)$.

This shows that the converse of Theorem 3.4 does not hold in general. However, with the additional requirement of continuity of the partial derivatives, the result is true, as we will now prove.

Definition 3.6. Let $U \subset \mathbb{R}^n$ be open. We say that $f : U \rightarrow \mathbb{R}^m$ is **continuously differentiable** at x_0 if each of the partial derivatives $\partial_{x_j} f^i(x)$ exists in a ball $B_\epsilon(x_0)$ and is continuous at x_0 .

Theorem 3.7. Let $U \subset \mathbb{R}^n$ be open and $x_0 \in U$. If $f : U \rightarrow \mathbb{R}^m$ is continuously differentiable at x_0 , then f is differentiable at x_0 and its Jacobian is given by

$$(Jf(x_0))_i = \partial_{x_i} f(x_0).$$

Proof. Suppose that the partial derivatives $\partial_{x_i} f(x)$ exist at each $x \in B_\epsilon(x_0)$. If h is such that $x_0 + h \in B_\epsilon(x_0)$, then by Proposition 3.6,

$$\begin{aligned} & f(x_0^1 + h_1, x_0^2 + h_2, \dots, x_0^n + h_n) - f(x_0^1, x_0^2, \dots, x_0^n) \\ &= f(x_0^1 + h_1, x_0^2 + h_2, \dots, x_0^n + h_n) - f(x_0^1, x_0^2 + h_2, \dots, x_0^n + h_n) \\ &\quad + f(x_0^1, x_0^2 + h_2, \dots, x_0^n + h_n) - f(x_0^1, x_0^2, x_0^3 + h_3, \dots, x_0^n + h_n) \\ &\quad + \dots + f(x_0^1, x_0^2, \dots, x_0^n + h_n) - f(x_0^1, x_0^2, \dots, x_0^n) \\ &= \partial_{x_1} f(c_1, x_0^2 + h_2, \dots, x_0^n + h_n)h_1 + \dots + \partial_{x_n} f(x_0^1, x_0^2, \dots, c_n)h_n, \end{aligned}$$

where each $c_i \in \mathbb{R}$ is between x_0^i and $x_0^i + h_i$. When $|h| \rightarrow 0$, each $c_i \rightarrow x_0^i$. Then,

$$\frac{|f(x_0 + h) - f(x_0) - \sum_i \partial_{x_i} f(x_0)h_i|}{|h|} \leq \sum_i |\partial_{x_i} f(c_1, \dots, x_0^i + h_2, \dots, x_0^n + h_n) - \partial_{x_i} f(x_0)| \frac{|h_i|}{|h|}$$

and therefore,

$$\lim_{h \rightarrow 0} \frac{|f(x_0 + h) - f(x_0) - \sum_i \partial_{x_i} f(x_0)h_i|}{|h|} = 0,$$

because the $\partial_{x_i} f$ are continuous at x_0 . □

By applying the result above to every component function of $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ we obtain the following

Corollary 3.8. If $f : U \rightarrow \mathbb{R}^m$ is continuously differentiable at x_0 , then it is differentiable at x_0 .

The inverse of this corollary is false, as shown by the following well-known example.

Example 3.2. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$f(x) = \begin{cases} 0, & x = 0, \\ x^2 \sin \frac{1}{x}, & x \neq 0. \end{cases}$$

The function is depicted in the center panel of Figure 3.2. We claim that f is differentiable in \mathbb{R} , but its derivative is not continuous at 0. Clearly the function is differentiable for $x \neq 0$. We will show that $f'(0) = 0$. Indeed, from the definition of the derivative we have that

$$|f'(0)| = \left| \lim_{h \rightarrow 0} \frac{f(0+h) - f(0)}{h} \right| = \left| \lim_{h \rightarrow 0} \frac{h^2 \sin(1/h) - 0}{h} \right| = \left| \lim_{h \rightarrow 0} h \sin(1/h) \right| \leq \left| \lim_{h \rightarrow 0} h \right| = 0.$$

Therefore f is differentiable at 0 and $f'(0) = 0$. On the other hand, for $x \neq 0$ we have that

$$f'(x) = 2x \sin(1/x) - \cos(1/x),$$




Figure 3.2: Left: The function $\sin(1/x)$, known as the topologist's sine is undefined at $x = 0$, and as its argument approaches 0, the frequency of the oscillation tends to infinity. Center: Modulating the amplitude of the topologist's sine by the factor x^2 and defining the value of the resulting function to be 0 for $x = 0$ makes the function differentiable everywhere. Right: Despite the fact that the derivative of the function is defined everywhere (and in particular equals 0 for $x = 0$), it is discontinuous at the origin.

that is not defined for $x = 0$. In particular, the limit as $x \rightarrow 0$ of the above expression does not exist. From this we conclude that

$$0 = f'(0) \neq \lim_{x \rightarrow 0} f'(x),$$

and f' is not continuous at $x = 0$.

To conclude this section, we will establish the classical version of the chain rule

Proposition 3.7. Let $U \subset \mathbb{R}^n$, $V \subset \mathbb{R}^m$ be open sets, $g : U \rightarrow V$ be continuously differentiable at $x_0 \in U$, and $f : V \rightarrow \mathbb{R}$ be differentiable at $g(x_0)$. Then,

$$\partial x_i (f \circ g)(x_0) = \sum_{j=1}^m \partial x_j f(g(x_0)) \partial x_i g^j(x_0). \quad (3.10)$$

Proof. Since g is continuously differentiable at x_0 , it is differentiable at x_0 , and by the chain rule, $f \circ g$ is differentiable at x_0 and

$$D(f \circ g)(x_0) = Df(g(x_0)) \circ Dg(x_0).$$

Hence, using the canonical basis we can represent each of the derivatives appearing above by their Jacobians, leading to

$$J(f \circ g)(x_0) = Jf(g(x_0)) \cdot Jg(x_0) = \begin{bmatrix} \partial_{x_1} f(g(x_0)) & \dots & \partial_{x_m} f(g(x_0)) \end{bmatrix} \cdot \begin{bmatrix} \partial_{x_1} g^1(x_0) & \dots & \partial_{x_n} g^1(x_0) \\ \partial_{x_1} g^2(x_0) & \dots & \partial_{x_n} g^2(x_0) \\ \vdots & \ddots & \vdots \\ \partial_{x_1} g^m(x_0) & \dots & \partial_{x_n} g^m(x_0) \end{bmatrix}.$$

Finally, recalling the definition of the partial derivatives and using the matrix representation above we see

that

$$\begin{aligned}
 \partial_{x_i}(f \circ g)(x_0) &= D(f \circ g)(x_0)(e_i) \\
 &= Jf(g(x_0)) \cdot Jg(x_0) \cdot e_i \\
 &= \underbrace{[\partial_{x_1}f(g(x_0)) \quad \dots \quad \partial_{x_m}f(g(x_0))]}_{Jf(g(x_0))} \cdot \underbrace{\begin{bmatrix} \partial_{x_1}g^1(x_0) & \dots & \partial_{x_n}g^1(x_0) \\ \partial_{x_1}g^2(x_0) & \dots & \partial_{x_n}g^2(x_0) \\ \vdots & \ddots & \vdots \\ \partial_{x_1}g^m(x_0) & \dots & \partial_{x_n}g^m(x_0) \end{bmatrix}}_{Jg(x_0)} \cdot \underbrace{\begin{bmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{bmatrix}}_{\text{i-th row}} \\
 &\quad \text{This product picks the i-th column of } Jg(x_0) \\
 &= [\partial_{x_1}f(g(x_0)) \quad \dots \quad \partial_{x_m}f(g(x_0))] \cdot \begin{bmatrix} \partial_{x_i}g^1(x_0) \\ \partial_{x_i}g^2(x_0) \\ \vdots \\ \partial_{x_i}g^m(x_0) \end{bmatrix} \\
 &= \sum_{j=1}^m \partial_{x_j}f(g(x_0))\partial_{x_i}g^j(x_0),
 \end{aligned}$$

as we were trying to prove. □

3.4 The inverse function theorem

We know that the derivative of a function f at a point x_0 is the best linear approximation to f around x_0 , that is,

$$f(x) \approx f(x_0) + Df(x_0)(x - x_0).$$

If the expression above were the **exact** equality, (i.e. if we could drop the term $o(|x - x_0|)$ in the linear approximation)

$$f(x) = f(x_0) + Df(x_0)(x - x_0),$$

and the linear transformation $Df(x_0)$ were invertible, then we could solve explicitly for x , leading to

$$\begin{aligned}
 x &= x_0 + [Df(x_0)]^{-1}(f(x_0)) + [Df(x_0)]^{-1}(f(x)) \\
 &= b + [Df(x_0)]^{-1}(f(x)),
 \end{aligned}$$

where we first grouped the first two *constant* terms in the right hand side into the vector b . Therefore, the inverse function of f would also be linear. Since linear functions are differentiable and the derivative of a linear function is equal to the linear function, this would then imply that f^{-1} is differentiable at $f(x_0)$ and

$$[Df(x_0)]^{-1}(f(x_0)) = [Df(x_0)]^{-1}(f(x)).$$

More generally, if A is an $n \times n$ matrix associated to the linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$, then the function T is invertible if and only if the determinant of A is nonzero. Moreover, if A is an $m \times n$ matrix with $m < n$, then we can solve the homogeneous equation

$$Ax = 0$$

for m of its variables in terms of the remaining $n - m$ (i.e. m variables are implicit functions of the remaining $n - m$), provided that m of the columns of A (or equivalently, its m rows) are linearly independent, i.e., A has full rank.

Thus, we want to know under what circumstances these invertibility and differentiability properties of the linear transformation $Df(x_0)$ are transferred to f in a neighborhood of x_0 . For functions of a single variable, we have an answer: if f is continuously differentiable in a neighborhood of $x_0 \in \mathbb{R}$ and $f'(x_0) \neq 0$, we know that there exists a local inverse function. In a neighborhood of x_0 , say U , such that f is invertible in U , f^{-1} is differentiable in $f(U)$, and

$$(f^{-1})'(f(x)) = \frac{1}{f'(x)}.$$

In what follows, we will answer these questions. We will start by answering the question about invertibility, establishing the multi variable version of this result, known as the **Inverse function theorem**. The proof of this theorem is quite involved, and we will need to prove three useful results first. The first one is a special case of an incredibly powerful and useful theorem both for theory and applications. It is at the heart of many numerical methods for the solution of nonlinear equations (algebraic, differential, partial differential, integral, etc.) and variants of it are used to prove many important theoretical results. For instance, Picard's classic proof for the existence and uniqueness of solutions for a wide class of ordinary differential equations. The theorem can be stated and proven in a more general setting, without much difficulty, but we will focus here on its version for \mathbb{R}^n .

Theorem 3.9 (Contraction mapping theorem). *Let $X \subseteq \mathbb{R}^n$ be closed, $0 < c < 1$, and $f : X \rightarrow X$ be such that*

$$|f(x) - f(y)| < c|x - y| \quad \text{for all } \forall x, y \in X. \quad (3.11)$$

Then there exists a unique $x_ \in X$ such that $f(x_*) = x_*$.*

The proof of this result will be left as an exercise, but a few remarks are in order. A function satisfying the condition (3.11) with $0 < c \leq 1$ is called a **contraction**; if the constant c is *strictly* less than one (as in the statement of the theorem), the function is called a **strict contraction**. If the constant is equal to 1, it is easy to build examples of functions that have infinitely many fixed points, and also that have no fixed points.

The second result will allow us to use the information from the derivatives of a function to ensure that the function is Lipschitz continuous.

Lemma 3.2. *Let $\epsilon > 0$, $f : B_\epsilon(x_0) \rightarrow \mathbb{R}^n$ be differentiable, and $M > 0$ such that*

$$|\partial_j f^i(x)| \leq M \quad \text{for } i, j = 1, \dots, n, \quad \text{and } x \in B_\epsilon(x_0).$$

Then, f is Lipschitz on the ball $B_\epsilon(x_0)$. More precisely

$$|f(x) - f(y)| \leq n^2 M |x - y|, \quad x, y \in B_\epsilon(x_0).$$

Proof. The proof of this lemma involves an argument similar to the one used in the proof of Theorem 3.7, where a bunch of zeros were added in order to express the difference $f(x) - f(y)$ by changing only one component of the argument at a time and be able to use the mean value theorem. By the mean value theorem

3.6, there exist $z_1, \dots, z_n \in B_\epsilon(x_0)$ such that

$$\begin{aligned} f^i(x) - f(y) &= f^i(x_1, x_2, \dots, x_n) - f^i(y_1, y_2, \dots, y_n) \\ &= f^i(x_1, x_2, \dots, x_n) - f^i(y_1, x_2, \dots, x_n) \\ &\quad + f^i(y_1, x_2, \dots, x_n) - f^i(y_1, y_2, \dots, x_n) \\ &\quad \vdots \qquad \qquad \qquad \vdots \\ &\quad + f^i(y_1, y_2, \dots, y_{n-1}, x_n) - f^i(y_1, y_2, \dots, y_n) \\ &= \sum_{j=1}^n \partial_{x_j} f^i(z_j)(x_j - y_j), \end{aligned}$$

so that

$$|f^i(x) - f^i(y)| \leq \sum_{j=1}^n |\partial_{x_j} f^i(z_j)| |x_j - y_j| \leq nM|x - y|.$$

Thus, repeating the argument for each component we obtain

$$|f(x) - f(y)| \leq \sum_{i=1}^n |f^i(x) - f^i(y)| \leq \sum_{i=1}^n nM|x - y| = n^2M|x - y|.$$

□

Finally, the third lemma below will be the core of the proof of the inverse function theorem.

Lemma 3.3. *Let $r > 0$ and $g : B_r(0) \subset \mathbb{R}^n \rightarrow \mathbb{R}^n$ be such that*

$$g(0) = 0 \quad \text{and} \quad |g(x) - g(y)| \leq \frac{1}{2}|x - y| \quad \forall x, y \in B_r(0). \quad (3.12)$$

Define

$$f : B_r(0) \rightarrow \mathbb{R}^n \quad \text{by} \quad f(x) := g(x) + x. \quad (3.13)$$

Then:

(a) f is injective.

(b) $B_{r/2}(0) \subset f(B_r(0))$.

The two conditions (a) and (b) above imply that if we restrict the co-domain of f to consist only of $B_{r/2}(0)$, the resulting function $f : B_r(0) \rightarrow B_{r/2}(0)$ is invertible.

Proof. We start by proving (a) by contradiction. If f were not injective, then there would exist $x \neq y \in B_r(0)$ such that $f(x) = f(y)$. This would imply that

$$|x - y| = |f(x) - x - f(y) + y| = |g(x) - g(y)| \leq \frac{1}{2}|x - y|,$$

which can only hold if $|x - y| = 0$, in contradiction with the hypothesis $x \neq y$.

We now prove (b). We want to show that for every $y \in B_{r/2}(0)$ there exists $x \in B_r(0)$ such that $y = f(x)$. Using the definition of f and reordering terms, we see that this is equivalent to proving that for every $y \in B_{r/2}(0)$ there exists $x \in B_r(0)$ such that

$$x = y - g(x) =: G_y(x). \quad (3.14)$$




Figure 3.3: The image of the ball $B_r(0)$ under the function f defined in (3.13) contains the ball $B_{r/2}(0)$. Therefore, every point $y \in B_{r/2}(0)$ is the image under f of a point in $B_r(0)$, which guarantees surjectivity of f when the range is restricted to $B_{r/2}(0)$.

In other words, we want to show that for every y , the function $G_y(x)$ has a fixed point. We will do that using the contraction mapping theorem 3.9. Take $0 < \epsilon < r$ and consider a point

$$y \in \overline{B_{(r-\epsilon)/2}(0)} := \{y \in \mathbb{R}^n : |y| \leq (r - \epsilon)/2\},$$

and define $G_y : \overline{B_{(r-\epsilon)}(0)} \rightarrow \mathbb{R}^n$ as in (3.14). Taking a point

$$x \in \overline{B_{(r-\epsilon)}(0)} := \{y \in \mathbb{R}^n : |y| \leq (r - \epsilon)\},$$

and applying G_y we see that:

$$\begin{aligned} |G_y(x)| &= |y - g(x)| \leq |y| + |g(x)| \\ &= |y| + |g(x) - g(0)| && \text{(Since } g(0) = 0\text{)} \\ &\leq \frac{1}{2}(r - \epsilon) + |g(x) - g(0)| && \text{(Since } y \in \overline{B_{(r-\epsilon)/2}(0)}) \\ &\leq \frac{1}{2}(r - \epsilon) + \frac{1}{2}|x| \\ &\leq \frac{1}{2}(r - \epsilon) + \frac{1}{2}(r - \epsilon) && \text{(Since } x \in \overline{B_{(r-\epsilon)}(0)}) \\ &= r - \epsilon. \end{aligned}$$

Therefore, if $x \in \overline{B_{(r-\epsilon)}(0)}$ it follows that $G_y(x) \in \overline{B_{(r-\epsilon)}(0)}$, which proves that

$$G_y : \overline{B_{(r-\epsilon)}(0)} \longrightarrow \overline{B_{(r-\epsilon)}(0)}.$$

It only remains to show that G_y is a contraction. To do that, we take $x_1, x_2 \in \overline{B_{(r-\epsilon)}(0)}$ and compute

$$|G_y(x_1) - G_y(x_2)| = |y - g(x_1) - y + g(x_2)| = |g(x_1) - g(x_2)| \leq \frac{1}{2}|x_1 - x_2|,$$

which verifies that G_y is a contraction. We can therefore apply the contraction mapping theorem (3.9) to conclude that for all $0 < \epsilon < r$ and all $y \in \overline{B_{(r-\epsilon)/2}(0)}$, there exists $x_* \in \overline{B_{(r-\epsilon)}(0)}$ such that $f(x_*) = y$. Finally, since this argument can be repeated verbatim for all $0 < \epsilon < r$, and f is injective on $B_r(0)$, we conclude that $f : B_r(0) \rightarrow B_{r/2}(0)$ is invertible. This is depicted in Figure 3.3. \square

We are finally in the position to state and prove the inverse function theorem.

Theorem 3.10 (Inverse function theorem). Let $U \subset \mathbb{R}^n$, $x_0 \in U$, and $f : U \rightarrow \mathbb{R}^n$ be differentiable in U , continuously differentiable at x_0 , and such that

$$\det(f'(x_0)) \neq 0.$$

Then, there exist a neighborhood V of x_0 and a neighborhood W of $f(x_0)$ such that $f : V \rightarrow W$ has an inverse $f^{-1} : W \rightarrow V$, f^{-1} is differentiable in W , and for each $y \in W$,

$$D(f^{-1})(y) = [Df(f^{-1}(y))]^{-1}.$$

Or, in terms of matrices,

$$Jf^{-1}(y) = [Jf(f^{-1}(y))]^{-1}.$$

In words: the inverse f^{-1} is differentiable, and its Jacobian matrix is the inverse of the Jacobian matrix of f .

Proof. We will make three simplifying assumptions:

1. We will assume that the function produces the zero vector at the point of differentiability. Explicitly, that $f(x_0) = 0$.
2. We will assume that the point of differentiability is the origin, i.e. $x_0 = 0$.
3. We will assume that the derivative of f at the point of differentiability is the identity; i.e. $Df(x_0) = I$.

The fully general result stated above can be recovered from this simplified statement easily, and will be left as an exercise.

Our first goal is to use Lemma 3.3 to prove that the function f is invertible in a neighborhood of x_0 and $f(x_0)$. To that avail, we define $g : U \rightarrow \mathbb{R}^n$ by

$$g(x) := f(x) - x.$$

By definition, f and g satisfy condition (3.13) from Lemma 3.3, while the first and second simplifying assumptions yield

$$g(0) = f(0) - 0 = 0,$$

and this the first condition in (3.12) is satisfied.

