MATH-404 FUNCTIONAL ANALYSIS II, SPRING 2025
LECTURE NOTES

MATHIAS BRAUN

Abstract. This is a shortened summary of the above course taught at EPFL and *no* complete script. Therefore, it cannot replace the careful study of a textbook, e.g. one of those listed in the bibliography. In particular, these notes contain only selected proofs.

These notes will be updated as the lecture moves on. References tagged by **??** refer to results that will be introducted later. Parts that have been added later (after the respective lecture) are colored in gray. Please send remarks or comments about mistakes, typos, lack of clarity, suggestions for improvements, etc. to mathias.braun@epfl.ch.

This summary is only intended for students visiting the above course.

Contents

0. A short compendium of topology 2
1. Locally convex vector spaces 4
2. Test functions and distributions 15
3. Calculus on Banach spaces 32
4. A selection of fixed point theorems 44
5. Gradient flows in Hilbert spaces 50
Appendices 60
Appendix A. Weak topologies induced by families of functions 60
Appendix B. Weak topologies and infinite dimensional Banach spaces 63
Appendix C. Weak compactness, separability, and reflexivity 67
Appendix D. Schauder–Tychonoff in Banach spaces 72
Appendix E. Bochner integration on Banach spaces 73
References 75

Lecture 1. Throughout these notes, we use indifferently $A \subset B$ and $A \subseteq B$ to state that A is a subset of B, with possibly $A = B$. We would use instead $A \subsetneq B$ to stress that A is a proper or strict subset of B, namely $A \subset B$ yet $A \neq B$.

The set $\mathbf{N}$ of natural numbers does not contain 0. We set $\mathbf{N}_0 := \mathbf{N} \cup \{0\}$.

The terms "function", "map", and "mapping" will be used synonymously. Given a function $f\colon X \to \mathbf{R}$ on a set X, given $t \in \mathbf{R}$ we abbreviate $\{x \in X : f(x) < t\}$ by $\{f < t\}$. The sets $\{f \leq t\}$, $\{f = t\}$, $\{f \geq t\}$, $\{f > t\}$, etc. are defined analogously.

These lecture notes are strongly based on summaries of previous versions of this course by Matthias Ruf, Michele Dolce, and Lucio Galeati, who thankfully shared their notes with me.

0. A short compendium of topology[1]

Let X be an arbitrary set. Then we have the following basic definitions. All implicit statements are left to the reader as an exercise.

- A **topology** $\tau \subset \mathcal{P}(X)$ on X is a family of sets such that $\emptyset, X \subset \tau$ that is stable under finite intersections and arbitrary unions. The sets in τ are called **open sets** and the pair (X, τ) is called a **topological space**. The **closed sets** of (X, τ) are exactly those sets whose complement is an open set. The collection of closed sets is closed under arbitrary intersections and finite unions.
- We say $\mathcal{B} \subset \tau$ is a **basis of the topology** τ if every open set can be written as union of elements in $\mathcal{B}$.
- If τ_1 and τ_2 are two topologies on X with $\tau_1 \subset \tau_2$, we say τ_1 is **coarser** than τ_2 and τ_2 is **finer** than τ_1[2].
- If (X, τ) and (Y, ρ) are topological spaces, then a function $f\colon X \to Y$ is called **continuous** if $f^{-1}(A) \in \tau$ for every $A \in \rho$, where f^{-1} denotes the usual preimage. (In general, images of open sets under continuous maps need not be open.) If we want to specify the respective reference topologies explicitly, we also write $f\colon (X, \tau) \to (Y, \rho)$.
- If $f\colon (X_1, \tau_1) \to (X_2, \tau_2)$ and $g\colon (X_2, \tau_2) \to (X_3, \tau_3)$ are continuous maps between the three given topological spaces, so is their composition $g \circ f\colon (X_1, \tau_1) \to (X_3, \tau_3)$ defined by $g \circ f(x) := g(f(x))$.
- If $B \subset X$, its **closure** $\overline{B}$ of B is defined as the smallest closed set containing B. In particular, a set is closed if and only if it coincides with its closure. The **interior** $\operatorname{int} B$ of B is defined as the largest open set contained in B. The **boundary** of B is defined as $\partial B = \overline{B} \setminus \operatorname{int} B$.
- For a topological space (X, τ), a set $K \subset X$ is **compact** is every open cover of K admits a finite subcover. That is, given any (not necessarily countable) collection $\{U_i : i \in I\} \subset \tau$ whose union contains K, there exists a finite index set $J \subset I$ such that $K \subset \bigcup_{i \in J} U_i$.
- The continuous image of compact sets is compact. That is, if $f\colon (X, \tau) \to (Y, \rho)$ is continuous and $K \subset X$ is compact, then $f(K) \subset Y$ is compact.
- A topological space (X, τ) is **Hausdorff** (or **T2**) if for all $x, y \in X$ with $x \neq y$ there exist $U_x, U_y \in \tau$ with $x \in U_x$ and $y \in U_y$ yet $U_x \cap U_y = \emptyset$. On a Hausdorff topological space, singletons are closed sets. Additionally, you should try to check in this case that compact sets are closed. (Note that singletons are compact in any topology.)
- Given a topological space (X, τ) and a sequence $(x_n)_{n \in \mathbf{N}}$ in X, we say that $(x_n)_{n \in \mathbf{N}}$ **converges** to a point $x \in X$ (with respect to τ) if for every $U \in \tau$ containing x there exists $n_0 \in \mathbb{N}$ such that $x_n \in U$ for all $n \geq n_0$.
- Continuous functions preserve convergence of sequences. More precisely, let $f\colon (X, \tau) \to (Y, \rho)$ be a continuous function and $(x_n)_{n \in \mathbf{N}}$ be a sequence converging to $x \in X$ with respect to τ. Then the sequence $(f(x_n))_{n \in \mathbf{N}}$ in Y converges to $f(x)$ with respect to ρ. This is an instructive exercise only requiring to apply the definitions.

Remark 0.1 (Caveat). Special care is required when working with sequences in abstract topological spaces! While they enjoy the natural property described above (which we will often use throughout the course), the correct notion to use in order

[1] The content of this preliminary chapter is not examinable. Yet, we will frequently use some of the facts provided here in the lectures and exercises.

[2] In general, one accepts the slight linguistic mismatch that "coarser" or "finer" can also mean the two topologies coincide.

to characterize the topology τ would be **nets**, i.e. families labeled by a directed but in general uncountable index set. We will not treat nets in detail and use them very sparingly, which is why in most arguments we will rather consider mostly open sets, and most importantly neighborhoods.

As a practical example, let us point out the following facts. On any topological space (X,τ), you should try to check closed sets are **sequentially closed**. In other words, if $E \subset X$ is closed and $(x_n)_{n\in\mathbf{N}}$ is a sequence in E converging to $x \in X$, then $x \in E$. However, the converse is not true: there exist topological spaces with sequentially closed sets which are not closed. ■

Although the basic objects in a topology are open sets, it is often convenient (particularly in the first part of this course) to think in terms of neighborhoods.

- A set $N \subset X$ is called a **neighborhood** of a point $x \in X$ if there exists $U \in \tau$ contained in N such that $x \in U$. In particular, an open set U is a neighborhood of all the points it contains. On the other hand, note that neighborhoods need not be open in general.
- A family $\mathcal{N}_x \subset \mathcal{P}(X)$ of subsets of X is called a **neighborhood basis** of $x \in X$ if every $N \in \mathcal{N}_x$ is a neighborhood of x and for every neighborhood $W \subset X$ of x there exists $N \in \mathcal{N}_x$ such that $N \subset W$. In particular, this applies when $W \in \tau$.
- Specifying a topology τ is equivalent to specifying a neighborhood basis $\mathcal{N}_x$ of every point $x \in X$. Indeed, any open set can then be reconstructed as a union of elements of $\mathcal{N}_x$: given $U \in \tau$, for any $y \in U$ there must exist $N_y \in \mathcal{N}_x$ such that $y \in N_y \subset U$, which implies $U = \bigcup_{y\in U} N_y$. In particular, the collection $\mathcal{N} := \bigcup_{x\in X} \mathcal{N}_x$ is a basis of the topology τ.
- Continuity of maps is characterized by neighborhoods. More precisely, a map $f\colon (X,\tau) \to (Y,\rho)$ is continuous if and only if the preimage of any neighborhood under f is a neighborhood. In fact, it suffices to verify this property on bases of neighborhood at every given point in Y.
- Let τ_1 and τ_2 be topologies on X. Then showing $\tau_1 \subset \tau_2$ is equivalent to verifying the following: for any $x \in X$ and any neighborhood $N_1 \subset X$ of x with respect to τ_1, there exists a neighborhood $N_2 \subset X$ with respect to τ_2 with $X_2 \subset N_1$ and $x \in N_2$. It is actually enough to verify this whenever N_1 belongs to a neighborhood basis of x with respect to τ_1.
- Closed sets can be defined intrinsically by neighborhoods. More precisely, given $B \subset X$, then $x \in \overline{B}$ if and only if for any neighborhood $N \subset X$ of x we have $N \cap B \neq \emptyset$.

As a basic example, let us recall how the above concepts readapt consistently to the case of metric spaces.

- Given any set X, a **semimetric** on X is a map $\mathsf{d}\colon X^2 \to \mathbf{R}_+$ such that $\mathsf{d}(x,x) = 0$ for every $x \in X$, $\mathsf{d}(x,y) = \mathsf{d}(y,x)$ for every $x, y \in X$, and $\mathsf{d}(x,z) \leq \mathsf{d}(x,y) + \mathsf{d}(y,z)$ for every $x,y,z \in X$. If additionally d does not vanish outside the diagonal of M^2, then d is called a **metric** on X. In the first case, the couple (X,d) is called a **semimetric space**, in the second a **metric space**.
- On a semimetric space (X,d), the **open ball** with center $x \in X$ and radius $r > 0$ are defined by $B_r(x) := \{y \in X : \mathsf{d}(x,y) < r\}$.
- The **topology τ^{d} induced by the semi-metric** d on X is defined as follows: a set $U \subset X$ is open if and only if for every $x \in U$ there exists $r > 0$ such that $B_r(x) \subset U$. (Note $B_r(x)$ always contains x.) In other words, the collection of open balls $\{B_r(x) : x \in X,\ r > 0\}$ forms a basis of the topology τ^{d}. Moreover, for any fixed $x \in X$, $\{B_r(x) : r > 0\}$ forms a neighborhood

basis of x. This is not the only option: for instance, $\{B_{1/n}(x) : n \in \mathbf{N}\}$ is also a (countable) neighborhood basis of x.

- Given a semimetric space (X, d), the topology τ^{d} is Hausdorff if and only if $\{\mathsf{d}(x, \cdot) = 0\} = \{x\}$ for every $x \in X$. In other words, τ^{d} is Hausdorff if and only if d is a metric. For this reason, in these notes we will only consider metric spaces instead of semimetric spaces.

We finally recall two basic operations on topological spaces.

- If (X, τ_X) is a topological space and $Y \subset X$, then the **subspace topology** (or **relative topology**) of Y is given by $\tau_Y := \{U \cap Y : U \in \tau_X\}$. Then (Y, τ_Y) is a topological space and τ_Y is the smallest topology which makes the **inclusion map** $\iota\colon Y \to X$ given by $\iota(y) := y$ continuous. Furthermore, a set $N' \subset Y$ is a neighborhood of y in τ_Y if and only if it is of the form $N' = N \cap Y$, where $N \subset X$ is a neighborhood of y in τ_X. Similarly, bases and neighborhood bases of τ_Y and τ_X can be shown to be in analogous direct correspondences.
- For an arbitrary family of sets $\{X_i : i \in I\}$, the product set $X := \prod_{i \in I} X_i$ can be identified as the collection of all possible indexed tuples $(x_i)_{i \in I}$ such that $x_i \in X_i$ for all $i \in I$. Given an arbitrary family of topological spaces $\{(X_i, \tau_i) : i \in I\}$, we define the **product topology** $\prod_{i \in I} \tau_i$ on their product X as the coarsest topology such that for every $i \in I$, the projection $\pi_i\colon X \to X_i$ defined by $\pi_i(x) := x_i$ is continuous.

Often, the reference topologies are clear from the context. In this case, we will name topological properties and notions without explicit reference to the topology. However, note carefully that sometimes topologies need specification to rule out ambiguities or errors. For instance, the clause "$[0, 1)$ is open" is not true for the standard topology on $\mathbf{R}$, but it is true for the relative topology on $[0, 1]$.

1. Locally convex vector spaces

This first chapter introduces the basic objects we will consider throughout the course in a rather abstract fashion. It might take a bit of time to develop the right intuition about them. The basic idea is to set the properties of convergence, compactness, etc. of test functions and distributions on a topological ground. For relevant and helpful applications and examples, see the exercise sheets and the upcoming chapters.

Remark 1.1 (Convention). In principle, vector spaces can be defined over arbitrary fields $\mathbf{K}$; throughout the course, however, we will restrict ourselves to vector fields over $\mathbf{R}$. In particular, whenever we say "let X be a vector space", we mean X is an $\mathbf{R}$-vector space. However, note that all stated results would still hold (with minor modifications) for complex vector spaces. ■

1.1. Basic notions.

Given $n \in \mathbf{N}$, Euclidean space $\mathbf{R}^n$ will always be endowed with the Euclidean topology τ_{Eucl} induced by the Euclidean metric d_2 given by

$$\mathsf{d}_2(x, y) := \sqrt{|x_1 - y_1|^2 + \cdots + |x_n - y_n|^2}.$$

Definition 1.2 (Topological vector space). *Let X constitute a vector space with given addition $+\colon X^2 \to X$ and scalar multiplication $\cdot\colon \mathbf{R} \times X \to X$. Moreover, let τ be a topology on X. The couple (X, τ) will be called a* **topological vector space**, *briefly TVS, if the mappings $+$ and $\cdot$ are continuous. Here, X^2 and $\mathbf{R} \times X$ are endowed with the evident product topologies τ^2 and $\tau_{\mathrm{Eucl}} \times \tau$, respectively.*

Remark 1.3 (Separate continuity). It follows from the definition of the product topology that addition and scalar multiplication in a TVS (X, τ) are separately continuous. Given any $\bar{x} \in X$, define the translation $f\colon X \to X$ by $f(x) := x + \bar{x}$. Then f is continuous. Moreover, it is straightforward to check f is invertible with continuous inverse. Therefore, translations are homeomorphisms on topological vector spaces. Analogously, given $\bar{x} \in X$ the map $g\colon \mathbf{R} \to X$ given by $g(\lambda) := \lambda x$ is continuous. For a fixed $\bar{\lambda} \in \mathbf{R}$ the map $h\colon X \to X$ with $h(x) := \lambda x$ is continuous. If $\bar{\lambda} \neq 0$, you should check h is invertible with continuous inverse. In other words, nontrivial dilations are homeomorphisms of X. ■

Remark 1.4 (Neighborhoods)*.* The topology of a TVS (X, τ) is characterized by bases of neighborhoods of the origin. Indeed, as translations are homeomorphisms by the previous remark, they map open sets to open sets and neighborhoods to neighborhoods. In particular, a set $N \subset X$ is a neighborhood of 0 if and only if $x + N$ is a neighborhood of x, where we set $x + N := \{x + n : n \in N\}$.

In addition, one can "shrink" (or — less relevant in applications — "enlarge") neighborhoods of the origin (and hence of arbitrary points in X) as much as needed by multiplication with an arbitrary parameter $\lambda > 0$. Indeed, a set $N \subset X$ is a neighborhood of 0 if and only if λN is, where we set $\lambda N := \{\lambda n : n \in N\}$. ■

Example 1.5 (Normed spaces)*.* Every normed space is a topological vector space. However, we are mainly interested in topologies that are not normable, possibly not even metrizable (like weak topologies on Banach spaces, see later). ■

In order to define "local convexity" of a TVS, we need the following concept.

Definition 1.6 (Seminorm)**.** *Let X be a vector space. A function $p\colon X \to \mathbf{R}_+$ is called a* **seminorm** *if*

a. *$p(\lambda x) = |\lambda|\, p(x)$ for every $\lambda \in \mathbf{R}$ and every $x \in X$ and*
b. *$p(x + y) \leq p(x) + p(y)$ for every $x, y \in X$.*

Every seminorm p on X satisfies $p(0) = 0$ and $p(x) = p(-x)$ for every $x \in X$. Note that every norm on X is a seminorm, but we do not require seminorms to only vanish at 0 (for instance, the map sending all of X to 0 is clearly a seminorm).

The reason why we will speak about "locally convex" topologies in Definition 1.12 is because its generating sets from (1.1) below are convex (cf. Definition 1.16). In turn, this follows from the following simple observation.

Lemma 1.7 (Convexity)**.** *Every seminorm p on a vector space X is convex.*

Proof. Given any $x, y \in X$ and any $\lambda \in [0, 1]$, simply note

$$p((1-\lambda)x + \lambda y) \leq p((1-\lambda)x) + p(\lambda y) = (1-\lambda)\, p(x) + \lambda\, p(y). \qquad \square$$

Vector spaces equipped with given families of seminorms are naturally endowed with an associated topology.

Definition 1.8 (Seminorm topology)**.** *Let X be a vector space and let $\mathcal{P} := \{p_i : i \in I\}$ be a family of seminorms on X. Let us define τ, the* **topology induced by** *$\mathcal{P}$ on X, as follows: a subset $U \subset X$ belongs to τ if and only if for every $x \in U$ there exist a finite set $J \subset I$ and $\varepsilon > 0$ such that $B_{\varepsilon,J}(x) \subset U$, where*

$$B_{\varepsilon,J}(x) := \{y \in X : p_i(y - x) < \varepsilon \text{ for every } i \in J\}. \tag{1.1}$$

In other words, the above seminorm topology τ on X is characterized by the requirement that, for every point $x \in X$, the set $\{B_{\varepsilon,J}(x) : \varepsilon > 0,\ J \subset I \text{ finite}\}$ is a basis of neighborhoods of x.

Since $B_{\varepsilon,\{i\}}(x) = \{p_i(x - \cdot) < \varepsilon\}$ for every $i \in I$, (1.1) equivalently becomes

$$B_{\varepsilon,J}(x) = \bigcap_{i \in J} B_{\varepsilon,\{i\}}(x). \tag{1.2}$$

Remark 1.9 (Seminorms vs. semimetrics). Definition 1.8 should be compared to the topology induced by a semimetric. Indeed, given $i \in I$ define $\mathsf{d}_i \colon X^2 \to \mathbf{R}_+$ by $\mathsf{d}_i(x,y) := p_i(x-y)$. It is straightforward to check d_i is a semimetric (and a metric if and only if p_i is a norm). Thus, $B_{\varepsilon,\{i\}}(x)$ is nothing but the open ε-ball with center $x \in X$ with respect to d_i. In turn, by (1.2) a basis of neighborhoods of x is given by the family of all finite intersections of such semimetric balls with a fixed yet arbitrary radius. ■

Lemma 1.10 (Seminorm topology yields TVS). *In the framework of Definition 1.8, the space (X, τ) is a TVS.*

Proof. Exercise 2.1. □

Example 1.11 (Continuous functions). A TVS with infinitely many seminorms is given by $C(\mathbf{R})$ with the family $\{p_n : n \in \mathbf{N}\}$, where

$$p_n(f) := \sup_{x \in [-n,n]} |f(x)|.$$ ■

We are in a position to state the first central notion of this course.

Definition 1.12 (Locally convex topological vector space). *We call a vector space X equipped with the topology induced by a family $\mathcal{P} := \{p_i : i \in I\}$ of seminorms a* **locally convex topological vector space**, *briefly LCTVS, if $X_0 = \{0\}$, where*

$$X_0 := \{x \in X : p_i(x) = 0 \text{ for every } i \in I\}.$$

Remark 1.13 (About Definition 1.12). The definition of LCTVSs appears rather restrictive, as it relies on a priori given seminorms. Given a topology τ on X, it is hard to identify the right seminorms it should be induced by per se. There is a more geometric characterization of LCTVSs based on the existence of a convex neighborhood basis of the origin (and the Hausdorff property). These definitions are equivalent, but in the lecture we will only prove a partial converse adding some geometric conditions to convexity, cf. Theorem 1.18. The full equivalence is established in Exercise 1.3. ■

Example 1.14 (Continuation of Example 1.11). The topology from Example 1.11 is easily seen to be locally convex: for every $f \in C(\mathbf{R}) \setminus \{0\}$ there exists $n \in \mathbf{N}$ such that $p_n(f) > 0$. ■

Recall the set X_0 above always contains 0. The point of the condition $X_0 = \{0\}$ defining LCTVSs is the Hausdorffness of their generating topology.

Lemma 1.15 (Hausdorff property). *Let τ be a topology on a vector space X as in Definition 1.8. Then τ is Hausdorff if and only if $X_0 = \{0\}$.*

Proof. We assume first $X_0 = \{0\}$. To show Hausdorffness of τ, let $x, y \in X$ satisfy $x \neq y$. Since $x - y \neq 0$, by assumption there exist $i \in I$ and a seminorm p_i such that $\delta > 0$, where $\delta := p_i(x-y)$. Then the open sets $B_{\delta/2,\{i\}}(x)$ and $B_{\delta/2,\{i\}}(y)$ are disjoint; indeed, otherwise a point $z \in B_{\delta/2,\{i\}}(x) \cap B_{\delta/2,\{i\}}(y)$ would satisfy

$$p_i(x-y) \leq p_i(x-z) + p_i(z-y) = p_i(z-x) + p_i(z-y) < \delta,$$

which contradicts the choice of δ.

On the other hand, suppose τ is Hausdorff. As noted above, it suffices to show $X_0 \subset \{0\}$. Let $x \in X_0$. If $x \neq 0$, by the hypothesized Hausdorffness there exists an open set $U \subset X$ containing x such that $0 \notin U$. Since U is open, we find $J \subset I$ finite such that $B_{\varepsilon,J}(x) \subset U$. But then $0 \notin B_{\varepsilon,J}(x)$, which implies there exists $i \in J$ such that $p_i(x) = p_i(-x) = p_i(0-x) \geq \varepsilon$, contradicting the inclusion $x \in X_0$. □

Definition 1.16 (Minkowski functional). *Let X be a vector space. A subset $A \subset X$ will be called*

a. **absorbing** *if for every* $x \in X$ *there exists* $\varepsilon > 0$ *such that for every* $t \in \mathbf{R}$ *with* $|t| \leq \varepsilon$ *we have* $tx \in A$,
b. **balanced** *if* $\lambda A \subset A$ *for all* $\lambda \in \mathbf{R}$ *with* $|\lambda| \leq 1$, *and*
c. **convex** *if* $(1-\lambda)x + \lambda y \in A$ *for all* $x, y \in A$ *and* $\lambda \in [0,1]$.

If A *is absorbing, balanced, and convex, we define the associated* **Minkowski functional** $p_A \colon X \to \mathbf{R}_+$ *by*

$$p_A(x) := \inf\{t > 0 : x \in tA\}.$$

The connection between geometry and seminorms becomes more evident in the following result. It also clarifies the conditions on the set A in the above definition used to set up the concept of Minkowski functionals.

Proposition 1.17 (Seminorms from Minkowski functionals)**.** *Let* X *be a vector space. Assume that* $A \subset X$ *is absorbing, balanced, and convex. Then the induced Minkowski functional* p_A *is a seminorm satisfying*

$$\{p_A < 1\} \subset A \subset \{p_A \leq 1\}. \tag{1.3}$$

Conversely, if q *is any seminorm on* X, *then* $q = p_{A_q}$ *with the absorbing, balanced, and convex set* $A_q := \{q < 1\}$.

Proof. To prove the first claim, first note since A is absorbing, p_A is finite.

To prove homogeneity of p_A, note first $p_A(0) = 0$ since A is balanced. Next, fix $\lambda \in \mathbf{R} \setminus \{0\}$ and $x \in X$. Again since A is balanced, we deduce $\lambda A = |\lambda| A$ and therefore $\lambda x \in tA$ if and only if $x \in tA/|\lambda|$. This entails

$$\begin{aligned}
p_A(\lambda x) &= \inf\{t > 0 : \lambda x \in tA\} \\
&= \inf\{t > 0 : x \in |\lambda|^{-1} tA\} \\
&= \inf\{|\lambda|\,|\lambda|^{-1} t > 0 : x \in |\lambda|^{-1} tA\} \\
&= |\lambda| \inf\{\mu > 0 : x \in \mu A\} \\
&= |\lambda|\, p_A(x).
\end{aligned}$$

Next, we prove subadditivity of p_A. Let $x, y \in X$. Since p_A is finite, by definition for every $\varepsilon > 0$ there exist $\lambda, \mu > 0$ such that

- $\lambda \leq p_A(x) + \varepsilon$ and $x \in \lambda A$ as well as
- $\mu \leq p_A(y) + \varepsilon$ and $y \in \mu A$.