Since f and x are both continuously differentiable at x_0 , it follows that g is continuously differentiable at x_0 , and its derivative is given by

$$Dg(x_0) = Df(x_0) - I = I - I = 0,$$

where we made use of the third simplifying assumption $Df(x_0) = I$. This provides condition (3.12) from Lemma 3.3. It only remains to show that g is a contraction with constant equal to $\frac{1}{2}$.

We then observe that, since $Dg(x_0) = 0$, it follows that $\partial_{x_j} g^i(x_0) = 0$ for every $i, j \in \{1, \dots, n\}$. Moreover, since g is continuously differentiable at x_0 , then every partial derivative is continuous at x_0 and thus, there exists a ball $B_r(x_0)$ such that

$$|\partial_{x_j} g^i(x)| \leq \frac{1}{2n^2} \quad \text{for every } x \in B_r(x_0) \quad \text{and } i, j \in \{1, \dots, n\}.$$

This bound allows us to use Lemma 3.2 with $M := \frac{1}{2n^2}$ to conclude that g satisfies the Lipschitz condition

$$|g(x) - g(y)| \leq \left(\frac{1}{2n^2} \right) n^2 |x - y| = \frac{1}{2} |x - y| \tag{3.15}$$

for all $x, y \in B_r(x_0)$. This contraction bound provides us with the final condition required to apply Lemma 3.3. We can therefore conclude that the function $f : B_r(x_0) \rightarrow B_{r/2}(f(x_0))$ is invertible.

Before moving on to prove differentiability of the inverse function, we will extract one more useful inequality from the result above. Substituting the definition of g in terms of f in the inequality (3.15) and using the two possibilities afforded by the reverse triangle inequality yields

$$\left. \begin{array}{l} |x - y| - |f(x) - f(y)| \\ |f(x) - f(y)| - |x - y| \end{array} \right\} \leq \frac{1}{2}|x - y|.$$

From the top line above we can infer the first inequality below, while from the bottom line we can obtain the second one, yielding

$$\frac{1}{2}|x - y| \leq |f(x) - f(y)| \leq \frac{3}{2}|x - y|. \quad (3.16)$$

This inequality will prove useful shortly.

We must now prove that the inverse function f^{-1} is differentiable, and we will do that by exhibiting the candidate for the derivative and proving that it satisfies the linear approximation property (3.2).

By definition, the functions f and f^{-1} satisfy the property

$$(f^{-1} \circ f)(x) = x.$$

If the inverse function were differentiable, we could differentiate both sides of the expression above and use the chain rule to obtain

$$Df^{-1}(f(x_0)) \circ Df(x_0) = I.$$

Since we assumed that $Df(x_0) = I$, the expression above simplifies to

$$Df^{-1}(f(x_0)) = I,$$

which gives us a candidate for the derivative of the inverse function.

We will now prove, by verifying the linear approximation property, that the identity is indeed the derivative of f^{-1} . Concretely, our goal is to prove that

$$\lim_{x \rightarrow x_0} \frac{|f^{-1}(x) - f^{-1}(x_0) - I(x - x_0)|}{|x - x_0|} = 0.$$

From the first two simplifying assumptions, namely that $f(x_0) = 0$ and $x_0 = 0$, together with the invertibility of $f : B_r(0) \rightarrow B_{r/2}(0)$ it follows that

$$f^{-1}(0) = 0.$$

Therefore, the quotient above simplifies into

$$\lim_{x \rightarrow 0} \frac{|f^{-1}(x) - x|}{|x|}.$$

The invertibility of f also guarantees that for every $x \in B_{r/2}(0)$ there is a $y \in B_r(0)$ such that $y = f^{-1}(x)$. We can thus express the quotient in terms of f as

$$\frac{|y - f(y)|}{|f(y)|}.$$

By letting $x = 0$ in the sequence of inequalities (3.16) and observing the leftmost inequality we conclude that:

$$\frac{1}{|f(y)|} \leq \frac{2}{|y|} \quad (3.17a)$$

$$\text{If } f(y) \rightarrow 0, \text{ then } y \rightarrow 0. \quad (3.17b)$$

We therefore have that

$$0 \leq \lim_{x \rightarrow 0} \frac{|f^{-1}(x) - x|}{|x|} = \lim_{f(y) \rightarrow 0} \frac{|y - f(y)|}{|f(y)|} \stackrel{(3.17a)}{\leq} 2 \lim_{f(y) \rightarrow 0} \frac{|y - f(y)|}{|y|} \stackrel{(3.17b)}{=} 2 \lim_{y \rightarrow 0} \frac{|f(y) - y|}{|y|}.$$

However, since f is differentiable at $y = 0$ and $Df(0) = I$, the limit on the right must be zero. This proves that f^{-1} is differentiable at $x = 0$ and $Df^{-1}(0) = I$. \square

Remark 3.3. From the proof of the inverse function theorem, we can conclude that if f is continuously differentiable in a neighborhood of x_0 and $\det Jf(x_0) \neq 0$, then we can choose V and W such that $f : V \rightarrow W$ and $f^{-1} : W \rightarrow V$ are continuously differentiable.

This follows from the expression for the Jacobian of f^{-1} and from Cramer's rule, which implies that the entries of $(f^{-1})'(f(x))$ are rational functions of the entries of $f'(x)$, and therefore continuous.

3.5 Implicit function theorem

In this section, we consider the second part of the question posed at the beginning of the previous one. Recall that if A is a matrix of size $m \times (n + m)$, and z is a vector in $\mathbb{R}^{(n+m)}$ then the equation

$$Az = 0$$

has a solution if A has m columns that are linearly independent—we say that **the rank of A** is equal to m . To see why this is the case, let us assume that the columns of A have been ordered so that the last m columns are indeed linearly independent, and let us write

$$A = [A_n | A_m] \quad \text{and} \quad z = (x, y)^\top$$

where $A_n \in \mathbb{R}^{m \times n}$ and $A_m \in \mathbb{R}^{m \times m}$ are the matrices formed by taking the first n and last m columns of A respectively, and $x \in \mathbb{R}^n$, $y \in \mathbb{R}^m$.

Doing so, we can then express the equation as

$$Az = [A_n | A_m](x, y)^\top = A_n x + A_m y = 0.$$

From here, and using the fact that the columns of A_m are linearly independent, we can solve for y , obtaining

$$y = -A_m^{-1} A_n x.$$

Therefore, knowledge of the values of the n components of x determines the value of y , and we could write

$$y = g(x) \quad \text{where} \quad g(x) := -A_m^{-1} A_n x.$$

For this reason we say that y is implicitly determined as a function of x by the equation

$$A(x, y)^\top = 0.$$

It then makes sense for some function $f : \mathbb{R}^{(m+n)} \rightarrow \mathbb{R}^m$ to ask whether the equation

$$f(x, y) = 0,$$

implicitly defines $y \in \mathbb{R}^m$ as a function of $x \in \mathbb{R}^n$ in some neighborhood of a point $(x_0, y_0) \in \mathbb{R}^{(n+m)}$ that satisfies the equation. Moreover, if implicitly $y = g(x)$, then we are also interested in the differentiability of g . The answer is given by the implicit function theorem.

Theorem 3.11 (Implicit function theorem). *Let $U \subset \mathbb{R}^{n+m}$ be open, $f : U \rightarrow \mathbb{R}^m$ differentiable, and continuously differentiable at $(x_0, y_0) \in U$, with $f(x_0, y_0) = 0$. Let M be the $m \times m$ matrix whose i, j -th component is given by*

$$M_{i,j} := \partial_{x_{n+j}} f^i(x_0, y_0), \quad \text{for } i, j = 1, \dots, m,$$

i.e. the Jacobian of f with respect to the last m variables, and suppose that $\det M \neq 0$. Then there exist an open set $V \subset \mathbb{R}^n$, $x_0 \in V$, and an open set $W \subset \mathbb{R}^m$, $y_0 \in W$, such that for each $x \in V$, there exists a unique differentiable $g(x) \in W$ such that $f(x, g(x)) = 0$. That is, the equation

$$f(x, y) = 0$$

implicitly defines y as a function of x , provided that the derivatives in y form a nonsingular matrix. Thus, we indeed have the analogous result to the linear case.

Proof. Let $F : U \rightarrow \mathbb{R}^{n+m}$ be the function $F(x, y) = (x, f(x, y))$. Then its Jacobian $F'(x_0, y_0)$ is given by

$$\begin{bmatrix} 1 & \dots & 0 & 0 & \dots & 0 \\ 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 1 & 0 & \dots & 0 \\ \partial_{x_1} f^1(x_0, y_0) & \dots & \partial_{x_n} f^1(x_0, y_0) & \partial_{x_{n+1}} f^1(x_0, y_0) & \dots & \partial_{x_{n+m}} f^1(x_0, y_0) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \partial_{x_1} f^m(x_0, y_0) & \dots & \partial_{x_n} f^m(x_0, y_0) & \partial_{x_{n+1}} f^m(x_0, y_0) & \dots & \partial_{x_{n+m}} f^m(x_0, y_0) \end{bmatrix}.$$

That is, it is of the form

$$F'(x_0, y_0) = \begin{bmatrix} I & 0 \\ * & M \end{bmatrix},$$

so that $\det F'(x_0, y_0) = \det M \neq 0$.

By the inverse function theorem, there exists a neighborhood $V \times W$ of (x_0, y_0) and a neighborhood W' of $F(x_0, y_0)$ such that $F : V \times W \rightarrow W'$ has an inverse $F^{-1} : W' \rightarrow V \times W$ that is differentiable.

Since $F(x, y) = (x, f(x, y))$, the inverse function is of the form

$$F^{-1}(x, v) = (x, h(x, v)),$$

for each $(x, v) \in W'$, where $h : W' \rightarrow W$ is differentiable. From here it follows that

$$(x, v) = F(x, h(x, v)) = (x, f(x, h(x, v))),$$

so that

$$f(x, h(x, v)) = v.$$

Thus, $f(x, h(x, 0)) = 0$, and therefore we can take

$$y = g(x) := h(x, 0).$$

□

3.6 Higher order derivatives

If a function $f : U \rightarrow \mathbb{R}$ is differentiable, it induces functions in U given by the partial derivatives $\partial_{x_i} f(x)$. These functions can themselves be differentiable (in which case all of their partial derivatives exist) or at least have *some* partial derivatives well defined. If the partial derivatives of $\partial_{x_i} f$ exist, they are called second-order partial derivatives of f and are denoted by

$$\partial_{x_i x_j} f(x) = \partial_{x_j} (\partial_{x_i} f)(x).$$

In what follows, we will try to make the (already cumbersome) notation a little simpler by suppressing the x in the notation for partial derivatives, so that

$$\partial_i f(x) := \partial_{x_i} f(x) \quad \text{and} \quad \partial_{ij} f(x) := \partial_{x_i x_j} f(x).$$

Similarly, higher-order partial derivatives of order k are denoted by

$$\partial_{i_1 i_2 \dots i_k} f(x) = \partial_{i_k} (\dots \partial_{i_2} (\partial_{i_1} f) \dots)(x).$$

In general, $D_{ij} f(x) \neq D_{ji} f(x)$, for instance in the case of the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$f(x, y) = \begin{cases} \frac{xy(x^2 - y^2)}{x^2 + y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0), \end{cases}$$

you will show as an exercise that $\partial_{12} f(0, 0) \neq \partial_{21} f(0, 0)$.

In the previous example, $\partial_{12} f$ and $\partial_{21} f$ are not continuous at $(0, 0)$. However, as the following theorem shows, continuity is enough to guarantee that the “mixed” derivatives are equal. This result has a long and interesting story (including a failed proof attempt by Euler himself), and is often referred to as either **Clairut’s theorem** (although his proof was reportedly flawed) or **Young’s theorem**.

Theorem 3.12 (Symmetry of mixed partial derivatives). *If $\partial_{ij} f$ and $\partial_{ji} f$ exist in a neighborhood U of x_0 and are continuous at x_0 , then*

$$\partial_{ij} f(x_0) = \partial_{ji} f(x_0).$$

Proof. We first consider the special case where $U \subset \mathbb{R}^2$ and $i = 1, j = 2$. In the general case, we define a function of two variables in a suitable neighborhood of (x_0) in \mathbb{R}^2 by freezing all the remaining components of x_0 with indices different from i and j , and apply the result for \mathbb{R}^2 . Namely, let ϕ be defined by

$$\phi(u, v) = f(x_0^1, \dots, \underset{i\text{-th}}{u}, \dots, \underset{j\text{-th}}{v}, \dots, x_0^n),$$

and observe that

$$\partial_{ij} f(x_0) = \partial_{12} \phi(x_0^i, x_0^j) \quad \text{and} \quad \partial_{ji} f(x_0) = \partial_{21} \phi(x_0^i, x_0^j).$$

The fully general theorem then follows from the case of functions of two variables applied to ϕ .

Therefore, we will focus on the two dimensional case. Let $r > 0$ be such that $B_r(x_0) \subset U$ and define $F : B_r^0(0) \rightarrow \mathbb{R}$ by

$$F(x) := f(x_0 + x).$$

Then,

$$\begin{aligned}\partial_1 F(0) &= \partial_1 f(x_0), & \partial_2 F(0) &= \partial_2 f(x_0), \\ \partial_{12} F(0) &= \partial_{12} f(x_0), & \partial_{21} F(0) &= \partial_{21} f(x_0).\end{aligned}$$

If we define $G : B_r^0(0) \rightarrow \mathbb{R}$ by $G(x, y) = F(y, x)$, and $x, y \in B_r^0(0)$ we have

$$D_1 G(x, y) = \lim_{h \rightarrow 0} \frac{G(x + h, y) - G(x, y)}{h} = \lim_{h \rightarrow 0} \frac{F(y, x + h) - F(y, x)}{h} = \partial_2 F(y, x), \quad (3.18)$$

as well as

$$\partial_{12} G(x, y) = \lim_{h \rightarrow 0} \frac{\partial_1 G(x, y + h) - \partial_1 G(x, y)}{h} \underset{\text{By (3.18)}}{=} \lim_{h \rightarrow 0} \frac{\partial_2 F(y + h, x) - \partial_2 F(y, x)}{h} = \partial_{21} F(y, x).$$

We want to prove that

$$D_{12} G(0) = D_{12} F(0).$$

We will proceed by contradiction and, without loss of generality, assume that

$$\partial_{12} G(0) < \partial_{12} F(0).$$

Since both derivatives are continuous at $0 \in \mathbb{R}^2$, there exists a rectangle

$$R := [-\epsilon, \epsilon] \times [-\epsilon, \epsilon] \subset B_r^0(0)$$

such that

$$0 < \partial_{12} F(x) - \partial_{12} G(x) \quad \text{for all } x \in R.$$

Then, for all $(x, y) \in [-\epsilon, \epsilon] \times [-\epsilon, \epsilon]$, we have

$$0 < \partial_2(\partial_1(F - G))(x, y)$$

This implies that, freezing the first argument, the function $y \mapsto \partial_1(F - G)(x, y)$ is strictly increasing in $[-\epsilon, \epsilon]$ for each $x \in [-\epsilon, \epsilon]$. In particular,

$$\begin{aligned}0 &< \partial_1(F - G)(x, \epsilon) - \partial_1(F - G)(x, -\epsilon) \\ &= \partial_1\left((F - G)(x, \epsilon) - (F - G)(x, -\epsilon)\right).\end{aligned} \quad (3.19)$$

Thus, if we define the function $H : [-\epsilon, \epsilon] \rightarrow \mathbb{R}$ by

$$H(t) := (F - G)(t, \epsilon) - (F - G)(t, -\epsilon),$$

the inequality (3.19) states that $H'(t) > 0$ in $[-\epsilon, \epsilon]$, so that H is strictly increasing in $[-\epsilon, \epsilon]$. Therefore, evaluating at the endpoints of the interval we must have

$$H(-\epsilon) < H(\epsilon). \quad (3.20)$$

However, computing the value of both sides of the inequality above, and recalling that $G(x, y) = F(y, x)$, we see that

$$\begin{aligned} H(\epsilon) &= F(\epsilon, \epsilon) - G(\epsilon, \epsilon) - F(\epsilon, -\epsilon) + G(\epsilon, -\epsilon), \\ &= F(\epsilon, \epsilon) - F(\epsilon, \epsilon) - F(\epsilon, -\epsilon) + F(-\epsilon, \epsilon), \\ &= F(-\epsilon, \epsilon) - F(\epsilon, -\epsilon) \end{aligned}$$

and

$$\begin{aligned} H(-\epsilon) &= F(-\epsilon, \epsilon) - G(-\epsilon, \epsilon) - F(-\epsilon, -\epsilon) + G(-\epsilon, -\epsilon) \\ &= F(-\epsilon, \epsilon) - F(\epsilon, -\epsilon) - F(-\epsilon, -\epsilon) + F(-\epsilon, -\epsilon) \\ &= F(-\epsilon, \epsilon) - F(\epsilon, -\epsilon). \end{aligned}$$

which implies $H(\epsilon) = H(-\epsilon)$, contradicting (3.20). □

Definition 3.7. We say that $f : U \rightarrow \mathbb{R}$ is of class \mathcal{C}^k , for $k = 1, 2, \dots$, and we write $f \in \mathcal{C}^k(U)$, if the partial derivatives of order k

$$\partial_{i_1 i_2 \dots i_k} f(x)$$

exist for each $x \in U$ and are continuous. We say that f is of class \mathcal{C}^0 , or simply of class \mathcal{C} , and write $f \in \mathcal{C}^0(U)$ (or $f \in \mathcal{C}(U)$, respectively), if f is continuous.

We say that f is of class \mathcal{C}^∞ , and denote it as $f \in \mathcal{C}^\infty(U)$, if all partial derivatives of any order exist. That is, $f \in \mathcal{C}^k(U)$ for all $k \geq 1$. Typically the term **smooth function** is used when referring to a \mathcal{C}^∞ function.

Theorem 3.12 can be extended to higher-order derivatives.

Corollary 3.13. If $f \in \mathcal{C}^k(U)$ and $\sigma : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$ is a permutation, then

$$\partial_{i_1 i_2 \dots i_k} f(x) = \partial_{i_{\sigma(1)} i_{\sigma(2)} \dots i_{\sigma(k)}} f(x)$$

for all $x \in U$ and any k -multi-index (i_1, i_2, \dots, i_k) .

3.7 Taylor approximation

By definition, a differentiable function can be well approximated locally by a linear function. When derivatives of higher order are available, the linear approximation can be improved by introducing a polynomial of higher order. Approximations of this sort are at the heart of many algorithms in both pure and applied mathematics and being able to determine when such an approximation is available and also to estimate the error associated is very relevant. We will start with a version of Taylor's theorem for real-valued functions of a single variable and will then use this result to extend the approximation to functions from \mathbb{R}^n to \mathbb{R} .

Theorem 3.14 (Taylor's theorem). Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function such that every derivative $f^{(k)}(x_0)$ of order $1 \leq k \leq n$ exists at a point $x_0 \in \mathbb{R}$. Define

$$a_k := \frac{f^{(k)}(x_0)}{k!} \quad \text{and} \quad P_{n, x_0}(x) := \sum_{k=0}^n a_k (x - x_0)^k.$$

The polynomial P_{n, x_0} is known as **Taylor's polynomial** of f around x_0 and satisfies the following properties

1. $P_{n,x_0}^{(k)}(x_0) = f^{(k)}(x_0)$ for every $1 \leq k \leq n$.
2. $f(x) - P_{n,x_0}(x) = o((x - x_0)^n)$.

In addition, if every derivative of order $1 \leq k \leq n + 1$ exists on an interval $[x_0, x]$, then

3. $f(x) - P_{n,x_0}(x) = \frac{f^{(n+1)}(t)}{(n+1)!}(x - x_0)^{n+1}$ for some $t \in (x_0, x)$.

Proof. We start by computing the k -th derivative of Taylor's polynomial, and observing that the first $k - 1$ terms of the sum (being all polynomials of degree strictly less than k) will yield 0 after k successive differentiations, therefore for any x ,

$$P_{n,x_0}^{(k)}(x) = \sum_{i=k}^n a_i \cdot (i) \cdot (i-1) \cdots (i-k)(x - x_0)^{i-k} = \sum_{i=k}^n a_i \frac{i!}{(i-k)!} (x - x_0)^{i-k}. \quad (3.21)$$

Substituting $x = x_0$ in the expression above sees all but the very first term in the sum vanishing (since the term with $i = k$ is a constant), leading to

$$P_{n,x_0}^{(k)}(x_0) = a_i \cdot i! = \frac{f^{(k)}(x_0)}{i!} i! = f^{(k)}(x_0), \quad (3.22)$$

which proves point 1.

For the second point, we must show that

$$\begin{aligned} \lim_{x \rightarrow x_0} \frac{f(x) - P_{n,x_0}(x)}{(x - x_0)^n} &= \lim_{x \rightarrow x_0} \frac{f(x) - \sum_{k=0}^n a_k (x - x_0)^k}{(x - x_0)^n} \\ &= \lim_{x \rightarrow x_0} \frac{f(x) - \sum_{k=0}^{n-1} a_k (x - x_0)^k}{(x - x_0)^n} - \frac{f^n(x_0)}{n!} \\ &= 0. \end{aligned}$$

To avoid very long expressions, we will define

$$Q(x) := \sum_{k=0}^{n-1} a_k (x - x_0)^k \quad \text{and} \quad G(x) := (x - x_0)^n.$$

With this notation, our goal can be restated as proving that

$$\lim_{x \rightarrow x_0} \frac{f(x) - Q(x)}{G(x)} = \frac{f^n(x_0)}{n!}.$$

Since all three functions f , Q , and G are continuous, and $Q(x_0) = a_0 = f(x_0)$, while $G(x_0) = 0$, it follows that

$$\lim_{x \rightarrow x_0} \frac{f(x) - Q(x)}{G(x)} \quad \text{is indeterminate of the form} \quad \frac{0}{0},$$

therefore we will use L'Hôpital's rule to compute the limit. However, since Q consists of the first $n - 1$ terms of the Taylor polynomial of f at the point x_0 , the discussion right before equation (3.22) implies that

$$Q^{(k)}(x_0) = f^{(k)}(x_0) \quad \text{for all } k \leq n - 1,$$

which in turn implies that

$$\lim_{x \rightarrow x_0} f^{(k)}(x) - Q^{(k)}(x) = 0 \quad \text{for all } k \leq n - 1.$$

At the same time

$$G^k(x) = \frac{n!}{(n-k)!} (x-x_0)^{n-k} \implies \lim_{x \rightarrow x_0} G^k(x) = 0 \quad \text{for all } k \leq n-2.$$

Hence, for all $0 \leq k \leq n-2$ the limits

$$\lim_{x \rightarrow x_0} \frac{f^{(k)}(x) - Q^{(k)}(x)}{G^{(k)}(x)}$$

are all also indeterminate of the form $0/0$. Yet, for $k = n-1$ we have

$$G^{(n-1)}(x) = n!(x-x_0) \quad \text{and} \quad Q^{(n-1)}(x) = a_{n-1}(n-1)! = \frac{f^{(n-1)}(x_0)}{(n-1)!} (n-1)! = f^{(n-1)}(x_0).$$

Therefore

$$\begin{aligned} \lim_{x \rightarrow x_0} \frac{f(x) - Q(x)}{G(x)} &= \lim_{x \rightarrow x_0} \frac{f^{(n-1)}(x) - Q^{(n-1)}(x)}{G^{(n-1)}(x)} \\ &= \lim_{x \rightarrow x_0} \frac{f^{(n-1)}(x) - f^{(n-1)}(x_0)}{n!(x-x_0)} \\ &= \frac{f^{(n)}(x_0)}{n!}, \end{aligned}$$

where in the final step we used the fact that f has n derivatives at x_0 . This proves point 2.

Finally, the statement for point 3 gives us the additional assumptions of f being $n+1$ times differentiable on an interval $[x_0, x]$. These assumptions will allow us to perturb the point around which the polynomial is constructed. Since now the hypotheses required for points 1 and 2 are valid for any $t \in [x_0, x]$, we construct the Taylor polynomial of f for an arbitrary $t \in [x_0, x]$ and define the discrepancy between the function f and its Taylor polynomial at t when evaluated at any point $x \in [x_0, x]$ by

$$R_n(x, t) := f(x) - P_{n,t}(x).$$

We will refer to this difference as the **remainder**. With this notation, and writing $P_{n,t}(x)$ explicitly we have

$$f(x) = f(t) + f'(t)(x-t) + \frac{f''(t)}{2}(x-t)^2 + \dots + \frac{f^{(n)}(t)}{n!}(x-t)^n + R_n(x, t).$$

Differentiating both sides with respect to t yields

$$0 = f'(t) - f'(t) + f''(t)(x-t) - \frac{f''(t)}{2}2(x-t) + \frac{f'''(t)}{2}(x-t)^2 - \dots + \frac{f^{(n+1)}(t)}{n!}(x-t)^n + R'_n(x, t),$$

which, due to the telescopic cancellation, leads to

$$R'_n(x, t) = -\frac{f^{(n+1)}(t)}{n!}(x-t)^n. \quad (3.23)$$

We now observe that both $R_n(x, t)$ and $g(x, t) := (x - t)^{n+1}$, as functions of t are continuous and differentiable over $[x_0, x]$, so we take apply the generalized mean value theorem (Theorem A.6 in the appendices) over the interval $[x_0, x]$ to this pair of functions to obtain the existence of a point $t \in (x_0, x)$ such that

$$\frac{R_n(x, x_0) - R(x, x)}{g(x, x_0) - g(x, x)} = \frac{R'_n(x, t)}{g'(x, t)} = \frac{-\frac{f^{(n+1)}(t)}{n!}(x-t)^n}{-(n+1)(x-t)^n} = \frac{f^{(n+1)}(t)}{(n+1)!}.$$

Observing that $R(x, x) = 0 = g(x, x)$, the expression above leads to

$$R_n(x, x_0) = \frac{f^{(n+1)}(t)}{(n+1)!}(x-x_0)^{n+1}.$$

□

A few remarks are in order:

- Taylor's polynomial is in fact **the only** polynomial of degree n that satisfies property 1.
- Property 2 is the basis for polynomial approximation of nonlinear functions: whenever a function has derivatives of very high order, the error made when replacing it for its associated Taylor polynomial decays very rapidly as the evaluation point approaches x_0 .
- Note that the hypotheses for point 3 are much stronger: the previous two points require n derivatives **at the single point** x_0 , while point 3 requires $n + 1$ derivatives **on a neighborhood** of x_0 .
- The difference $f(x) - P_{n,x_0}(x)$ is called the remainder, and there are several different representations for it. The one given in point 3 is known as **Lagrange's** form of the remainder.
- If $f \in C^\infty$, we can compute as many terms as we like, and the series thus obtained is called the Taylor series around x_0 . If this series converges to $f(x)$, i.e., if $R_k(x) \rightarrow 0$ around x_0 , then we say that f is **real analytic** at x_0 .

Not every C^∞ function is analytic. If a function is of class C^∞ , it is not necessarily analytic. In fact, it is possible that the expansion converges to a limit different from $f(x)$, as shown by the following example.

Consider the function $f : \mathbb{R} \rightarrow \mathbb{R}$ given by

$$f(x) = \begin{cases} e^{-1/x^2}, & x \neq 0, \\ 0, & x = 0. \end{cases}$$

It is not that hard to prove that f is of class C^∞ , but $f^{(k)}(0) = 0$ for all k . Thus, the Taylor expansion of f around 0 is identically 0, but clearly $f(x) \neq 0$ for all $x \neq 0$. This shows that its Taylor series does not converge to $f(x)$ for any point $x \neq 0$.

We can now address the multi-variable version of Taylor's theorem.

Theorem 3.15 (Multivariable Taylor's theorem). *Let $U \subset \mathbb{R}^n$ be open, $x_0 \in U$, $\epsilon > 0$ be such that $B_\epsilon(x_0) \subset U$, and let the function $f : U \rightarrow \mathbb{R}$ be of class $C^k(U)$. Then, for $x \in B_\epsilon(x_0)$, there exists y along the line segment connecting x_0 and x such that*

$$\begin{aligned} f(x) = & f(x_0) + \sum_{i=1}^n \partial_i f(x_0)(x^i - x_0^i) + \frac{1}{2} \sum_{j=1}^n \sum_{i=1}^n \partial_{ij} f(x_0)(x^i - x_0^i)(x^j - x_0^j) + \dots \\ & \dots + \frac{1}{(k-1)!} \sum_{i_1, i_2, \dots, i_{k-1}=1}^n \partial_{i_1 i_2 \dots i_{k-1}} f(x_0) \prod_{l=1}^{k-1} (x^{i_l} - x_0^{i_l}) + R_k(x, x_0), \end{aligned} \quad (3.24)$$

where the remainder is given by

$$R_k(x, x_0) := \frac{1}{k!} \sum_{i_1, i_2, \dots, i_k=1}^n \partial_{i_1 i_2 \dots i_k} f(y) \prod_{l=1}^k (x^{i_l} - x_0^{i_l}).$$

The polynomial in the right hand side of equation (3.24) is called the **Taylor polynomial** of f around x_0 .

Proof. We start by defining $\phi : [0, 1] \rightarrow \mathbb{R}$ as

$$\phi(t) := f(x_0 + t(x - x_0)) = f(x_0^1 + t(x^1 - x_0^1), \dots, x_0^n + t(x^n - x_0^n)).$$

Then $\phi(0) = f(x_0)$ and $\phi(1) = f(x)$. Moreover, since $f \in \mathcal{C}^k(U)$, we have $\phi \in \mathcal{C}^k((0, 1))$. Therefore, building a Taylor polynomial for ϕ around 0 and evaluating it at 1 is equivalent to building a Taylor polynomial for f around x_0 and evaluating it at x .

By the chain rule,

$$\phi^{(k)}(t) = \sum_{i_1, i_2, \dots, i_k=1}^n \partial_{i_1 i_2 \dots i_k} f(x_0 + t(x - x_0)) \prod_{l=1}^k (x^{i_l} - x_0^{i_l}). \quad (3.25)$$

By Taylor's theorem, there exists $c \in (0, 1)$ such that

$$\phi(1) = \phi(0) + \phi'(0) + \dots + \frac{\phi^{(k-1)}(0)}{(k-1)!} + R_k, \quad (3.26)$$

where

$$R_k = R_k(1, 0) = \frac{\phi^{(k)}(c)}{k!}.$$

The result follows from (3.26) by simply substituting $\phi(1) = f(x)$, $\phi(0) = f(x_0)$ and $\phi^{(k)}(0)$ in terms of f as given by (3.25). \square

Just like in the one dimensional case, if the hypothesis of $f \in \mathcal{C}^k(U)$ is relaxed into f having k derivatives at the point x_0 , only the explicit expression for the remainder is lost. The polynomial and its first k derivatives will still coincide with the function and its first k derivatives at the point x_0 , and the difference $f(x) - P(x)$ will still behave as $o(|x - x_0|^n)$.

Similarly, the requirement of the evaluation point being contained in a ball that is completely contained in U is only needed for the explicit expression of the remainder. If the domain U is not convex, then the point y in the expression of the remainder (located along the segment connecting x_0 to x) may fall outside of U (see Figure 1.3). Since U is the domain of differentiability of f and the expression requires the $k+1$ derivative, the explicit form is lost in this case. This is the reason why the multivariable Taylor's theorem is sometimes stated with the hypothesis of U being convex instead of the requirement of x belonging to a ball contained in U . However, even when the domain is non convex or x is not inside of such a ball, as long as the expansion and evaluation points x_0, x lie within U , the error in the polynomial approximation will remain $o(|x - x_0|^n)$.

The theorem can also be extended in a very similar way to functions from \mathbb{R}^n to \mathbb{R}^m . The idea is to apply the result above to each component of f . Notation is absolutely nightmarish, but the concepts remain identical.

The first three terms in the Taylor polynomial (which are the most commonly used) are often denoted using the slightly different (and perhaps more familiar) notation

$$f(x) = f(x_0) + \nabla f(x_0)^\top (x - x_0) + \frac{1}{2} (x - x_0)^\top H f(x_0) (x - x_0) + \dots + R_k(x, x_0),$$

where the matrix of second partial derivatives

$$Hf(x_0) = \begin{bmatrix} \partial_{11}f(x_0) & \partial_{12}f(x_0) & \cdots & \partial_{1n}f(x_0) \\ \partial_{21}f(x_0) & \partial_{22}f(x_0) & \cdots & \partial_{2n}f(x_0) \\ \vdots & \ddots & \ddots & \vdots \\ \partial_{n1}f(x_0) & \partial_{n2}f(x_0) & \cdots & \partial_{nn}f(x_0) \end{bmatrix} \quad (3.27)$$

is known as the **Hessian**. It is in fact the matrix representation of the derivative of the derivative of f , so that $Hf(x_0) = D(Df)(x_0)$ and is sometimes also denoted by $\nabla^2 f(x_0)$, especially in the mathematical literature—in physics and engineering the symbol ∇^2 is reserved for a different second order differential operator: the Laplacian.

3.8 Exercises

1. Prove proposition 3.1.
2. Prove Proposition 3.3.
3. Let $U \subset \mathbb{R}^n$ be open and let $f, g : U \rightarrow \mathbb{R}$ be such that f is continuous at $x_0 \in U$, g is differentiable at x_0 , and $g(x_0) = 0$. Show that fg is differentiable at x_0 .
4. Compute the derivative and the Jacobian of each of the following functions using the chain rule:
 - a) $(x, y) \mapsto (x^2 - y^2, 2xy)$, at each point $(x_0, y_0) \in \mathbb{R}^2$;
 - b) $(x, y) \mapsto (\sin(x^2 + xy + y^2), e^{xy})$, at each point $(x_0, y_0) \in \mathbb{R}^2$.
5. Repeat the previous exercise using partial derivatives.
6. If we extend the definition of directional derivatives to vectors u, v that are not necessarily unitary, prove that if f is differentiable at x_0 and $t \in \mathbb{R}$ it holds
 - (a) $D_{tu}f(x_0) = tD_u f(x_0)$,
 - (b) $D_{u+v}f(x_0) = D_u f(x_0) + D_v f(x_0)$.
7. We say that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is homogeneous of degree α if $f(tx) = t^\alpha f(x)$, for $x \in \mathbb{R}^n$, $t > 0$.
If f is also differentiable, show Euler's formula

$$\sum_{i=1}^n x_i D_i f(x) = \alpha f(x).$$

8. If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is differentiable and $f(0) = 0$, prove that there exist functions $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$f(x) = \sum_{i=1}^n x_i g_i(x).$$

9. Prove Proposition 3.5.
10. Prove Proposition 3.6.
11. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear function that is invertible. Prove that the inverse function f^{-1} is also linear.

12. We will prove the version of the contraction mapping theorem stated in Theorem 3.9. Let $0 < c < 1$ and $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a function such that

$$|f(x) - f(y)| < c|x - y| \quad \text{and} \quad \forall x, y \in X.$$

- (a) Pick an arbitrary $x_1 \in \mathbb{R}^n$ and define a sequence $\{x_k\} \subset \mathbb{R}^n$ iteratively by

$$x_{k+1} = f(x_k).$$

Prove that this sequence converges to some $x_* \in \mathbb{R}^n$.

- (b) Prove that $f(x_*) = x_*$. (The point x_* is known as a fixed point of f)
 (c) Prove that the fixed point is unique.

13. We recall the simplified version of the inverse function theorem proven in the notes (Theorem 3.10):

Let $U \subset \mathbb{R}^n$ be open and $x_0 \in U$. If the function $f : U \rightarrow \mathbb{R}^m$ is differentiable on U , continuously differentiable at $x_0 \in \mathbb{R}^n$, and such that the following hypotheses are satisfied

- (i) $x_0 = 0$,
 (ii) $f(x_0) = 0$,
 (iii) $Df(x_0) = I$ (the identity function).

Then there exist open sets $V \subset \mathbb{R}^n$ and $W \subset \mathbb{R}^m$ such that $x_0 \in V$, $f(x_0) \in W$ and $f : V \rightarrow W$ is invertible. Moreover, the inverse function $f^{-1} : W \rightarrow V$ is differentiable on W and $Df^{-1}(y) = I$ for all $y \in W$.

In this problem we will use the result above to prove the general version of the theorem.

- (a) Use the simplified result above, to prove that the hypothesis (i) can be dropped; i.e. the point x_0 can be arbitrary. Do that by defining $g(x) := f(x + x_0)$ and applying to this function. Verify carefully that the hypothesis of are verified for g .
 (b) Use the more general result proven in the previous step, to show that the hypothesis (ii) can be dropped; i.e. the value $f(x_0)$ can be arbitrary. Do that by defining $g(x) := f(x) - f(x_0)$, verifying that the hypotheses of the previous case are verified and then applying the result from the previous step.
 (c) Use the result from the previous step to show that the hypothesis (iii) can be replaced by “ $Df(x_0)$ is invertible”. Do that by defining a function $g(x) := [Df(x_0)]^{-1} \circ f(x)$, and applying the result from the previous step. Prove that in this case the expression for the derivative of the inverse must be changed to

$$Df^{-1}(f(x_0)) = [Df(x_0)]^{-1}.$$

14. Let $U \subset \mathbb{R}^n$ be an open set and let $f : U \rightarrow \mathbb{R}^n$ be injective and continuously differentiable such that $\det Df(x) \neq 0$ for all $x \in U$.
- (a) Show that $f(U)$ is open and that $f^{-1} : f(U) \rightarrow U$ is differentiable.
 (b) Show that $f(V)$ is open for every open $V \subset U$.
15. (a) Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be continuously differentiable. Show that f is not injective. (Hint: Consider the function $g(x, y) = (f(x, y), y)$.)
 (b) Generalize this result to continuously differentiable functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, with $m < n$.

16. (a) Show that if $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfies $f'(x) \neq 0$ for all $x \in \mathbb{R}$, then f is injective.

(b) However, show that the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by

$$f(x, y) = (e^x \cos y, e^x \sin y)$$

satisfies $\det f'(x, y) \neq 0$ for all $(x, y) \in \mathbb{R}^2$, but is not injective.

17. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be continuously differentiable such that there exists $c > 0$ such that

$$|f(x) - f(y)| \geq c|x - y|$$

for all $x, y \in \mathbb{R}^n$. Show that:

(a) f is injective;

(b) $\det f'(x) \neq 0$ for all $x \in \mathbb{R}^n$; and

(c) $f(\mathbb{R}^n) = \mathbb{R}^n$. (Hint: consider the function $g(x) = |y - f(x)|^2$).

18. Use the implicit function theorem 3.11 to prove the inverse function theorem.

19. Consider the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$f(x, y) := \begin{cases} \frac{xy(x^2 - y^2)}{x^2 + y^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0). \end{cases}$$

(a) Prove that its first order partial derivatives are given by

$$\partial_x f(x, y) = \begin{cases} \frac{x^4 y + 4x^2 y^3 - y^5}{(x^2 + y^2)^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0), \end{cases}$$

$$\partial_y f(x, y) = \begin{cases} \frac{x^5 - 4x^3 y^2 - xy^4}{(x^2 + y^2)^2}, & (x, y) \neq (0, 0), \\ 0, & (x, y) = (0, 0). \end{cases}$$

You will have to use the limit definition for the values at the origin.

(b) Prove that $\partial_{x,y} f(x, y) = \partial_{y,x} f(x, y)$ for all $(x, y) \neq (0, 0)$, but $\partial_{x,y} f(0, 0) \neq \partial_{y,x} f(0, 0)$.

You will have to use the limit definition for the values at the origin.

20. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be given by

$$f(x) = \begin{cases} e^{-1/x^2}, & x \neq 0, \\ 0, & x = 0. \end{cases}$$

Show by induction that, for $x \neq 0$ and $k = 1, 2, \dots$,

$$f^{(k)}(x) = \frac{P_k(x)}{Q_k(x)} e^{-1/x^2},$$

where P_k and Q_k are polynomials. Conclude that $f \in C^\infty$ and that $f^{(k)}(0) = 0$ for all k .