By convexity of A we know $\lambda A + \mu A \subset (\lambda + \mu)A$, so that $x + y \in (\lambda + \mu)A$. Hence by definition and the arbitrariness of ε we infer $p_A(x+y) \leq p_A(x) + p_A(y)$. This shows p_A is a seminorm.

Finally, we establish the inclusions (1.3). Since $a \in 1 \cdot A$ for every $a \in A$, we obtain $p_A(a) \leq 1$, which shows $A \subset \{p_A \leq 1\}$. Furthermore, given any $x \in X$ with $p_A(x) < 1$ there exists $\lambda \in [0,1)$ such that $x \in \lambda A \subset A$, where we used that A is balanced. Hence $\{p_A < 1\} \subset A$.

Now we show the representation of arbitrary seminorms claimed in the second statement. We start by showing the sublevel set $\{q < 1\}$ is absorbing, balanced, and convex. If $x \in X$ satisfies $q(x) = 0$, then $tx \in \{q < 1\}$ for every $t \in \mathbf{R}$ by homogeneity of q. If $q(x) > 0$, then for every $\varepsilon \in (0, 1/q(x))$ and every $t \in \mathbf{R}$ with $|t| \leq \varepsilon$ we have $q(tx) = |t|\, q(x) < 1$ and therefore $tx \in \{q < 1\}$. This shows $\{q < 1\}$ is absorbing. In order to show $\{q < 1\}$ is balanced, let $x \in \{q < 1\}$ and $\lambda \in \mathbf{R}$ be such that $|\lambda| \leq 1$. Then $q(\lambda x) = |\lambda|\, q(x) < 1$, so that $\lambda x \in \{q < 1\}$. Finally, since seminorms are convex by Lemma 1.7, the set $\{q < 1\}$ is convex as the sublevel set of a convex function.

Moreover, note that every $x \in X$ satisfies

$$p_{A_q}(x) = \inf\{t > 0 :\, x \in tA_q\}$$

$$
\begin{aligned}
&= \inf\{t > 0 : t^{-1}x \in A_q\} \\
&= \inf\{t > 0 : q(t^{-1}x) < 1\} \\
&= \inf\{t > 0 : q(x) < t\} \\
&= q(x).
\end{aligned}
$$

□

Lecture 2. Now we can formulate the main theorem on the equivalent definitions of LCTVSs indicated in Remark 1.13.

Theorem 1.18 (Characterization of local convexity)**.** *Let X be an LCTVS in the sense of Definition 1.12. Then 0 has a neighborhood basis consisting of absorbing, balanced, convex, and open sets.*

Conversely, if (X, τ) is a Hausdorff topological vector space such that 0 has a neighborhood basis consisting of convex sets, then its topology is induced by a family of seminorms on X according to Definition 1.12.

Note that any absorbing or balanced set contains 0 (see the proof of Proposition 1.17). In a topological vector space, translations are homeomorphisms. Thus, the origin plays no special role in the previous theorem from a topological point of view. Theorem 1.18 can be adapted in that every point $x \in X$ has a neighborhood basis consisting of convex sets that are x-translates of absorbing and balanced sets.

The proof of Theorem 1.18 relies on the following fact.

Lemma 1.19 (Existence of good neighborhoods)**.** *Let (X, τ) constitute a TVS. Then every convex neighborhood of 0 contains an absorbing, balanced, convex, and open neighborhood of 0.*

Proof. Exercise 1.3. □

Proof of Theorem 1.18. Assume (X, τ) satisfies Definition 1.12. We consider the following family of sets:

$$\mathcal{N} := \{B_{\varepsilon,I_0}(0) : \varepsilon > 0,\ I_0 \subset I \text{ finite}\}.$$

We claim $\mathcal{N}$ is a neighborhood basis of the origin of absorbing, balanced, convex, and open sets. Clearly, $0 \in B_{\varepsilon,I_0}(0)$ for every $\varepsilon > 0$ and every $I_0 \subset I$ finite. Next we show every $B \in \mathcal{N}$ is open with respect to the topology inherited by the seminorms. Write $B = B_{\varepsilon,I_0}(0)$ and let $y \in B_{\varepsilon,I_0}(0)$, define $\delta := \max_{i \in I_0} p_i(y) < \varepsilon$, and choose $z \in B_{\varepsilon-\delta,I_0}(y)$. Then

$$p_i(z) \leq p_i(z - y) + p_i(y) < (\varepsilon - \delta) + \delta = \varepsilon.$$

This shows $B_{\varepsilon-\delta,I_0}(y) \subset B_{\varepsilon,I_0}(0)$ and by definition the set $B_{\varepsilon,I_0}(0)$ is open. In order to prove the remaining properties, it suffices to note that the function $\max_{i\in I_0} p_i$ is still a seminorm. Then we can repeat the proof of Proposition 1.17 to show $B_{\varepsilon,I_0}(0) = \{\max_{i\in I_0} p_i < \varepsilon\}$ is absorbing, balanced and convex. We conclude the first part of the proof by noting the definition of the seminorm topology implies $\mathcal{N}$ is a neighborhood basis of 0.

Now we prove the converse. Let $\mathcal{N}'$ be a family of convex sets that forms a neighborhood basis of the origin. By Lemma 1.19, there exists a family $\mathcal{N}''$ of absorbing, balanced, convex, and open sets that still forms a neighborhood basis of the origin. Given any $U \in \mathcal{N}''$, consider the Minkowski functional $p_U : X \to \mathbf{R}_+$ from Definition 1.16, which defines a seminorm by Proposition 1.17.

We first claim $U = \{p_U < 1\}$. We already know from Proposition 1.17 that $\{p_U < 1\} \subset U$. Now let $x \in U$. Since the scalar multiplication is continuous from $\mathbf{R} \times X$ to X and $1 \cdot x \in U$, there exists $s > 1$ such that $sx \in U$. In particular, this gives $p_U(x) < 1/s < 1$. Hence $x \in \{p_U < 1\}$, which shows $U \subset \{p_U < 1\}$.

Next, let us consider the topology induced by the seminorms $\{p_U : U \in \mathcal{N}''\}$ according to Definition 1.8. Since $\{p_U < \varepsilon\} = \varepsilon\,\{p_U < 1\}$ for every $\varepsilon > 0$, by continuity of the scalar multiplication the set $\{p_U < \varepsilon\}$ is open. Moreover, given any $x \in X$, continuity of the addition implies

$$\{p_U(\cdot - x) < \varepsilon\} = \{x + z \in X : p_U(z) < \varepsilon\} = x + \{p_U < \varepsilon\} \in \tau.$$

As open sets are stable under finite intersections, we conclude $B_{\varepsilon,I_0}(x) \in \tau$ for every $\varepsilon > 0$ and every $I_0 \subset I$ finite. This means the topology induced by the seminorms is coarser than τ. Conversely let $O \in \tau$ be open. Since addition is continuous, we know $x + \mathcal{N}''$ defines a neighborhood basis for any $x \in X$. Hence, we can write

$$O = \bigcup_{x\in O}(x + U_x) = \bigcup_{x\in O}(x + \{p_{U_x} < 1\}) = \bigcup_{x\in O}\{p_{U_x}(\cdot - x) < 1\}$$

with $U_x \in \mathcal{N}''$. Hence O is an open set also with respect to the topology induced by the above seminorms.

We finally need to show the Hausdorff property implies that if an element of X vanishes on all seminorms, it has to be zero. Assume $x \in X \setminus \{0\}$. Since (X, τ) is Hausdorff, there exists $U \in \mathcal{N}''$ such that $x \notin U$. Then the above reasoning implies $p_U(x) \geq 1$, which entails the claim. □

Given an LCTVS, from now on we can choose between a representation of the topology with seminorms or a convex neighborhood basis of the origin.

1.2. **Metrization and normability.** Given an LCTVS endowed with a family of seminorms $\mathcal{P} := \{p_i : i \in I\}$, it is a natural question to understand which conditions prevent the topology from being compatible with a metric or a norm. In this section we characterize these two notions. To this aim, we first need to discuss if one can get rid of some of the seminorms p_i while keeping the same induced topology τ — as we shall see, an LCTVS is metrizable if and only if the topology is induced by a countable family of seminorms.

Definition 1.20 (Seminorm basis)**.** *Let X be a vector space. Let $\mathcal{P}$ be a family of seminorms on X. A subfamily $\mathcal{Q} \subset \mathcal{P}$ is called a* **basis of seminorms** *for $\mathcal{P}$ if for every $p \in \mathcal{P}$ there exist $s > 0$ and $q \in \mathcal{P}$ such that $p(x) \leq s\,q(x)$ for every $x \in X$.*

You should compare this to the Lipschitz continuity of a norm $\|\cdot\|_1$ with respect to another norm $\|\cdot\|_2$ on X. Suppose there is $C > 0$ with $\|x\|_1 \leq C\,\|x\|_2$. What does this entail about the respectively induced topologies on X?

The following shows no topological information is lost when restricting ourselves to the topology induced by a basis of seminorms.

Lemma 1.21 (Reduction lemma)**.** *Let X be a vector space endowed with a family $\mathcal{P}$ of seminorms inducing the seminorm topology τ. Then any basis $\mathcal{Q}$ of seminorms for $\mathcal{P}$ induces the same topology.*

Proof. For clarity, we write $\tau_{\mathcal{P}} := \tau$. Let $\tau_{\mathcal{Q}}$ denote the topology induced by $\mathcal{Q}$.

By the inclusion $\mathcal{Q} \subset \mathcal{P}$, we obtain $\tau_{\mathcal{Q}} \subset \tau_{\mathcal{P}}$.

Conversely, let $U \in \tau_{\mathcal{P}}$ be given. By definition of $\tau_{\mathcal{P}}$, given any $x \in U$ there exist $p_1, \dots, p_n \in \mathcal{P}$ and $\varepsilon > 0$ such that $\bigcap_{i=1}^n \{p_i(\cdot - x) < \varepsilon\} \subset U$. Since $\mathcal{Q}$ is a basis of seminorms for $\mathcal{P}$, given any $i \in \{1 \dots, n\}$ there exist $q_i \in \mathcal{Q}$ and $s_i > 0$ such that $p_i \leq s_i\,q_i$. Since we are dealing with finitely many seminorms, replacing s_i by $s := \max\{s_1, \dots, s_n\}$, we may and will assume the constants do not depend on the index i. By the consequential inequality $p_i \leq s\,q_i$ for every $i \in \{1, \dots, n\}$,

$$\bigcap_{i=1}^n \{q_i(\cdot - x) < s^{-1}\,\varepsilon\} \subset \bigcap_{i=1}^n \{p_i(\cdot - x) < \varepsilon\} \subset U.$$

This shows $U \in \tau_{\mathcal{Q}}$ to conclude the inclusion $\tau_{\mathcal{P}} \supset \tau_{\mathcal{Q}}$. □

Example 1.22 (Continuation of Example 1.11). In the framework of Example 1.11, given any sequence $(n_k)_{k\in\mathbf{N}}$ diverging to ∞, $\mathcal{Q} := \{p_{n_k} : k \in \mathbf{N}\}$ is a basis for $\mathcal{P}$. Note the family $\mathcal{Q}$ is countable. In view of the upcoming Theorem 1.23, this implies the seminorm topology on $C(\mathbf{R})$ induced by $\mathcal{P}$ is metrizable. ■

Recall a topological space (X, τ) is **metrizable** if there exists a metric d on X that induces τ, viz. $\tau = \tau^{\mathsf{d}}$.

Theorem 1.23 (When is an LCTVS metrizable?)**.** *Let X be an LCTVS with a topology τ induced by a family of seminorms $\{p_i : i \in I\}$. Then (X, τ) is metrizable if and only if there exists a countable set $I' \subset I$ such that $\{p_i : i \in I'\}$ induces τ*[3]*.*

In particular, a metrizable LCTVS has some useful properties that all metric spaces have. For instance, it is second-countable, paracompact, Lindelöf, normal, and completely regular. We will not enter details as we do not need these properties and refer the interested reader to any standard textbook on topology.

Proof of Theorem 1.23. Suppose there exists a countable subfamily $\{p_n : n \in \mathbf{N}\}$ of seminorms generating τ. Define $\mathsf{d}_\tau \colon X^2 \to \mathbf{R}_+$ by

$$\mathsf{d}_\tau(x, y) = \sum_{n\in\mathbf{N}} 2^{-n} \frac{p_n(x - y)}{1 + p_n(x - y)}.$$

As shown in Exercise 2.4, d_τ is a metric generating τ. (A similar construction works if the subfamily is merely finite).

Conversely, assume τ is generated by a metric d. Then there exists a countable neighborhood basis of the origin given by the balls $\{B_{1/n}(0) : n \in \mathbf{N}\}$. From the definition of the topology τ, we infer for every $n \in \mathbf{N}$ there exists $\varepsilon_n > 0$ and $I_0^n \subset I$ finite with the property

$$B_{\varepsilon_n, I_0^n}(0) \subset B_{1/n}(0). \tag{1.4}$$

Define the countable set

$$I' := \bigcup_{n\in\mathbf{N}} I_0^n.$$

Since the family $\{B_{1/n}(0) : n \in \mathbf{N}\}$ of metric balls is a neighborhood basis of the origin, (1.4) implies the seminorms $\{p_i : i \in I'\}$ induce a finer topology than τ. Since $I' \subset I$, these topologies actually agree. □

Recall a topological space (X, τ) is **normable** if there exists a norm on X whose induced metric d obeys $\tau = \tau^{\mathsf{d}}$. A normable topological space is clearly metrizable, but the converse is not true (the discrete metric is not induced by a norm). Hence, normability of a topology is strictly stronger than metrizability. It is therefore no surprise that compared to Theorem 1.23, normability of locally convex topologies deals with a more restrictive realm than countable families of seminorms.

Proposition 1.24 (When is an LCTVS normable?)**.** *Let X be an LCTVS with a topology τ induced by a family $\{p_i : i \in I\}$ of seminorms. Then (X, τ) is normable if and only if there exists a finite set $I' \subset I$ such that $\{p_i : i \in I'\}$ induces τ.*

---

[3]One would be tempted to claim (X, τ) is metrizable if and only if there exists a countable basis for the family of seminorms $\{p_i : i \in I\}$. However, this is not true. For a counterexample suggested by Matthias Ruf, consider $X := C([0,1])$ with the seminorms $p_0(f) := \sup_{x\in[0,1/2]} |f(x)|$ and $p_{1/2}(f) := \sup_{[1/2,1]} |f(x)|$. Moreover, given any point $x \in (0, 1/2)$ consider the seminorm $p_x(f) := |f(x)| + |f(x + 1/2)|$. Then the space is even normable because $\max\{p_1, p_2\}$ generates the uniform convergence. However, by construction the family of seminorms $\{p_x : x \in [0, 1/2]\}$ does not contain a countable base.

Proof. Assume $\{p_i : i \in I'\}$ induces τ, where $I' \subset I$ is finite. Define the seminorm $\|x\|_\tau = \max_{i\in I'} p_i(x)$, which becomes a norm when the seminorms separate points in the sense of Definition 1.12. Is not difficult to show that this norm induces the same topology as τ[4].

To prove the converse, assume $\|\cdot\|$ is a norm on X that generates τ. Then there exist $\varepsilon, \delta > 0$ and $I_0 \subset I$ finite such that the induced balls satisfy

$$B_\delta(0) \subset \{x \in X : p_i(x) < \varepsilon \text{ for every } i \in I_0\} \subset B_1(0).$$

This shows the seminorms $\max_{i\in I_0} p_i$ and $\|\cdot\|$ are equivalent, which implies they generate the same topology. In particular, $\{p_i : i \in I_0\}$ is a finite subfamily of seminorms inducing the topology on X. □

Remark 1.25 (Comments on Theorems 1.18 and 1.23). The proofs of Theorems 1.18 and 1.23 reveal a LCTVS is metrizable if and only if there exists a countable convex neighborhood basis of the origin. ■

Lecture 3. This remark provides an intrinsic, more geometric characterization of metrizability for LCTVS which does not rely on seminorms. Similarly, normability can be characterized by a more geometric condition.

Definition 1.26 (Boundedness)**.** *Let (X, τ) be a TVS. A subset $E \subset X$ is called* **bounded** *if for every neighborhood V of 0 there exists $s > 0$ such that $E \subset sV$.*

Theorem 1.27 (Kolmogorov's criterion)**.** *Let X be a LCTVS. Then X is normable if and only if the origin has a bounded neighborhood.*

Proof. If $\|\cdot\|$ is a norm on X that generates the reference topology τ, the induced open unit ball is a bounded neighborhood of the origin since every neighborhood of the origin contains a set of the form $B_\delta(0)$ for some $\delta > 0$.

Conversely, we let U be a bounded neighborhood of the origin and denote by $\{p_i : i \in I\}$ a family of seminorms generating τ. By definition, there exist $\varepsilon > 0$ and $I_0 \subset I$ finite such that $\{p < \varepsilon\} \subset U$, where $p := \max_{i\in I_0} p_i$. Fix any other seminorm p_i with $i \in I$. Since $\{p_i < R\} = R\{p_i < 1\}$ for every $R > 0$ by homogeneity and $\{p_i < 1\}$ is a neighborhood of the origin, the boundedness of U implies that exists $s > 0$ such that

$$\varepsilon\{p < 1\} = \{p < \varepsilon\} \subset U \subset \{p_i < s\} \quad \Longrightarrow \quad \{p < 1\} \subset \{p_i < \varepsilon^{-1} s\}. \tag{1.5}$$

We claim (1.5) implies that

$$p_i \leq \frac{2s}{\varepsilon} p. \tag{1.6}$$

Once (1.6) is shown, since the argument works for any p_i, we conclude p induces the topology τ by arguing as in the proof of Lemma 1.21. Furthermore, since p is a seminorm inducing an Hausdorff topology, it must hold $p(x) = 0$ if and only if $x = 0$, implying p is the desired norm.

It remains to prove (1.6). Fix $x \in X$ and let us assume first $p(x) \neq 0$; then by homogeneity and (1.5), we have

$$p\Big[\frac{x}{2p(x)}\Big] = \frac{1}{2} < 1 \quad \Longrightarrow \quad p_i\Big[\frac{x}{2p(x)}\Big] = \frac{p_i(x)}{2p(x)} < \frac{s}{\varepsilon} \quad \Longrightarrow \quad p_i(x) < \frac{2s}{\varepsilon} p(x).$$

If $p(x) = 0$, by homogeneity and (1.5), for every $\delta > 0$ we have

$$p(x) < \delta \quad \Longrightarrow \quad x \in \delta\{p < 1\} \subset \delta\{p_i < \varepsilon^{-1} s\} \quad \Longrightarrow \quad p_i(x) < \frac{\delta s}{\varepsilon};$$

by the arbitrariness of δ, this implies $p_i(x) = 0$. Therefore we have verified (1.6), which concludes the proof. □

[4]In fact, given any $x \in X$ and any $\varepsilon > 0$ one has $B_{\varepsilon,I'}(x) = \{\|\cdot - x\|_\tau < \varepsilon\}$.

On TVS there exists a notion of Cauchy sequences — in fact, a related definition of Cauchy sequences makes sense in arbitrary topological spaces, not necessarily endowed with a metric.

Definition 1.28 (Cauchy sequence)**.** *Let X be a TVS. A sequence $(x_n)_{n\in\mathbf{N}}$ is called a* **Cauchy sequence** *if for every neighborhood U of the origin there exists $n_0 \in \mathbf{N}$ such that for every $n, m \in \mathbf{N}$ with $n, m \geq n_0$, we have $x_m - x_n \in U$.*

Remark 1.29 (About Definition 1.28). If there exists a translation-invariant metric d that generates the topology on X, then Definition 1.28 coincides with the metric definition of Cauchy sequences, since $\mathsf{d}(x, y) = \mathsf{d}(x - y, 0)$ for every $x, y \in X$. In particular, this is the case for a LCTVS (X, τ) with the metric induced by countably many seminorms (cf. the construction from Theorem 1.23).

However, translation-invariance is crucial, otherwise these two notions do not coincide. This can be shown for example by endowing $\mathbf{R}$ with the metric d given by $\mathsf{d}(x, y) := |\arctan(x) - \arctan(y)|$; d still induces the Euclidean topology on $\mathbf{R}$, but $(\mathbf{R}, \mathsf{d})$ is not sequentially complete, since the sequence $(x_n)_{n\in\mathbf{N}}$ given by $x_n := n$ is a nonconvergent Cauchy sequence. ■

Definition 1.30 (Fréchet space)**.** *A metric space (X, d) is* **sequentially complete** *if every Cauchy sequence $(x_n)_{n\in\mathbf{N}}$ converges.*

A LCTVS (X, τ) is called a **Fréchet space** *if τ is induced by a translation-invariant metric d which is sequentially complete.*

1.3. **Linear maps and the dual space.** In what follows, we let (X, τ) and (Y, ρ) be Hausdorff TVS. Let $\mathcal{L}(X, Y)$ denote the space of all continuous linear mappings from X to Y. Since Y is a TVS, $\mathcal{L}(X, Y)$ is itself a vector space by the usual addition and scalar multiplication for maps between vector spaces. A particular case is given by the choice $Y = \mathbf{R}$ equipped as usual with the Euclidean topology. In this case, we use the shorter notation $X' = \mathcal{L}(X, \mathbb{R})$, which is called the (topological) **dual space** of X. Before discussing the continuity of linear maps, we collect a similar result concerning the continuity of seminorms.

Lemma 1.31 (Continuity of seminorms)**.** *Let X be an LCTVS with seminorms $\{p_i : i \in I\}$ generating the reference topology. Consider another seminorm q on X. Then the following are equivalent.*

(i) *The seminorm q is continuous.*
(ii) *There exist $c > 0$ and $I_0 \subset I$ finite such that*

$$q \leq c \sum_{i\in I_0} p_i.$$

Proof. Exercise 3.3. □

Proposition 1.32 (Characterization of continuous linear maps)**.** *Let (X, τ) and (Y, ρ) be locally convex topological vector spaces. Let $\mathcal{P} := \{p_i : i \in I\}$ and $\mathcal{Q} := \{q_j : j \in J\}$ two families of seminorms generating τ and ρ, respectively. Let $T \colon X \to Y$ be linear. Then the following are equivalent.*

(i) *$T \in \mathcal{L}(X, Y)$.*
(ii) *T is continuous at zero. That is, for any $V \in \rho$ containing 0, $T^{-1}(V) \in \tau$.*
(iii) *For all $j \in J$ there exist $c_j > 0$ and $I_j \subset I$ finite such that the seminorm $P_j = \max_{i\in I_j} p_i$ satisfies*

$$q_j \circ T \leq c_j P_j.$$

In particular, for every $f \in X'$ there exist $c > 0$ and $I_0 \subset I$ finite such that

$$|f| \leq c \sum_{i\in I_0} p_i.$$

Proof. We will show (ii) $\implies$ (i), (i) $\implies$ (iii), and (iii) $\implies$ (ii).

(ii) $\implies$ (i). Let $V \subset Y$ be open and $x \in T^{-1}(V)$. Then the translation $V - T(x)$ is a neighborhood of 0 in Y. By the hypothesized continuity of T in 0, there exists a neighborhood U of 0 in X such that $T(U) \subset V - T(x)$. The set $x + U$ forms a neighborhood of x, and for all $z \in x + U$ one has

$$T(z) \subset T(x) + T(U) \subset T(x) + V - T(x) = V.$$

Thus $x + U \in T^{-1}(V)$ implying openness of $T^{-1}(V)$, verifying the continuity of T.