21. Use the previous exercise to find a function $F : \mathbb{R} \rightarrow \mathbb{R}$, $F \in C^\infty$, such that $F(x) = 0$ for $x \leq 0$ and $F(x) > 0$ for $x > 0$.

22. Find a function $G : \mathbb{R} \rightarrow \mathbb{R}$ such that $G(x) > 0$ for $x \in (-1, 1)$ and $G(x) = 0$ for $|x| \geq 1$.

23. **Newton's method.**

Let $U \subset \mathbb{R}^n$ be open, $x_* \in U$ and $B_q(x_*) \subset U$. Let $f : U \rightarrow \mathbb{R}^m$ be twice continuously differentiable on U and such that $Df(x)$ is invertible for all $x \in U$. Let $y_* := f(x_*)$ and define a function $N : B_q(x_*) \rightarrow \mathbb{R}^n$ by

$$N(x) := x + [Jf(x)]^{-1} (y_* - f(x)).$$

(a) Show that there exists $0 < \epsilon$ and $0 < C < \infty$ such that for all $0 \leq r < \epsilon$,

$$|x - x_*| \leq r \implies |N(x) - x_*| < Cr^2.$$

(b) Show that if

$$|x_1 - x_*| < \min\{\epsilon, \frac{1}{2C}\},$$

then the sequence $x_{n+1} := N(x_n)$ converges very rapidly to x_* .

The function $N(x)$ is sometimes referred to as a *Newton update* and is used iteratively to obtain a solution to the equation $f(x) = y_*$ (whose solution is x_*). Part (a) asks you to prove that if a point x is *close enough* to the solution x_* , then the point $N(x)$ obtained when *feeding* x to the Newton update will be *much closer* to the solution x_* —recall that if $r < 1$ then $r^2 \ll 1$. Part (b) shows that if the *initial guess* x_1 is *close enough* to the solution, then the approximation obtained by applying $N(\cdot)$ iteratively will not only converge to the solution, but will in fact do so very quickly. The ball centered at x_* and with radius equal to $\min\{\epsilon, \frac{1}{2C}\}$ is often called the *basin of convergence* of the method.

When it works (i.e. when the hypothesis of the theorem are satisfied), Newton's method is essentially unbeatable as a tool for finding roots of nonlinear equations. However, the hypothesis of the theorem are quite restrictive: the function must have two continuous derivatives, the Jacobian must be invertible **and** the initial guess must be already quite close to the solution (this last point is often overlooked when singing the praises of the method). There are other algorithms for root finding that are much less restrictive, at the cost of a slower rate of convergence. There is no free lunch.

As an example of the efficiency of Newton's method when the function f is nice enough:

(c) Show that if f is linear, Newton's method will converge to the solution of the equation $f(x) = y_*$ after a single iteration regardless of the initial guess.

Chapter 4

Maxima and minima

4.1 A refresher of proven results

We will now focus on the extreme values of real-valued functions of several variables of the form $f : \mathbb{R}^n \rightarrow \mathbb{R}$. This function is often referred to as the **objective function**. We start by recalling the Definition 3.4 on **local extrema**

Definition (Extrema). Let A be a set in \mathbb{R}^n and $x_0 \in A$. We say that $f : A \rightarrow \mathbb{R}$ has a **minimum** at x_0 if

$$f(x_0) \leq f(x) \quad \text{for all } x \in A.$$

We say that f has a **local minimum** at x_0 if there exists a ball $B_\epsilon(x_0) \subset U$ such that

$$f(x_0) \leq f(x) \quad \text{for all } x \in B_\epsilon(x_0) \cap A.$$

Similarly, we say that f has a **maximum** at x_0 if

$$f(x_0) \geq f(x) \quad \text{for all } x \in A,$$

and we say that f has **local maximum** at x_0 if there exists a ball $B_\epsilon(x_0) \subset U$ such that

$$f(x_0) \geq f(x) \quad \text{for all } x \in B_\epsilon(x_0) \cap A.$$

Since every subset $A \subset \mathbb{R}^n$ can be decomposed (see Proposition B.6) as

$$A = \partial A \cup A^\circ$$

where the boundary ∂A is closed and the interior A° is open, and $\partial A \cap A^\circ = \emptyset$, we can focus our efforts separately on the cases where the extrema belong to the interior and when they belong to the boundary.

Regarding the case when the extrema belong to the open interior, in Proposition 3.5 we have also proven that:

Proposition: If $f : U \rightarrow \mathbb{R}$ has a local minimum or maximum at x_0 and its partial derivatives exist, then $\partial_{x_i} f(x_0) = 0$, for $i = 1, \dots, n$.

So that, if the extrema of a differentiable function belong to the interior of a set, they must necessarily be critical points of the function. We would like to have also a sufficiency condition for a point to be an extremum of a function.

4.2 Second derivative criterion

Lets consider a function $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ belonging to the class $\mathcal{C}^2(U)$. By Taylor's Theorem 3.14, for any $x, x_0 \in U$ we can write

$$f(x) - f(x_0) = \nabla f(x_0) \cdot (x - x_0) + \frac{1}{2}(x - x_0)^\top Hf(x_0)(x - x_0) + o(|x - x_0|^2),$$

where the matrix $Hf(x_0)$ is the *Hessian* matrix defined in (3.27). If the point x_0 is a critical point of f then, from 3.5, the expression above simplifies into

$$f(x) - f(x_0) = \frac{1}{2}(x - x_0)^\top Hf(x_0)(x - x_0) + o(|x - x_0|^2).$$

This equality makes it evident that the interplay between the Hessian and the remainder of the Taylor polynomial will determine the sign of the difference $f(x) - f(x_0)$. Thus, the Hessian must contains information regarding whether x_0 is a local extremum. We will now explore its properties.

Definition 4.1. A function $F : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be a **quadratic form** if there exists a matrix $A \in \mathbb{R}^{n \times n}$ such that

$$F(x) = x^\top Ax \quad \text{for all } x \in \mathbb{R}^n.$$

In addition, we say that F is **positive semi-definite**, **positive definite**, **negative semi-definite**, or **negative definite** if the matrix associated to F is either (See Definition C.1 in the appendix).

Quadratic forms receive their name because they are in fact second degree polynomials in the components of their arguments. Some of their properties are summarized in the following proposition. Property 4 will be of particular relevance when developing a multi-variable analogue to the *second derivative criterion*.

Proposition 4.1. Let $F_A : \mathbb{R}^n \rightarrow \mathbb{R}$ be the quadratic form defined by

$$F_A(x) := x^\top Ax \quad \text{for all } x \in \mathbb{R}^n.$$

Then

1. F_A is continuous at every $x_0 \in \mathbb{R}^n$.
2. If $B := \frac{1}{2}(A + A^\top)$, then $F_A(x) = F_B(x)$ for all $x \in \mathbb{R}^n$.
3. The matrix B defined above, is the only symmetric matrix satisfying point 2.
4. If F_A is positive definite (resp. negative definite), then there exists a constant $m > 0$ (resp. $M > 0$) such that

$$F_A(x) \geq m|x|^2 \quad (\text{resp. } F_A(x) \leq M|x|^2) \quad \text{for all } x \in \mathbb{R}^n. \quad (4.1)$$

Proof. The proof of the first three points is left as an exercise, For the fourth point we will prove only the case of a positive definite quadratic form, as the negative definite case follows an analogous argument.

We first observe that if $x = 0$ then the inequality $F_A(x) \geq m|x| = m \cdot 0 = 0$ holds trivially for any positive m . Hence, we will assume that $x \neq 0$ and will consider the unit vector $\hat{x} = x/|x|$. The unit spherical shell

$$S_{n-1} := \{x \in \mathbb{R}^n : |x| = 1\}$$

is both bounded and closed, and thus is compact by Theorem 1.3. This in turn implies, by the extreme value theorem and the continuity of F_A , proven in point 1, that there exists a minimum value $m > 0$ such that

$$F_A(\hat{x}) = \hat{x}^\top A\hat{x} \geq m \quad \text{for all } \hat{x} \in S_{n-1}.$$

The fact that the minimum value is strictly positive for $x \neq 0$ follows from the positive definiteness of A . Keeping this observation in mind, we take compute

$$m \leq F_A(\hat{x}) = \hat{x}^\top A \hat{x} = \left(\frac{x^\top}{|x|} \right) A \left(\frac{x}{|x|} \right) = \frac{1}{|x|^2} F_A(x),$$

from which the desired inequality 4.1 follows. \square

We can now use the *Hessian* to determine whether a critical point is a local extremum.

Theorem 4.1. *Let $U \subset \mathbb{R}^n$ be open and $f : U \rightarrow \mathbb{R}$ be of class $\mathcal{C}^2(U)$. If $x_0 \in U$ is a critical point of f such that $Hf(x_0)$ is positive definite (resp. negative definite), then x_0 is a local minimum of f .*

Proof. We will prove only the case where the Hessian $Hf(x_0)$ is positive definite, as the argument for the negative definite case is completely analogous. We go back to the second order Taylor expansion of f around the critical point x_0

$$f(x) - f(x_0) = \frac{1}{2}(x - x_0)^\top Hf(x_0)(x - x_0) + o(|x - x_0|^2),$$

and recall that the fact that the remainder is *little oh* of $|x - x_0|^2$ implies that there exists $\delta > 0$ such that, if $|x - x_0| < \delta$, then

$$|o(|x - x_0|^2)| < \frac{1}{2}m|x - x_0|, \quad (4.2)$$

where $m > 0$ is the constant such that $x^\top Hf(x_0)x > m|x|^2$ guaranteed by the positive definiteness of the Hessian and Proposition 4.1. We then pick x such that $|x - x_0| < \delta$ and compute

$$\begin{aligned} f(x) - f(x_0) &= \frac{1}{2}(x - x_0)^\top Hf(x_0)(x - x_0) + o(|x - x_0|^2) \\ &\geq \frac{1}{2}m|x - x_0|^2 + o(|x - x_0|^2) && \text{By (4.1)} \\ &\geq \frac{1}{2}m|x - x_0|^2 - |o(|x - x_0|^2)| \\ &> \frac{1}{2}m|x - x_0|^2 - \frac{1}{2}m|x - x_0|^2 && \text{By (4.2)} \\ &> 0. \end{aligned}$$

Proving that x_0 is a local minimum in the ball $B_\delta(x_0)$. \square

The converse of this proposition is not true (i.e. a \mathcal{C}^2 function may have a local extremum at a point where the Hessian quadratic form vanishes) as is exemplified by the function

$$f : \mathbb{R}^2 \rightarrow \mathbb{R} \quad \text{given by} \quad f(x, y) = x^4 + y^4.$$

This means that at a local minimum (resp. maximum) the Hessian of a \mathcal{C}^2 function is not necessarily positive definite (resp. negative definite), however, as we will now prove at a local extremum the Hessian must be sign semi-definite.

Theorem 4.2. *If $U \subset \mathbb{R}^n$ is open, $f : U \rightarrow \mathbb{R}$ is of class $\mathcal{C}^2(U)$ and f has a local minimum (resp. maximum) at a point $x_0 \in U$, then the Hessian $Hf(x_0)$ is positive semi-definite (resp. negative semi-definite).*

Proof. We will only provide the full details for the case where x_0 is a local minimizer, as the argument for a local maximizer is completely analogous. Since x_0 is a local minimizer, there exists some $\epsilon > 0$ such that for every $x \in B_\epsilon(x_0)$ we have

$$f(x_0) \leq f(x).$$

Let \hat{x} be a unit vector; we define the function $\phi : (-\epsilon, \epsilon) \rightarrow \mathbb{R}$ as

$$\phi(t) := f(x_0 + t\hat{x}).$$

A simple application of the chain rule shows that

$$\phi'(t) = \nabla f(x_0 + t\hat{x}) \cdot \hat{x} \quad \text{and} \quad \phi''(t) = \hat{x}^\top Hf(x_0 + t\hat{x}) \hat{x}.$$

On the other hand, from the definition of ϕ and the fact that if $t \in (-\epsilon, \epsilon)$ then $x_0 + t\hat{x} \in B_\epsilon(x_0)$ it follows that

$$\phi(t) = f(x_0 + t\hat{x}) \geq f(x_0) = \phi(0)$$

and therefore ϕ has a local minimum at $t = 0$. We know from analysis of a single variable that, if a function $f : \mathbb{R} \rightarrow \mathbb{R}$ has a local minimum at a point x_0 and it is twice continuously differentiable at x_0 , then $f''(x_0) \geq 0$. Therefore, since ϕ has a local minimum at 0 it follows that

$$0 \leq \phi''(0) = \hat{x}^\top Hf(x_0) \hat{x} \implies 0 \leq x^\top Hf(x_0) x \quad \forall x = |x|\hat{x} \in \mathbb{R}^n.$$

This proves that the Hessian at x_0 is positive semi-definite. □

4.3 Equality constraints and level sets

Often we are interested in finding the extreme values of a function $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ restricted to a subset of its domain. For instance, we may want the sum of the arguments to be a given constant value c , or we may want the argument to have a prescribed Euclidean norm c , etc. The additional conditions imposed on the arguments are referred to as **constraints** and can often be expressed as an equation of the form

$$g(x) = 0,$$

where the function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ expresses the condition that needs to be satisfied by the arguments. For instance

$$g(x) := c - \sum_{i=1}^n x_i, \quad \text{or} \quad g(x) := c^2 - \sum_{i=1}^n x_i^2$$

in the case of a constant sum or a constant Euclidean norm, respectively. It is possible to have more complex constraints (for instance expressed as inequalities or inclusions) but we will focus here only on the simple case of equality constraints.

For a function $g : \mathbb{R}^n \rightarrow \mathbb{R}$, the set of all points in the domain that get mapped by g to a given, fixed value c is denoted by

$$L_c(g) := \{x \in \mathbb{R}^n : g(x) = c\}$$

and is called the **c -level set** of g . Note that this is simply the inverse image of c under g ; i.e.

$$L_c(g) = g^{-1}(\{c\}).$$

The notation $\{c\}$ has been used above to emphasize that, in the definition of a level set, the value c is considered *as a set*. Note that, even if the level set $L_c(g)$ *always* exists, it may be empty or may be a set of disconnected points. For instance, if $g(x, y) = (\sin x)(\sin y)$ then

$$L_2(g) = \emptyset$$

and

$$L_1(g) = \left\{ (x, y) \in \mathbb{R}^2 : x = \frac{1}{2}(4n + 1)\pi \text{ and } y = \frac{1}{2}(4m + 1)\pi \text{ for } n, m \in \mathbb{Z} \right\}.$$

With this consideration, we see that the problem of finding the extreme points of f subject to the constraint g , is equivalent to finding the extreme points of f over the zero-level set of g .

By now we are very familiar with the process needed to search for extrema of functions *defined over open sets*: the use of Proposition 3.5 which tells us that at a local maximum or minimum, the gradient of a differentiable function must vanish. Hence, if the objective function is differentiable, we compute the gradient and set out to build a list of candidate points by find the locations where the gradient vanishes. Therefore, the reader might wonder what is the big deal about constrained optimization. The answer is simple: *the set where the constraints are satisfied may not be open, and therefore we can not rely on Proposition 3.5 to search for candidate points*.

An effective and elegant way to deal with constrained optimization problems is the **method of Lagrange multipliers**. Before dealing with it rigorously, we will start by introducing the method intuitively and discussing the ideas behind it.

4.4 Lagrange multipliers: an informal analytic motivation

We will start with an informal, non rigorous motivation of the method. We are interested in locating the extreme values (maximum/minimum) of a function $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ of class $\mathcal{C}^1(U)$ subject to $c < n$ equality constraints encoded by c equations of the form

$$g_1(x) = 0, \quad g_2(x) = 0, \quad \dots, \quad g_c(x) = 0,$$

where every one of the constraining functions $g_i : \mathbb{R}^n \rightarrow \mathbb{R}$ is of class $\mathcal{C}^1(U)$. Using our previous knowledge about the extrema of differentiable functions (in particular Proposition 3.5), we would be tempted to look for the points satisfying the conditions

$$\nabla f(x, y) = 0 \quad \text{and} \quad g_1(x) = 0, g_2(x) = 0, \dots, g_c(x) = 0.$$

However, the vanishing gradient condition on the left is equivalent to n equations (one for each component of the gradient) which, together with the c conditions imposed by the constraints, add up to a total of $n + c$ equations. Our unknowns, on the other hand, are the n coordinates of a possible extremum of the problem. Therefore, this approach would lead to an overdetermined system. Moreover, the vanishing gradient condition stemming from Proposition 3.5 reflects the geometric fact that **when allowed to take values over its entire domain of definition** U a differentiable function f *turns around* on every direction at a local extremum and thus all of its partial derivatives must vanish. However, when dealing with constrained problems, **the constraints reduce the accessible regions of the domain** from U to the set

$$U \cap S \quad \text{where} \quad S := \bigcap_{i=1}^c L_0(g_i),$$

and there is no reason to expect that the unconstrained extrema (where the gradient must vanish) will belong to this region. Put informally into words: the gradient ∇f knows only about the geometry of the function f but does not contain information about the constraints.

The discussion above suggests considering an extended function that will be privy to the geometric properties of both the objective function and the constraints. One natural choice is to define a function $F : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ as

$$F(x) := f(x) - \sum_{i=1}^c g_i(x).$$

This function certainly includes the information from f and all the constraints g_i . Moreover, if $x^* \in \mathbb{R}^n$ is a point satisfying the constraints, we see that

$$F(x^*) = f(x^*) - \sum_{i=1}^c g_i(x^*) = f(x^*),$$

so that $F(x) = f(x)$ for all $x \in U \cap S$ and therefore if F attains an extremum on the region $U \cap S$ so will f . We would then be tempted again to make use of Proposition 3.5 and search for the points satisfying

$$\nabla F(x, y) = \nabla f(x) - \sum_{i=1}^c \nabla g_i(x) = 0 \quad \text{and} \quad g_1(x) = 0, g_2(x) = 0, \dots, g_c(x) = 0.$$

Now at least the vanishing gradient condition $\nabla F(x) = 0$ includes geometric information from f and all constraints. However once again we would be facing an overdetermined system, with n equations coming from the vanishing gradient condition, c additional equations coming from the constraints, and only n unknowns: one for each coordinate of the extremum.

The previous attempt seems to be geared in the right direction, as the vanishing gradient condition now includes all the relevant analytic information, and the only problem seems to be that we have too few unknowns, so we try a naïve approach: how about we add just enough additional unknowns to solve the problem of overdetermination? We then define the function $\mathcal{L}(x, \lambda) : U \times V \subset \mathbb{R}^n \times \mathbb{R}^c \rightarrow \mathbb{R}$ by

$$\mathcal{L}(x, \lambda) := f(x) - \sum_{i=1}^c \lambda_i g_i(x), \tag{4.3}$$

where the numbers $\lambda_1, \dots, \lambda_c$ are c unknowns to be determined along with the n coordinates of the extremum point x^* . Just as before, if $x^* \in U \cap S$, so that the constraints are all satisfied, we have that $\mathcal{L}(x^*, \lambda) := f(x^*)$ and thus constrained extrema of f will also be constrained extrema of \mathcal{L} . Moreover, considering the quantities $\lambda_1, \dots, \lambda_c$ as additional variables, the vanishing gradient condition applied to \mathcal{L} yields:

$$\partial_{x_j} \mathcal{L}(x, \lambda) = \partial_{x_j} f(x) - \sum_{i=1}^c \lambda_i \partial_{x_j} g_i(x) = 0 \quad \forall 1 \leq j \leq n, \quad \text{and} \quad \partial_{\lambda_j} \mathcal{L}(x, \lambda) = g_j(x) = 0.$$

These conditions can be written more succinctly in the form

$$\nabla f(x) = \sum_{i=1}^c \lambda_i \nabla g_i(x) \quad \text{and} \quad g_1(x) = 0, g_2(x) = 0, \dots, g_c(x) = 0. \tag{4.4}$$

The first vector equation on the left contains n scalar equations, while the constraint equation on the right contribute with c additional conditions for a total of $n + c$ equations. However, the addition of the new $\lambda_1, \dots, \lambda_c$ unknowns to the original n unknowns from the coordinates of the optimizer means that we now have also $n + c$ unknowns and the system is neither over nor under determined.

The **method of Lagrange multipliers**, says that if the system of equations (4.4) has one or more solutions, then the local extrema of the constrained problem will be amongst them. Thus, the algorithmic process it

so compute the gradients necessary to set up the system (4.4), and then proceeding to solve it. The function $\mathcal{L}(x, \lambda)$ defined in (4.3) is known as the **Lagrangian** while each of the λ_i 's is known as a **Lagrange multiplier**. The system of equations (4.4) is often referred to as the **first order constrained optimality condition**.

Before proceeding to prove under what conditions this process will indeed produce a local extremum to the constrained problem, we will further motivate the validity of the algorithm by looking at it from a geometric perspective.

4.5 A geometric interlude

We say that a set $\Gamma \subset \mathbb{R}^n$ is a **parametrizable curve** if there exists a function $\gamma : \mathbb{R} \rightarrow \mathbb{R}^n$ and an interval $I \subset \mathbb{R}$ such that

$$\gamma(I) = \Gamma.$$

Conversely, if the equality above is satisfied, we say that the function γ is a **parametric curve** or a **parametrization** of the curve Γ . If the parametric curve $\gamma : \mathbb{R} \rightarrow \mathbb{R}^n$ is differentiable at a point t_0 , then the vector $\gamma'(t_0) \in \mathbb{R}^n$ will be tangent to the curve at the point $\gamma(t_0)$. The vector γ' is often referred to as the **velocity** of the parametrization.

It is important to remark that parametrizations of a curve are not unique. For instance, the functions $\gamma_1, \gamma_2, \gamma_3 : [0, 1) \rightarrow \mathbb{R}^2$ given by

$$\gamma_1(t) := (\cos(2\pi t), \sin(2\pi t)), \quad \gamma_2(t) := (\sin(2\pi t), \cos(2\pi t)) \quad \text{and} \quad \gamma_3(t) := (\sin(2\pi t^2), \cos(2\pi t^2))$$

are all parametrizations of the unit circle. However, even though they all describe the same geometric object, they differ—amongst other things—in their velocities.

Since the straight line that locally resembles a differentiable curve the most at any given point is the one whose slope is given by the derivative of the curve at that point, the velocity vector enables us to generalize the notion of orthogonality to curves.

Definition 4.2 (Orthogonality to a curve). Let $n \in \mathbb{R}^n$ be a non-zero vector and $\gamma : I \subset \mathbb{R} \rightarrow \mathbb{R}^n$ be a parametric curve such that $\gamma'(t_0) \neq 0$. We say that the vector n is orthogonal to γ at the point $\gamma(t_0)$ if

$$\gamma'(t_0) \cdot n = 0.$$

The condition of $\gamma'(t_0) \neq 0$ is simply a technical requirement. Since all vectors are orthogonal to the zero vector, attempting to define geometric orthogonality in terms of the inner product with zero is a little futile. The fact that the velocity vector vanishes may not be a problem of the curve itself, but rather of our choice of parametrization. If a curve Γ (as a geometric object) is not discontinuous or have a corner at a point, it is always possible to define a tangent vector there. However, our choice of parametrization γ (as an analytic object, i.e. a function) might be poor in the sense that its velocity vector may vanish at a location where the associated geometric object Γ has a perfectly valid tangent. For instance, as depicted schematically in Figure 4.1, the parametrization γ_3 of the unit circle presented above has vanishing velocity vector for $t = 0$. However, the circle has a perfectly valid tangent line at the point $\gamma_3(0) = (1, 0)$, we simply need to take a better parametrization, such as γ_1 or γ_2 , of the circle to define a tangent at the point $(1, 0)$.

Definition 4.3 (Orthogonality between curves). Let $\gamma_1, \gamma_2 : I \subset \mathbb{R} \rightarrow \mathbb{R}^n$ be parametric curves. We say that γ_1, γ_2 are orthogonal at a point $p \in \mathbb{R}^n$ if:




Figure 4.1: Three different parametrizations of the unit circle, denoted by Γ ; the point $P = (0, 1)$ is marked by a red circle and the velocity vector γ' given by each parametrization is represented by a red arrow. At the point P , the velocity vector for the parametrizations in the left and center panels points left and right respectively. However, for the parametrization on the right panel, the velocity vector at P vanishes.

1. $\gamma_1(t_1) = p$ for some $t_1 \in I$ and $\gamma_2(t_2) = p$ for some $t_2 \in I$.
2. $\gamma_1'(t_1) \neq 0$ and $\gamma_2'(t_2) \neq 0$.
3. $\gamma_1'(t_1) \cdot \gamma_2'(t_2) = 0$.

In words, we say that two curves are orthogonal to each other at a point if they intersect at that point (this is the essence of point 1) and their tangent vectors are orthogonal at the intersection (this is the essence of point 3). Just as in the previous definition, point 2 is a purely technical requirement that forces us to choose good parametrizations for both curves.

We can now prove an interesting geometric connection between the gradient of a function and its level sets.

Proposition 4.2. *Let $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be of class $\mathcal{C}^1(U)$ and consider that the level set $L_c(f)$ admits a parametrization $\gamma : I \subset \mathbb{R} \rightarrow \mathbb{R}^n$. If we let $x_0 = \gamma(t_0)$ and the parametrization γ is differentiable at a point t_0 with $\gamma'(t_0) \neq 0$, then $\nabla f(x_0)$ is orthogonal to the level set $L_c(f)$ at the point x_0 .*

Proof. Since γ is a parametrization of the c -level set of f , it follows that

$$(f \circ \gamma)(t) = c \quad \forall t \in I.$$

Therefore, an application of the chain rule at t_0 yields

$$0 = D(f \circ \gamma)(t_0) = \nabla f(\gamma(t_0)) \cdot \gamma'(t_0) = \nabla f(x_0) \cdot \gamma'(t_0).$$

It then follows from the definition 4.2 that the gradient $\nabla f(x_0)$ is orthogonal to the level set $L_c(f) = \gamma$ at the point x_0 . \square

4.6 Lagrange multipliers: An informal geometric motivation

(To be completed)

4.7 Proof of the Lagrange multipliers theorem

We can now state and prove rigorously the **Lagrange multiplier** theorem.

Theorem 4.3 (Lagrange multipliers). *Let c and ℓ be positive integers so that $n = c + \ell$ and split $\mathbb{R}^n = \mathbb{R}^c \times \mathbb{R}^\ell$. Therefore, every $z \in \mathbb{R}^n$ can be expressed as $z = (x, y)$ with $x \in \mathbb{R}^c$ and $y \in \mathbb{R}^\ell$. Let $U \subset \mathbb{R}^n$ be open and consider an objective function $f : U \rightarrow \mathbb{R}$ such that $f \in \mathcal{C}^1(U)$, and c constraint functions $g_1, \dots, g_c : U \rightarrow \mathbb{R}$ such that $g_1, \dots, g_c \in \mathcal{C}^1(U)$. We will define the subset of \mathbb{R}^n where all the constraints are satisfied as*

$$\mathcal{S} := \{(x, y) \in \mathbb{R}^c \times \mathbb{R}^\ell : g_i(x, y) = 0 \text{ for all } 1 \leq i \leq c < n\}.$$

Let $(x_0, y_0) \in U \cap \mathcal{S} \subset \mathbb{R}^n = \mathbb{R}^c \times \mathbb{R}^\ell$ be such that the set of vectors $\nabla g_1(x_0, y_0), \dots, \nabla g_c(x_0, y_0)$ is linearly independent and $f : U \cap \mathcal{S} \rightarrow \mathbb{R}$ has a local extremum on (x_0, y_0) . Then there exist c constants $\lambda_1, \dots, \lambda_c$ known as **Lagrange multipliers** such that

$$\nabla f(x_0, y_0) = \sum_{i=1}^c \lambda_i \nabla g_i(x_0, y_0).$$

Intuitive idea: The proof can seem technical, but the idea behind it is simple and intuitive, so let us state it here and use it as a beacon throughout the argument. Since the gradients $\nabla g_1(x_0, y_0), \dots, \nabla g_c(x_0, y_0)$ are linearly independent, they constitute a basis for a certain subspace $W^\perp \subset \mathbb{R}^n$ whose dimension is equal to the number of constraints. We will then prove that, at the point (x_0, y_0) , the gradient $\nabla f(x_0, y_0)$ also belongs to the subspace W^\perp and therefore it must be a linear combination of $\nabla g_1(x_0, y_0), \dots, \nabla g_c(x_0, y_0)$. The coefficients of $\nabla f(x_0, y_0)$ with respect to this basis are precisely the Lagrange multipliers. The argument will use several concepts from linear algebra.

Proof. .

Step 1:

We will prove that the set $\{\nabla g_1(x_0, y_0), \dots, \nabla g_c(x_0, y_0)\}$ is a basis for a very particular vector subspace.

Consider the function $G : U \subset \mathbb{R}^n = \mathbb{R}^c \times \mathbb{R}^\ell \rightarrow \mathbb{R}^c$ given by

$$G(x, y) := (g_1(x, y), \dots, g_c(x, y)).$$

Since every g_i is continuously differentiable over U , then $G \in \mathcal{C}^1(U)$ and its Jacobian is given by the matrix

$$JG(x_0, y_0) := \begin{bmatrix} \partial_{x_1} g_1(x_0, y_0) & \cdots & \partial_{x_c} g_1(x_0, y_0) & \partial_{y_1} g_1(x_0, y_0) & \cdots & \partial_{y_\ell} g_1(x_0, y_0) \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \partial_{x_1} g_c(x_0, y_0) & \cdots & \partial_{x_c} g_c(x_0, y_0) & \partial_{y_1} g_c(x_0, y_0) & \cdots & \partial_{y_\ell} g_c(x_0, y_0) \end{bmatrix}$$

which has c rows and $c + \ell = n$ columns. The entries of the i -th row consist of the components of the gradient $\nabla g_i(x_0, y_0)$. Since the gradients $\nabla g_i(x_0, y_0)$ for $1 \leq i \leq c < n$ are linearly independent, the c rows of the matrix are linearly independent. Since the column rank and the row rank of a matrix are equal, it follows that the matrix must have c linearly independent columns. Let us assume that the first c columns are linearly independent (the fact that we can rearrange the columns and relabel variables in this way follows from the Rank Theorem).

Therefore, given that $G \in \mathcal{C}^1(U)$ and that the first c columns of its Jacobian $JG(x, y)$ are linearly independent at the point (x_0, y_0) , the Implicit function theorem guarantees the existence of

- an open domain $V \subset \mathbb{R}^\ell$,
- a $\mathcal{C}^1(V)$ function $h(y) = (h_1(y), \dots, h_c(y)) : V \rightarrow \mathbb{R}^c$,
- and a ball $B_\delta(x_0, y_0) \subset U$,

such that

$$(1) y_0 \in V, \quad (2) h(y_0) = x_0 \quad \text{and} \quad (3) (h(y), y) \in B_\delta(x_0, y_0) \cap \mathcal{S} \text{ for all } y \in V.$$

We define $H : V \subset \mathbb{R}^\ell \rightarrow \mathbb{R}^n = \mathbb{R}^c \times \mathbb{R}^\ell$ by

$$H(y) = (h(y), y),$$

and note that $H \in \mathcal{C}^1(V)$. Therefore, for any $y = (y_1, \dots, y_\ell) \in V$ the derivative $DH(y)$ can be represented by the Jacobian

$$JH(y) = \begin{bmatrix} \partial_{y_1} h_1(y) & \cdots & \partial_{y_\ell} h_1(y) \\ \vdots & \ddots & \vdots \\ \partial_{y_1} h_c(y) & \cdots & \partial_{y_\ell} h_c(y) \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \ddots & 1 & \ddots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & \cdots & 0 & 1 \end{bmatrix}.$$

Observe that the bottom ℓ rows of the Jacobian above correspond to the $\ell \times \ell$ identity matrix, and therefore the columns of $JH(y)$ are linearly independent. Let $W \subset \mathbb{R}^n$ be the vector subspace spanned by the columns of $JH(y)$ and note that the dimension of W is equal to ℓ . Thus, the columns of $JH(y)$ are a basis of W , and the dimension of the orthogonal complement

$$W^\perp := \{z \in \mathbb{R}^n : z \cdot w = 0 \text{ for all } w \in W\}$$

must be equal to c , for \mathbb{R}^n is the direct sum of W and W^\perp and $n = c + \ell$.

On the other hand, for all $y \in V$ the point $(h(y), y) \in \mathbb{R}^n$ belongs to $B_\delta(x_0, y_0) \cap \mathcal{S}$, which implies that **all the constraints are satisfied** for all $y \in V$ and therefore

$$g_i \circ H(y) = g_i(h(y), y) = 0 \quad \text{for all } y \in V \text{ and } 1 \leq i \leq c.$$

In other words, every g_i is constant over V and its derivative must vanish. This implies that

$$\begin{aligned} \overbrace{(0, \dots, 0)}^{\ell \text{ components}} &= D(g_i \circ H)(y_0) && (g_i \circ H \text{ is constant on } V) \\ &= Dg_i(H(y_0)) DH(y_0) && (\text{Chain rule}) \\ &= \nabla g_i(x_0, y_0)^\top JH(y_0). \end{aligned}$$

The expression above tells us that for every $1 \leq i \leq c$, the gradient $\nabla g_i(x_0, y_0)$ is orthogonal to each and every one of the columns of $JH(y_0)$. Since the columns are a basis of W , it follows that every $\nabla g_i(x_0, y_0)$ must be orthogonal to every $x \in W$ and therefore

$$\nabla g_i(x_0, y_0) \in W^\perp \quad \text{for all } 1 \leq i \leq c.$$

On the other hand, we assumed that all the gradients $\nabla g_i(x_0, y_0)$ are linearly independent, and since there are c of them and the dimension of W^\perp is equal to c , we must conclude that the set

$$\{\nabla g_1(x_0, y_0), \dots, \nabla g_c(x_0, y_0)\}$$

is a basis for the subspace W^\perp .

Step 2:

We will now show that the gradient $\nabla f(x_0, y_0)$ also belongs to W^\perp .

We start by defining the function $F : V \subset \mathbb{R}^\ell \rightarrow \mathbb{R}$ as

$$F(y) = (f \circ H)(y) = f(h(y), y).$$

Recalling that if $y \in V$ then $(h(y), y) \in B_\delta(x_0, y_0) \cap \mathcal{S}$, and that f restricted to $U \cap \mathcal{S}$ has a local extremum at (x_0, y_0) , and that $h(y_0) = x_0$, it follows that F has a local extremum at (y_0) . Since V is an unconstrained open set, then by Proposition 3.5 the gradient of F must vanish at y_0 , leading to

$$\begin{aligned} \overbrace{(0, \dots, 0)}^{\ell \text{ components}} &= \nabla F(y_0) && \text{(By 3.5)} \\ &= D(f \circ H)(y_0) && \text{(Definition of } F) \\ &= Df(H(y_0)) DH(y_0) && \text{(Chain rule)} \\ &= \nabla f(x_0, y_0)^\top JH(y_0). \end{aligned}$$

Similarly to the argument in Step 1, the expression above shows that the gradient $\nabla f(x_0, y_0)$ is orthogonal to every one of the columns of $JH(y_0)$. Since the columns of this Jacobian are a basis for the subspace W , it follows that $\nabla f(x_0, y_0)$ is orthogonal to every $w \in W$ and therefore $\nabla f(x_0, y_0) \in W^\perp$.

Finally, since the set $\{\nabla g_1(x_0, y_0), \dots, \nabla g_c(x_0, y_0)\}$ is a basis for W^\perp and $\nabla f(x_0, y_0) \in W^\perp$, there must exist constants $\lambda_1, \dots, \lambda_c$ such that

$$\nabla f(x_0, y_0) = \sum_{i=1}^c \lambda_i \nabla g_i(x_0, y_0),$$

which is what we set out to prove. □

4.8 Exercises

1. Prove the first three points in Proposition 4.1.
2. Consider $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by $f(x, y) := x^2 - y^2$. Find its critical points and prove that they are neither local maxima nor local minima.
3. Let A be an $n \times n$ matrix with real coefficients and $F_A : \mathbb{R}^n \rightarrow \mathbb{R}$ be the quadratic form defined by

$$F_A(x) := x^\top Ax.$$

- (a) Prove that F is continuous for any $x_0 \in \mathbb{R}^n$.
- (b) Consider the matrix $B := \frac{1}{2}(A + A^\top)$. Prove that $F_B(x) = F_A(x)$ for all $x \in \mathbb{R}^n$.

- (c) Prove that B is the only symmetric matrix whose quadratic form coincides with that of A .
4. Let $U \subset \mathbb{R}^n$ be an open set such that $x_0 \in U$, and $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be of class $\mathcal{C}^2(U)$. Prove that if $f(x_0)$ is a local minimum, then

$$\partial_{ii}f(x_0) \geq 0 \quad \text{for all } 1 \leq i \leq n.$$

5. **Least squares fitting.** Let $(x_1, y_1), \dots, (x_n, y_n)$ be a set of points in \mathbb{R}^2 such that $x_1 < x_2 < \dots < x_n$.
- (a) Prove that the function $d : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$d(m, b) := \sum_{i=1}^n (y_i - (mx_i + b))^2$$

attains a minimum value on \mathbb{R}^2 . Find the explicit expression for the point (m^*, b^*) that minimizes d .

- (b) Prove that if (m^*, b^*) is the minimizer found in the previous step, then

$$\frac{1}{n} \sum_{i=1}^n (y_i - (m^*x_i + b^*)) = 0.$$

In other words, the average distance from the points (x_i, y_i) to the line $y = m^*x + b^*$ is zero. And thus this line is the *best fit* to the data, in the sense that the average error is zero.

6. Three important inequalities

- (a) Consider a set of k positive real numbers $\{a_n\}_{n=1}^k$. Use Lagrange multipliers to prove that

$$(a_1 \cdot a_2 \cdot \dots \cdot a_k)^{1/k} \leq \frac{1}{k} \sum_{n=1}^k a_n.$$

This inequality is known as the **geometric mean / arithmetic mean inequality**.

- (b) Consider two sets of real numbers $\{a_i\}_{i=1}^n$ and $\{b_i\}_{i=1}^n$ such that $a_i b_i \geq 0$ for all $i \leq n$, and two numbers p and q such that $\frac{1}{p} + \frac{1}{q} = 1$. Use Lagrange multipliers to prove that

$$\sum_{i=1}^n a_i b_i \leq \left(\sum_{i=1}^n a_i^p \right)^{1/p} \left(\sum_{i=1}^n b_i^q \right)^{1/q}.$$

This inequality is known as **Hölder's inequality** and the numbers p, q are said to be **Hölder conjugates**.

- (c) Use the previous inequality to prove that

$$\left(\sum_{i=1}^n (a_i + b_i)^p \right)^{1/p} \leq \left(\sum_{i=1}^n a_i^p \right)^{1/p} + \left(\sum_{i=1}^n b_i^p \right)^{1/p}.$$

The inequality above is known as **Minkowski's inequality**. The quantity

$$\|(a_1, \dots, a_n)\|_p := \left(\sum_{i=1}^n a_i^p \right)^{1/p}$$

is known as the **p-norm** of the vector (a_1, \dots, a_n) , and Minkowski's inequality is nothing but the equivalent of the triangle inequality for the p-norm.