(i) $\implies$ (iii). By Lemma 1.31, q_j is continuous on Y. This implies continuity of the composition $q_j \circ T$. On the other hand, since q_j is a seminorm and T is linear, $q_j \circ T$ is a continuous seminorm. Applying Lemma 1.31 again, we deduce there exist $C_j > 0$ and $I_j \subset I$ finite such that

$$q_j \circ T \le C_j \sum_{i \in I_j} p_i \le C_j \, \#I_j \, \max_{i \in I_j} p_i.$$

(iii) $\implies$ (ii). It suffices to show that for any neighborhood of V of 0, $T^{-1}(V)$ is also a neighborhood of 0. Given such V, we can find $\varepsilon > 0$ and $J_0 \subset J$ finite such that $0 \in B_{\varepsilon, J_0}(0) \subset V$. For all $j \in J_0$, by assumption there exist c_j and a seminorm P_j — which is continuous by Lemma 1.31 — with $q_j \circ T \le c_j P_j$. Since P_j is continuous in 0, we can find an open neighborhood $W_j \subset X$ of the origin such that $P_j(W_j) < \varepsilon / c_j$. Defining $W := \bigcap_{j \in J_0} W_j$, by construction W is an open neighborhood of the origin satisfying

$$q_j \circ T \le c_j \, P_j < \varepsilon.$$

Hence $T(W) \subset U$ and $0 \in W \subset T^{-1}(U)$. Since $W \in \tau$, this shows $T^{-1}(U)$ is a neighborhood of 0, implying continuity of T in 0. □

Next we will discuss the finite-dimensional case, which is quite special in that continuity becomes a redundant assumption.

As usual, we equip Euclidean space with the usual Euclidean topology. We note that the second statement below does not simply follow from the first and basic linear algebra, since Y — albeit being a vector space as the image of a linear map — comes with its own subspace topology.

Proposition 1.33 (Continuity redundancy)**.** *Let X constitute a Hausdorff TVS and $Y \subset X$.*

(i) *If $f : \mathbf{R}^n \to X$ is linear, then f is continuous.*
(ii) *If $f : \mathbf{R}^n \to Y$ is linear and bijective, then $f^{-1} : Y \to \mathbf{R}^n$ is continuous.*
(iii) *If Y is a subspace of finite dimension, then Y is closed.*

Proof. (i) Let $\{e_1, \dots, e_n\}$ denote the canonical basis of $\mathbf{R}^n$. Then for every $x \in \mathbf{R}^n$,

$$f(x) = f\Big[\sum_{i=1}^n x_i \, e_i\Big] = \sum_{i=1}^n x_i \, f(e_i).$$

Since $f(e_1), \dots, f(e_n)$ are fixed vectors in X and all coordinate projections are continuous, the continuity of f follows from the hypothesized continuity of the vector space operations.

(ii) As f is bijective and linear, the map $f^{-1} : Y \to \mathbf{R}^n$ is well-defined and linear. By Proposition 1.32, to verify its continuity it suffices to do so at 0. That is, we need to show for any neighborhood U of 0 on $\mathbf{R}^n$, there exists a neighborhood $\tilde{W}$ of 0 on Y with $(f^{-1})(y) \in U$ for every $y \in \tilde{W}$. In other words, since $\{B_\varepsilon(0) : \varepsilon > 0\}$ is a neighborhood basis of 0 on $\mathbf{R}^n$, and the open sets on Y are defined by intersecting open sets of X with Y, we must verify the following: given any $\varepsilon > 0$, there is a neighborhood W of 0 in X such that $\|f^{-1}(y)\| < \varepsilon$ for all $y \in W \cap Y$.

To this end, set $S := \{x \in \mathbf{R}^n : \|x\| = 1\}$. Since S is compact, the continuity of f implied by (i) yields $f(S) \subset X$ is compact. Since $f(0) = 0$, the bijectivity of f implies $0 \notin f(S)$. Since X is a Hausdorff space, compact sets are closed and their complements are open. Therefore, there exists an open neighborhood V of 0 such that $V \cap f(S) = \emptyset$; by Exercise 1.3, we can take V to be balanced.

Define $E := f^{-1}(V)$. We claim $E \subset B_1(0)$. Suppose by contradiction there exists $x \in \mathbf{R}^n$ with $\|x\| \geq 1$ and $x \in E$; then $x/\|x\| \in S$ and the balancedness of V implies the relations

$$f\Big[\frac{x}{\|x\|}\Big] = \frac{f(x)}{\|x\|} \in \frac{1}{\|x\|} f(E) \subset \frac{1}{\|x\|} V \subset V,$$

which is a contradiction. Hence $\|y\| < 1$ for all $y \in f^{-1}(V)$. In particular, given any $\varepsilon > 0$, the set $W = \varepsilon V \cap Y$ is a neighborhood of the origin in Y. Since Y is a linear subspace and f^{-1} is linear,

$$f^{-1}(\varepsilon V \cap Y) = f^{-1}(\varepsilon(V \cap Y)) = \varepsilon f^{-1}(V \cap Y) = \varepsilon E \subset \varepsilon B_1(0) = B_\varepsilon(0).$$

Hence f^{-1} is continuous in 0.

(iii) Take $y \in \overline{Y}$. Set $d := \dim Y$, then by standard results from linear algebra there exists a bijective, linear map $f\colon \mathbf{R}^d \to Y$. Let $V \subset X$ be constructed as in (ii) and $E = f^{-1}(V)$. As V is an open neighborhood of the origin, it is absorbing, therefore there exists $s > 0$ such that $y \in sV$. We claim that openness of sV and the inclusion $y \in \overline{Y}$ imply $y \in \overline{Y \cap sV}$. To see this, we use the characterization of the closure by neighborhoods, viz. $y \in \overline{Y \cap sV}$ if and only if, for any neighborhood W of y, we have $W \cap (Y \cap sV) \neq \emptyset$.

Since sV is open, $W \cap sV$ is a neighborhood of y; since $y \in \overline{Y}$, applying again the aforementioned characterization (for $\tilde{W} = W \cap sV$), we must have

$$W \cap (Y \cap sV) = (W \cap sV) \cap Y = \tilde{W} \cap Y \neq \emptyset.$$

Since this holds for any W, we deduce $y \in \overline{Y \cap sV}$.

Next, the homogeneity and bijectivity of f from $\mathbf{R}^d$ to Y (and the fact $sY = Y$) imply $Y \cap sV = f(sE)$. Hence we have the inclusions

$$y \in \overline{Y \cap sV} \subset \overline{f(sE)} \subset \overline{f(\overline{B_s(0)})} = f(\overline{B_s(0)}) \subset Y,$$

where we used that since $\overline{B_s(0)}$ is compact, so is $f(\overline{B_s(0)})$, which is therefore also closed since X is Hausdorff. Thus $\overline{Y} = Y$, which concludes the proof. □

Remark 1.34 (Consequences). The above result has several key implications.

- If (X, τ) is a finite-dimensional Hausdorff TVS we can construct a linear isomorphism $f\colon X \to \mathbf{R}^{\dim X}$, the latter being endowed with the Euclidean topology. In particular, f is continuous with continuous inverse, i.e. a linear homeomorphism. This yields open sets on X and on $\mathbf{R}^{\dim X}$ are in a one-to-one correspondence.
- Consider now $\mathbf{R}^d$ with locally convex topologies τ_1 and τ_2 induced by two norms $\|\cdot\|_1$ and $\|\cdot\|_2$ respectively. Then the identity map from $(\mathbf{R}^d, \tau_1)$ to $(\mathbf{R}^d, \tau_2)$ is a homeomorphism, whence $\tau_1 = \tau_2$. Since $\|\cdot\|_1$ and $\|\cdot\|_2$ induce the same topology, applying Lemma 1.31 twice they must be **equivalent norms** (in the usual sense). In other words, there exists only one topology on $\mathbf{R}^d$ which makes it a Hausdorff TVS, which is the Euclidean one. As a byproduct, we have shown the basic fact that **all norms on $\mathbf{R}^d$ are equivalent**. By (i), similar considerations apply to any other finite dimensional Hausdorff TVS (X, τ).
- Let (X, τ) be a Hausdorff TVS. Then it can be shown the origin has a neighborhood basis of precompact sets if and only if X is finite-dimensional. For one implication, apply (i) above. The other will studied in Exercise 4.4.

In particular, on infinite-dimensional normed spaces, the closed unit ball $\overline{B}_1(0)$ is **never compact** — a fact that is usually already shown in a first course on functional analysis. ■

Lecture 4. One of the most important properties of locally convex topological vector spaces is the presence of a Hahn–Banach theorem, which we state below without proof. There exist more versions that distinguish whether X is a real or complex vector space, cf. e.g. [10, §3].

Theorem 1.35 (Hahn–Banach theorem, analytical version). *Let X constitute a vector space and let $p\colon X \to \mathbf{R}_+$ be a seminorm. Suppose $M \subset X$ is a linear subspace and $f\colon M \to \mathbf{R}$ is a linear functional such that $|f(x)| \le p(x)$ for every $x \in M$. Then f can be extended to a (nonrelabeled) linear functional $f : X \to \mathbb{R}$ satisfying the property $|f(x)| \le p(x)$ for all $x \in X$.*

Theorem 1.36 (Hahn–Banach theorem, geometric version). *Let (X, τ) be a LCTVS and let $A, B \subset X$ be convex sets such that A is compact and B closed. If $A \cap B = \emptyset$, then there exist $x' \in X'$ and $\alpha, \beta \in \mathbf{R}$ such that for every $x \in A$ and every $y \in B$,*

$$x'(x) < \alpha < \beta < x'(y).$$

From these theorems, we deduce in particular that the dual space of a locally convex topological vector space separates points.

Corollary 1.37 (Point separation in the dual space). *Let (X, τ) be a LCTVS and $x \in X \setminus \{0\}$. Then there exists $x' \in X'$ such that $x'(x) \neq 0$.*

Proof. Apply Theorem 1.36 to the compact sets $\{0\}$ and $\{x\}$, which are closed as (X, τ) is Hausdorff. Clearly $x'(0) = 0$ for every $x' \in X'$, so that $x'(x) \neq 0$. □

Definition 1.38 (Weak topology). *Let (X, τ) be a LCTVS. We define the* **weak topology** *on X as the coarsest topology τ_w such that all elements $x' \in X'$ are continuous with respect to τ_w.*

Corollary 1.39 (Local convexity). *If (X, τ) is a LCTVS, so is (X, τ_w).*

Proof. Thanks to Corollary 1.37, the proof is identical to Exercise 2.1. □

Similarly to the weak topology and in analogy to the case of Banach spaces, we define the so-called weak* topology on X'.

Definition 1.40 (Weak* topology). *Let (X, τ) be a TVS with dual X'. The* **weak* topology** *on X' is the coarsest topology such that the mappings $x' \mapsto x'(x)$ are continuous on X' for every $x \in X$.*

In the above situation, X' always becomes a LCTVS when equipped with the weak* topology. We will study this and further elementary properties in the exercises. One could fill a whole course with locally convex spaces, but we shall stop here with the abstract considerations as they suffice for the next chapter. The interested reader should consult [10, §§1–3] for a complete treatment about basic linear functional analysis on topological vector spaces. See also §A for more details about weak and weak* topologies on LCTVS.

2. Test functions and distributions

We start by introducing a locally convex topology on the function space $C^\infty(\Omega)$[5] of functions on Ω which are differentiable infinitely often, where $\Omega \subset \mathbf{R}^d$ is open.

[5]In all function spaces, unless stated otherwise the target domain will be $\mathbf{R}$. The theory that we develop also holds for $\mathbf{C}$-valued functions; in that case, note however that differentiability does never refer to complex differentiability.

Remark 2.1 (Disclaimer). This space and its topology discussed below are not going to be the usual ones used in the theory of distributions! Those will appear later in Definition 2.10. The relevant space will be the set $C_c^\infty(\Omega)$ of *compactly supported* elements of $C^\infty(\Omega)$, cf. Definition 2.2 below. To avoid confusion of these two spaces, we will write $\mathcal{D}(\Omega)$ in place of $C_c^\infty(\Omega)$. ■

Definition 2.2 (Test functions)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and $\varphi \in C(\Omega)$.*

a. *The* **support** *of φ is defined by*

$$\operatorname{spt}\varphi := \overline{\{x \in \Omega : \varphi(x) \neq 0\}}^{\Omega} = \overline{\{x \in \Omega : \varphi(x) \neq 0\}}^{\mathbf{R}^d} \cap \Omega.$$

b. *We define the set of* **test functions** *by*

$$\mathcal{D}(\Omega) := \{\varphi \in C^\infty(\mathbf{R}^d) : \operatorname{spt}\varphi \textit{ is a compact subset of } \Omega\}.$$

c. *Given a compact set $K \subset \mathbf{R}^d$, we further define the space*

$$\mathcal{D}_K = \{\varphi \in C^\infty(\mathbf{R}^d) : \operatorname{spt}\varphi \subset K\}.$$

Remark 2.3 (Basic observations). Since $\operatorname{spt}(f+g) \subset \operatorname{spt} f \cup \operatorname{spt} g$ and $\operatorname{spt}(\lambda f) = \mathit{supp} f$ for every continuous functions f and g and every $\lambda \in \mathbf{R} \setminus \{0\}$, the spaces $\mathcal{D}(\Omega)$ and $\mathcal{D}_K$ are vector spaces with respcect to the usual pointwise addition and scalar multiplication of functions.

Moreover, as shown in the exercises, the function $f\colon \mathbf{R} \to \mathbf{R}$ given by

$$f(t) := \begin{cases} \mathrm{e}^{-1/t} & \text{if } t > 0, \\ 0 & \text{otherwise} \end{cases}$$

lies in $C^\infty(\mathbf{R})$. By the chain rule, the function

$$\varphi(x) = \begin{cases} \mathrm{e}^{1/(\|x\|^2-1)} & \text{if } \|x\| < 1, \\ 0 & \text{otherwise,} \end{cases}$$

belongs to $C^\infty(\mathbf{R}^d)$ and satisfies $\operatorname{spt}\varphi = \overline{B}_1(0)$. Given any $x_0 \in \mathbf{R}^d$ and any $\varepsilon > 0$, the function $\varphi_{x_0,\varepsilon}\colon \mathbf{R}^d \to \mathbf{R}$ defined through $\varphi_{x_0,\varepsilon}(x) := \varphi((x-x_0)/\varepsilon)$ belongs to $C^\infty(\mathbf{R}^d)$ with $\operatorname{spt}\varphi_{x_0,\varepsilon} = \overline{B}_\varepsilon(x_0)$[6]. In particular, $\mathcal{D}(\Omega) \neq \{0\}$; the same holds for $\mathcal{D}_K$ if K has non-empty interior[7]. Since both spaces are stable under multiplication with smooth functions, it follows that both spaces are infinite dimensional (again when K has nonempty interior). ■

In Definition 2.4 below, we construct a locally convex topology on $C^\infty(\Omega)$; this induces a similar topology on $\mathcal{D}_K$, which is a linear subspace of $C^\infty(\Omega)$ provided $K \subset \Omega$ is compact. Later in Definition 2.10, we will see a different locally convex topology on $\mathcal{D}(\Omega)$ that induces the same relative topology on $\mathcal{D}_K$.

We recall some basic notation from vector calculus. A **multiindex** is an element of $\mathbf{N}_0^d$. The **norm** of a multiindex $\alpha \in \mathbf{N}_0^d$ is defined by $|\alpha| := \alpha_1 + \cdots + \alpha_n$. Moreover, for such an α we define the differential operator (the "α'th partial derivative")

$$D^\alpha := \Big[\frac{\partial}{\partial x_1}\Big]^{\alpha_1} \cdots \Big[\frac{\partial}{\partial x_d}\Big]^{\alpha_d},$$

where the above powers mean so-and-so-fold application of the respective partial derivative. It will always act on smooth functions, so by Schwarz' theorem, we can group the partial derivatives by directions and not worry about their order.

[6] The flock of functions thus defined is a very useful tool to regularize functions in the study of PDEs often called (standard) mollifier (up to a dimensional "normalization"). We refer to [6, §C] for details.

[7] When K has empty interior, we have $\mathcal{D}_K = \{0\}$.

Definition 2.4 (A seminorm topology on $C^\infty(\Omega)$)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and $(K_n)_{n\in\mathbf{N}}$ be a sequence of compact sets with $K_n \subset \operatorname{int} K_{n+1}$ for every $n \in \mathbf{N}$ and $\Omega = \bigcup_{n\in\mathbf{N}} K_n$*[8] *For $N \in \mathbf{N}$, we define the seminorm $p_N\colon C^\infty(\Omega) \to \mathbf{R}_+$ by*

$$p_N(\varphi) := \max\{|D^\alpha\varphi(x)| : \alpha \in \mathbf{N}_0^d,\ |\alpha| \le N,\ x \in K_N\}.$$

Let ρ denote the induced seminorm topology.

Let τ_K denote the relative topology induced by the family of seminorms $\{p_N : N \in \mathbf{N}\}$ on the space $\mathcal{D}_K$ whenever $K \subset \Omega$ is compact.

Definition 2.5 (Heine–Borel property)**.** *We say a TVS (X,τ) has the* **Heine–Borel property** *if every bounded and closed subset of X is compact.*

Of course, this terminology stems from the Heine–Borel theorem, which states $\mathbf{R}^d$ has the Heine–Borel property.

Proposition 2.6 (Fundamental properties)**.** *The LCTVS $(C^\infty(\Omega),\rho)$ is a Fréchet space and has the Heine–Borel property.*

Moreover, if $K \subset \Omega$ is compact, then $\mathcal{D}_K$ is a closed subset of $C^\infty(\Omega)$.

Proof. The separation property from Definition 1.12 is satisfied, implying the seminorms $\{p_N : N \in \mathbf{N}\}$ induce a locally convex topology.

By construction, this family of seminorms is countable, implying metrizability of the space in question by Theorem 1.23. By the proof of that theorem, the metric inducing the seminorm topology in question may and will be chosen to be translation-invariant. To be a Fréchet space, we are left to show completeness. Let $(\varphi_n)_{n\in\mathbf{N}}$ be a Cauchy sequence in $C^\infty(\Omega)$. This means that for every $\varepsilon > 0$ and every $N \in \mathbf{N}$ there exists $M \in \mathbf{N}$ such that

$$\sup_{x\in K_N} |D^\alpha\varphi_n(x) - D^\alpha\varphi_m(x)| < \varepsilon$$

for all multiindices $\alpha \in \mathbf{N}_0^d$ with $|\alpha| \le N$ and every $n, m \ge M$. In particular, for any fixed such α and a fixed compact set $K \subset \Omega$ the sequence $(D^\alpha\varphi_n)_{n\in\mathbf{N}}$ is a Cauchy sequence with respect to the uniform topology on K. Since the latter is complete, $(D^\alpha\varphi_n)_{n\in\mathbf{N}}$ converges uniformly on K to a continuous function $g_\alpha\colon K \to \mathbf{R}$. Since the sets $(K_n)_n$ exhaust Ω, the limit g_α does not depend on K, $D^\alpha\varphi_n(x) \to g_\alpha(x)$ as $n \to \infty$ for every $x \in \Omega$, and $(D^\alpha\varphi_n)_{n\in\mathbf{N}}$ converges locally uniformly (i.e. uniformly on each compact subset of Ω) to the function g_α. By standard results from analysis, we obtain $D^\alpha g_0(x) = g_\alpha(x)$ for every $x \in \Omega$, so that $g_0 \in C^\infty(\Omega)$. Moreover, by construction $\varphi_n \to g_0$ as $n \to \infty$ with respect to the topology of $C^\infty(\Omega)$, as the convergence reduces to the uniform convergence on compact subsets for any partial derivative. This shows the claimed completeness.

We turn to the Heine–Borel property. Let $E \subset C^\infty(\Omega)$ be bounded and closed. Since the topology in question is metrizable, it suffices to prove sequential compactness of E. This will be a consequence of the Arzelà–Ascoli Theorem 2.7 below (applied to each partial derivative), whose hypotheses we now verify. By Exercise 3.1, boundedness of E implies for every $N \in \mathbf{N}$ there exists $M_N > 0$ such that $\sup\{p_N(\varphi) : \varphi \in E\} \le M_N$. Note that the inequality $\sup_{x\in K_N} |D^\alpha\varphi(x)| \le M_N$ whenever $|\alpha| \le N$ implies the sequence $(D^\beta\varphi_n)_{n\in\mathbf{N}}$ is equicontinuous — in fact, equi-Lipschitz, i.e. Lipschitz continuous with uniformly bounded Lipschitz constants — on K_{N-1} whenever $|\beta| \le N-1$. Indeed, since $K_{N-1} \subset \operatorname{int} K_N$, compactness implies $r_N := \operatorname{dist}(K_{N-1}, \partial K_N) > 0$. Hence for any $x \in K_{N-1}$ one has $B_{r_N}(x) \subset K_N$; the mean value theorem implies for every $y \in B_{r_N}(x)$ that

$$|D^\beta\varphi(y) - D^\beta\varphi(x)| \le \sup_{z\in B_{r_N}(x)} |\nabla D^\beta\varphi(z)||y-x| \le d\, p_N(\varphi)\, |y-x|.$$

[8]In this case, we say the sequence $(K_n)_{n\in\mathbf{N}}$ **exhausts** Ω. For an explicit construction, one can take $K_n := \{x \in \Omega : |x| \le n,\ \mathsf{d}(x,\partial\Omega) \ge 1/n\}$; cf. Exercise 4.2.

This yields the desired equi-Lipschitz continuity. For equiboundedness, note that $\sup_{x\in K_{N-1}} |D^\beta\varphi(x)| \leq p_N(\varphi) \leq M_N$. Therefore, the Arzelà–Ascoli theorem applies which, together with a classical diagonal argument, implies for every sequence $(\varphi_n)_{n\in\mathbf{N}}$ in E there exists a subsequence that converges in $C^\infty(\Omega)$. As E is closed, this limit belongs to E.

Finally, we show closedness of $\mathcal{D}_K$ in $C^\infty(\Omega)$. Note that for every $x \in \Omega$ the function $\delta_x\colon C^\infty(\Omega) \to \mathbf{R}$[9] defined by $\delta_x(\varphi) := \varphi(x)$ is linear and continuous, so that its kernel is closed. Since $\varphi \in \mathcal{D}_K$ if and only if $\varphi \in C^\infty(\Omega)$ and $\delta_x(\varphi) = 0$ for every $x \in \Omega \setminus K$ (noting K is closed), the set $\mathcal{D}_K$ is closed as the intersection of closed sets. □

The following result, used to show the Heine–Borel property in the above proof and repeated for the convenience of the reader, forms *the* central characteristic of (pre)compactness result in spaces of continuous functions.

To formulate it, given a compact metric space (K, d), let $C(K)$ denote the space of real-valued continuous functions defined on K. When endowed with the uniform topology (i.e. the topology induced by the usual supremum norm on K), this becomes a Banach space. Lastly, recall a subset of a Hausdorff topological space is called precompact if its closure is compact.

Theorem 2.7 (Arzelà–Ascoli theorem)**.** *Let (K, d) be a compact metric space. Then a set $\mathcal{F} \subset C(K)$ is precompact with respect to the uniform topology if and only if the following two conditions hold simultaneously.*

(i) **Uniform equicontinuity.** *For every $\varepsilon > 0$ there exists $\delta > 0$ such that for every $x \in K$, we have the implication*
$$\mathsf{d}(x,y) \leq \delta \quad\Longrightarrow\quad \sup_{f\in\mathcal{F}} |f(y) - f(x)| \leq \varepsilon.$$

(ii) **Pointwise boundedness.** *For every $x \in K$, the subset $\{f(x) : f \in \mathcal{F}\}$ is a bounded subset of $\mathbf{R}$.*

Uniform equicontinuity might seem like an odd condition at first glance. We hope the following example (that you might want to comprehend as an instructive exercise) clarifies the two conditions above in a more practical way.

Example 2.8 (Lipschitz functions). Let (K, d) be a compact metric space. A function $f\colon K \to \mathbf{R}$ is called Lipschitz continuous if there exists a constant $L > 0$ such that $|f(y) - f(x)| \leq L\,\mathsf{d}(x,y)$ for every $x, y \in K$. The smallest possible choice of L with this property is denoted $\operatorname{Lip} f$ and called *Lipschitz constant* of f.

Let $\mathcal{F}$ be any given set of functions with $\sup_{f\in\mathcal{F}} \operatorname{Lip} f < \infty$. Then $\mathcal{F}$ is a subset of $C(K)$; it is in fact precompact therein. ■

Remark 2.9 (Further considerations)*.* We have the following

- Since $\mathcal{D}_K$ is closed in $C^\infty(\Omega)$, $\mathcal{D}_K$ is a Fréchet space as well. It also has the Heine–Borel property: bounded sets in $\mathcal{D}_K$ are also bounded in $C^\infty(\Omega)$.
- The above topology on $C^\infty(\Omega)$ does not depend on the sequence $(K_n)_{n\in\mathbf{N}}$ of compact sets exhausting Ω. Conversely, note that a neighborhood basis of the origin in $(\mathcal{D}_K, \tau_K)$ is given by the sets
$$V_{\varepsilon,N} = \big\{\varphi \in \mathcal{D}_K : \sup_{x\in K} |D^\alpha\varphi(x)| < \varepsilon \text{ whenever } |\alpha| \leq N\big\},$$
so that the topology τ_K of $\mathcal{D}_K$ does not depend on the ambient open set Ω containing K. In fact, τ_K is induced by the family $\{\tilde{p}_N : N \in \mathbf{N}\}$ of

[9] Later and in the literature, this is called Dirac δ-distribution.

seminorms given as follows for every $\varphi \in \mathcal{D}_K$:

$$\tilde{p}_N(\varphi) := \max\{|D^\alpha \varphi(x)| : x \in K,\, \alpha \in \mathbf{N}_0^d,\, |\alpha| \le N\}. \tag{2.1}$$

- In light of Proposition 2.6, $C^\infty(\Omega)$ might seem like a great space to work with. However, it has the problem that an element $\varphi \in C^\infty(\Omega)$ might behave badly as it approaches the boundary of Ω. As a practical example, consider $\Omega := (0, +\infty)$ and $\varphi \in C^\infty(\Omega)$ defined by $\varphi(x) := \sin(1/x)$. This function has no limit as $x \to 0+$; even worse, all its derivatives explode near zero. The spaces $\mathcal{D}_K$ and $\mathcal{D}(\Omega)$ do not have this problem.
- One could try to equip the space $\mathcal{D}(\Omega)$ with the relative topology of $C^\infty(\Omega)$. In this way we would produce a metrizable, locally convex topology. However, **this space fails to be complete**! Intuitively, this is clear: the limit of a sequence $(\varphi_n)_{n\in\mathbf{N}}$ in $\mathcal{D}(\Omega)$ may fail to have compact support in Ω, since the supports $\operatorname{spt}\varphi_n$ can reach the boundary as $n \to \infty$. See Exercise 5.3 for a counterexample. ■

Lecture 5.

2.1. **A locally convex topology on $\mathcal{D}(\Omega)$.** Let us heuristically discuss what we have seen thus far. Let $\Omega \subset \mathbf{R}^d$ open. On one hand, given any compact $K \subset \Omega$, we have the Fréchet spaces $(\mathcal{D}_K, \tau_K)$. On the other, we have $(C^\infty(\Omega), \rho)$, which is also a Fréchet space, but allows for functions with pathological behavior at the boundary. The space $\mathcal{D}(\Omega)$ lies somewhat in between: given an exhaustion $(K_n)_{n\in\mathbf{N}}$ of Ω,

$$\mathcal{D}_K \subset \mathcal{D}(\Omega) = \bigcup_{n\in\mathbf{N}} \mathcal{D}_{K_n} \subset C^\infty(\Omega).$$

We would like to endow $\mathcal{D}(\Omega)$ with an appropriate topology. However, by Remark 2.9, the topology inherited from its inclusion into $C^\infty(\Omega)$ does not appear right. Ideally, the topology τ we are looking for should satisfy the following.

- It should turn $(\mathcal{D}(\Omega), \tau)$ into a LCTVS. In this way, we have the Hahn–Banach theorems at our disposal and we can turn its dual space $\mathcal{D}'(\Omega)$, endowed with the weak* topology, into a LCTVS as well, cf. the discussion right after Definition 1.40.
- It should "respect" the inclusion $\mathcal{D}_K \subset \mathcal{D}(\Omega)$, where $K \subset \Omega$ is compact. Namely, the topology induced by τ on $\mathcal{D}_K$ as a subspace of $\mathcal{D}(\Omega)$ should coincide exactly with τ_K. We would then hope $(\mathcal{D}(\Omega), \tau)$ inherits some of the nice topological properties of $(\mathcal{D}_K, \tau_K)$.
- The topology τ should be as large as possible. In this way, "many" linear operators will become continuous, and we can expect a rich structure for the dual $\mathcal{D}'(\Omega)$.

We are now in a position to introduce a topology on $\mathcal{D}(\Omega)$ that fullfills the above requirements. The price we pay is we lose metrizability of $\mathcal{D}(\Omega)$ (see Corollary 2.15 below). On the other hand, we gain many useful properties inherited from the spaces $\mathcal{D}_{K_n}$ coming from the exhaustion sequence $(K_n)_{n\in\mathbf{N}}$ (see Theorem 2.14).

Definition 2.10 (Topology on $\mathcal{D}(\Omega)$)**.** *Let $\Omega \subset \mathbf{R}^d$ be open. We set*

$$\begin{aligned}\mathcal{N}(\Omega) := \{U \subset \mathcal{D}(\Omega) : U \text{ balanced and convex},\\ U \cap \mathcal{D}_K \in \tau_K \text{ for every compact } K \subset \Omega\}.\end{aligned}$$

We define a collection τ of subsets of $\mathcal{D}(\Omega)$ as follows. We say $E \in \tau$ if and only if it can be written as the (possibly empty) union of sets of the form $\varphi + U$, where $\varphi \in \mathcal{D}(\Omega)$ and $U \in \mathcal{N}(\Omega)$.

Definition 2.10 is rather abstract, but we will gradually see in the next results why it is the right choice. If it does not appear intuitive, this is because indeed it is not. Historically, Schwartz[10] actually first defined the notion of *convergence* in $\mathcal{D}(\Omega)$ (see Theorem 2.14), which in practical applications to PDEs is often enough, and only later understood how to construct the topology τ inducing this convergence.

We first have to show τ actually defines a locally convex topology.

Proposition 2.11 (Local convexity)**.** *Consider the class τ from Definition 2.10. Then $(\mathcal{D}(\Omega), \tau)$ is a LCTVS and $\mathcal{N}(\Omega)$ defines a neighborhood basis of the origin.*

Before entering the proof, we notice $\mathcal{N}(\Omega)$ is stable under finite intersections and multiplication by nonzero scalars. The first fact can be verified using the intersection of balanced and convex sets is still balanced and convex, and τ_K is stable under finite intersection; the second one is similar.

Proof. We start by showing τ defines a topology. By definition, $\emptyset, \mathcal{D}(\Omega) \in \tau$ and τ is stable under arbitrary unions. To show τ is closed under finite intersections, by induction it suffices to show that if $V_1, V_2 \in \tau$ then $V_1 \cap V_2 \in \tau$. To this end, it suffices to show that for any $\varphi \in V_1 \cap V_2$, there exists $U \in \mathcal{N}(\Omega)$ such that

$$\varphi + U \subset V_1 \cap V_2. \tag{2.2}$$

Once (2.2) is shown, we can conclude τ is a topology. Since $V_1, V_2 \in \tau$, there exist $\phi_1, \phi_2 \in \mathcal{D}(\Omega)$ and $U_1, U_2 \in \mathcal{N}(\Omega)$ with $\varphi \in \phi_1 + U_1 \subset V_1$ and $\varphi \in \phi_2 + U_2 \subset V_2$. Next, let $K \subset \Omega$ be compact such that $\varphi, \phi_1, \phi_2 \in \mathcal{D}_K$. Note $U_1 \cap \mathcal{D}_K, U_2 \cap \mathcal{D}_K \in \tau_K$. By continuity of the scalar multiplication there exists $\delta > 0$ such that $\varphi - \phi_1 \in (1-\delta)U_1$ and $\varphi - \phi_2 \in (1-\delta)U_2$. By convexity of U_1 and U_2, for every $i \in \{1,2\}$,

$$\varphi - \phi_i + \delta U_i \subset (1-\delta)U_i + \delta U_i = U_i.$$

Hence, setting $U = \delta U_1 \cap \delta U_2$, we deduce that for all i as above,

$$\varphi + U \subset \varphi + \delta U_i \subset \phi_i + U_i \subset V_i.$$

Therefore $\varphi + U \subset V_1 \cap V_2$, which proves the first claim since the intersection of balanced and convex sets remains balanced and convex.

The claim that $\mathcal{N}(\Omega)$ defines a neighborhood basis of 0 follows immediately from the definition of $\mathcal{N}(\Omega)$.

It remains to show the local convexity of the topology. By Theorem 1.18, it suffices to show $\mathcal{D}(\Omega)$ is a Hausdorff topological vector space with a convex neighborhood basis of the origin (which in this case is naturally chosen to be $\mathcal{N}(\Omega)$, which is already made of convex sets by definition).

We start with the Hausdorff property. Fix distinct $\varphi_1, \varphi_2 \in \mathcal{D}(\Omega)$ and set

$$W := \{\varphi \in \mathcal{D}(\Omega) : 2\|\varphi\|_\infty < \|\varphi_1 - \varphi_2\|_\infty\}, \tag{2.3}$$

where $\|\cdot\|_\infty$ is the usual supremum norm. We claim $W \in \mathcal{N}(\Omega)$. Indeed, since it is a ball with respect to a norm, W is balanced and convex. Moreover, for every compact $K \subset \Omega$,

$$W \cap \mathcal{D}_K = \big\{\varphi \in \mathcal{D}_K : 2 \sup_{x \in K} |\varphi(x)| < \|\varphi_1 - \varphi_2\|_\infty\big\}, \tag{2.4}$$

which belongs to τ_K by Remark 2.9. This shows $W \in \mathcal{N}(\Omega)$. Therefore, the sets $W_i := \varphi_i + W$ are open and contain φ_i, where $i \in \{1,2\}$. Moreover, if $\varphi \in W_1 \cap W_2$,

$$\|\varphi_1 - \varphi_2\|_\infty \le \|\varphi_1 - \varphi\|_\infty + \|\varphi - \varphi_2\|_\infty < \|\varphi_1 - \varphi_2\|_\infty,$$

which yields a contradiction. Hence τ has the Hausdorff property.

[10] Laurent Schwartz (1915–2002), French mathematician, received a Fields Medal in 1950 for his invention of the theory of distributions.

We turn to continuity of the vector space operations. To this aim, we shall systematically use that sets of the form $\psi + \mathcal{N}(\Omega)$ form a basis of the topology. Let $\varphi_1, \varphi_2 \in \mathcal{D}(\Omega)$ and $W \in \mathcal{N}(\Omega)$. By convexity of W,

$$\varphi_1 + \frac{1}{2}W + \varphi_2 + \frac{1}{2}W \subset (\varphi_1 + \varphi_2) + W,$$

which implies continuity, since $W/2 \in \mathcal{N}(\Omega)$. To treat the scalar multiplication, fix $\lambda_0 \in \mathbf{R}$, $\varphi_0 \in \mathcal{D}(\Omega)$, and $W \in \mathcal{N}(\Omega)$. It suffices to show there exist $U \in \mathcal{N}(\Omega)$ and $\delta > 0$ such that for every $\lambda \in (\lambda_0 - \delta, \lambda_0 + \delta)$ and every $\varphi \in \varphi_0 + U$,

$$\lambda\varphi \in \lambda_0\varphi_0 + W. \tag{2.5}$$

We first claim for $W \in \mathcal{N}(\Omega)$ there exists $\delta > 0$ such that $\delta\varphi_0 \in W/2$. Indeed, there exists $K \subset \Omega$ compact such that $\varphi \in \mathcal{D}_K$. Then by continuity of the scalar multiplication in $\mathcal{D}_K$, there exists $\delta > 0$ such that $\delta\varphi_0 \in W/2 \cap \mathcal{D}_K$, as the latter set is an open neighborhood of 0 in $\mathcal{D}_K$. Given such $\delta > 0$, fix $c > 0$ such that $2c(|\lambda_0| + \delta) = 1$ and set $U := cW$. By balancedness and convexity of W, for all $\lambda \in \mathbf{R}$ with $|\lambda - \lambda_0| < \delta$ and $\varphi \in \varphi_0 + U = \varphi_0 + cW$ we obtain

$$\lambda\varphi - \lambda_0\varphi_0 = \lambda(\varphi - \varphi_0) + (\lambda - \lambda_0)\varphi_0 \subset \frac{1}{2}W + \frac{1}{2}W \subset W.$$

which proves the desired relation (2.5). □

Remark 2.12 (Explicit generating seminorms[11]). It is not immediate to find an explicit form of the seminorms generating the topology. We know from Theorem 1.18 that we can take the Minkowski functionals p_U, where $U \in \mathcal{N}(\Omega)$. With some effort, one can show a neighborhood basis of the origin is given by sets of the form

$$V_{A,\varepsilon} := \{\varphi \in \mathcal{D}(\Omega) : |D^\alpha\varphi(x)| < \varepsilon(x) \text{ for every } x \in \Omega,\, |\alpha| \leq A(x)\},$$

where $\varepsilon : \Omega \to (0, \infty)$ and $A \colon \Omega \to (0, \infty)$ are continuous functions. Equivalently, a corresponding seminorm is given by

$$p_{A,\varepsilon}(\varphi) := \sup\{|\varepsilon^{-1}(x)\, D^\alpha\varphi(x)| : x \in \Omega,\, |\alpha| \leq A(x)\}.$$

As we will not use these seminorms, we leave the proof — partitions of unity are helpful — to the motivated reader. ■

Remark 2.13 (Algebraic considerations[12]). For the readers familiar with categories, $(\mathcal{D}(\Omega), \tau)$ is the locally convex direct inductive limit of the Fréchet spaces $(\mathcal{D}_{K_n}, \tau_{K_n})$, for any sequence $(K_n)_{n\in\mathbf{N}}$ of compacts exhausting Ω. In particular, $(\mathcal{D}(\Omega), \tau)$ is a so-called LF-space.

Alternatively, the topology τ can be characterized as the largest topology such that for every $n \in \mathbf{N}$, the inclusion $\iota_n \colon (\mathcal{D}_{K_n}, \tau_{K_n}) \to (\mathcal{D}(\Omega), \tau)$ is continuous. ■

In what follows, we tacitly endow $\mathcal{D}(\Omega)$ with the topology τ given in Definition 2.10. **All topological results refer to this topology**, unless explicitly stated otherwise.

Theorem 2.14 below provides many important structural properties of τ. Before stating and proving it, we need to recall some basic facts.

- Given a LCTVS (X, τ_X) and a closed linear subspace $Y \subset X$, endowing Y with the subspace topology τ_Y, many properties of τ_Y and τ_X are in a one-to-one correspondence, cf. Exercise 6.1.

[11]This remark is not examinable.
[12]This remark is not examinable.

- If X is a LCTVS with topology induced by a family of seminorms $\{p_i : i \in I\}$, then a sequence $(x_n)_{n\in\mathbf{N}}$ in X is a Cauchy sequence if and only if, for every $i \in I$ and for every $\varepsilon > 0$, there exists $n_0 \in \mathbf{N}$ such that $p_i(x_n - x_m) < \varepsilon$ for every $n, m \geq n_0$. The same property can be schematically stated as follows: for every $i \in I$,

$$\lim_{n,m\to\infty} p_i(x_n - x_m) = 0. \tag{2.6}$$

The proof relating Cauchy sequences to (2.6) is similar to Exercise 3.1 and is left to the reader.

Theorem 2.14 (Fundamental properties of τ)**.** *The following statements hold.*

(i) *A convex and balanced set $U \subset \mathcal{D}(\Omega)$ is open if and only if $U \in \mathcal{N}(\Omega)$.*
(ii) *For every compact $K \subset \Omega$, $\mathcal{D}_K$ is a closed linear subspace of $\mathcal{D}(\Omega)$.*
(iii) *Given any compact set $K \subset \Omega$, the topology τ_K of $\mathcal{D}_K$ coincides with the subspace topology inherited from $\mathcal{D}(\Omega)$.*
(iv) *If $E \subset \mathcal{D}(\Omega)$ is bounded, there exists $K \subset \Omega$ compact such that $E \subset \mathcal{D}_K$ and for every $N \in \mathbf{N}$ there exists $M_N > 0$ such that for every $\varphi \in E$ and every $\alpha \in \mathbf{N}_0^d$ with $|\alpha| \leq N$,*

$$\sup_{x\in K} |D^\alpha \varphi(x)| \leq M_N. \tag{2.7}$$

(v) *The class $\mathcal{D}(\Omega)$ has the Heine–Borel property from Definition 2.5.*
(vi) *If $(\varphi_n)_{n\in\mathbf{N}}$ is a Cauchy sequence, then there exists $K \subset \Omega$ compact such that $\varphi_n \in \mathcal{D}_K$ for every $n \in \mathbf{N}$ and for every $\alpha \in \mathbf{N}_0^d$,*

$$\lim_{n,m\to+\infty} \sup_{x\in K} |D^\alpha \varphi_n(x) - D^\alpha \varphi_m(x)| = 0.$$

(vii) *A sequence $(\varphi_n)_{n\in\mathbf{N}}$ in $\mathcal{D}(\Omega)$ converges to $\varphi \in \mathcal{D}(\Omega)$ if and only if there is $K \subset \Omega$ compact such that $\operatorname{spt} \varphi_n \subset K$ for every $n \in \mathbf{N}$ and $D^\alpha \varphi_n \to D^\alpha \varphi$ uniformly in K as $n \to \infty$ for every $\alpha \in \mathbf{N}_0^d$.*
(viii) *The space $\mathcal{D}(\Omega)$ is sequentially complete. That is, every Cauchy sequence in $\mathcal{D}(\Omega)$ converges in that space.*

The main "thumb rule" takeaways of this theorem are the following.

- Convergence effectively takes place in compact subsets of Ω.
- Cauchy sequences in the space $\mathcal{D}(\Omega)$ are characterized by all derivatives of the function sequence in question being a Cauchy sequence with respect to uniform convergence on the above compact subset.
- Convergence in $\mathcal{D}(\Omega)$ is characterized by uniform convergence of all derivatives of the function sequence in question on the above compact subset.

Proof of Theorem 2.14. (i) Given any $U \in \tau$, we claim $U \cap \mathcal{D}_K \in \tau_K$ for every compact $K \subset \Omega$. Let $\varphi \in U \cap \mathcal{D}_K$. By Proposition 2.11, there exists $V \in \mathcal{N}(\Omega)$ such that $\varphi + V \subset U$. Since $\varphi \in \mathcal{D}_K$,

$$\varphi + (V \cap \mathcal{D}_K) = (\varphi + V) \cap \mathcal{D}_K \subset U \cap \mathcal{D}_K.$$

Since $\varphi + (V \cap \mathcal{D}_K)$ is a neighborhood of φ in τ_K, it follows $\mathcal{D}_K \cap U \in \tau_K$. If in addition U is balanced and convex, it follows from Definition 2.10 that $U \in \mathcal{N}(\Omega)$.

The converse implication is trivial since $\mathcal{N}(\Omega) \subset \tau$.

(ii) Fix $K \subset \Omega$ compact. It is clear that $\mathcal{D}_K$ is a linear subspace. Thus, we only need to check it is closed. Note that $\varphi \in \mathcal{D}(\Omega)$ belongs to $\mathcal{D}_K$ if and only if $\varphi(x) = 0$ for all $x \in \Omega \setminus K$. In other words, setting $Z_x := \{\varphi \in \mathcal{D}(\Omega) : \varphi(x) = 0\}$, we have

$$\mathcal{D}_K = \bigcap_{x\in\Omega\setminus K} Z_x;$$

in order to show $\mathcal{D}_K$ is closed, it suffices to show the subsets Z_x are closed for every $x \in \Omega$. Fix $x \in \Omega$ and let $\varphi \notin Z_x$, which means $|\varphi(x)| > 0$. Define

$$W_\varphi := \{\tilde{\varphi} \in \mathcal{D}(\Omega) : 2\|\tilde{\varphi} - \varphi\|_\infty < |\varphi(x)|\}.$$

Going through similar arguments to those developed for W in (2.3) and (2.4), one checks $W_\varphi - \varphi \in \mathcal{N}(\Omega)$, so that $W_\varphi \in \tau$. Moreover, by the triangle inequality, given any $\tilde{\varphi} \in W_\varphi$ we have $|\tilde{\varphi}(x)| \geq |\varphi(x)| - |\varphi(x)|/2 > 0$, which means $W_\varphi \subset Z_x^c$. As the argument holds for any $\varphi \in Z_x^c$, we deduce Z_x^c is open in τ, thus Z_x is closed.

(iii) The proof of the first part of (i) shows the subspace topology $\mathcal{D}_K$ inherits from $\mathcal{D}(\Omega)$ is contained in τ_K.

Hence, it suffices to show the converse inclusion. Namely, given any $E \in \tau_K$, we need to show there exists $U \in \tau$ with $E = \mathcal{D}_K \cap U$. By Remark 2.9, given any $\varphi \in E$ we can find $N_\varphi \in \mathbf{N}$ and $\delta_\varphi > 0$ such that

$$\big\{\psi \in \mathcal{D}_K : \sup_{x \in K} |D^\alpha \psi(x) - D^\alpha \varphi(x)| < \delta_\varphi \text{ whenever } |\alpha| \leq N_\varphi\big\} \subset E.$$

Now we define

$$U_\varphi := \big\{\psi \in \mathcal{D}(\Omega) : \sup_{x \in \Omega} |D^\alpha \psi(x)| < \delta_\varphi \text{ whenever } |\alpha| \leq N_\varphi\big\}.$$

Going through the same argument we developed for W in (2.3) and (2.4), one checks $U_\varphi \in \mathcal{N}(\Omega)$, therefore $\varphi + U_\varphi \in \tau$. Hence $U := \bigcup_{\varphi \in E}(\varphi + U_\varphi) \in \tau$ and by construction $E = \mathcal{D}_K \cap U$.

(iv) Let $E \subset \mathcal{D}(\Omega)$ be bounded and assume by contradiction $E \setminus \mathcal{D}_K \neq \emptyset$ for all compact sets $K \subset \Omega$. Then there exists a sequence $(x_n)_{n \in \mathbf{N}}$ in Ω that has no accumulation point in Ω and $\varphi_n \in \mathcal{D}(\Omega) \cap E$ for every $n \in \mathbf{N}$ such that $\varphi_n(x_n) \neq 0$ for every $n \in \mathbf{N}$. Now we define

$$W := \{\varphi \in \mathcal{D}(\Omega) : n\,|\varphi(x_n)| < |\varphi_n(x_n)| \text{ for every } n \in \mathbf{N}\}.$$

Given any compact set $K \subset \Omega$, the number of $n \in \mathbf{N}$ such that $x_n \in K$ is finite. Since the evaluation mapping φ to $\varphi(x_n)$ is continuous on $\mathcal{D}_K$, it follows $W \cap \mathcal{D}_K \in \tau_K$. Since W is also balanced and convex, we deduce $W \in \mathcal{N}(\Omega)$. Noting $\varphi_n \notin sW$ for all $s \leq n$, we deduce there exists no finite $s > 0$ such that $E \subset sW$, which yields a contradiction.

Therefore, if E is bounded, there exists a compact set $K \subset \Omega$ such that $E \subset \mathcal{D}_K$. By (iii), the set E is also bounded in $\mathcal{D}_K$. Property (2.7) then follows from the topology τ_K being induced by the seminorms (2.1) and Exercise 3.1.

(v) By Remark 2.9, the space $\mathcal{D}_K$ has the Heine–Borel property. Hence the claim follows from properties (iii) and (iv). Indeed, by (iv) we know if $E \subset \mathcal{D}(\Omega)$ is bounded and closed, then $E \subset \mathcal{D}_K$ for some compact set $K \subset \Omega$. Then (iii) implies E is also bounded and closed in $\mathcal{D}_K$ and therefore compact in $\mathcal{D}_K$ by Remark 2.9, which together with (iii) implies compactness in $\mathcal{D}(\Omega)$.

(vi) Let $(\varphi_n)_{n \in \mathbf{N}}$ be a Cauchy sequence in $\mathcal{D}(\Omega)$. By Exercise 3.1, the sequence constitutes a bounded set in $\mathcal{D}(\Omega)$, hence by (iv) there exists $K \subset \Omega$ compact such that $\varphi_n \in \mathcal{D}_K$ for every $n \in \mathbf{N}$. By (iii), $(\varphi_n)_{n \in \mathbf{N}}$ is also a Cauchy sequence in $\mathcal{D}_K$, whose topology is induced by the seminorms from (2.1) by Remark 2.9. The conclusion then follows from (2.6).