Chapter 5

Integration

5.1 The Riemann integral in \mathbb{R}^n

Let $R = [a_1, b_1] \times \cdots \times [a_n, b_n] \subset \mathbb{R}^n$ be a closed rectangle. We refer to each of the one-dimensional intervals in the Cartesian product as a *component interval* and in particular to $[a_i, b_i]$ as the *i -th component interval*. The **volume** of R , denoted $v(R)$, is defined as

$$v(R) = (b_1 - a_1)(b_2 - a_2) \cdots (b_n - a_n).$$

To define a partition of a rectangle R , we will appeal to the definition of a partition of an interval. We take a partition P_i for each of the component intervals of R , and say that P is a **partition** of R if P is the set of all rectangles of the form

$$Q = [y_1, z_1] \times [y_2, z_2] \times \cdots \times [y_n, z_n],$$

where $y_i, z_i \in P_i$ are consecutive points in the partition P_i of the i -th component interval of R . Each rectangle $Q \in P$ is called a subrectangle of R . We will denote P in the form

$$P = (P_1, P_2, \dots, P_n),$$

where each P_i is a partition of $[a_i, b_i]$.

It is a simple exercise to see that if P is a partition of the closed rectangle R , then

$$\sum_{Q \in P} v(Q) = v(R).$$

Let $f : R \rightarrow \mathbb{R}$ be bounded, where $R \subset \mathbb{R}^n$ is a closed rectangle, and let P be a partition of R . Analogously to the one-dimensional case, the **lower sum** of f with respect to P is given by

$$L(f, P) = \sum_{Q \in P} m_Q(f) \cdot v(Q),$$

where $m_Q(f) := \inf\{f(x) : x \in Q\}$. Similarly, the **upper sum** of f with respect to P is given by

$$U(f, P) = \sum_{Q \in P} M_Q(f) \cdot v(Q),$$

where $M_Q(f) = \sup\{f(x) : x \in Q\}$.

It is clear that from their definition as infima and suprema that for every partition P of R , we have

$$L(f, P) \leq U(f, P).$$

A similar inequality holds for any two partitions P and T . However, to prove it, we need to introduce the concept of refinement.

Definition 5.1 (Refinement). Let T and P be partitions of a closed rectangle $R \subset \mathbb{R}^n$. We say that $T = (T_1, T_2, \dots, T_n)$ is a **refinement** of P if $P_i \subset T_i$ for each i . That is, each rectangle in T is a subrectangle of some rectangle in P .

Proposition 5.1. *If T is a refinement of P , then*

$$L(f, P) \leq L(f, T) \quad \text{and} \quad U(f, T) \leq U(f, P).$$

Proof. If $Q \in T$, there exists a rectangle $S \in P$ such that $T \subset S$. Then we have

$$m_Q(f) \geq m_S(f) \quad \text{and} \quad M_Q(f) \leq M_S(f).$$

Furthermore, each $S \in P$ is subdivided into rectangles $Q_1, Q_2, \dots, Q_k \in T$, and

$$v(S) = v(Q_1) + v(Q_2) + \dots + v(Q_k).$$

Then

$$m_S(f)v(S) = m_S(f) \sum_{j=1}^k v(Q_j) = \sum_{j=1}^k m_S(f)v(Q_j) \leq \sum_{j=1}^k m_{Q_j}(f)v(Q_j),$$

so that $L(f, P) \leq L(f, T)$.

Similarly,

$$M_S(f)v(S) = \sum_{j=1}^k M_S(f)v(Q_j) \geq \sum_{j=1}^k M_{Q_j}(f)v(Q_j),$$

so $U(f, P) \geq U(f, T)$. □

Corollary 5.1. *If P and T are partitions of R , then*

$$L(f, P) \leq U(f, Q).$$

Proof. Let \mathcal{R} be the common partition

$$\mathcal{R} = (P_1 \cup Q_1, P_2 \cup Q_2, \dots, P_n \cup Q_n).$$

Then \mathcal{R} is a refinement of both P and Q , and by the previous proposition,

$$L(f, P) \leq L(f, \mathcal{R}) \leq U(f, \mathcal{R}) \leq U(f, Q).$$

□

Since for any bounded function $f : R \rightarrow \mathbb{R}$ and for any partition P of R it holds that

$$-\infty < \inf_{x \in R} f(x) \cdot v(R) \leq L(f, P) \quad \text{and} \quad U(f, P) \leq \sup_{x \in R} f(x) \cdot v(R) < \infty.$$

We see that the set of all possible lower sums is bounded from below, while the set of all possible upper sums is bounded from above. This fact enables us to define the **lower integral** of f as

$$\int_{\underline{R}} f = \sup\{L(f, P) : P \text{ is a partition of } R\},$$

and the **upper integral** of f as

$$\int_{\overline{R}} f = \inf\{U(f, P) : P \text{ is a partition of } R\}.$$

Whenever the integration domain R is clear from the context, it is common practice to drop it from the notation and simply use the symbols

$$\int_{\underline{\quad}} \quad \text{for the lower integral and } \int_{\overline{\quad}} \quad \text{for the upper integrals.}$$

From the definition it is clear that for any bounded f

$$\int_{\underline{\quad}} f \leq \int_{\overline{\quad}} f.$$

Definition 5.2 (Riemann integrability). We say that the bounded function $f : R \rightarrow \mathbb{R}$ is **Riemann-integrable** if

$$\int_{\underline{\quad}} f = \int_{\overline{\quad}} f.$$

The common value of the upper and lower integrals is called the **integral** of f and is denoted by

$$\int f.$$

or by

$$\int_R f,$$

whenever the domain of integration needs to be specified explicitly.

Proposition 5.2 (Constant functions are Riemann-integrable). Let $f : R \rightarrow \mathbb{R}$ be given by $f(x) = c$ for some $c \in \mathbb{R}$. Then f is Riemann-integrable.

Proof. If P is a partition of R and $S \in P$, then

$$m_S(f) = M_S(f) = c,$$

so that

$$L(f, P) = U(f, P) = c \cdot v(R).$$

Clearly, $L(f) = U(f)$, and thus

$$\int f = c \cdot v(R).$$

□

Sometimes it is not easy to verify the equality between the upper and lower integrals of a function. The following theorem provides an alternative characterization of whether a function is Riemann-integrable that can be easier to verify. The proof of this fact will be left as an exercise.

Theorem 5.2. *Let $f : R \rightarrow \mathbb{R}$ be bounded. Then f is Riemann-integrable if and only if for every $\varepsilon > 0$, there exists a partition P such that*

$$U(f, P) - L(f, P) < \varepsilon.$$

The following proposition summarizes the basic properties of the Riemann integral.

Proposition 5.3. *Let $f, g : R \rightarrow \mathbb{R}$ be Riemann-integrable. Then:*

1. $f + g$ is Riemann-integrable and

$$\int (f + g) = \int f + \int g;$$

2. if $f \leq g$, then

$$\int f \leq \int g;$$

3. $|f|$ is Riemann-integrable and

$$\left| \int f \right| \leq \int |f|.$$

Proof. 1. Let P be a partition of R and $S \in P$. Then

$$m_S(f) + m_S(g) \leq m_S(f + g) \quad \text{and} \quad M_S(f) + M_S(g) \geq M_S(f + g).$$

This implies that

$$L(f, P) + L(g, P) \leq L(f + g, P) \quad \text{and} \quad U(f, P) + U(g, P) \geq U(f + g, P).$$

Given $\varepsilon > 0$, Theorem 5.2 implies that there exists a partition P such that

$$U(f, P) - L(f, P) < \frac{\varepsilon}{2} \quad \text{and} \quad U(g, P) - L(g, P) < \frac{\varepsilon}{2}.$$

Then

$$\begin{aligned} U(f + g, P) - L(f + g, P) &\leq U(f, P) + U(g, P) - (L(f, P) + L(g, P)) \\ &= (U(f, P) - L(f, P)) + (U(g, P) - L(g, P)) \\ &< \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon, \end{aligned}$$

and $f + g$ is Riemann-integrable. Now, if P is a partition of R ,

$$\int (f + g) = L(f + g) \geq L(f + g, P) \geq L(f, P) + L(g, P),$$

so that

$$\int (f + g) \geq L(f) + L(g) = \int f + \int g.$$

Similarly,

$$\int (f + g) \leq \int f + \int g,$$

and therefore

$$\int (f + g) = \int f + \int g.$$

2. The proof of this property will be left as an exercise.
3. Let P be a partition of R and $S \in P$. If $x, y \in S$, the reverse triangle inequality implies that

$$||f(x)| - |f(y)|| \leq |f(x) - f(y)|.$$

Then $M_S(|f|) - m_S(|f|) \leq M_S(f) - m_S(f)$, and this implies that

$$U(|f|, P) - L(|f|, P) \leq U(f, P) - L(f, P).$$

So, given $\varepsilon > 0$, if P is such that $U(f, P) - L(f, P) < \varepsilon$, then

$$U(|f|, P) - L(|f|, P) < \varepsilon,$$

and $|f|$ is Riemann-integrable. The inequality follows from property 2 and the fact that

$$-|f| \leq f \leq |f|.$$

□

5.2 Sets of measure zero

Definition 5.3. Let $A \subset \mathbb{R}^n$. We say that A is of **measure zero** if for every $\varepsilon > 0$ there exist rectangles R_1, R_2, \dots such that

$$A \subset \bigcup_i R_i \quad \text{and} \quad \sum_i v(R_i) < \varepsilon.$$

In addition, if the family of rectangles R_i is finite for every ε , we say that A has **content zero**.

Proposition 5.4. Any finite set $\{x_1, \dots, x_k\} \subset \mathbb{R}^n$ is of measure zero.

Proof. Notice that if $p = (p_1, p_2, \dots, p_n)$ is a point in \mathbb{R}^n , then the set

$$S := [p_1, p_1] \times \dots \times [p_n, p_n] = \{p\}$$

is indeed a closed rectangle. Then, given $\varepsilon > 0$, for each x_i we can take a rectangle R_i , with $x_i \in R_i$, such that $v(R_i) < \varepsilon/k$ and

$$\{x_1, \dots, x_k\} \subset \bigcup_{i=1}^k R_i \quad \text{and} \quad \sum_{i=1}^k v(R_i) < \sum_{i=1}^k \frac{\varepsilon}{k} = \varepsilon.$$

□

Proposition 5.5. Let $A_1, A_2, \dots \subset \mathbb{R}^n$ be sets of measure zero. Then the countable union

$$\bigcup_{i=1}^{\infty} A_i$$

is of measure zero.

Proof. Let $\varepsilon > 0$. Since every A_i has measure zero, for any $\delta_i > 0$ it is possible to choose a family of rectangle $\{R_{ij}\}_{j=1}^{\infty}$ such that

$$A_i \subset \bigcup_{j=1}^{\infty} R_{ij} \quad \text{and} \quad \sum_{j=1}^{\infty} v(R_{ij}) < \delta_i. \quad (5.1)$$

If we then take the union over i of these rectangle it follows that

$$\bigcup_{i=1}^{\infty} A_i \subset \bigcup_{i=1}^{\infty} \bigcup_{j=1}^{\infty} R_{ij}.$$

The sum of the volumes of all of the rectangles will then be given by the series

$$\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} v(R_{ij}) < \sum_{i=1}^{\infty} \delta_i,$$

where the inequality comes from (5.1). We see then that it is enough to choose

$$\delta_i < \frac{\varepsilon}{2^i}$$

to obtain

$$\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} v(R_{ij}) < \sum_{i=1}^{\infty} \frac{\varepsilon}{2^i} < \varepsilon$$

Since all volumes $v(R_{ij})$ are non negative, the series independent of the order of summation because it converges absolutely. □

Corollary 5.3. All countable sets are of measure zero.

5.3 Riemann integrability and continuity

In this section we will classify Riemann-integrable functions according to their points of continuity. We begin by introducing the concept of oscillation of a function at a point.

Definition 5.4 (Oscillation). Let $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be bounded and $x_0 \in A$. Given $\delta > 0$, we will denote the supremum and infimum of the f in a neighborhood of size δ of x_0 respectively by

$$M(f, x_0, \delta) := \sup\{f(x) : x \in A \cap B_\delta(x_0)\}, \quad m(f, x_0, \delta) := \inf\{f(x) : x \in A \cap B_\delta(x_0)\}.$$

The quantities $M(f, x_0, \delta)$ and $m(f, x_0, \delta)$ exist due to the boundedness of f . Note that if $0 < \eta \leq \delta$,

$$M(f, x_0, \eta) \leq M(f, x_0, \delta) \quad \text{and} \quad m(f, x_0, \eta) \geq m(f, x_0, \delta). \quad (5.1)$$

The **oscillation** of f at x_0 is given by

$$\mathcal{O}(f, x_0) = \lim_{\delta \rightarrow 0} (M(f, x_0, \delta) - m(f, x_0, \delta)).$$

The fact that the limit in the definition of $\mathcal{O}(f, x_0)$ exists is due to the monotonicity of the difference $M(f, x_0, \delta) - m(f, x_0, \delta)$ as a function of δ . Thus, the limit is simply the infimum of all differences for $\delta > 0$. The proof of this fact will be left as an exercise.

The oscillation of a function is a measure of how much a function *jumps* at a point. Since continuous functions do not have sudden jumps, the following proposition—whose proof will be left as an exercise—should not come as a surprise.

Proposition 5.6. *Let $R \subset \mathbb{R}^n$ be a closed rectangle, $x_0 \in R$, and $f : R \rightarrow \mathbb{R}$ be bounded. Then, f is continuous at x_0 if and only if $\mathcal{O}(f, x_0) = 0$.*

The relationship between the oscillation of a function and the upper and lower sums of the function with respect to a partition is given by the following proposition.

Proposition 5.7. *Let $R \subset \mathbb{R}^n$ be a closed rectangle and $f : R \rightarrow \mathbb{R}$ a bounded function such that $\mathcal{O}(f, x) < \varepsilon$ for all $x \in R$. Then there exists a partition P of R such that*

$$U(f, P) - L(f, P) < v(R)\varepsilon.$$

Proof. For each $x \in R$, we take an open rectangle R_x such that $x \in R_x$ and

$$M_{\overline{R_x}}(f) - m_{\overline{R_x}}(f) < \varepsilon,$$

where $\overline{R_x}$ is the closure of the rectangle R_x . Such a rectangle exists because $\mathcal{O}(f, x) < \varepsilon$. Then the collection $\{R_x\}_{x \in R}$ is an open cover for R , and since R is compact, there exist R_{x_1}, \dots, R_{x_k} such that

$$R \subset R_{x_1} \cup \dots \cup R_{x_k}.$$

Let P be the partition induced by all the endpoints of the component intervals of all the rectangles R_{x_i} . This partition is such that if S is a subrectangle of P , then $S \subset \overline{R_{x_i}}$ for some i and thus

$$M_S(f) - m_S(f) \leq M_{\overline{R_{x_i}}}(f) - m_{\overline{R_{x_i}}}(f) < \varepsilon.$$

Multiplying by the volume of the corresponding rectangle and adding over P and taking the supremum over all partitions it follows that

$$U(f, P) - L(f, P) < \varepsilon \sum_{S \in P} v(S) = \varepsilon v(R).$$

□

As an immediate result, we can prove the fact that continuous functions are Riemann-integrable.

Corollary 5.4. *If $f : R \rightarrow \mathbb{R}$ is continuous, then it is Riemann-integrable.*

Proof. By Proposition 5.6, the oscillation of f satisfies

$$\mathcal{O}(f, x) = 0 \quad \text{for every } x \in R.$$

Therefore we can apply Proposition 5.7 for any $\varepsilon > 0$ and use the ε -principle (Theorem A.1) to conclude that

$$U(f, P) - L(f, P) = 0$$

□

If a function f is not continuous everywhere, then its oscillation will be positive at its points of discontinuity. Hence, Proposition 5.7 suggests that its integrability will depend on whether we can cover those points with rectangles of arbitrarily small volume. We will spend the remainder of the section proving that this is the case.

The following characterization of the set of discontinuities of a function will be useful. Let $f : R \rightarrow \mathbb{R}$ be bounded and let

$$F := \{x \in R : f \text{ is not continuous at } x\}.$$

For $\varepsilon > 0$, define F_ε as

$$F_\varepsilon := \{x \in R \mid \mathcal{O}(f, x) \geq \varepsilon\}$$

By Proposition 5.6 and the fact that the oscillation of a function is always non-negative, the set of discontinuities F can be characterized as

$$F = \bigcup_{k=1}^{\infty} F_{1/k}.$$

To prove the main result of this section we will need to use the following lemma.

Lemma 5.1. *Let $f : R \rightarrow \mathbb{R}$ be bounded and such that F_ε can be covered by open rectangles R_1, \dots, R_k satisfying*

$$\sum_{i=1}^k v(R_i) < \delta,$$

for some $\delta > 0$. Then, if $M := \sup_R |f|$, there exists a partition P such that

$$U(f, P) - L(f, P) < v(R)\varepsilon + 2M\delta.$$

Proof. As in the proof of Proposition 5.7, we can find a partition P such that, for every subrectangle $S \in P$, if $S \cap F_\varepsilon = \emptyset$, then

$$M_S(f) - m_S(f) < \varepsilon.$$

Thus, if we define

$$P' := \{S \in P : S \cap F_\varepsilon = \emptyset\}, \quad \text{and} \quad P'' = \{S \in P : S \cap F_\varepsilon \neq \emptyset\},$$

and write $P = P' \cup P''$, it follows that

$$\begin{aligned} U(f, P) - L(f, P) &= \sum_{S \in P'} (M_S(f) - m_S(f))v(S) + \sum_{S \in P''} (M_S(f) - m_S(f))v(S) \\ &< \varepsilon \sum_{S \in P'} v(S) + 2M \sum_{S \in P''} v(S) \\ &\leq \varepsilon v(R) + 2M \sum_{i=1}^k v(R_i) \\ &< \varepsilon v(R) + 2M\delta. \end{aligned}$$

□

We are now in position to state the main result of this section. There will be only one step in the argument that we will use without proof.

Theorem 5.5. Let $f : R \rightarrow \mathbb{R}$ be bounded and let

$$F = \{x \in R : f \text{ is not continuous at } x\}.$$

Then f is Riemann-integrable if and only if F is of measure zero.

Proof. For $\varepsilon > 0$, define F_ε as

$$F_\varepsilon := \{x \in R \mid \mathcal{O}(f, x) \geq \varepsilon\}$$

and note that, by Proposition 5.6 and the fact that the oscillation of a function is always non-negative, the set of discontinuities F can be characterized as

$$F = \bigcup_{k=1}^{\infty} F_{1/k}.$$

\Rightarrow Suppose first that f is Riemann-integrable. We will show that each $F_{1/k}$ is of measure zero and the result will follow from Proposition 5.5.

Let $\varepsilon > 0$, and let P be a partition such that

$$U(f, P) - L(f, P) < \frac{\varepsilon}{k}.$$

Let $P' = \{S \in P : S \cap F_{1/k} \neq \emptyset\}$, i.e. the collection of subrectangles where the oscillation of the function exceeds $1/k$. Then

$$F_{1/k} \subset \bigcup_{S \in P'} S,$$

and for each subrectangle $S \in P'$ it holds $M_S(f) - m_S(f) \geq \frac{1}{k}$. Thus,

$$\sum_{S \in P'} \frac{1}{k} v(S) \leq \sum_{S \in P'} (M_S(f) - m_S(f)) v(S) \leq U(f, P) - L(f, P) < \frac{\varepsilon}{k}.$$

Therefore

$$\sum_{S \in P'} v(S) < \varepsilon.$$

Since ε is arbitrary, $F_{1/k}$ is of content zero, and therefore F is of measure zero.

\Leftarrow Now suppose that F is of measure zero. Given $\varepsilon > 0$, define

$$\bar{\varepsilon} := \frac{\varepsilon}{2v(R)}.$$

We claim that the set

$$F_{\bar{\varepsilon}} := \{x \in R \mid \mathcal{O}(f, x) \geq \bar{\varepsilon}\}$$

is of content zero. Clearly, since $F_{\bar{\varepsilon}} \subset R$, it is bounded. If one can prove that the set is also closed, it would follow that the set can be covered by finitely many rectangles whose volume can be shown to be arbitrarily small. We won't prove the claims in the previous sentence, as the results are quite technical.