(vii) If $\varphi_n \to \varphi$ in $\mathcal{D}(\Omega)$ as $n \to \infty$, then $(\varphi_n)_{n \in \mathbf{N}}$ is a Cauchy sequence in $\mathcal{D}(\Omega)$, which implies the conclusion by (vi).

For the converse implication, observe by assumption $\varphi \in \mathcal{D}_K$, and by Remark 2.9 — e.g. using the seminorms (2.1) — we have $\varphi_n \to \varphi$ with respect to τ_K as $n \to \infty$. By (iii), we have $\varphi_n \to \varphi$ with respect to τ as $n \to \infty$.

(viii) This is now a consequence of (vi), completeness of $\mathcal{D}_K$, and (vii). □

Lecture 6.

Corollary 2.15 (Nonmetrizability)**.** *The space $\mathcal{D}(\Omega)$ is not metrizable.*

Proof. Exercise 6.3. □

In the next subsection we will study the dual space of $\mathcal{D}(\Omega)$. To this aim, we formulate a very useful characterization of continuous linear maps defined on the space of test functions.

Proposition 2.16 (Characterizations of continuity)**.** *Let Y be an LCTVS and $T\colon \mathcal{D}(\Omega) \to Y$ be linear. Then the following properties are equivalent.*

(i) *T is continuous.*
(ii) *If $(\varphi_n)_{n\in\mathbf{N}}$ converges to 0 in $\mathcal{D}(\Omega)$, then $(T\varphi_n)_{n\in\mathbf{N}}$ converges to 0 in Y*[13].
(iii) *The restriction of T to $\mathcal{D}_K$ is continuous for every compact $K \subset \Omega$.*

Proof. (i) $\Longrightarrow$ (ii). This is trivial.

(ii) $\Longrightarrow$ (iii). Since $\mathcal{D}_K$ is metrizable and T is linear, sequential continuity in 0 with respect to τ_K is equivalent to continuity of T. Now let $(\varphi_n)_{n\in\mathbf{N}}$ be a sequence in $\mathcal{D}_K$ which converges to zero in $\mathcal{D}_K$. By Theorem 2.14, this convergence also happens in $\mathcal{D}(\Omega)$, which by (ii) and the above implies the continuity of T.

(iii) $\Longrightarrow$ (i). Let $V \subset Y$ be a balanced and convex neighborhood of 0 and set $U := T^{-1}(V)$. Since T is linear, the set U is also balanced and convex. By (iii), the set $U \cap \mathcal{D}_K$ is open in τ_K for every compact set $K \subset \Omega$, which implies U is open in $\mathcal{D}(\Omega)$ by Definition 2.10. This proves the continuity of T in 0; the linearity of T implies its continuity by Proposition 1.32. □

Corollary 2.17 (Some continuous maps on $\mathcal{D}(\Omega)$)**.** *The following maps are continuous from $\mathcal{D}(\Omega)$ to $\mathcal{D}(\Omega)$.*

(i) *The derivative assignment $\varphi \mapsto D^\alpha\varphi$, where $\alpha \in \mathbf{N}_0^d$ is given.*
(ii) *The multiplication assignment $\varphi \mapsto \psi\,\varphi$, where $\psi \in C^\infty(\Omega)$ is given.*
(iii) *For $\Omega = \mathbf{R}^d$ the affine transformation assignment $\varphi \mapsto \varphi(\lambda\cdot - z)$, where $\lambda \in \mathbf{R}\setminus\{0\}$ and $z \in \mathbf{R}^d$ are fixed.*

Proof. Exercise 6.2. □

2.2. **Distributions from a topological point of view.** Next we study the dual space of $\mathcal{D}(\Omega)$. Observe that since the convergence on $\mathcal{D}(\Omega)$ is quite strong, one might expect many linear functionals to be continuous with respect to the topology of $\mathcal{D}(\Omega)$, so that its topological dual should be quite large.

Definition 2.18 (Distributions)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and let $\mathcal{D}(\Omega)$ be the space of test functions introduced in Definition 2.2. We denote by $\mathcal{D}'(\Omega)$ the topological dual space of $\mathcal{D}(\Omega)$.*

Elements of $\mathcal{D}'(\Omega)$ are called **distributions***.*

Example 2.19 (Concrete distributions). The following are examples of distributions $T \in \mathcal{D}'(\mathbf{R}^d)$.

- **Dirac delta distribution.** The assignment $T\varphi := \varphi(0)$. This distribution is often denoted by δ_0. Analogously δ_{x_0} for the Dirac delta distribution at an arbitrary point $x_0 \in \mathbf{R}^d$.
- **Evaluation of derivatives.** The assignment $T\varphi := D^\alpha\varphi(x_0)$ for some fixed multiindex $\alpha \in \mathbf{N}_0^d$ and some fixed $x_0 \in \mathbf{R}^d$. In short, $T = \delta_{x_0} \circ D^\alpha$.

[13] By Corollary 2.15, the space $\mathcal{D}(\Omega)$ is not metrizable, so a priori continuity is not equivalent to sequential continuity. Nevertheless, Proposition 2.16 guarantees that for linear functionals on $\mathcal{D}(\Omega)$, this is actually the case!

- **Integration.** The assignment $T_f(\varphi) := \int_{\mathbf{R}^d} f\,\varphi\,\mathrm{d}\mathcal{L}^d$, where $f \in L^1_{\mathrm{loc}}(\mathbf{R}^d, \mathcal{L}^d)$; the latter means f is integrable on every compact subset $K \subset \mathbf{R}^d$;
- **Integration of derivatives.** The assignment $T\varphi := \int_{\mathbf{R}^d} f\,D^\alpha\varphi\,\mathrm{d}\mathcal{L}^d$ for some fixed multiindex $\alpha \in \mathbf{N}_0^d$ and f as above.
- **Borel measures.** The assignment $T\varphi := \int_{\mathbf{R}^d} \varphi\,\mathrm{d}\mu$, where μ is a locally finite Borel measure[14] on $\mathbf{R}^d$.

All these examples are special cases of Exercise 6.2, so we omit the proofs. ■

Remark 2.20 (Towards fundamental lemmata of distribution theory). In the case $f \in L^1_{\mathrm{loc}}(\Omega, \mathcal{L}^d)$, the distribution T_f defined in the previous Example 2.19 uniquely determines f $\mathcal{L}^d$-a.e[15]. In this manner, we can regard $L^1_{\mathrm{loc}}(\Omega, \mathcal{L}^d)$ as a "subset" of distributions on Ω. One often informally says that a distribution T is (represented by) a function if there exists $f \in L^1_{\mathrm{loc}}(\Omega, \mathcal{L}^d)$ such that $T\varphi = \int_\Omega f\,\varphi\,\mathrm{d}\mathcal{L}^d$ for every $\varphi \in \mathcal{D}(\Omega)$ and one tacitly identifies T with f.

Similarly one says a distribution is (represented by) a Radon measure in the sense of Example 2.19 above. There are quite general representation results for distributions by measures which do not even require to work with smooth functions, cf. e.g. the **Riesz–Markov–Kakutani representation theorem**. We do not state it rigorously here, but informally it says a distribution which remains nonnegative when evaluated at a nonnegative function is necessarily given by a measure. ■

We have the following characterization of distributions.

Lemma 2.21 (Distributions by seminorms)**.** *Let $\Omega \subset \mathbf{R}^d$ open. Given any $N \in \mathbf{N}$ and $K \subset \Omega$ compact, define a seminorm $p_{N,K}$ on $\mathcal{D}(\Omega)$ by*

$$p_{N,K}(\varphi) := \max\{|D^\alpha\varphi(x)| : \alpha \in \mathbf{N}_0^d,\ |\alpha| \le N,\ x \in K\}.$$

Let $T\colon \mathcal{D}(\Omega) \to \mathbf{R}$ be linear. Then $T \in \mathcal{D}'(\Omega)$ if and only if for every compact set $K \subset \Omega$ there exist $C_K > 0$ and $N_K \in \mathbf{N}_0$ such that for every $\varphi \in \mathcal{D}_K$,

$$|T\varphi| \le C_K\, p_{N_K,K}(\varphi). \tag{2.8}$$

Proof. By Proposition 2.16 continuity of T is equivalent to continuity of the restriction $T\big|_{\mathcal{D}_K}$ for every compact set $K \subset \Omega$. Since the seminorms $p_{N,K}$ are increasing in N, the statement follows from Proposition 1.32 and the fact that the given family $\{p_{N,K} : N \in \mathbf{N}_0\}$ generates the topology τ_K, cf. Remark 2.9. □

Definition 2.22 (Order of distributions)**.** *If there exists a common number $N \in \mathbf{N}_0$ such that inequality (2.8) holds true for all compact sets $K \subset \Omega$ — but possibly with a varying constant C_K —, we say the distribution T has* **finite order***.*

In this case, its **order** *is the smallest number $\bar{N} \in \mathbf{N}_0$ with this property.*

Otherwise, the distribution is said to have **infinite order***.*

Remark 2.23 (About Example 2.19). All distributions appearing in Example 2.19 have finite order (try to compute it).

As specified in the exercise sheets, an example in $\mathcal{D}'(\mathbf{R})$ of infinite order is the following assignment:

$$T\varphi := \sum_{n\in\mathbf{N}} D^n\varphi(n).$$

[14]Recall a Borel measure on a topological space is a measure defined on the Borel σ-algebra of that topological space, i.e. the smallest σ-algebra containing all open sets. Such a measure is called locally finite if it is finite on each compact set.

[15]This is actually a quite important result, usually referred to as the fundamental lemma of the calculus of variations, cf. Lemma 2.28 later. This insight is also a basis for the definition of **Sobolev spaces** in PDE theory.

2.3. **A short introduction to distributional calculus.** On distributions, one can define many operations by duality, that is, by "moving the operation on the argument $\varphi \in \mathcal{D}(\Omega)$". This is a commonly used method that in PDE theory, one usually sees in action when "sufficiently many derivatives are moved to the test function in question in an integration by parts formula".

In functional analytic terms, whenever we have a continuous linear functional $A\colon \mathcal{D}(\Omega) \to \mathcal{D}(\Omega)$, we may define its **adjoint**[16] $A^* : \mathcal{D}'(\Omega) \to \mathcal{D}'(\Omega)$ by the following formula for every $\varphi \in \mathcal{D}(\Omega)$:

$$(A^*T)(\varphi) := T(A\varphi).$$

One can then verify that, given any $T \in \mathcal{D}'(\Omega)$, A^*T is again an element of $\mathcal{D}'(\Omega)$ and that A^* is continuous with respect to the weak* topology of $\mathcal{D}'(\Omega)$.

In the next definition, we present three important examples of this procedure.

Definition 2.24 (Differential calculus on distributions)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and let $T \in \mathcal{D}'(\Omega)$ be a distribution.*

a. *Given any $\alpha \in \mathbf{N}_0^d$, the* **partial derivative** *D^α of T is defined by*
$$(D^\alpha T)(\varphi) := (-1)^{|\alpha|}\, T(D^\alpha \varphi).$$
b. *Given any $\psi \in C^\infty(\Omega)$, the* **product** *of T with ψ is defined by*
$$(\psi\, T)(\varphi) := T(\psi\, \varphi).$$
c. *If $\Omega = \mathbf{R}^d$ and $\psi \in C_c^\infty(\mathbf{R}^d)$, the* **convolution** *of ψ with T is a function defined by the assignment*
$$(\psi * T)(x) := T(\psi(x - \cdot))$$

Remark 2.25 (About Definition 2.24). By Corollary 2.17, for a distribution T the partial derivatives $D^\alpha T$ and the product ψT with a smooth function are again distributions. Instead, $\psi * T$ is a C^∞-function on $\mathbf{R}^d$.

The idea behind these definitions is they coincide with the classical definitions in the situation when T is replaced by T_f from Example 2.19. For instance, when $f \in C^m(\mathbf{R}^d)$ with $m \in \mathbf{N}$, then for any multiindex $\alpha \in \mathbf{N}_0^d$ such that $|\alpha| \leq m$, by integration by parts[17] we have

$$(D^\alpha T_f)(\varphi) = (-1)^{|\alpha|} \int_{\mathbf{R}^d} f\, D^\alpha \varphi \,\mathrm{d}\mathcal{L}^d = \int_{\mathbf{R}^d} D^\alpha f\, \varphi \,\mathrm{d}\mathcal{L}^d = T_{D^\alpha f}(\varphi).$$

Similarly, we obtain $\psi\, T_f = T_{\psi f}$ and $\psi * T_f = \psi * f$.

Studying PDEs often leads to considering distributional derivatives of integrable functions (e.g. by considering Sobolev spaces, cf. Remark 2.20). This means exactly the quantity $D^\alpha T_f$, which makes sense even when the function f in question is not classically differentiable. For instance, consider the **Heaviside function** $f = 1_{\mathbf{R}_+}$. Then for any $\varphi \in \mathcal{D}(\mathbf{R})$,

$$T_f'(\varphi) = -\int_{\mathbf{R}} f\, \varphi' \,\mathrm{d}\mathcal{L}^1 = -\int_0^\infty \varphi' \,\mathrm{d}\mathcal{L}^1 = \varphi(0) = \delta_0(\varphi),$$

where we used φ has compact support. This means the distributional derivative of f is the Dirac delta centered in 0; this is often stated as "$f' = \delta_0$ in the sense of distributions". Observe that f is differentiable at every $x \in \mathbf{R} \setminus \{0\}$ with $f'(x) = 0$, so that in particular $f' = 0$ $\mathcal{L}^d$-a.e. However, its distributional derivative is not the function identically equal to 0![18] Some care must taken when trying to identify

[16] Think of the pairing $T\varphi$ of a linear map $T \in \mathcal{D}'(\Omega)$ and $\varphi \in \mathcal{D}(\Omega)$ as a "scalar product" $\langle T, \varphi\rangle$; in this notation, the adjoint of A satisfies $\langle T, A\varphi\rangle = \langle A^*T, \varphi\rangle$.

[17] No boundary terms appear since φ has compact support.

[18] Yet, intuitively this recovers the "correct" derivative of the Heaviside function f. Indeed, while $f' = 0$ outside zero, f makes a jump of height 1 at zero. In comparison, note $T_{2f}' = 2\,\delta_0$.

derivatives defined a.e. point with distributional derivatives, in general they might not coincide.

To be more precise, in the above example the distributional derivative of the Heaviside step function is a *measure* with Lebesgue decomposition $0 \cdot \mathcal{L}^1 + \delta_0$[19]. Informally, this explains what is going on here: even when a function has vanishing derivative $\mathcal{L}^1$-a.e. (which would contribute to the $\mathcal{L}^1$-absolutely continuous part of the distributional derivative), its distributional derivative may still have an $\mathcal{L}^1$-singular part. As a disclaimer, we note that not all distributional derivatives are measures, hence the above analogy should be taken with care; the class of functions for which this is the case (in an appropriate sense) are termed to have **bounded variation**. ■

We can also speak about the convergence of distributions.

Definition 2.26 (Convergence)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and $(T_n)_{n\in\mathbf{N}}$ be a sequence in $\mathcal{D}'(\Omega)$. We say it* **converges** *to $T \in \mathcal{D}'(\Omega)$ if $T_n\varphi \to T\varphi$ for every $\varphi \in \mathcal{D}(\Omega)$.*

This is exactly the convergence in the weak*-topology, cf. Definition 1.40. We interchangeably use the notations $T_n \rightharpoonup^* T$ as $n \to \infty$ or $T_n \to T$ as $n \to \infty$.

The following result shows why calculus with distributions is often very simple.

Lemma 2.27 (Convergence vs. partial derivatives)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and let $(T_n)_{n\in\mathbf{N}}$ be a sequence in $\mathcal{D}'(\Omega)$ converging to $T \in \mathcal{D}'(\Omega)$. Then for all $\alpha \in \mathbf{N}_0^d$, the sequence $(D^\alpha T_n)_{n\in\mathbf{N}}$ converges to $D^\alpha T$. In other words, partial derivatives respect convergence in $\mathcal{D}'(\Omega)$.*

Proof. Fix $\varphi \in \mathcal{D}(\Omega)$. Then $D^\alpha\varphi \in \mathcal{D}(\Omega)$ as well and therefore the convergence of T_n implies

$$\lim_{n\to\infty}(D^\alpha T_n)(\varphi) = (-1)^{|\alpha|}\lim_{n\to\infty} T_n(D^\alpha\varphi) = (-1)^{|\alpha|}\,T(D^\alpha\varphi) = (D^\alpha T)(\varphi).$$

This terminates the proof. □

Using convolutions, it can be shown for every $T \in D'(\Omega)$ there exists a sequence of functions $(f_n)_{n\in\mathbf{N}}$ in $\mathcal{D}(\Omega)$ such that $T_{f_n} \to T$ as $n \to \infty$. Roughly speaking, this means $\mathcal{D}(\Omega)$ is sequentially dense in $\mathcal{D}'(\Omega)$, with respect to the weak* topology, up to identifying $f \in \mathcal{D}(\Omega)$ with $T_f \in \mathcal{D}'(\Omega)$ with a slight abuse of notation.

Lecture 7. In Example 2.19 we have introduced examples of distributions $T \in \mathcal{D}'(\mathbf{R}^d)$. The example in its third bullet point is particularly interesting since, as written in Remark 2.20, the distribution T_f associated to an $L^1_{\mathrm{loc}}(\Omega, \mathcal{L}^d)$ function uniquely determines f $\mathcal{L}^d$-a.e. This is a highly nontrivial statement, implying the injectivity of $T_\cdot$ on $L^1_{\mathrm{loc}}(\Omega, \mathcal{L}^d)$. The statement is the following.

Lemma 2.28 (Fundamental lemma of calculus of variations)**.** *Let $\Omega \subset \mathbf{R}^d$ be open and let $f \in L^1_{\mathrm{loc}}(\Omega, \mathcal{L}^d)$. If every $\varphi \in \mathcal{D}(\Omega)$ satisfies*

$$\int_\Omega f\,\varphi\,\mathrm{d}\mathcal{L}^d = 0,$$

then $f = 0$ $\mathcal{L}^d$-a.e.

For an application of this lemma that justifies its name, see Exercise 7.4.

The result is trivial if $f \in \mathcal{D}(\Omega)$, since it would be enough to take $\varphi = f$ to conclude the proof. The key point in the naive ansatz $\varphi = f$ is that we are choosing a test function that is able to 'detect' the sign of f. Looking at the sign of f is a much less restrictive property which does not require any smoothness. We only

[19]Here, δ_0 denotes the Dirac mass at zero.

need to sample the values of f on some set. This is the key idea on which we will build the proof by contradiction.

In the proof below, we assume the reader knows the inner and outer regularity properties of the Lebesgue measure.

Proof of Lemma 2.28. We first explain the proof strategy. If the conclusion of the lemma fails, then $f \neq 0$ on a set of positive measure. Hence, there are positive measure sets where $f > 0$ or $f < 0$. We then need to construct suitable test functions supported on either set depending on the situation. However, this is nontrivial because the positivity and negativity sets can be complicated. For instance, they could be a Cantor set; per se, they neither need to be closed nor open.

We divide the proof into the following six steps.

1. For every ball $B_r(x_0) \subset \mathbf{R}^d$ there exists a cutoff function $g \in C_c(\mathbf{R}^d)$ with values in $[0,1]$ such that $g = 1$ on $B_r(x_0)$ and $\operatorname{spt} g \subset B_{2r}(x_0)$.
2. Given any compact set $K \subset \mathbf{R}^d$ and any open set $U \subset \mathbf{R}^d$ containing K, there exists $g \in C_c(\mathbf{R}^d)$ which takes values in $[0,1]$ such that $g = 1$ on K and $\operatorname{spt} g \subset U$.
3. Let $A \subset \mathbf{R}^d$ be a measurable set with finite measure. Then there exists a sequence $(g_n)_{n\in\mathbf{N}}$ in $C_c(\mathbf{R}^d)$ with values in $[0,1]$ such that $g_n \to 1_A$ in $L^1(\mathbf{R}^d)$ as $n \to \infty$. Moreover, if $A \subset V$ for some open set $V \neq \mathbf{R}^d$, the above sequence can be constructed to satisfy $\operatorname{spt} g_n \subset V$ for every $n \in \mathbf{N}$.
4. Upgrade Step 3 from continuous to smooth functions, i.e. ensure $g_n \in \mathcal{D}(V)$ for every $n \in \mathbf{N}$.
5. Take A to be an appropriate subset of $\{f > 0\}$ or $\{f < 0\}$ to achieve the desired contradiction.

Step 1. We define $g(x) := \max\{1 - 2\,\mathsf{d}_2(x, B_r(x_0))/r, 0\}$, where $x \in \mathbf{R}^d$. Here $\mathsf{d}_2(\cdot, B_r(x_0))$ denotes the customary distance function to $B_r(x_0)$ induced by the Euclidean norm on $\mathbf{R}^d$.

Step 2. Set $r := \mathsf{d}_2(K, \partial U)/3$. By compactness of K and since $\partial U \cap K = \emptyset$, we have $r > 0$. Cover K by finitely many balls $B_r(x_1), \dots, B_r(x_m)$ with $x_1, \dots, x_m \in K$. Then $B_{2r}(x_i) \subset U$ for all $i \in \{1, \dots, m\}$ by our choice of r. Let $g_1, \dots, g_m$ be as constructed in Step 1 and set $\widetilde{g} = g_1 + \dots + g_m$ as well as $g = \min\{1, \widetilde{g}\}$. Then g satisfies all the claimed properties thanks to the inclusions $\operatorname{spt} g_i \subset B_{2r}(x_i) \subset U$ for every $i \in \{1, \dots, m\}$.

Step 3. Since the Lebesgue measure is regular[20] and A has finite measure, for every $n \in \mathbf{N}$ there exists a compact set $K_n \subset A$ and an open set $U_n \supset A$ such that $\mathcal{L}^d[U_n \setminus K_n] \leq 1/n$. Applying Step 2 we find a sequence $(g_n)_{n\in\mathbf{N}}$ in $C_c(\mathbf{R}^d)$ with values in $[0,1]$ such that $g_n = 1$ on K_n and $g_n = 0$ outside U_n. In particular, $g_n = 1_A$ on $K_n \cup (\mathbf{R}^d \setminus U_n)$ and $|g_n - 1_A| \leq 1$ on $U_n \setminus K_n$. Hence

$$\int_{\mathbf{R}^d} |g_n - 1_A| \,\mathrm{d}\mathcal{L}^d \leq \mathcal{L}^d[U_n \setminus K_n] \leq \frac{1}{n}.$$

Therefore, $g_n \to 1_A$ in $L^1(\mathbf{R}^d, \mathcal{L}^d)$ as $n \to \infty$. Finally, if $A \subset V$ for some open set V, we can always replace the set U_n above by the open set $U_n \cap V$ and hence the claim follows from the properties of the sequence $(g_n)_{n\in\mathbf{N}}$ from Step 2.

Step 4. Let $A \subset \mathbf{R}^d$ be a set of finite measure and $V \subset \mathbf{R}^d$ be an open set containing A. We will regularize the sequence $(g_n)_{n\in\mathbf{N}}$ in $C_c(V)$ found in Step 3 by

[20] This means that $\mathcal{L}^d$ is simultaneously

- inner regular, i.e. $\mathcal{L}^d[A] = \sup\{\mathcal{L}^d[K] : K \subset A \text{ compact}\}$ for every measurable $A \subset \mathcal{L}^d$,
- outer regular, i.e. $\mathcal{L}^d[A] = \inf\{\mathcal{L}^d[U] : U \supset A \text{ open}\}$ for every measurable $A \subset \mathcal{L}^d$.

convolution, a standard method to regularize functions. Define $\eta \in \mathcal{D}(\mathbf{R}^d)$ by

$$\eta(x) = \begin{cases} C\,\mathrm{e}^{1/(\|x\|^2-1)} & \text{if } \|x\| < 1, \\ 0 & \text{otherwise,} \end{cases}$$

where $C > 0$ is chosen such that η integrates up to one. Furthermore, given any $\varepsilon > 0$ define $\eta_\varepsilon \in \mathcal{D}(\mathbf{R}^d)$ by $\eta_\varepsilon(x) := \eta(x/\varepsilon)/\varepsilon^d$. Observe $\operatorname{spt}\eta_\varepsilon \subset \overline{B}_\varepsilon(0)$ as well as $\int_{\mathbf{R}^d} \eta_\varepsilon \,\mathrm{d}\mathcal{L}^d = 1$. These functions are often referred to as standard mollifiers. We then define the convolution

$$g_n \star \eta_\varepsilon(x) := \int_{\mathbf{R}^d} g_n(x-y)\,\eta_\varepsilon(y)\;\mathrm{d}y = \int_{B_\varepsilon(0)} g_n(x-y)\,\eta_\varepsilon(y)\;\mathrm{d}y.$$

It is a classical result from analysis that $g_n \star \eta_\varepsilon \in C^\infty(\mathbf{R}^d)$, cf. e.g. [6, §C] for details. Moreover, by general properties of the convolution we have

$$\operatorname{spt} g_n \star \eta_\varepsilon \subset \operatorname{spt} g_n + \operatorname{spt}\eta_\varepsilon \subset \operatorname{spt} g_n + \overline{B}_\varepsilon(0). \tag{2.9}$$

Since g_n has compact support, it follows that $g_n \star \eta_\varepsilon \in \mathcal{D}(\mathbf{R}^d)$. Moreover, for ε small enough (depending on n) we can ensure $\operatorname{spt} g_n + \overline{B}_\varepsilon(0) \subset V$, so that $g_n \star \eta_\varepsilon \in \mathcal{D}(V)$. With n fixed, the uniform continuity of g_n and the fact that the standard mollifiers integrate up to one imply $g_n \star \eta_\varepsilon \to g_n$ uniformly on $\mathbf{R}^d$ and in $L^1(\mathbf{R}^d, \mathcal{L}^d)$ as $\varepsilon \to 0$ due to (2.9). Choose $\varepsilon_n > 0$ such that $\|g_n \star \eta_{\varepsilon_n} - g_n\|_{L^1(\mathbf{R}^d,\mathcal{L}^d)} \le 1/n$. Then by the triangle inequality $g_n \star \eta_{\varepsilon_n} \to 1_A$ in $L^1(\mathbf{R}^d, \mathcal{L}^d)$ as $n \to \infty$. Finally, $0 \le g_n \star \eta_\varepsilon \le 1$ since η_ε is non-negative and has integral one, and g_n takes values in $[0,1]$.

Step 5. We are in a position to prove the statement of the lemma. Assume to the contrary that $f \neq 0$ on a set of positive $\mathcal{L}^d$-measure. Without loss of generality (up to switching the sign of f) we may and will assume $\mathcal{L}^d[\{f > 0\}] > 0$. Using an exhaustion of Ω by compact subsets and Levi's monotone convergence theorem, we find $A \subset \Omega$ $\mathcal{L}^d$-measurable such that $\mathcal{L}^d[A] \in (0,\infty)$, $f > 0$ on A, and A has compact closure in Ω. Again by monotone convergence and possibly shrinking A, we may and will assume there exists $\delta > 0$ such that $f \ge \delta$ on A.

Since A has compact closure in Ω, there exists an open, bounded set V such that $\overline{V} \subset \Omega$ and $A \subset V$. By Step 4, there exists a sequence $(\varphi_n)_{n\in\mathbf{N}}$ in $\mathcal{D}(V)$ with values in $[0,1]$ such that $\varphi_n \to 1_A$ in $L^1(\Omega, \mathcal{L}^d)$ as $n \to \infty$. Choose a nonrelabeled subsequence such that $\varphi_n \to 1_A$ $\mathcal{L}^d$-a.e. as $n \to \infty$[21]. Then clearly $f\,\varphi_n \to f\,1_A$ $\mathcal{L}^d$-a.e. and $|f\,\varphi_n| \le 1_V\,|f|$ on $\mathbf{R}^d$, which is integrable since V is bounded. Hence by assumption and Lebesgue's dominated convergence theorem,

$$0 = \lim_{n\to\infty} \int_\Omega f\,\varphi_n \,\mathrm{d}\mathcal{L}^d = \int_{\mathbf{R}^d} f\,1_A \,\mathrm{d}\mathcal{L}^d \ge \delta\,\mathcal{L}^d[A],$$

which is a contradiction. □

2.4. **The Schwartz space and tempered distributions.** In our short introduction of distributions, we have seen some examples and basic elements of distributional calculus. As both distributions and the Fourier transform are important in the theory of PDEs, there is a final natural question we address: can we also define the **Fourier transform** acting on distributions?

Again our ansatz is by duality, extending the standard Fourier transform acting on integrable functions. As an example, let us try to define the Fourier transform of the distribution associated to T_f with $f \in L^1(\mathbf{R}^d, \mathcal{L}^d)$. Following the paradigm

[21] L^1-convergence implies subsequential a.e. convergence.

from Definition 2.24, the Fourier transform of T_f should be given by $T_{\mathcal{F}[f]}$, where $\mathcal{F}$ denotes the usual Fourier transform on $\mathbf{R}^d$. Using Fubini's theorem, we compute

$$\begin{aligned} T_{\mathcal{F}[f]}(\varphi) &= \int_{\mathbf{R}^d} \varphi(k)\,\mathcal{F}[f](k)\,\mathrm{d}k \\ &= \int_{\mathbf{R}^d} \varphi(k) \int_{\mathbf{R}^d} \mathrm{e}^{-\mathrm{i}k\cdot x}\, f(x)\,\mathrm{d}x\,\mathrm{d}k \\ &= \int_{\mathbf{R}^d} f(x)\,\mathcal{F}[\varphi](x)\,\mathrm{d}x \\ &\text{“=”}\ T_f(\mathcal{F}[\varphi]). \end{aligned}$$

Is the last term on the right-hand side a well defined object? Can we extend it to any distribution? The answer is no, since the Fourier transform of a nontrivial $\mathcal{D}(\mathbf{R}^d)$-function never has compact support (cf. Exercise 6.4), which suggests that pairing it with a distribution is not defined in general. To overcome this issue, one needs to slightly enlarge the space of test functions and therefore we introduce the **Schwartz space**[22].

Definition 2.29 (Schwartz space)**.** *The Schwartz space $\mathcal{S}(\mathbf{R}^d)$ is defined as*

$$\mathcal{S}(\mathbf{R}^d) := \{f \in C^\infty(\mathbf{R}^d;\mathbf{C}) : \sup_{x\in\mathbf{R}^d} |x^\alpha D^\beta f(x)| < \infty \text{ for every } \alpha,\beta \in \mathbf{N}_0^d\},$$

where $x^\alpha := x_1^{\alpha_1}\cdots x_d^{\alpha_d}$.

We endow this Schwartz space by the locally convex topology given by the countable family of seminorms $\{p_{\alpha,\beta} : \alpha,\beta \in \mathbf{N}_0^d\}$ given by

$$p_{\alpha,\beta}(f) := \sup_{x\in\mathbf{R}^d} |x^\alpha D^\beta f(x)|.$$

Clearly, the LCTVS $\mathcal{S}(\mathbf{R}^d)$ is metrizable, cf. Theorem 1.23.

Roughly speaking, Schwartz functions are smooth functions that decay rapidly at infinity, as quantified by the previous definition[23]. More precisely, by definition all derivatives of a Schwartz functions decay superpolynomially at infinity.

For instance, the function $f\colon \mathbf{R}^d \to \mathbf{R}$ given by $f(x) := \mathrm{e}^{-x^\top Ax}$, where $A \in \mathbf{R}^{d\times d}$ is symmetric and positive definite, belongs to $\mathcal{S}(\mathbf{R}^d)$ yet has noncompact support. In the next lemma we collect some elementary properties of $\mathcal{S}(\mathbf{R}^d)$ whose proof is left to the interested reader.

Lemma 2.30 (Involutions of Schwartz space)**.** *Let $f, g \in \mathcal{S}(\mathbf{R}^d)$ and $\alpha \in \mathbf{N}_0^d$. Then the following functions belong also to $\mathcal{S}(\mathbf{R}^d)$.*

(i) *The complex conjugation $x \mapsto \overline{f}(x)$.*
(ii) *The product $x \mapsto f(x)\,g(x)$.*
(iii) *The product with arbitrary monomials $x \mapsto x^\alpha f(x)$.*
(iv) *The derivative map $x \mapsto D^\alpha f(x)$.*

We also recall less obvious properties of the Schwartz space without proof.

Proposition 2.31 (Basic properties of Schwartz space)**.** *The following properties hold true.*

(i) *The space $\mathcal{D}(\mathbf{R}^d)$ is densely contained in $\mathcal{S}(\mathbf{R}^d)$.*

---

[22]The space is named after Laurent Schwartz, who pioneered the theory of distributions and his work was awarded with a Fields medal in the 1950. See his biography at https://mathshistory.st-andrews.ac.uk/Biographies/Schwartz/.

[23]Since our aim is to study the Fourier transform (and antitransform) of functions, it makes sense in the following to always take them $\mathbf{C}$-valued. In particular, $\overline{f}(x)$ will denote the complex conjugate of $f(x)$. Recall that all results from the previous lectures on $C^\infty(\Omega)$ seemlessy transfer to $C^\infty(\Omega;\mathbf{C})$ (e.g. by splitting f into its real and imaginary parts) and that by differentiability we always mean real differentiability, not complex differentiability.

(ii) *We have* $\mathcal{S}(\mathbf{R}^d) \subset L^p(\mathbf{R}^d, \mathcal{L}^d)$ *for every* $p \in [1,\infty]$.
(iii) $\mathcal{S}(\mathbf{R}^d)$ *is complete and has the Heine–Borel property.*
(iv) *The Fourier transform* $\mathcal{F}\colon \mathcal{S}(\mathbf{R}^d) \to \mathcal{S}(\mathbf{R}^d)$ *defined by*

$$\mathcal{F}[f](k) := \int_{\mathbf{R}^d} f(x)\, e^{-\mathrm{i}k\cdot x}\,\mathrm{d}x$$

is a linear homeomorphism.

Let us briefly explain the last property. As known from classical Fourier analysis, the decay of a function and its derivatives at infinity is related to the smoothness of its Fourier transform, while its smoothness is related to the decay of the Fourier transform at infinity. For this reason, the seminorms from Definition 2.29 encode both decay at infinity and smoothness and it is natural to expect that the Fourier transform inherits these properties. The mathematical proof of (iv) is merely a rigorous formulation of the above thoughts.

By the above proposition the Fourier transform maps $\mathcal{D}(\mathbf{R}^d)$ into $\mathcal{S}(\mathbf{R}^d)$. Hence we could define the Fourier transform on "distributions" that are defined on $\mathcal{S}(\mathbf{R}^d)$. Since $\mathcal{D}(\mathbf{R}^d) \subset \mathcal{S}(\mathbf{R}^d)$, this means those are special family of distributions.

Definition 2.32 (Tempered distribution)**.** *A tempered distribution is a linear functional* $T\colon \mathcal{S}(\mathbf{R}^d) \to \mathbf{C}$ *that is continuous with respect to the convergence introduced in Definition 2.29, symbolically* $T \in \mathcal{S}'(\mathbf{R}^d)$.

Note every tempered distribution is a "classical" distribution on $\mathbf{R}^d$ since the convergence on $\mathcal{D}(\mathbf{R}^d)$ implies the convergence in $\mathcal{S}(\mathbf{R}^d)$. Moreover, every L^p-function can be interpreted as tempered distribution in the sense of Example 2.19. However, not all cases in Example 2.19 define tempered distributions. For instance, the function $f \in L^1_{\mathrm{loc}}(\mathbf{R}, \mathcal{L}^1)$ given by $f(x) := \mathrm{e}^{x^2}$ grows so quickly such that its product with the Schwartz function $1/f$ has no finite integral on $\mathbf{R}$. With some small restrictions the operations on distributions also make sense for tempered distributions[24].

Theorem 2.33 (Fourier transform on tempered distributions)**.** *Given* $T \in \mathcal{S}'(\mathbf{R}^d)$*, we define its Fourier transform* $\widehat{T} \in \mathcal{S}'(\mathbf{R}^d)$ *by* $\widehat{T}(\varphi) = T(\mathcal{F}[\varphi])$*. This Fourier transform is a linear bijective map from* $\mathcal{S}'(\mathbf{R}^d)$ *to* $\mathcal{S}'(\mathbf{R}^d)$.

Proof. First we show $\widehat{T}$ is again a tempered distribution. First of all, we know $\mathcal{F}[\varphi] \in \mathcal{S}(\mathbf{R}^d)$ for every $\varphi \in \mathcal{S}(\mathbf{R}^d)$, meaning $\widehat{T}$ is well-defined. Linearity follows from linearity of the Fourier transform and linearity of T. Thus, it remains to show the Fourier transform defined above is continuous. By linearity, it suffices to prove if $(\varphi_n)_{n\in\mathbf{N}}$ is a sequence in $\mathcal{S}(\mathbf{R}^d)$ converging to zero, then $\widehat{T}(\varphi_n) \to 0$ as $n \to \infty$. By definition and continuity of T this reduces to prove $\mathcal{F}[\varphi_n] \to 0$ in $\mathcal{S}(\mathbf{R}^d)$ as $n \to \infty$. This is the last item of Proposition 2.31.

To conclude the proof, we need to show the Fourier transform is bijective. We first prove injectivity. By linearity it suffices to prove $\widehat{T} = 0$ implies $T = 0$. Let $\varphi \in \mathcal{S}(\mathbf{R}^d)$ be given. Then $\mathcal{F}^{-1}[\varphi] \in \mathcal{S}(\mathbf{R}^d)$ by the Fourier inversion theorem on the Schwartz space. This yields

$$0 = \widehat{T}(\mathcal{F}^{-1}[\varphi]) = T(\mathcal{F}[\mathcal{F}^{-1}[\varphi]]) = T(\varphi).$$

Since $\varphi \in \mathcal{S}(\mathbf{R}^d)$ was arbitrary, the above identity implies $T = 0$. To prove surjectivity, let $T \in \mathcal{S}'(\mathbf{R}^d)$ and define a tempered distribution $S \in \mathcal{S}'(\mathbf{R}^d)$ by the formula $S(\varphi) := T(\mathcal{F}^{-1}[\varphi])$. Again this is a tempered distribution since the inverse Fourier

[24]One has to be careful with the definition of the product with smooth functions. The product of a Schwartz function with $\psi \in C^\infty(\mathbf{R}^d)$ is not a Schwartz function in general. However, it suffices to require ψ is smooth and ψ and all its derivatives grow at most polynomially.

transform is also linear and continuous with respect to the convergence in $\mathcal{S}(\mathbf{R}^d)$ by Proposition 2.31. The definition entails

$$\widehat{S}(\varphi) = S(\mathcal{F}[\varphi]) = T(\mathcal{F}^{-1}[\mathcal{F}[\varphi]]) = T(\varphi).$$

This terminates the proof. □

In the previous proof, we obtained the formula $\widehat{T}^{-1} = T \circ \mathcal{F}^{-1}$. In particular, it is elementary to verify the Fourier transform and its inverse are sequentially weakly*-continuous on $\mathcal{S}'(\mathbf{R}^d)$. (The weak*-continuity is true as well, but the proof is slightly more involved.)

This concludes our concise survey on (tempered) distributions.

Lecture 8.

3. Calculus on Banach spaces

If not specified otherwise, in this section X, Y, and Z always denote Banach spaces over a field $\mathbf{K}$ equal to $\mathbf{R}$ or $\mathbf{C}$. We introduce basic calculus for maps $F\colon X \to Y$ and deduce several properties well-known from the finite-dimensional case. Since many proofs hardly differ from the case of maps from $\mathbf{R}^n$ to $\mathbf{R}^m$, where $n, m \in \mathbf{N}$, we often just sketch them.

3.1. Basic definitions and elementary results.

In a basic undergraduate course in analysis, one defines the "derivative" of a given map $F\colon \mathbf{R}^n \to \mathbf{R}^m$, where $n, m \in \mathbf{N}$, at a point $x_0 \in \mathbf{R}^n$ in two related yet different ways.

- If there exists a linear map $A \in \mathcal{L}(\mathbf{R}^n, \mathbf{R}^m)$ such that
$$\lim_{h\to 0} \frac{|F(x_0+h) - F(x_0) - Ah|}{|h|} = 0$$
then this map is unique, termed the (**total**) **derivative** of F at x_0, and denoted by $DF(x_0)$.
- If for a given $v \in \mathbf{R}^n$, the limit
$$\lim_{\varepsilon\to 0} \frac{F(x_0+\varepsilon v) - F(x_0)}{\varepsilon}$$
exists in $\mathbf{R}^m$, it is (clearly) unique, termed the **directional derivative** of F at x_0 in the direction of v, and denoted by $D_v F(x_0)$.

Recall a differentiable map is differentiable in every direction, but the converse is not true in general. Hence, the first notion is stronger, while in applications directional derivatives are frequently easier to compute.

We will generalize both notions to Banach spaces. The results to follow valid when we replace Banach spaces by general normed spaces.

The first element of calculus on such spaces is an appropriate notion of (total) derivative. It is defined in complete analogy to the finite-dimensional case that has already been treated in an undergraduate analysis course.

Definition 3.1 (Fréchet differentiability). *Let $F\colon X \to Y$ and let $U \subset X$ be open. We say F is* **Fréchet differentiable** (*or simply* **differentiable**) *at a point $x_0 \in U$ if there exists a bounded linear map $A \in \mathcal{L}(X, Y)$ such that*

$$\lim_{h\to 0} \frac{\|F(x_0+h) - F(x_0) - Ah\|_Y}{\|h\|_X} = 0.$$

We say F is Fréchet differentiable (*or simply* differentiable) *on U if it is differentiable at every point $x_0 \in U$.*

The operator A above is easily seen to be unique if it exists. Depending on the context, we will thus write $F'(x_0) := A$ or $DF(x_0) := A$.

Note that a differentiable function is in particular continuous (understood in the evident way, since we are dealing with normed spaces).

If a map $F\colon U \to Y$ is differentiable on some open set $U \subset X$, then the map $F'\colon U \to \mathcal{L}(X,Y)$ is also a map with values in a Banach space. By iterating this observation[25], we can define higher order derivatives as follows.

To simplify the notation, given any $n \in \mathbf{N}$ we inductively introduce the target space T_n of the n-th derivative of the map F in question by $T_1 := \mathcal{L}(X,Y)$ and $T_{n+1} := \mathcal{L}(X,T_n)$.

Definition 3.2 (Higher order derivatives)**.** *Let $U \subset X$ be open and $F\colon U \to Y$. Fix $n \in \mathbf{N}$. Inductively, if F is n-times differentiable in a neighborhood V of x_0 and its n-th Fréchet derivative $F^{(n)}\colon V \to T_n$ is differentiable at x_0, we say that F is $(n+1)$-times differentiable at x_0.*

We say F is $(n+1)$-times differentiable in U if it is $(n+1)$-times differentiable at every $x_0 \in U$.

Finally, we define

$$C^n(U;Y) = \{F\colon U \to Y : F \text{ is } n\text{-times differentiable in } U \text{ and } F^{(n)} \text{ is continuous on } U\}.$$

Example 3.3 (Fréchet differentiable maps).

- Every bounded linear map is Fréchet differentiable with constant derivative. Indeed, if $A \in \mathcal{L}(X,Y)$ then we have $A'(x) = A$ for every $x \in X$. This follows from the simple observation $A(x+h) - Ax - Ah = 0$ for every $x, h \in X$.
- Let $X := L^2([0,1],\mathcal{L}^1)$ and $Y := \mathbf{R}$. Define $F\colon X \to Y$ by

$$F(u) = \int_0^1 |u|^2 \,\mathrm{d}\mathcal{L}^1.$$

Then F is differentiable and

$$F'(u)v = 2\int_0^1 u\,v \,\mathrm{d}\mathcal{L}^1.$$

This is a consequence of the following more abstract result. Let H be a real Hilbert space with scalar product $\langle \cdot,\cdot \rangle$. Then the map $F\colon X \to \mathbf{R}$ defined by $F(x) := \langle x,x \rangle$ is differentiable with $F'(x)y = 2\langle x,y \rangle$ for every $x, y \in H$. Indeed, by bilinearity we easily see

$$\langle x+h, x+h \rangle - \langle x,x \rangle - 2\langle x,h \rangle = \|h\|_H^2,$$

for every $h \in H$, implying

$$\lim_{h\to 0} \frac{|F(x+h) - F(x) - 2\langle x,h \rangle|}{\|h\|_H} = \lim_{h\to 0} \|h\|_H = 0. \qquad \blacksquare$$

We now turn to extensions of classical calculus rules.

Lemma 3.4 (Chain rule)**.** *Let $U \subset X$ and $V \subset Y$ be open. Let $F\colon U \to Y$ and $G\colon V \to Z$. Assume both F is differentiable at a point $x_0 \in U$ with $F(x_0) \in V$ and G is differentiable at $F(x_0)$. Then $G \circ F\colon U \to Z$ is differentiable at x_0 and*

$$(G \circ F)'(x_0) = G'(F(x_0))F'(x_0).$$

Proof. Exercise 9.1. □

[25] Note that, for instance, $F''\colon U \to \mathcal{L}(X, \mathcal{L}(X,Y))$. Compare this with the customary Hessian of maps between finite-dimensional vector spaces.

Be careful with the notation in the lemma above, in the sense that we are not taking products of derivatives, but compositions. More precisely, $F'(x_0)$ constitutes a bounded linear map sending $x \in X$ to $F'(x_0)x \in Y$, after which the bounded linear map $G'(F(x_0))$ sends $y := F'(x_0)x \in Y$ to $G'(F(x_0))y$.

Another important tool in analysis is the mean value theorem. First, we recall the classical example of the circle $f\colon [0, 2\pi] \to \mathbf{R}^2$ given by $f(t) := (\cos(t), \sin(t))$ for which the equality $f(x) - f(y) = f'(\xi)(x - y)$ — that one proves for real-valued functions in a first analysis course — cannot hold for any $\xi \in [x, y]$ (take e.g. $x = 2\pi$ and $y = 0$). In other words, whenever the target domain of f has more than one dimension we cannot expect to extend the standard mean value theorem. The natural partial generalization to higher dimensions (which is often enough in applications) is the inequality $|f(x) - f(y)| \leq |f'(\xi)||x - y|$ for some $\xi \in [x, y]$. This is what we are going to generalize on Banach spaces.

To simplify the notation, in what follows we do not specify the norms whenever confusion is excluded. If $x \in X$ then $\|x\|$ refers to its norm in X, while terms like $\|F'(x)\|$ refer to the operator-norm of $F'(x)$ in $\mathcal{L}(X, Y)$.

Lemma 3.5 (Mean value inequality)**.** *Let $[a, b] \subset \mathbf{R}$ be an interval and let the map $F\colon [a, b] \to X$ be continuous on all of $[a, b]$ and differentiable on (a, b). Then there is $\xi \in (a, b)$ such that*

$$\|F(b) - F(a)\| \leq \|F'(\xi)\| \, |b - a|.$$

Proof. Without loss of generality, we may and will assume $F(a) \neq F(b)$, otherwise the claim is clear. Then by the Hahn–Banach theorem, we know there is $g \in X'$ such that $\|g\| = 1$ and $g(F(b) - F(a)) = \|F(b) - F(a)\|$[26]. We intend to apply the standard mean value theorem[27] to the mapping $g \circ F$. In particular, we know there exists $\xi \in [a, b]$ such that

$$(g \circ F)(b) - (g \circ F)(a) = (g \circ F)'(\xi)(b - a).$$

Using the chain rule and linearity of g,

$$(g \circ F)'(\xi)(b - a) = g'(F(\xi))(F'(\xi)(b - a)) = g(F'(\xi)(b - a)).$$

We finally get

$$\begin{aligned} \|F(b) - F(a)\| &= g(F(b) - F(a)) \\ &= g(F'(\xi)(b - a)) \\ &\leq \|g\| \, \|F'(\xi)\| \, |b - a| \\ &= \|F'(\xi)\| \, |b - a|. \end{aligned}$$

This terminates the proof. □

The above lemma allows us to prove local Lipschitz continuity of differentiable functions whose derivatives are locally[28] uniformly bounded. Indeed, it suffices to consider the restriction to one-dimensional segments.

We also have the following version of the Schwarz theorem on the symmetry of second derivatives.