Let $M = \sup_R |f|$, and let R_1, \dots, R_k be closed rectangles such that

$$F_{\bar{\varepsilon}} \subset \bigcup_{i=1}^k R_i \quad \text{and} \quad \sum_{i=1}^k v(R_i) < \frac{\varepsilon}{4M}.$$

By Lemma 5.1, there exists a partition P such that

$$U(f, P) - L(f, P) < \varepsilon v(R) + 2M \cdot \frac{\varepsilon}{4M} = \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Since ε is arbitrary, f is Riemann-integrable. \square

The theorem above offers a powerful criterion for verifying whether a function is Riemann-integrable. For example, the following corollary establishes that the product of Riemann-integrable functions is Riemann-integrable.

Corollary 5.6. *Let $f, g : R \rightarrow \mathbb{R}$ be Riemann-integrable. Then fg is Riemann-integrable.*

Proof. Let

$$F := \{x \in R : f \text{ is discontinuous at } x\}, \quad \text{and} \quad G := \{x \in R : g \text{ is discontinuous at } x\}.$$

Since the product of continuous functions is continuous, if

$$H := \{x \in R : fg \text{ is discontinuous at } x\},$$

then $H \subset F \cup G$. If f and g are Riemann-integrable, then F and G are of measure zero. Therefore, H is of measure zero and fg is Riemann-integrable. \square

5.4 Exercises

1. Let $R \subset \mathbb{R}^n$ be a closed rectangle and $f : R \rightarrow \mathbb{R}$ be a bounded function. Prove that f is Riemann integrable if and only if for every $\varepsilon > 0$ there exists a partition P of R such that

$$U(f, P) - L(f, P) < \varepsilon.$$

[Hint: the argument is very similar to the equivalent statement for functions of a single variable.]

2. Let $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be a bounded function and $x_0 \in A$. The oscillation of f at x_0 is defined as the limit

$$O(f, x_0) := \lim_{\delta \rightarrow 0} (\sup\{f(x) : x \in B_\delta(x_0) \cap A\} - \inf\{f(x) : x \in B_\delta(x_0) \cap A\}).$$

Prove that this limit exists. [Hint: consider a family of balls with radius $\delta = 1/n$.]

3. Let $f : A \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be a bounded function. Prove that f is continuous at $x_0 \in A$ if and only if the oscillation $O(f, x_0) = 0$.
4. (a) Let $f : [a, b] \rightarrow \mathbb{R}$ be increasing and $x_1 < x_2 < \dots < x_n \in [a, b]$. Prove that

$$\sum_{i=1}^n O(f, x_i) \leq f(b) - f(a).$$

- (b) Let $f : [a, b] \rightarrow \mathbb{R}$ be increasing. Prove that the set

$$S_n := \{x \in [a, b] : O(f, x) > 1/n\}$$

is finite for each n . [Assume that this is not true and contradict part (a).]

(c) Use part b) to prove that the set

$$S := \{x \in [a, b] : f \text{ is discontinuous at } x\}$$

has measure 0.

5. Let $R \subset \mathbb{R}^n$ be a closed rectangle and $f : R \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be bounded and continuous except for a finite number of points. Prove that f is Riemann integrable.

Appendix A

Useful facts from analysis in one variable

We present, without proof, some important facts from real analysis of a single variable.

Theorem A.1 (The ϵ -principle [5]). Consider a real number x such that $0 \leq x < \epsilon$ for all $\epsilon > 0$. Then $x = 0$.

Proof. We will proceed by contradiction and assume that there exists a real number $x \neq 0$ and such that $0 \leq x < \epsilon$ for all $\epsilon > 0$. This would imply that $1 < \epsilon/x$ for all $\epsilon > 0$. Therefore, letting $\epsilon = x$ it must hold that $1 < \epsilon/x = x/x = 1$, which is a contradiction. Therefore $x = 0$. \square

Definition A.1. Let $E \subset \mathbb{R}$ be a non empty set. A number $x \in \mathbb{R}$ is called:

- An **upper bound** of E if $x \geq e$ for all $e \in E$.
- A **lower bound** of E if $x \leq e$ for all $e \in E$.

A set is said to be **bounded from above** if it has at least one upper bound, **bounded from below** if it has at least one lower bound and simply **bounded** if it has both upper and lower bounds.

The axiom of the infimum/supremum. The following properties will be taken as an axiom of the real number system:

- If $E \subset \mathbb{R}$ is non empty and bounded from above, then there exists $S \in \mathbb{R}$ such that if U is an upper bound of E , then $S \leq U$. This number is called the **supremum** or the **least upper bound** of E , and is denoted as $\sup(E)$.
- If $E \subset \mathbb{R}$ is non empty and bounded from below, then there exists $I \in \mathbb{R}$ such that if L is a lower bound of E , then $I \geq L$. This number is called the **infimum** or the **greatest lower bound** of E , and is denoted as $\inf(E)$.

An alternate characterization of the infimum and supremum that is particularly useful to build convergent sequences is given in the following Theorem.

Theorem A.2. Let $E \subset \mathbb{R}$ be a non empty set.

1. If E is bounded from above and is S an upper bound of E . Then $S = \sup(E)$ if and only if for all $\epsilon > 0$ there exists $e \in E$ such that $S - \epsilon < e$.




Figure A.1: Nested interval theorem: The uncountable intersection of closed nested intervals whose length decreases to zero contains a single point.

2. If E is bounded from below and I is a lower bound of E . Then $I = \inf(E)$ if and only if for all $\epsilon > 0$ there exists $e \in E$ such that $I + \epsilon > e$.

The supremum and infimum have several important properties associated with the sum and multiplication of sets.

Lemma A.1. If $A \subset \mathbb{R}$ is non-empty and bounded and $c \in \mathbb{R}$, then

1. If $c > 0$ then

$$i) \inf(cA) = c \inf(A) \quad \text{and} \quad ii) \sup(cA) = c \sup(A).$$

2. If $c < 0$ then

$$iii) \inf(cA) = c \sup(A) \quad \text{and} \quad iv) \sup(cA) = c \inf(A).$$

The existence of the supremum and infimum can be used to show the **completeness of \mathbb{R}** i.e., the fact that every Cauchy sequence in \mathbb{R} converges. The proof of this fact starts from the assumption that the infimum and supremum exist, and one then proves the three following results in succession.

Theorem A.3 (Nested interval theorem). Let $I_n := [a_n, b_n]$ be a family of closed and nested intervals

$$I_1 \supset I_2 \supset \dots \supset I_n \supset I_{n+1} \supset \dots$$

such that $|I_n| := |b_n - a_n| \rightarrow 0$ as $n \rightarrow \infty$. Then, the intersection $\bigcap_{n=1}^{\infty} I_n$ contains exactly one point (See Figure A.1).

Theorem A.4 (Bolzano-Weierstrass theorem). Let $\{x_n\} \subset \mathbb{R}$ be a bounded sequence. Then there exists a subsequence $\{x_{n_k}\} \subset \{x_n\}$ that converges to a point $x \in \mathbb{R}$.

Theorem A.5 (Completeness of the real line). Let $\{a_n\}_{n=1}^{\infty} \subset \mathbb{R}$ be a Cauchy sequence. Then, there exists $a \in \mathbb{R}$ such that

$$a = \lim_{n \rightarrow \infty} a_n.$$

In words: every Cauchy sequence of real numbers is convergent.

Lemma A.2 (Fundamental lemma of differentiation). Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be differentiable at a point x_0 . Then, there exists a real-valued function η , defined in an interval around zero such that η is continuous at zero and

$$\eta(0) = 0, \tag{A.1a}$$

$$f(x) = f(x_0) + (x - x_0)(f'(x_0) + \eta(x - x_0)). \tag{A.1b}$$

Proof. We will simply define a function that satisfies these two properties by construction, and will then show that it is indeed continuous at zero. Define

$$\eta(h) := \begin{cases} \frac{f(x_0+h)-f(x_0)}{h} - f'(x_0) & \text{if } h \neq 0 \\ 0 & \text{if } h = 0. \end{cases} \quad (\text{A.2})$$

Letting $h = x - x_0$ in the definition above, it is easy to see that properties (A.1a) and (A.1b) are satisfied by construction, while from the definition of η and the fact that f is differentiable at x_0 we have

$$\lim_{h \rightarrow 0} \eta(h) = \lim_{h \rightarrow 0} \left(\frac{f(x_0+h) - f(x_0)}{h} - f'(x_0) \right) = f'(x_0) - f'(x_0) = 0,$$

and thus η is continuous at 0 as desired. □

Remark A.1. The value of the function η can be interpreted as a measure of how far from the derivative the Newton quotient is as h approaches zero. Clearly, if the function is differentiable, the “distance” must decrease smoothly as $h \rightarrow 0$ and should vanish on the limit. This is the intuitive interpretation of (A.1a) and the continuity of η at zero. Moreover, (A.1b) states that, if a function is differentiable at a point x_0 , one can approximate it in the vicinity of x_0 by the straight line going through the point $(x_0, f(x_0))$ and with slope given by $f'(x_0)$ as

$$f(x) \approx f(x_0) + (x - x_0)f'(x_0)$$

incurring an error that vanishes faster than the distance between the evaluation point x and the approximation point x_0 , as measured by the magnitude of the term $(x - x_0)\eta(x - x_0)$. Due to this geometric interpretation, the lemma is also known as the **linear approximation lemma**.

The following result is sometimes called the **generalized mean value theorem** and some other times the **Cauchy’s mean value theorem**.

Theorem A.6 (Generalized mean value theorem). If the functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$ are both continuous over the interval $[a, b]$ and differentiable over (a, b) , then there exists a point $c \in (a, b)$ such that

$$\frac{f'(c)}{g'(c)} = \frac{f(b) - f(a)}{g(b) - g(a)}.$$

Proof. Since both f and g are continuous over the interval $[a, b]$ and differentiable over (a, b) , the function $h : \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$h(x) := (g(b) - g(a))f(x) - (f(b) - f(a))g(x)$$

is also continuous over $[a, b]$ and differentiable over (a, b) . Moreover, simple algebraic manipulations show that

$$\begin{aligned} h(a) &= (g(b) - g(a))f(a) - (f(b) - f(a))g(a) \\ &= g(b)f(a) - g(a)f(b) \\ &= g(b)f(a) - g(b)f(b) + f(b)g(b) - g(a)f(b) \\ &= (g(b) - g(a))f(b) - (f(b) - f(a))g(b) \\ &= h(b). \end{aligned}$$

We can therefore apply Rolle’s theorem to obtain a point $c \in (a, b)$ such that

$$0 = h'(c) = (g(b) - g(a))f'(c) - (f(b) - f(a))g'(c),$$

from which the theorem follows. □

Appendix B

Essential point-set topology

The topology of a space allows us to study basic concepts of analysis, such as convergence and compactness. In this section, we will introduce the basic topological properties of the Euclidean space.

Definition B.1. We say that $U \subset \mathbb{R}^n$ is an **open set** if, for every $x \in U$, there exists $\epsilon > 0$ such that $B_\epsilon(x) \subset U$.

Example B.1. The sets \emptyset and \mathbb{R}^n are open. The case of \mathbb{R}^n is clear; however, the fact that \emptyset is open follows from a technical convention in formal logic: in a conditional statement of the form $A \Rightarrow B$, when the hypothesis A is false, then the conclusion B is considered to be true. Therefore, “ $x \in \emptyset$ ” is false for any x , the statement “If $x \in \emptyset$, then there exists $\epsilon > 0$ such that $B_\epsilon(x) \subset \emptyset$,” is true.

Example B.2. An open ball is an open set (See Figure B.1). To show this, consider the ball:

$$B_r(x) = \{y \in \mathbb{R}^n : |x - y| < r\},$$

and let $y \in B_r(x)$. Let $\delta = \frac{r - |x - y|}{2}$ and $z \in B_\delta(y)$. Then, by the triangle inequality:

$$|z - x| \leq |z - y| + |y - x| \leq \delta + |x - y| = \frac{1}{2}r + \frac{1}{2}|x - y| < r,$$

so $z \in B_r(x)$, and therefore $B_\delta(y) \subset B_r(x)$.

Example B.3. An open rectangle

$$R = (a_1, b_1) \times (a_2, b_2) \times \cdots \times (a_n, b_n),$$

is an open set. Let $x \in R$, and let ϵ be the smallest one-dimensional distance between the components of x and the endpoints of their respective intervals; i.e.

$$\epsilon := \frac{1}{2} \min\{x_1 - a_1, b_1 - x_1, \dots, x_n - a_n, b_n - x_n\}.$$

With this choice for ϵ it follows that $B_\epsilon(x) \subset R$.

The previous example allows us to conclude the following proposition, which provides an equivalent definition of an open set that can be useful when working on the Cartesian setting.




Figure B.1: Open balls and open rectangles in \mathbb{R}^n are open sets.

Proposition B.1. $U \subset \mathbb{R}^n$ is open if and only if, for every $x \in U$, there exists an open rectangle R such that $x \in R$ and $R \subset U$.

Proof. Let U be open and $x \in U$. Then there exists $\epsilon > 0$ such that $B_\epsilon(x) \subset U$. Let

$$R = \left(x_1 - \frac{\epsilon}{\sqrt{n}}, x_1 + \frac{\epsilon}{\sqrt{n}}\right) \times \left(x_2 - \frac{\epsilon}{\sqrt{n}}, x_2 + \frac{\epsilon}{\sqrt{n}}\right) \times \cdots \times \left(x_n - \frac{\epsilon}{\sqrt{n}}, x_n + \frac{\epsilon}{\sqrt{n}}\right).$$

Then $x \in R$ and $R \subset B_\epsilon(x) \subset U$.

Now suppose that for every $x \in U$, we can find an open rectangle $R = (a_1, b_1) \times (a_2, b_2) \times \cdots \times (a_n, b_n)$ such that $x \in R$ and $R \subset U$. Let

$$\epsilon = \frac{1}{2} \min\{x_1 - a_1, b_1 - x_1, \dots, x_n - a_n, b_n - x_n\}.$$

Then $B_\epsilon(x) \subset R \subset U$, and therefore U is open. \square

If $x \in \mathbb{R}^n$, a neighborhood of x is an open set $U \subset \mathbb{R}^n$ that contains x , i.e., $x \in U$. \square

Definition B.2 (accumulation point). Let $A \subset \mathbb{R}^n$ and $x \in \mathbb{R}^n$. We say that x is an accumulation point of A if, for every $r > 0$, the set of $B_r(x) \cap A$ is infinite¹.

Analogous to the definition of an open set, we can show that x is an accumulation point of A if and only if, for every open rectangle R such that $x \in R$, $R \cap A$ is infinite. In general, x is an accumulation point of A if every neighborhood of x contains an infinite number of points in A .

If the set A has any accumulation points, then A , by the previous definition, is infinite. Additionally, if x is an accumulation point of A , it does not necessarily follow that $x \in A$. However, if x is an accumulation point of A and $x \notin A$, then we can “approach” x arbitrarily closely from A ; that is, for every $r > 0$, there exists $y \in A$ such that $|x - y| < r$.

Proposition B.2. Let $A \subset \mathbb{R}^n$ and $x \in \mathbb{R}^n$. x is an accumulation point of A if and only if, for every $r > 0$,

$$B_r(x) \cap (A \setminus \{x\}) \neq \emptyset.$$

That is, x is an accumulation point of A if and only if every ball around x contains points of A distinct from x .

¹The expression “the set A is infinite” should be understood as “the cardinality of A is infinite”

Clearly, if x is an accumulation point of A , $B_r(x) \cap (A \setminus \{x\}) \neq \emptyset$ because $B_r(x) \cap A$ is infinite.

To show the converse, suppose x is not an accumulation point of A . Then there exists $r > 0$ such that $B_r(x) \cap A$ is finite. If $B_r(x) \cap A$ is equal to $\{x\}$ or empty, then $B_r(x) \cap (A \setminus \{x\}) = \emptyset$. Suppose instead that:

$$B_r(x) \cap A = \{x_1, \dots, x_k\} \neq \{x\},$$

and let $\delta = \frac{1}{2} \min\{|x_i - x| : x_i \neq x\}$. Then $B_\delta(x) \cap (A \setminus \{x\}) = \emptyset$. \square Of course, this proposition can also be equivalently stated using open rectangles.

Definition B.3 (Closed set). We say that $A \subset \mathbb{R}^n$ is closed if it contains all its accumulation points.

This definition suggests that a closed set does not have "nearby external points," hence the name "closed." In particular, if A is closed and $x \notin A$, by Proposition 1.16 there exists $r > 0$ such that $B_r(x) \cap A$ contains at most x . Since $x \notin A$, $B_r(x) \cap A$ is empty.

Example B.4. The sets \emptyset and \mathbb{R}^n are closed. It is clear that \mathbb{R}^n contains all its accumulation points because it contains all its points, while \emptyset has no accumulation points.

Example B.5. Like the empty set, any set without accumulation points is closed. This includes finite sets and the lattice:

$$\mathbb{Z}^n = \{(k_1, k_2, \dots, k_n) : k_i \in \mathbb{Z}\}.$$

Example B.6. A closed rectangle $R = [a_1, b_1] \times [a_2, b_2] \times \dots \times [a_n, b_n]$ is a closed set. If $x \notin R$, then some coordinate $x_i \notin [a_i, b_i]$, meaning $x_i < a_i$ or $x_i > b_i$. If we define:

$$\epsilon = \frac{1}{2} \max\{a_i - x_i, x_i - b_i\},$$

then $\epsilon > 0$ and $B_\epsilon(x) \cap R = \emptyset$, so x is not an accumulation point of R . Thus, R contains all its accumulation points and is closed.

Now let us examine the relationship between closed and open sets.

Proposition B.3. $A \subset \mathbb{R}^n$ is closed if and only if $\mathbb{R}^n \setminus A$ is open².

Proof. $\boxed{\Rightarrow}$ Suppose A is closed and $x \in \mathbb{R}^n \setminus A$. Since $x \notin A$, x is not an accumulation point of A , so there exists $\epsilon > 0$ such that $B_\epsilon(x) \cap A = \emptyset$. That is, $B_\epsilon(x) \subset \mathbb{R}^n \setminus A$. Hence, $\mathbb{R}^n \setminus A$ is open.

$\boxed{\Leftarrow}$ Now suppose $\mathbb{R}^n \setminus A$ is open and $x \notin A$. Then $x \in \mathbb{R}^n \setminus A$. Since $\mathbb{R}^n \setminus A$ is open, there exists $\epsilon > 0$ such that $B_\epsilon(x) \subset \mathbb{R}^n \setminus A$. Thus, $B_\epsilon(x) \cap A = \emptyset$, meaning x is not an accumulation point of A . Therefore, A is closed. \square

Definition B.4 (Boundary of a set). Let $A \subset \mathbb{R}^n$. The boundary of A , denoted ∂A , is the set of $x \in \mathbb{R}^n$ such that, for every $\epsilon > 0$:

$$B_\epsilon(x) \cap A \neq \emptyset \quad \text{and} \quad B_\epsilon(x) \cap (\mathbb{R}^n \setminus A) \neq \emptyset.$$

²This proposition allows us to define a closed set equivalently as the complement of an open set, without reference to accumulation points. Conversely, it also provides an alternative: we can first define closed sets through their accumulation points and then define an open set as the complement of a closed set. Either approach is valid for defining the topology on \mathbb{R}^n , and both are used in different analysis texts depending on the author's preference.

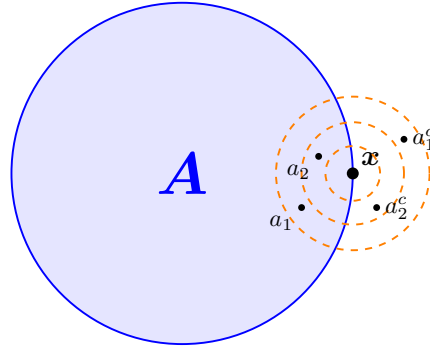


Figure B.2: Every neighborhood around boundary point $x \in \partial A$ contains points from the set A (other than itself) and from its complement A^c .