---

[26]The Hahn–Banach theorem has the following consequence: for each non-zero $x_0 \in X$ there exists $g \in X'$ such that $g(x_0) = \|x_0\|$. To prove this, consider the subspace M spanned by x_0 and define $g(\lambda x_0) = \lambda\|x_0\|$ where $g \in M'$ and $\|g\| = 1$. Then apply the Hahn–Banach theorem to extend g to all of X'. Note that this is another proof of Corollary 1.37.

[27]We can assume g is real-valued since we can always interpret X as a vector space over $\mathbf{R}$.

[28]Attention: in infinite dimensions "locally" has to be understood on open neighborhoods instead of compact subsets.

Lemma 3.6 (Schwarz theorem). *Let $U \subset X$ be open. Assume that $F \colon U \to Y$ is two-times differentiable on U. Then for every $x \in U$ and every $v, w \in X$,*

$$(F''(x)v)w = (F''(x)w)v.$$

Proof. Fix $\varepsilon > 0$ and $v, w \in X$ with $\|v\|, \|w\| < \eta$, where $\eta > 0$ is sufficiently small such that for every $s \in [0,1]$,

$$\|F'(x+v+sw) - F'(x) - F''(x)(v+sw)\|_{\mathcal{L}(X,Y)} \leq \varepsilon\,(\|v\| + s\|w\|),$$
$$\|F'(x+sw) - F'(x) - F''(x)sw\|_{\mathcal{L}(X,Y)} \leq \varepsilon\, s\|w\|.$$

By the triangle inequality and the definition of the operator norm,

$$\|F'(x+v+sw)w - F'(x+sw)w - (F''(x)v)w\|_Y \leq \varepsilon\,(\|v\| + 2s\|w\|)\,\|w\|,$$

We define $g \colon [0,1] \to Y$ through $g(s) = F(x+v+sw) - F(x+sw) - s(F''(x)v)w$. Then the mean value inequality implies

$$\begin{aligned}
&\|F(x+v+w) - F(x+w) - F(x+v) + F(x) - (F''(x)v)w\| \\
&\quad\leq \sup_{s\in(0,1)} \Big\|\frac{\mathrm{d}}{\mathrm{d}s} g(s)\Big\| \\
&\quad\leq \sup_{s\in(0,1)} \|F'(x+v+sw)w - F'(x+sw)w - (F''(x)v)w\|_Y \\
&\quad\leq \varepsilon\,(\|v\| + 2\|w\|)\,\|w\|.
\end{aligned}$$

Exchanging the roles of v and w we conclude

$$\|F(x+v+w) - F(x+w) - F(x+v) + F(x) - (F''(x)w)v\| \leq \varepsilon\,(2\|v\| + \|w\|)\,\|v\|.$$

This entails

$$\|(F''(x)v)w - (F''(x)w)v\| \leq 2\,\varepsilon(\|v\| + \|w\|)^2$$

Since both sides are positively homogeneous of degree two, the above restriction on the norm of $\|v\|$ and $\|w\|$ can be dropped. Sending $\varepsilon \to 0$ yields the claim. ☐

We have defined a notion of differentiability. Next, we will introduce directional derivatives. They have several advantages. For instance, they are somewhat easier to compute. Moreover, differentiability is a very strong property in infinite dimensions[29] and one often needs to use weaker notions of differentiability. This is motivated by the following example.

Example 3.7 (A nowhere differentiable map). Let $X := L^2((0,1), \mathcal{L}^1)$ and define $F \colon X \to X$ as $F(u) := \cos \circ u$. Then F is not differentiable at 0. Indeed, assume F is differentiable at 0. Then for every $v \in \mathcal{D}((0,1))$ with unit L^2-norm,

$$F'(0)v = F'(0)v + \lim_{\varepsilon\to 0} \frac{F(\varepsilon v) - F(0) - F'(0)\varepsilon v}{\varepsilon} = \lim_{\varepsilon\to 0} \frac{F(\varepsilon v) - F(0)}{\varepsilon},$$

where the limits are understood in $L^2((0,1), \mathcal{L}^1)$. We compute

$$\lim_{\varepsilon\to 0} \Big\|\frac{F(\varepsilon v) - F(0)}{\varepsilon}\Big\|^2 = \lim_{\varepsilon\to 0} \frac{1}{\varepsilon^2} \int_0^1 |\cos(\varepsilon v(x)) - 1|^2 \,\mathrm{d}x = 0,$$

where the latter equality follows from the global bound $|\cos(x) - 1| \leq |x|^2/2$ for every $x \in \mathbf{R}_+$, which can be seen by Taylor expansion, and the boundedness of v. This implies $F'(0)v = 0$ for every $v \in \mathcal{D}((0,1))$; by density of $\mathcal{D}((0,1))$ in $L^2((0,1), \mathcal{L}^1)$, we deduce $F'(0) = 0$. Hence for every $h \in L^2((0,1), \mathcal{L}^1)$,

$$\int_0^1 |\cos(h(x)) - 1|^2 \,\mathrm{d}x = \mathrm{o}(\|h\|^2).$$

[29] Roughly speaking, in finite dimensions one only needs to control finitely many directions to control the entire derivative.

Now choose $h = 1_{[0,\varepsilon]}$ with $\varepsilon \in (0,1)$. The above estimate implies

$$\varepsilon\,|\cos(1) - 1|^2 = \mathrm{o}(\varepsilon^2).$$

Dividing by ε, we obtain a contradiction since the left-hand side is different from zero. With more effort one can show F is in fact nowhere differentiable. ■

Lecture 9. Now we turn to the analog of directional differentiation.

Definition 3.8 (Gâteaux differentiability). *Let $U \subset X$ be open and $F\colon U \to Y$. We say F is* **Gâteaux differentiable** *at a point $x_0 \in U$ if there exists $A \in \mathcal{L}(X,Y)$ such that for every $v \in X$,*

$$Av = \lim_{\varepsilon \to 0} \frac{F(x_0 + \varepsilon v) - F(x_0)}{\varepsilon}.$$

We say F is Gâteaux differentiable in U if it is Gâteaux differentiable at every $x_0 \in U$.

In other words, we require the above limit to exist and to constitute a bounded linear map from X to Y.

In the above case, we write $\delta F(x_0) := A$ (which is trivially unique if existent) to distinguish the Gâteaux derivative from the Fréchet derivative.

Note that in contrast to the classical definition of differentiability, in the above limit the convergence can be nonuniform with respect to the direction v. The relationship between the two concepts is similar to the finite-dimensional case. Hence we omit the proof of the lemma below.

Lemma 3.9 (Gâteaux vs. Fréchet differentiation). *Let $U \subset X$ be open and $F\colon U \to Y$. If F is differentiable at $x_0 \in U$, it is Gâteaux differentiable and $\delta F(x_0) = F'(x_0)$.*

Conversely, if F is Gâteaux differentiable in a neighborhood V of x_0 and the Gâteaux derivative $\delta F\colon V \to \mathcal{L}(X,Y)$ is continuous, then F is differentiable at x_0 with $\delta F(x_0) = F'(x_0)$.

As customary in the Euclidean case, the Gâteaux derivative provides a useful necessary condition for optimization problems.

Lemma 3.10 (Necessary conditions for extremizers). *Let $U \subset X$ be open and $F\colon U \to \mathbf{R}$. If $x_0 \in U$ is a local extremizer of F and F is Gâteaux differentiable at x_0, then $\delta F(x_0) = 0$.*

Proof. By replacing F with $-F$, it suffices to establish the claim when x_0 is a local minimizer. Given any $v \in X$, the small perturbation $x_0 + \varepsilon v$ lies in U for every sufficiently small $\varepsilon > 0$. Given any such ε, local minimality of F at x_0 implies $F(x_0) \leq F(x_0 + \varepsilon v)$. Therefore,

$$\delta F(x_0)v = \lim_{\varepsilon \to 0+} \frac{F(x_0 + \varepsilon v) - F(x_0)}{\varepsilon} \geq 0.$$

Replacing v by $-v$ we deduce $\delta F(x_0)v = 0$. □

Similar to the finite-dimensional situation one can also derive necessary and sufficient conditions for the second derivatives in optimization problems. We will briefly discuss them in the exercises.

3.2. **Partial derivatives and the implicit function theorem.** Now we consider functions defined on Cartesian products of Banach spaces, i.e., we assume that given any $n \in \mathbf{N}$, $X_1, \dots, X_n$ are Banach spaces. We equip $X_1 \times \dots \times X_n$ with the complete maximum norm given by $\|(x_1, \dots, x_n)\| = \max\{\|x_i\|_{X_i} : i \in \{1, \dots, n\}\}$. If $U \subset X_1 \times \dots \times X_n$ is an open set and $F\colon U \to Y$ is a map, then for a fixed vector

$(x_1, \dots, \hat{x}_i, \dots, x_n) \in X_1 \times \dots \times \hat{X}_i \times \dots \times X_n$[30] we can consider the restriction $x_i \mapsto F(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n)$ defined on the open set

$$U_i = \{x_i \in X_i : (x_1, \dots, x_i, \dots, x_n) \in U\},$$

where $i \in \{1, \dots, n\}$. The i-th partial derivative $\partial_i F \colon U \to \mathcal{L}(X_i, Y)$ is defined as the derivative of this restriction (always intended in the sense of Fréchet).

The following result shows partial derivatives enable us to recover the entire derivative provided f is differentiable on the product space.

Lemma 3.11 (Partial vs. total derivative)**.** *Let $U \subset X_1 \times \dots \times X_n$ be open and let $F \colon U \to Y$. If F is differentiable, then all its partial derivatives exist and satisfy the following for every $(x_1, \dots, x_n) \in U$ and every $(h_1, \dots, h_n) \in X_1 \times \dots \times X_n$:*

$$F'(x_1, \dots, x_n)(h_1, \dots, h_n) = \sum_{j=1}^{n} \partial_j F(x_1, \dots, x_n) h_j. \tag{3.1}$$

Moreover, if $F \in C^1(U; Y)$, then $\partial_j F \in C(U; \mathcal{L}(X_j, Y))$ for every $j \in \{1, \dots, n\}$.

Conversely, if all partial derivatives exist and obey $\partial_j F \in C(U; \mathcal{L}(X_j, Y))$, then $F \in C^1(U; Y)$ and its derivative computes as (3.1).

Proof. The argument for the first two claims is left as an exercise to the reader.
The last statement is shown in Exercise 10.1. □

Now let us address the implicit function theorem. We first collect some consequences of Banach's fixed point theorem. Note that the first statement below is a nonlinear, abstract version of the Neumann series.

Lemma 3.12 (Contractions)**.**

(i) *Let $T \colon X \to X$. If there exists $\theta \in (0, 1)$ such that $\|T(x) - T(y)\| \leq \theta \|x - y\|$ for every $x, y \in X$, then $\mathrm{Id} - T$ is a homeomorphism from X to X.*

(ii) *Let $S \colon \overline{B}_\delta(0) \to X$, where $\overline{B}_\delta(0)$ denotes the closed δ-ball in X for $\delta > 0$, and assume there exists $\theta \in (0, 1)$ such that $\|S(x) - S(y)\| \leq \theta \|x - y\|$ for every $x, y \in \overline{B}_\delta(0)$. If $\|S(0)\| < \delta(1 - \theta)$ then the map $\mathrm{Id} + S$ has a unique zero. Moreover, setting $\rho := (1 - \theta)\delta - \|S(0)\|$ we have*

$$B_\rho(0) \subset (\mathrm{Id} + S)(B_\delta(0)).$$

Proof. Exercise 10.2. □

As we will see, the second part of the above lemma is tailor-made for the proof of the implicit function theorem.

Theorem 3.13 (Implicit function theorem)**.** *Let $x_0 \in X$ and $y_0 \in Y$. Let $U \subset X$ and $V \subset Y$ be open neighborhoods of x_0 and y_0, respectively. Assume $F : U \times V \to Z$ is continuous, partially differentiable in its second component, and that $\partial_2 F : U \times V \to \mathcal{L}(Y, Z)$ is continuous. Suppose $F(x_0, y_0) = 0$. If the linear map $\partial_2 F(x_0, y_0)$ is invertible, there exist closed balls $\overline{B}_r(x_0) \subset U$ and $\overline{B}_\delta(y_0) \subset V$ and exactly one map $T \colon \overline{B}_r(x_0) \to \overline{B}_\delta(y_0)$ such that $F(x, T(x)) = 0$ for all $x \in \overline{B}_r(x_0)$.*

Moreover, the map T is continuous and satisfies $T(x_0) = y_0$.

As usual, the implicit function theorem gives a local parametrization of a subset of the zero level set of F around (x_0, y_0) by the first component.

Proof of Theorem 3.13. Up to a translation (which does not cost generality), we may and will assume $x_0 = 0$ and $y_0 = 01$. We set $L := \partial_2 F(0, 0)$. Since for given $x \in U$ and $y \in V$ the equation $F(x, y) = 0$ holds if and only if $y + L^{-1} F(x, y) - y = 0$ by the hypothesized invertibility of L, we verify the assumptions of the second part

[30] We use the usual hat notation from multilinear algebra to indicate the i-th slot is omitted.

of Lemma 3.12 for $S(x,\cdot) := L^{-1}F(x,\cdot) - \mathrm{Id}_Y$, where Id_Y denotes the identity on Y. Indeed, a zero of $\mathrm{Id}_Y + S(x,\cdot)$ will depend on x and we use this zero to construct the desidered map T.

Since $\partial_2 S(0,0) = 0$ and $\partial_2 S$ is continuous on $U \times V$, given any $\theta \in (0,1)$ there exists $\delta > 0$ such that $\|\partial_2 S(x,y)\| \leq \theta$ on $\overline{B}_\delta(0) \times \overline{B}_\delta(0)$. The mean value theorem on the convex set $\overline{B}_\delta(0) \subset Y$ implies

$$\|S(x,y_1) - S(x,y_2)\| \leq \theta \|y_1 - y_2\|.$$

Moreover, the continuity in the first variable implies that for some $r \in (0,\delta)$, we have $\|S(\cdot,0)\| < \delta(1-\theta)$ on $\overline{B}_r(0)$. Hence the second statement of Lemma 3.12 yields there exists a unique zero $T(x) \in \overline{B}_\delta(0)$ of the map $y \mapsto y + S(x,y)$ for every $x \in \overline{B}_r(0)$. By uniqueness, we have $T(0) = 0$.

In order to prove continuity of T, note that for every $x, x_0 \in \overline{B}_r(0)$ we have

$$\begin{aligned}\|T(x) - T(x_0)\| &= \|S(x,T(x)) - S(x_0,T(x_0))\| \\ &\leq \|S(x,T(x)) - S(x,T(x_0))\| + \|S(x,T(x_0)) - S(x_0,T(x_0))\| \\ &\leq \theta\,\|T(x) - T(x_0)\| + \|S(x,T(x_0)) - S(x_0,T(x_0))\|\end{aligned}$$

Absorbing the θ-dependent term in the left-hand side, the continuity follows from the continuity of S in the first variable. □

Next we prove the inverse function theorem on Banach spaces.

Theorem 3.14 (Inverse function theorem)**.** *Let $U_0 \subset X$ be open and $F \in C^1(U_0;Y)$. Let $x_0 \in U_0$ be such that $F'(x_0)$ is invertible. Then there exists an open neighborhood $U \subset U_0$ of x_0 such that $F\big|_U \colon U \to F(U)$ is a homeomorphism onto the open neighborhood $F(U)$ of $y_0 := F(x_0)$.*

Moreover, there is a possibly smaller open neighborhood $V \subset U$ of x_0 such that $F\big|_U^{-1} \in C^1(F(V);X)$, and the following identity holds for every $x \in V$:

$$(F\big|_U^{-1})'(F(x)) = F'(x)^{-1}.$$

Proof. We apply the implicit function theorem to the assignment defined by $H(x,y) := F(x) - y$ with $Z = Y$ and the roles of X and Y reversed. We thus find $W = B_r(y_0)$ and $\overline{B}_\delta(x_0) \subset U_0$ and a unique continuous map $T : W \to B_\delta(x_0)$[31] such that $T(y_0) = x_0$ and $H(T(y),y) = 0$, that is, $F(T(y)) = y$ for every $y \in W$. In particular, W is in the image of F and the set $U := F^{-1}(W) \cap B_\delta(x_0)$ containing x_0 is open in X and $F(U) = W$ is an open neighborhood of y_0. Moreover, F has to be injective on U (otherwise there are at least two possibilities to construct the map T). Note that $T(W) \subset F^{-1}(W) \cap B_\delta(x_0) = U$. Hence $T \colon W \to U$, but we still need to show that $T(F(x)) = x$ for every $x \in U$. Clearly $F \colon U \to W$ is bijective (by injectivity and since $W = F(U)$). Hence as a general fact left and right inverse agree and we found $F\big|_U^{-1} = T$, so that F is indeed a homeomorphism.

It remains to show the differentiability properties of T. As the set of invertible linear maps is open[32], for some possibly smaller open neighborhood $W_0 \subset W$ the bounded, linear operator $F'(T(y))$ is still invertible. Set $V = F^{-1}(W_0) \cap B_\delta(x_0)$. Let us show that T is differentiable at $y \in W_0 = F(V)$ with derivative $F'(T(y))^{-1}$. Observe that

$$\begin{aligned}&\|T(y+h) - T(y) - F'(T(y))^{-1}h\| \\ &\qquad \leq \|F'(T(y))^{-1}\|_{\mathcal{L}(Y,X)}\,\|F'(T(y))T(y+h) - F'(T(y))T(y) - h\|\end{aligned}$$

[31] To obtain the open ball in the image, it suffices to decrease the radius.

[32] This is a well-known consequence of the von Neumann series when $Y = X$. In the general case, given L_0 invertible, write $L = L_0(I + L_0^{-1}(L - L_0))$ and then L is invertible whenever $I + L_0^{-1}(L - L_0)$ is invertible. The latter operator is again an operator from X to X and hence it is invertible when the operator norm $\|L - L_0\|$ is small enough.

$$= c\,\|F'(x)(x_h - x) - (F(x_h) - F(x))\|,$$

where $x := T(y)$, $x_h := T(y+h)$, and $c := \|F'(T(y))^{-1}\|_{\mathcal{L}(Y,X)}$. Note that by the continuity of T it follows that $x_h \to x$ as $h \to 0$. Hence, given $\varepsilon > 0$, the differentiability of F at x implies that for $\|h\|$ small enough,

$$\|T(y+h) - T(y) - F'(T(y))^{-1}h\| \le c\,\varepsilon\|x_h - x\| = c\,\varepsilon\|T(y+h) - T(y)\|.$$

In particular, for ε small enough we conclude $\|T(y+h) - T(y)\| \le c(1-\varepsilon c)^{-1}\|h\|$. Inserting this bound into the estimate above we obtain

$$\|T(y+h) - T(y) - F'(T(y))^{-1}h\| \le c\,\varepsilon\|x_h - x\| = c^2\,\varepsilon(1-c\varepsilon)^{-1}\|h\|,$$

which shows T is differentiable in W_0 with $T'(y) = F'(T(y))^{-1}$. Since T and F' are continuous, we obtain $T \in C^1(W_0; X)$. □

Remark 3.15 (Smoothness). One can show that the local inverse function inherits the smoothness of F, that is, if $F \in C^m(U_0, Y)$ for some $m \in \mathbf{N} \cup \{\infty\}$, then $F\big|_U^{-1} \in C^m(F(V); X)$ as well. A detailed proof of this fact is quite technical. Here is a sketch of the argument. Using $T'(y) = F'(T(y))^{-1}$ and the fact that taking the inverse of a bounded, linear operator is a smooth function, it suffices to show that the composition of smooth functions and C^{m-1}-functions belongs to C^{m-1}, which follows essentially from the chain rule. A similar statement holds for the function given by the implicit function theorem. ■

Lecture 10. As a further application of the implicit function theorem we discuss global diffeomorphisms[33]. Recall a diffeomorphism between two Banach spaces is a smooth bijective map with smooth inverse.

Corollary 3.16 (Sufficient criterion for diffeomorphy). *Let $F \in C^1(X;Y)$ be such that $F'(x)$ is invertible for all $x \in X$. If there exists a constant $c > 0$ such that $\|F'(x)^{-1}\| \le c\,\|x\| + c$ for every $x \in X$, then F is a diffeomorphism from X to Y.*

Proof. By the inverse function theorem, it suffices to show F is bijective. We first show F is surjective. Fix $y \in Y$ and $x_0 \in X$. Define $H\colon \mathbf{R} \times X \to Y$ by $H(t,x) := F(x) - ty - (1-t)F(x_0)$. Then $H(0,x_0) = 0$ and $\partial_2 H(0,x_0) = F'(x_0)$ is invertible. By the implicit function theorem there exists $x\colon [0,\delta] \to X$ such that

$$F(x(t)) = ty + (1-t)F(x_0).$$

If $\delta \ge 1$ we know $y = F(x(1))$. To this end, we derive an ODE satisfied by x.

We use the following trick to obtain the desired ODE by applying the inverse function theorem. Consider $G\colon \mathbf{R}\times X \to \mathbf{R}\times X$ defined by $G(t,x) = (t, F'(x_0)^{-1}H(t,x))$. Then $G \in C^1(\mathbf{R}\times X, \mathbf{R}\times X)$ and $G'(0,x_0)(t,h) = (t, h - yt + F(x_0)t)$, a map which is clearly invertible. Since $G^{-1}(t,0) = (t, x(t))$, we deduce x is indeed differentiable in an open interval around 0 and then the chain rule yields that

$$x'(t) = F'(x(t))^{-1}(y - F(x_0)).$$

Let T denote the maximal interval of the form $[0,T)$ where this differential equation can be solved with initial value $x(0) = x_0$. If $T \le 1$ then

$$\begin{aligned}\|x(t)\| &\le \|x_0\| + \int_0^t \big|F'(x(s))^{-1}(y - F(x_0))\big|\,\mathrm{d}s\\ &\le \|x_0\| + \int_0^t \big[c\|x(s)\| + c\big]\,\|y - F(x_0)\|\,\mathrm{d}s\\ &= c_1 + c_2\int_0^t \|x(s)\|\,\mathrm{d}s\end{aligned}$$

[33]The proof of this result was skipped in the lecture. The proof is not examinable.

for some appropriate constant $C > 0$. Gronwall's inequality then yields $\|x(t)\| \leq C\exp(Ct) \leq C'$ for all $t \in [0,T)$. In turn, using the ODE the quantity $\|x'\|$ is uniformly bounded on $[0,T)$. In particular, the limit $\lim_{t\to T} x(t)$ exists (here we use that X is complete). Denote the limit by x_T. Then we repeat the above argument with x_0 replaced by x_T and obtain a contradiction to the maximality of T. This shows $T > 1$ — actually $T = \infty$! — and therefore F is surjective.

For proving injectivity, assume $F(x_1) = F(x_2) = y$ yet $x_1 \neq x_2$. Without loss of generality we may and will assume $y = 0$. We will construct a continuous function $\varphi\colon [0,1]^2 \to X$ such that $F(\varphi(t,s)) = -sF(tx_1 + (1-t)x_2)$ for every $s,t \in [0,1]$. Then $F(\varphi(t,0)) = 0$ for all $t \in [0,1]$, which contradicts the local invertibility of F. To find φ, define $C_0([0,1];X) = \{u \in C([0,1];X) : x(0) = x(1) = 0\}$ (space of X-valued continuous functions on $[0,1]$ with Dirichlet boundary conditions) equipped with the maximum norm $\|\cdot\|_\infty$. Define $H\colon C_0^0([0,1];X) \to C([0,1];Y)$ by

$$H(u)(t) := F(u(t) + tx_1 + (1-t)x_2).$$

Note that $H(u)(0) = H(u)(1) = 0$, so that $H(u) \in C_0^0([0,1];Y))$. Moreover, one can show $H \in C^1(C_0[0,1];X);C([0,1];Y))$ with derivative

$$(H'(u)h)(t) = F'(u(t) + tx_1 + (1-t)x_2)h(t).$$

In particular,

$$(H'(u)^{-1}y)(t) = F'(u(t) + tx_1 + (1-t)x_2)^{-1}y(t).$$

Hence $\|H'(u)^{-1}\| \leq C\|u\|_\infty + C$ for some constant $C > 0$. Repeating the first part of the proof with F replaced by H and $x_0 = 0 \in C_0([0,1],X)$ and $y = 0 \in C([0,1];Y)$ to find a differentiable function $v : [0,1] \to C_0([0,1];X)$ such that $v(0) = 0$ and

$$v'(s) = -H'(v(s))^{-1}H(0).$$

Setting $\varphi(t,s) = tx_1 + (1-t)x_2 + v(s)(t)$ we get

$$\begin{aligned} F(\varphi(t,s)) &= H(v(s))(t) \\ &= \int_0^s \left(H'(v(r))v'(r)\right)(t)\,\mathrm{d}r \\ &= -sH(0)(t) \\ &= -sF(tx_1 + (1-t)x_2). \end{aligned}$$

This terminates the proof. □

3.3. **Taylor expansion on Banach spaces.** For function $f : \mathbf{R} \to \mathbf{R}$ the Taylor expansion around a point $x \in \mathbf{R}$ reads

$$f(x+h) = f(x) + f'(x)\,h + \ldots + \frac{f^{(n)}(x)}{n!}\,h^n + \mathrm{o}(|h|^n),$$

provided f is n-times differentiable, where $n \in \mathbf{N}$. In case one has more information, one can express the remainder $\mathrm{o}(|h|^n)$[34] either as an integral or an evaluation of a term involving the $(n+1)$-st derivative. We already know from the mean value theorem that the latter possibility cannot be expected for functions with target domain not being a subset of the real line.