That is, a point x is a **boundary point** if every ball around x intersects both A and its complement $\mathbb{R}^n \setminus A$ (Figure B.2). In this case we write $x \in \partial A$. Equivalently, $x \in \partial A$ if and only if, for every open rectangle R containing x :

$$R \cap A \neq \emptyset \quad \text{and} \quad R \cap (\mathbb{R}^n \setminus A) \neq \emptyset.$$

Note that if $x \in \partial A$, then x is an accumulation point of A or $\mathbb{R}^n \setminus A$. Moreover, if x is an accumulation point of A and $x \notin A$, then $x \in \partial A$ (Exercise 12).

Additionally, we observe that $\partial A = \partial(\mathbb{R}^n \setminus A)$.

Example B.7.

$$\partial \mathbb{R}^n = \partial \emptyset = \emptyset.$$

Example B.8. The boundary of a ball is the sphere that surrounds it. In fact:

$$\partial B_r(x) = \partial B_r^\circ(x) = S_r(x).$$

Moreover:

$$\partial S_r(x) = S_r(x).$$

Example B.9. If $R = (a_1, b_1) \times \cdots \times (a_n, b_n)$, then:

$$\partial R = \{a_1\} \times [a_2, b_2] \times \cdots \times [a_n, b_n] \cup \{b_1\} \times [a_2, b_2] \times \cdots \times [a_n, b_n] \cup \cdots \cup [a_1, b_1] \times \cdots \times \{b_n\}.$$

That is, ∂R is the union of the "faces" of R .

Example B.10. Let $Q = [0, 1] \cap \mathbb{Q}$ and consider $Q \times [0, 1] \subset \mathbb{R}^2$. If $x \in [0, 1] \times [0, 1]$ and $x \in (a, b) \times (c, d)$, then there exists:

$$q \in (a, b) \cap [0, 1] \cap \mathbb{Q},$$

so that $(q, x_2) \in Q \times [0, 1]$. Moreover, there exists:

$$\alpha \in (a, b) \cap [0, 1] \setminus \mathbb{Q},$$

so that $(\alpha, x_2) \in \mathbb{R}^2 \setminus (Q \times [0, 1])$. Therefore:

$$\partial(Q \times [0, 1]) = [0, 1] \times [0, 1].$$

Proposition B.4. Let $A \subset \mathbb{R}^n$. The boundary of A is closed.

Proof. We will show that the complement of the boundary is an open set. From the definition B.4 of the boundary we see that $x \in (\partial A)^c$ if and only if there exists $r_0 > 0$ such that either of the following conditions holds

$$B_{r_0}(x) \cap A = \emptyset \quad \text{or} \quad B_{r_0}(x) \cap A^c = \emptyset.$$

In either case, we see that there is a ball around x that contains points only of A or only of A^c . Thus neither of the points of this ball belongs to ∂A and so $B_{r_0}(x) \subset (\partial A)^c$. This proves that $(\partial A)^c$ is open and therefore ∂A is closed. \square

Definition B.5 (Closure of a set). Let $A \subset \mathbb{R}^n$. The closure of A , denoted by \overline{A} , is defined as the union of A and its accumulation points.

Alternatively, the closure of A is the union of itself with its boundary (Exercise 15). The following proposition establishes some properties of the closure.

Proposition B.5. Let $A \subset \mathbb{R}^n$.

1. \overline{A} is closed.
2. If E is closed and $A \subset E$, then $\overline{A} \subset E$.
3. If $A \subset B$, then $\overline{A} \subset \overline{B}$.
4. $\overline{\overline{A}} = \overline{A}$.

Proof. 1. Let x be an accumulation point of \overline{A} , and let R be a rectangle containing x . We want to show that $R \cap A$ is infinite, such that x is an accumulation point of A and hence $x \in \overline{A}$. Otherwise, since $R \cap \overline{A}$ is infinite, we can take $y \in R \cap (\overline{A} \setminus A)$. But then y is an accumulation point of A , and since $y \in R$, $R \cap A$ is infinite, leading to a contradiction.

2. If x is an accumulation point of A and $A \subset E$, then x is an accumulation point of E . Since E is closed, $x \in E$. Hence, $\overline{A} \subset E$.
3. The proof is similar to (2) (Exercise 16).
4. By definition, $\overline{A} \subset \overline{\overline{A}}$. Now, from (1) and (2), since \overline{A} is closed:

$$\overline{\overline{A}} \subset \overline{A}.$$

\square

Part (2) of Proposition B.5 implies that the closure of the set A is the "smallest" closed set that contains A .

Definition B.6. Let $A \subset \mathbb{R}^n$. The *interior of a set* A is defined as:

$$\text{int}(A) = A^\circ = \{x \in A : \text{there exists } \epsilon > 0 \text{ such that } B_\epsilon(x) \subset A\}.$$

The *exterior of a set* A is defined as the set:

$$\text{ext}(A) = \{x \in \mathbb{R}^n \setminus A : \text{there exists } \epsilon > 0 \text{ such that } B_\epsilon(x) \cap A = \emptyset\}.$$

Similar to the closure, the interior of A is now the "largest" open set contained in A (Exercise 11). Additionally, we note that:

$$\text{ext}(A) = \text{int}(\mathbb{R}^n \setminus A).$$

The following proposition follows easily from the definitions (Exercise 12).

Proposition B.6. Let $A \subset \mathbb{R}^n$.

1. $A^\circ = A \setminus \partial A$.
2. $\text{ext}(A) = \text{int}(\mathbb{R}^n \setminus \overline{A})$.

Example B.11.

$$Q^\circ = \emptyset \quad \text{and} \quad \overline{Q} = \mathbb{R}.$$

Note that in this case, the interior is empty, even though the closure is "large."

B.1 Exercises

1. Show that $U \subset \mathbb{R}^n$ is open if and only if, for every $x \in U$, there exists $\epsilon > 0$ such that:

$$B_\epsilon(x) \subset U.$$

In other words, open sets can be defined in terms of open balls.

2. Show that a half-space is open.
3. Show that if $\{U_\alpha\}$ is a collection of open sets in \mathbb{R}^n , then the union:

$$\bigcup_{\alpha} U_{\alpha}$$

is an open set.

4. Show that if U_1, U_2, \dots, U_k are open sets in \mathbb{R}^n , then the intersection:

$$\bigcap_{i=1}^k U_i$$

is an open set.

5. Show that x is an accumulation point of A if and only if, for every open rectangle R containing x ,

$$R \cap (A \setminus \{x\}) \neq \emptyset.$$

6. Show that, if $x \notin A$, then $x \in \partial A$ if and only if x is an accumulation point of A .
7. Show that if A is closed and $x \in \partial A$, then $x \in A$.
8. Show that, for any $A \subset \mathbb{R}^n$,

$$\partial A = \overline{A} \cap (\mathbb{R}^n \setminus A).$$

9. Show that:

$$\overline{A} = A \cup \partial A.$$

10. Prove the third part of Proposition 1.28.
11. Let $A \subset \mathbb{R}^n$ and $U \subset A$ be open. Show that $U \subset \text{int}A$.
12. Prove Proposition B.6.

Appendix C

Elements of linear algebra

We say that vectors $u_1, u_2, \dots, u_m \in \mathbb{R}^n$ **span** \mathbb{R}^n if for every $x \in \mathbb{R}^n$, there exist scalars $\alpha_1, \dots, \alpha_m \in \mathbb{R}$ such that:

$$x = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_m u_m.$$

When the equality above is verified, we say that x is a **linear combination** of the vectors u_1, u_2, \dots, u_m .

We say that u_1, u_2, \dots, u_m are **linearly independent** if:

$$\alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_m u_m = 0$$

implies that:

$$\alpha_1 = \alpha_2 = \dots = \alpha_m = 0.$$

We say that the set of vectors u_1, u_2, \dots, u_m is a **basis** of \mathbb{R}^n if they span \mathbb{R}^n and are linearly independent. The proof of the following theorem can be found in any text on linear algebra

Theorem C.1. *If the set $\{u_1, u_2, \dots, u_m\}$ forms a basis of \mathbb{R}^n , then $m = n$.*

If u_1, u_2, \dots, u_m form a basis of \mathbb{R}^n , then for every $x \in \mathbb{R}^n$, there exist unique scalars $\alpha_1, \alpha_2, \dots, \alpha_n$ such that:

$$x = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_n u_n.$$

The **canonical basis** of \mathbb{R}^n consists of the vectors e_1, e_2, \dots, e_n , where:

$$e_i = (0, 0, \dots, \overset{i\text{-th}}{1}, \dots, 0).$$

If $x \in \mathbb{R}^n$,

$$x = (x_1, x_2, \dots, x_n) = x_1 e_1 + x_2 e_2 + \dots + x_n e_n.$$

We say that u_1, u_2, \dots, u_n form an **orthonormal basis** they are a basis and if the vectors are mutually orthogonal and unitary, that is, $|u_i| = 1$ for all i .

Proposition C.1. *Let u_1, u_2, \dots, u_n be an orthonormal basis of \mathbb{R}^n and $x, y \in \mathbb{R}^n$. Then*

1. $x = (x \cdot u_1)u_1 + \dots + (x \cdot u_n)u_n.$
2. $|x| = \sqrt{\sum_i (x \cdot u_i)^2}.$
3. $x \cdot y = \sum_{i=1}^n (x \cdot u_i)(y \cdot u_i).$

Proof. Since u_1, u_2, \dots, u_n is a basis of \mathbb{R}^n , every $x \in \mathbb{R}^n$ can be expressed in the form

$$x = \alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_n u_n.$$

Then:

$$x \cdot u_i = (\alpha_1 u_1 + \alpha_2 u_2 + \dots + \alpha_n u_n) \cdot u_i = \alpha_i (u_i \cdot u_i) = \alpha_i,$$

where in the last two equalities we used the orthonormality of the set. This proves 1. Using the property 1, we have:

$$|x|^2 = x \cdot x = \left(\sum_{i=1}^n (x \cdot u_i) u_i \right) \cdot \left(\sum_{i=1}^n (x \cdot u_i) u_i \right) = \sum_{i,j=1}^n (x \cdot u_i) (x \cdot u_j) (u_i \cdot u_j).$$

Since the set u_1, u_2, \dots, u_n is orthonormal it follows that:

$$|x|^2 = \sum_{i=1}^n (x \cdot u_i)^2,$$

and the second property follows readily. Similarly, to prove 3 we observe that

$$x \cdot y = \left(\sum_{i=1}^n (x \cdot u_i) u_i \right) \cdot \left(\sum_{i=1}^n (y \cdot u_i) u_i \right) = \sum_{i,j=1}^n (x \cdot u_i) (y \cdot u_j) (u_i \cdot u_j) = \sum_{i=1}^n (x \cdot u_i) (y \cdot u_i),$$

where once again we used the orthonormality of the basis in the last step. \square

The space generated by the vectors v_1, v_2, \dots, v_r is the **subspace** of \mathbb{R}^n formed by all linear combinations of v_1, v_2, \dots, v_r , and we denote it by $\text{span}\{v_1, v_2, \dots, v_r\}$. We will denote the orthogonal projection of x onto the subspace V by $\text{Proj}_V x$. This is the unique vector $y \in V$ such that $x - y$ is orthogonal to every vector in V .

Proposition C.2. *If V is the subspace of \mathbb{R}^n generated by the orthonormal vectors v_1, v_2, \dots, v_r , then:*

$$\text{Proj}_V x = \sum_{i=1}^r (x \cdot v_i) v_i.$$

Proof. Since the set v_1, v_2, \dots, v_r is orthonormal, by Proposition C.1 we have that if $z \in V$, then:

$$z = \sum_{i=1}^r (z \cdot v_i) v_i.$$

Therefore, if $y = \sum_{i=1}^r (x \cdot v_i) v_i$, then $y \in V$ and, for every $z \in V$:

$$(x - y) \cdot z = \left(x - \sum_{i=1}^r (x \cdot v_i) v_i \right) \cdot \sum_{i=1}^r (z \cdot v_i) v_i = x \cdot \sum_{i=1}^r (z \cdot v_i) v_i - \sum_{i=1}^r (x \cdot v_i) (z \cdot v_i) = 0.$$

\square

The following theorem guarantees that, given a space spanned by vectors v_1, v_2, \dots, v_r , we can always choose an orthonormal basis within it. Its proof is constructive, and the resulting algorithm is known as **Gram-Schmidt orthonormalization**.

Theorem C.2 (Gram-Schmidt orthonormalization). *Let v_1, v_2, \dots, v_r be linearly independent vectors in \mathbb{R}^n . Then there exist orthonormal vectors u_1, u_2, \dots, u_r such that:*

$$\text{span}\{u_1, u_2, \dots, u_k\} = \text{span}\{v_1, v_2, \dots, v_k\} \quad \text{for } k = 1, \dots, r.$$

Proof. We take:

$$u_1 := \frac{v_1}{|v_1|},$$

and notice that $|u_1| = 1$. To construct u_2 , we recall that for any $x \in \mathbb{R}^n$ it holds that

$$x = \text{proj}_x u_1 u_1 + (x - \text{proj}_x u_1 u_1) = \text{proj}_x u_1 u_1 + (\text{proj}_x u_1 u_1)^\top$$

and therefore we let:

$$w_2 = v_2 - (v_2 \cdot u_1)u_1.$$

We see that w_2 is orthogonal to u_1 , so we take:

$$u_2 = \frac{w_2}{|w_2|}.$$

Since u_1 and u_2 are linear combinations of v_1 and v_2 :

$$\text{span}\{u_1, u_2\} \subset \text{span}\{v_1, v_2\}.$$

Similarly, v_1 and v_2 are linear combinations of u_1 and u_2 , so:

$$\text{span}\{v_1, v_2\} \subset \text{span}\{u_1, u_2\}.$$

□

Definition C.1. A matrix $A \in \mathbb{R}^{n \times n}$ is said to be **positive semi-definite** (resp. **negative semi-definite**) if

$$0 \leq x^\top A x \quad (\text{resp. } x^\top A x \leq 0) \quad \text{for all } x \in \mathbb{R}^n.$$

In addition, if $x^\top A x = 0$ if and only if $x = 0$, we say that A is **positive definite** (resp. **negative definite**).

Bibliography

- [1] J. Marsden and A. Weinstein. *Calculus III*. Springer US, 1985.
- [2] P. Mikusiński and M. D. Taylor. *An Introduction to Multivariable Analysis from Vector to Manifold*. Birkhäuser Boston, 2002.
- [3] J. Páez Cárdenas. *Cálculo integral de varias variables*. Las prensas de Ciencias, Universidad Nacional Autónoma de México, 2nd edition, 2014.
- [4] J. Páez Cárdenas. *Cálculo diferencial de varias variables*. Las prensas de Ciencias. Universidad Nacional Autónoma de México, 2nd edition, 2017.
- [5] C. C. Pugh. *Real Mathematical Analysis*. Springer New York, 2002.
- [6] W. Rudin. *Principles of mathematical analysis*. International series in pure and applied mathematics. McGraw-Hill Professional, New York, NY, 3 edition, Feb. 1976.
- [7] R. A. Sáenz. *Análisis de varias variables*, 2014.
- [8] T. Sánchez-Vizuet. *Minimal analysis I: A concise introduction to real analysis of a single variable*, 2025.
- [9] J. Shurman. *Calculus and Analysis in Euclidean Space*. Springer International Publishing, 2016.
- [10] M. Spivak. *Calculus on manifolds*. Westview Press, Philadelphia, PA, Jan. 1971.
- [11] M. Spivak. *Calculus*. Publish or Perish, fourth edition, 2008.
- [12] T. Tao. *Analysis II*. Springer Singapore, 2016.
- [13] M. E. Taylor. *Introduction to analysis in several variables*. Pure and Applied Undergraduate Texts. American Mathematical Society, Providence, RI, Sept. 2020.

Alphabetical Index

Symbols

$D_u f(x_0)$	32
$Df(x_0)$	27
$Jf(x_0)$	27
\mathcal{C}^0	49
\mathcal{C}^∞	49
\mathcal{C}^k	49
$Hf(x_0)$	54
$O(x)$	25
S^{n-1}	5
ϵ -principle	81
$\mathcal{O}(f, x_0)$	76
∇	35
$\nabla^2 f(x_0)$	54
$\nabla f(x_0)$	35
$\partial_{x_i} f(x_0)$	33
$f'(x_0)$	27
$o(x)$	25

Lagrange's 52

A

Absolute homogeneity 2	
accumulation point	85
affine function	19
angle between two vectors	3
asymptotic notation	25

B

basis	90
big O	25
bijjective	15
Bolzano-Weierstrass	7
Bolzano-Weierstrass theorem	7, 82
Boundary of a set	86

boundary point	87
bounded	81
bounded from above	81
bounded from below	81
bounded function	17
bounded sequence	6
bounded set	9

C

canonical basis	3, 90
Cauchy sequence	8
Cauchy's mean value theorem	83
Cauchy-Schwarz inequality	2
Chain Rule	30
Clairut's theorem	47
Closed set	86
Closure of a set	88
compact	9
completeness of \mathbb{R}	82
Completeness of the real line	82
component function	14
constraints	61
content zero	74
continuous	15
continuously differen- tiable	37
contraction	40
Contraction mapping theorem	40
convergence	6
convex	4
cover	9

D

derivative	27
----------------------	----

directional derivative	25, 32
dot product	1

E

Euclidean norm	1
Euclidean space	1
exterior of a set	88

F

first order constrained optimality condition	64
Fundamental lemma of differ- entiation	82

G

generalized mean value theorem	83
gradient	34
Gram-Schmidt orthonormal- ization	91
greatest lower bound	81

H

Heine-Borel theorem	10
Hessian	54, 59, 60
hypercube	5
hyperplane	4

I

image	14
Implicit function theorem	46
infimum	81
injective	14
inner product	1
integral	72
interior of a set	88

Inverse function
 theorem 40
 Inverse function theorem
 43
 inverse image 14
 invertible 15

J
 Jacobian 27, 34

L
 Lagrange multiplier 64,
 66
 Lagrangian 64
 least upper bound 81
 level set 61
 limit 15
 line 3
 linear approximation
 lemma 83
 linear combination 90
 linear function 18
 linearly independent 90
 Lipschitz continuity 18
 Lipschitz function 17
 little o 25
 local extrema 58
 local maximum 33, 58
 local minimum 33, 58
 lower bound 81
 lower integral 72
 lower sum 70

M
 maximum 33, 58
 mean value theorem 34
 measure zero 74
 method of Lagrange
 multipliers 62,
 63
 minimum 33, 58
 Multivariable Taylor's
 theorem 52

N
 nabla 35
 negative definite 59, 92
 negative semi-definite 59, 60,
 92
 Nested interval
 theorem 82
 normal vector 4

O
 objective function 58
 open ball 4
 open set 84
 orthogonal 1
 Orthogonality between
 curves 64
 Orthogonality to a
 curve 64
 orthonormal basis 90
 Oscillation 75

P
 parametric curve 64
 parametric equation of a
 hyperplane 4
 parametric equation of the
 line 4
 parametrizable curve 64
 parametrization 64
 partial derivative 33
 partition 70
 positive definite 59, 92
 Positive definiteness 2
 positive semi-definite 59, 60,
 92
 preimage 14
 projection 1
 projections 15

Q
 quadratic form 59

R
 real analytic 52

rectangle 5
 refinement 71
 rejection 1
 remainder 51
 Riemann-integrable 72

S
 scalar product 1
 sequentially compact 9
 sequentially
 continuous 16
 smooth function 49
 span 90
 sphere 4
 star-shaped set 4
 strict contraction 40
 subcover 9
 subsequence 7
 subspace 91
 supremum 81
 Symmetry of mixed partial
 derivatives 47

T
 Taylor polynomial 53
 Taylor's polynomial 49
 Taylor's theorem 49
 Triangle inequality 2

U
 uniformly continuous 20
 unit sphere 5
 upper bound 81
 upper integral 72
 upper sum 70

V
 velocity 64
 volume 70

Y
 Young's theorem 47