Before we prove any formulation of the Taylor expansion in Banach spaces, we need to understand the structure of higher order derivatives. It will be convenient to identify them with bounded, multilinear maps.

[34]Here is a little reminder on Landau notation. We will write $f \in \mathrm{o}(h)$ as $x \to a$ provided $\limsup_{x\to a} |f(x)/h(x)| = 0$. Equivalently, for every $\varepsilon > 0$ there exists $\delta > 0$ such that for every $x \in B_\delta(a)$, we have $|f(x)| \leq \varepsilon\,|h(x)|$. Whenever we write $f = g + \mathrm{o}(h)$ as $x \to a$ we intend to say $f - g \in \mathrm{o}(h)$ as $x \to a$.

Recall $\mathfrak{S}_n$ denotes the group of permutations of $\{1, \dots, n\}$, i.e. the set of all bijective maps $\sigma\colon \{1, \dots, n\} \to \{1, \dots, n\}$, where $n \in \mathbf{N}$.

Definition 3.17 (Multilinear maps)**.** *Given any $n \in \mathbf{N}$, let $\mathcal{M}^n(X, Y)$ denote the set of multilinear, bounded maps from X^n to Y, i.e., functions $m\colon X^n \to Y$ that are linear in each variable and such that the following quantity is finite:*

$$\|m\|_{\mathcal{M}^n(X,Y)} := \sup\{\|m(x_1, \dots, x_n)\|_Y : \|x_i\|_X \leq 1 \text{ for every } i \in \{1, \dots, n\}\}.$$

We call $m \in \mathcal{M}^n(X, Y)$ symmetric if for every permutation $\sigma \in \mathfrak{S}_n$ and every $(x_1, \dots, x_n) \in X^n$,

$$m(x_{\sigma(1)}, \dots, x_{\sigma(n)}) = m(x_1, \dots, x_n)$$

Remark 3.18 (Functional analytic properties). The space $\mathcal{M}^n(X, Y)$ with the above norm becomes a Banach space. Additionally, the space of symmetric maps is a closed subspace of it. ■

The identification of higher order derivatives derivative with an appropriate symmetric, bounded, multilinear map is justified in the next lemma.

Lemma 3.19 (Generalized Schwarz theorem)**.** *Let $U \subset X$ be open. Let $F\colon U \to Y$ be N-times differentiable, where $N \in \mathbf{N}$. Then for every $n \in \{1, \dots, N\}$, every permutation $\sigma \in \mathfrak{S}_n$, and every $x \in U$,*

$$F^{(n)}(x)(h_1, \dots, h_n) = F^{(n)}(x)(h_{\sigma(1)}, \dots, h_{\sigma(n)}). \tag{3.2}$$

Moreover, given any $x \in U$, we have $F^{(n)}(x) \in \mathcal{M}^n(X, Y)$ and

$$\|F^{(n)}(x)\| = \|F^{(n)}(x)\|_{\mathcal{M}^n(X,Y)}.$$

Recall the terms in (3.2) are understood as iterated applications. For instance,

- if $n = 2$ we have $F''(x)(h_1, h_2) = (F''(x)h_1)h_2$,
- if $n = 3$ we have $F'''(x)(h_1, h_2, h_3) = ((F''(x)h_1)h_2)h_3$,

and so on.

Proof of Lemma 3.19. We first prove symmetry by induction on n, starting with $n = 2$. In this case the statement is exactly Schwarz theorem, cf. Lemma 3.6.

To prove the general case, assume for the moment that σ is a permutation such that $\sigma(1) = 1$. The hypothesized differentiability of $F^{(n)}$ implies for every $h_1, \dots, h_n \in X$ the map $x \mapsto F^{(n)}(x)(h_1, \dots, h_n)$ from U to Y is differentiable with

$$\big[F^{(n)}(\cdot)(h_1, \dots, h_n)\big]'(x)\, h = F^{(n+1)}(x)(h, h_1, \dots, h_n)$$

for every $h \in X$. By our induction hypothesis,

$$\begin{aligned}
F^{(n+1)}(x)(h_1, \dots, h_{n+1}) &= \big[F^{(n)}(\cdot)(h_2, \dots, h_{n+1})\big]'(x)\, h_1 \\
&= \big[F^{(n)}(\cdot)(h_{\sigma(2)}, \dots, h_{\sigma(n+1)})\big]'(x)\, h_{\sigma(1)} \\
&= F^{(n+1)}(x)(h_{\sigma(1)}, \dots, h_{\sigma(n+1)})
\end{aligned}$$

Next, we consider $\sigma \in \mathfrak{S}_n$ such that $\sigma(1) = 2$ and $\sigma(2) = 1$ yet $\sigma(k) = k$ for every $k \in \{3, \dots, n+1\}$. Then the statement follows from Schwarz's theorem applied to the assignment $x \mapsto F^{(n-1)}(h_3, \dots, h_{n+1})$. The general case follows from the fact that any permutation can be written as a composition of the permutations considered above[35].

[35]Given any $\sigma \in \mathfrak{S}_n$, consider first a permutation that exchanges $\sigma(1)$ and 2 (first type when $\sigma(1) \neq 1$ or second type when $\sigma(1) = 1$). Then exchange the first and second element (second type). The remaining permutation will be of the first type.

Using the symmetry proven in the preceding part, it suffices to prove linearity in the first component. This directly follows from the definition of the n-th derivative and the linearity of point evaluations $f(x)$ in f.

To show boundedness, let $h_1, \dots, h_n \in X$ with $\|h_j\| \leq 1$ for all $j \in \{1, \dots, n\}$. Iterating the boundedness of the n-th derivative,

$$\begin{aligned}\big\|F^{(n)}(x)(h_1, \dots, h_n)\big\|_Y &\leq \big\|F^{(n)}(x)(h_1, \dots, h_{n-1})\big\|_{\mathcal{L}(X,Y)} \|h_n\| \\ &\leq \big\|F^{(n)}(x)(h_1, \dots, h_{n-2})\big\|_{\mathcal{L}(X,\mathcal{L}(X,Y))} \|h_{n-1}\| \, \|h_n\| \\ &\leq \dots \\ &\leq \|F^{(n)}(x)\| \, \|h_1\| \cdots \|h_n\|.\end{aligned}$$

This proves $F^{(n)}(x) \in \mathcal{M}^n(X, Y)$ and $\|F^{(n)}(x)\|_{\mathcal{M}^n(X,Y)} \leq \|F^{(n)}(x)\|$. To prove the reverse inequality, note that by definition

$$\begin{aligned}\|F^{(n)}(x)\| &= \sup_{\substack{h_1 \in X,\\ \|h_1\| \leq 1}} \cdots \sup_{\substack{h_n \in X,\\ \|h_n\| \leq 1}} \|F^{(n)}(x)(h_1, \dots, h_n)\|_Y \\ &\leq \sup\{\|F^{(n)}(x)(h_1, \dots, h_n) : \|h_j\| \leq 1 \text{ for every } j \in \{1, \dots, n\}\} \\ &= \|F^{(n)}(x)\|_{\mathcal{M}^n(X,Y)}.\end{aligned}$$

This terminates the proof. ☐

Armed with the multilinearity of the n-th-derivative, the Taylor expansion is an easy consequence of the mean value theorem combined with an induction argument.

Theorem 3.20 (Taylor expansion without quantitative remainder estimate)**.** *Let $U \subset X$ be open. Let $F: U \to Y$ be n-times differentiable, where $n \in \mathbf{N}$. Then for every $x_0 \in U$,*

$$F(x_0 + h) = \sum_{k=0}^{n} \frac{1}{k!} F^{(k)}(x_0) \underbrace{(h, \dots, h)}_{k \text{ times}} + \mathrm{o}(\|h\|^n) \quad \textit{as } h \to 0,$$

with the usual convention $F^{(0)} = F$.

Proof. We prove the statement by induction on n. For $n = 1$ it reduces to the definition of Fréchet-differentiability.

Now assume F is $(n+1)$-times differentiable. Note that the Taylor expansion is exact for $h = 0$. Hence by the mean value theorem applied to the assignment

$$\begin{aligned}G(t) &:= F(x_0 + th) - \sum_{k=0}^{n+1} \frac{1}{k!} F^{(k)}(x_0)(th, \dots, th) \\ &= F(x_0 + th) - t^k \sum_{k=0}^{n+1} F^{(k)}(x_0)(h, \dots, h)\end{aligned}$$

and an index shift we obtain

$$\begin{aligned}&\Big\|F(x_0 + h) - \sum_{k=0}^{n+1} \frac{1}{k!} F^{(k)}(x_0)(h, \dots, h)\Big\| \\ &\qquad \leq \sup_{t \in [0,1]} \Big\|F'(x_0 + th)h - \sum_{k=1}^{n+1} \frac{1}{(k-1)!} F^{(k)}(x_0)(th, \dots, th, h)\Big\| \\ &\qquad \leq \sup_{t \in [0,1]} \Big\|F'(x_0 + th) - \sum_{k=0}^{n} \frac{1}{k!} F^{(k+1)}(x_0)(th, \dots, th)\Big\|_{\mathcal{L}(X,Y)} \|h\|\end{aligned}$$

$$= \sup_{t\in[0,1]} \Big\| F'(x_0 + th) - \sum_{k=1}^{n} \frac{1}{k!} (F')^{(k)}(x_0)(th, \dots, th) \Big\|_{\mathcal{L}(X,Y)} \|h\|$$

Applying the induction hypothesis for F',

$$\Big\| F(x_0 + h) - \sum_{k=0}^{n+1} \frac{1}{k!} F^{(k)}(x_0)(h, \dots, h) \Big\| \\ \leq \sup_{t\in[0,1]} \mathrm{o}(\|th\|^n) \|h\| = \mathrm{o}(\|h\|^{n+1}) \quad \text{as } h \to 0,$$

which is the desired asymptotic. □

Next we present one version of the Taylor expansion with a more precise control of the error assuming higher order differentiability.

Theorem 3.21 (Taylor expansion with remainder estimate)**.** *Let $U \subset X$ be open. Let $F\colon U \to Y$ be $(n+1)$-times differentiable. Then for every $x \in U$ and every $h \in X$ such that the segment $[x, x+h]$ is contained in U, there is $\zeta \in [x, x+h]$ with*

$$\Big\| F(x + h) - \sum_{k=0}^{n} \frac{1}{k!} F^{(k)}(x)(h, \dots, h) \Big\| \leq \frac{1}{(n+1)!} \big\| F^{(n+1)}(\zeta)(h, \dots, h) \big\|.$$

Proof. By the Hahn–Banach theorem, there exists $g \in Y'$ with $\|g\| = 1$ and

$$g\Big[F(x + h) - \sum_{k=0}^{n} \frac{1}{k!} F^{(k)}(x)(h, \dots, h) \Big] = \Big\| F(x + h) - \sum_{k=0}^{n} \frac{1}{k!} F^{(k)}(x)(h, \dots, h) \Big\|.$$

Define the function $f\colon [0, 1] \to \mathbf{R}$ by

$$f(t) := g\Big[F(x + th) - \sum_{k=1}^{n} \frac{1}{k!} F^{(k)}(x)(th, \dots, th) \Big].$$

By the chain rule $f \in C^{n+1}((a, b))$, where (a, b) is an open interval containing $[0, 1]$. (Recall U is open.) By Taylor's theorem, there exists $t_0 \in [0, 1]$ such that

$$\Big\| F(x + h) - \sum_{k=0}^{n} \frac{1}{k!} F^{(k)}(x)(h, \dots, h) \Big\| \\ = f(1) - f(0) = \sum_{k=1}^{n} \frac{1}{k!} f^{(k)}(0) + \frac{1}{(n+1)!} f^{(n+1)}(t_0).$$

The derivatives of f are given by

$$f^{(k)}(t) = g\Big[F^{(k)}(x + th)(h, \dots, h) - \sum_{j=k}^{n} \frac{t^{j-k}}{(j-k)!} F^{(k)}(x)(h \dots, h) \Big].$$

In particular, for $t = 0$ we find $f^{(k)}(0) = 0$. This implies

$$\begin{aligned} \Big\| F(x + h) - \sum_{k=0}^{n} \frac{1}{k!} F^{(k)}(x)(h, \dots, h) \Big\| \\ &= \frac{1}{(n+1)!} f^{(n+1)}(t_0) \\ &= \frac{1}{(n+1)!} g(F^{(n+1)}(x + t_0 h)(h \dots, h)) \\ &\leq \frac{1}{(n+1)!} \|F^{(n+1)}(\zeta)(h, \dots, h)\|, \end{aligned}$$

where we have set $\zeta := x + t_0 h$. □

Remark 3.22 (Integral remainders). In the case $Y = \mathbf{R}$, the above proof yields $g \in \{-1, 1\}$ and therefore the Taylor expansion is exact with the intermediate value ζ. Remainder formulas with integral expression require integration theory in Banach spaces and therefore we omit them in this course. ■

Lecture 11.

4. A selection of fixed point theorems

Fixed point problems occur all over analysis. For instance, solving an equation of the form $G(x) = y$, where x and y belong to the same space, can easily be transformed into the fixed point equation $F(x) = x$, where $F(x) = x - G(x) - y$. The function F can be seen as a perturbation of the identity, even though the perturbation might not be small in any sense.

Another well-known domain where fixed point arguments occur is the theory of ODEs. Suppose $f\colon \mathbf{R}_+ \times \mathbf{R}^n \to \mathbf{R}^n$ is a (say) smooth function, where $n \in \mathbf{N}$. Solving the ODE $x'(t) = f(t, x(t))$ with $x(0) = x_0$ directly would amount to looking for solutions x in a space of C^1-functions, which is analytically challenging. Instead, we transform the ODE by integration into a fixed point equation $F(x) = x$, where

$$F(x)(t) := x_0 + \int_0^t f(s, x(s))\,\mathrm{d}s.$$

Observe we only have to solve this fixed point equation in a space of continuous functions, which is one derivative better! Indeed, by the fundamental theorem of calculus a continuous solution x of $F(x) = x$ is automatically continuously differentiable, hence solves the original ODE.

In this section we study the existence of fixed points of functions $F\colon K \to X$, where $K \subset X$ is an appropriate set.

4.1. Banach-style fixed point theorems.

A well-known result is Banach's fixed point theorem, which ensures existence and uniqueness whenever F is a strict contraction on a complete metric space. Moreover, the fixed points depend Lipschitz continuously on F.

Theorem 4.1 (Banach)**.** *Assume (X, d) is a complete metric space. Moreover, let $F\colon X \to X$ be a contraction, i.e. there exists $\lambda \in (0, 1)$ such that $\mathsf{d}(F(x), F(y)) \leq \lambda\,\mathsf{d}(x, y)$ for every $x, y \in X$. Then F has a unique fixed point.*

The restriction on λ cannot be dropped if one wants to retain unique existence of a fixed point. For instance, nontrivial translations in $\mathbf{R}^n$ are 1-Lipschitz continuous yet do not have any fixed points. Moreover, the identity map on $\mathbf{R}^n$ is again 1-Lipschitz yet *all* points are fixed points.

Proof. Define a sequence $(x_n)_{n \in \mathbf{N}_0}$ by $x_n := F^n(x_0)$, where $x_0 \in X$ is an arbitrarily chosen initial point. We claim $(x_n)_{n \in \mathbf{N}}$ converges to the unique fixed point of F.

Uniqueness is trivial. Indeed, if $x, x' \in X$ are two fixed points of F,

$$\mathsf{d}(x, x') = \mathsf{d}(F(x), F(x')) \leq \lambda\,\mathsf{d}(x, x'),$$

which forces $\mathsf{d}(x, x') = 0$.

To show the claimed convergence, first note for every $k \in \mathbf{N}$,

$$\mathsf{d}(x_{k+1}, x_k) = \mathsf{d}(T(x_k), T(x_{k-1})) \leq \lambda\,\mathsf{d}(x_k, x_{k-1}) \leq \cdots \leq \lambda^k\,\mathsf{d}(x_0, x_1).$$

Given $n, m \in \mathbf{N}$ with $n < m$, the triangle inequality and a geometric sum yield

$$\mathsf{d}(x_m, x_n) \leq \sum_{k=n}^{m-1} \mathsf{d}(x_{k+1}, x_k)$$

$$
\begin{aligned}
&\leq \sum_{k=n}^{m-1} \lambda^k \, \mathsf{d}(x_0, x_1) \\
&\leq \lambda^n \, \mathsf{d}(x_0, x_1) \sum_{k=0}^{m-n-1} \lambda^k \\
&\leq \frac{\lambda^n}{1-\lambda} \, \mathsf{d}(x_0, x_1).
\end{aligned}
$$

This shows $(x_n)_{n\in\mathbf{N}}$ is a Cauchy sequence. Since (X, d) is complete, it has a limit in X we denote by x. Given any $n \in \mathbf{N}$, we have

$$\mathsf{d}(x, F(x)) \leq \mathsf{d}(x, x_n) + \mathsf{d}(x_n, F(x)) \leq \mathsf{d}(x, x_n) + \lambda \, \mathsf{d}(x_{n-1}, x).$$

Sending $n \to \infty$ and using convergence of $(x_n)_{n\in\mathbf{N}}$ to x shows $x = F(x)$. □

The following simple consequences of the proof of Banach's fixed point theorem are left as an exercise for the interested reader.

Corollary 4.2 (Error bounds)**.** *Retain the hypotheses and the notation from the previous Theorem 4.1. Then for every $x_0 \in X$, the sequence $(x_n)_{n\in\mathbf{N}}$ defined by $x_n := T^n(x_0)$ converges to the unique fixed point x of F.*

Moreover, given any $n \in \mathbf{N}$ we have the **a priori estimate**

$$\mathsf{d}(x_n, x) \leq \frac{\lambda^n}{1-\lambda} \, \mathsf{d}(x_0, x_1)$$

and the **a posteriori estimate**

$$\mathsf{d}(x_n, x) \leq \frac{\lambda}{1-\lambda} \, \mathsf{d}(x_{n-1}, x_n).$$

However, the strict contraction property might be quite restrictive in applications. We thus provide a small selection of different fixed point theorems. Let us start with the following elementary generalization of Banach's fixed point Theorem 4.1 under the stronger assumption that the domain is compact.

Theorem 4.3 (Edelstein)**.** *Let (X, d) be a compact metric space and $F\colon X \to X$ be such that $\mathsf{d}(F(x), F(y)) < \mathsf{d}(x, y)$ for every distinct $x, y \in M$. Then F has a unique fixed point.*

Proof. Uniqueness is shown as in the proof of Theorem 4.1.

To show existence, consider the real-valued assignment $h(x) := \mathsf{d}(x, F(x))$. It is Lipschitz continuous since

$$
\begin{aligned}
\mathsf{d}(x, F(x)) &\leq \mathsf{d}(x, y) + \mathsf{d}(y, F(x)) \\
&\leq \mathsf{d}(x, y) + \mathsf{d}(y, F(y)) + \mathsf{d}(F(y), F(x)) \\
&\leq 2 \, \mathsf{d}(x, y) + \mathsf{d}(y, F(y)),
\end{aligned}
$$

which implies $h(x) - h(y) < 2 \, \mathsf{d}(x, y)$. Exchanging the roles of x and y then yields the desired Lipschitz continuity of h. By continuity the image of $h(X)$ is compact in $\mathbf{R}$. Hence h achieves its minimum on M. Let x_0 denote such a minimizer. If $x_0 = F(x_0)$, we are done. Otherwise, we would get

$$h(F(x_0)) = \mathsf{d}(F(x_0), F(F(x_0))) < \mathsf{d}(x_0, F(x_0)) = h(x_0),$$

which contradicts the minimality of x_0. □

4.2. **Schauder's fixed point theorem.** In this subsection we prove a quite general **existence theorem** for fixed points. Usually the theorem is stated on Banach spaces, but we show a version for LCTVS. We will need the following auxiliary result that relies on the Brouwer fixed point theorem, cf. Theorem 4.9 below, that we will prove later.

Recall given any subset $A \subset X$ of a vector space X, its convex hull is defined by

$$\operatorname{co} A := \Big\{\sum_{i=1}^n \lambda_i\, x_i : n \in \mathbf{N},\, x_1, \dots, x_n \in A,\, \lambda_1, \dots, \lambda_n \in [0,1],\, \sum_{i=1}^n \lambda_i = 1\Big\}.$$

Lemma 4.4 (Nonempty intersections)**.** *Let X be a Hausdorff TVS and $B \subset X$. For every $x \in B$ let $A(x) \subset X$ be a closed set such that*

$$\operatorname{co}\{x_1, \dots, x_n\} \subset \bigcup_{i=1}^n A(x_i)$$

for every finite subset $\{x_1, \dots, x_n\} \subset B$. Then for every finite set $\{y_1, \dots, y_k\} \subset B$, the intersection $\bigcap_{j=1}^k A(y_j)$ is nonempty.

Proof. Assume to the contrary there exist $y_1, \dots, y_k \in B$ such that $\bigcap_{j=1}^k A(y_j)$ is empty. Let W denote the linear span of $\{y_1, \dots, y_k\}$. By Proposition 1.33, the space W is closed. Hence, by assumption on the family A, also $W \cap A(x)$ is closed for every $x \in B$. In particular, given any $z \in W$, we have $z \notin W \cap A(y_j)$ if and only if $\mathsf{d}(z, W \cap A(y_j)) > 0$; here, d is the induced Euclidean norm on W (which generates the subspace topology of W by Proposition 1.33). Note that $\bigcap_{j=1}^k W \cap A(y_j) = \emptyset$, so that for all $c \in C$, where $C := \operatorname{co}\{y_1, \dots, y_k\}$, we have

$$\sum_{j=1}^k \mathsf{d}(c, W \cap A(y_j)) > 0.$$

Define $F \colon C \to C$ by

$$F(c) = \Big[\sum_{j=1}^k \mathsf{d}(c, W \cap A(y_j))\Big]^{-1} \sum_{j=1}^k \mathsf{d}(c, W \cap A(y_j))\, y_j.$$

Then F is continuous since the distance function is continuous. Moreover, C is compact, convex and nonempty. By Brouwer's fixed point theorem the map F has a fixed point $c_0 \in C$. Let $I \subset \{1, \dots, k\}$ denote the set of those indices j such that $\mathsf{d}(c_0, W \cap A(y_j)) > 0$. Then by the assumption on the sets $A(y_j)$ we have

$$c_0 = F(c_0) \in \operatorname{co}\{y_j : j \in I\} \subset \bigcup_{j \in I} A(y_j),$$

which yields a contradiction to the definition of the set I. □

Now we can state and prove Schauder's fixed point theorem on LCTVS, which was first established by Tychonoff. This theorem is usually presented in Banach spaces but it is quite interesting to see that it holds on any LCTVS. We refer to Remark 4.6 for more comments about the version in Banach spaces.

Theorem 4.5 (Schauder–Tychonoff)**.** *Let X be an LCTVS and $K \subset X$ be closed, convex, and nonempty. Let $F \colon K \to K$ be continuous such that $\overline{F(K)}$ is compact. Then F has a fixed point in K.*

Proof. Let U be a convex, balanced, and open neighborhood of the origin. Define $S := \overline{F(K)}$. Since S is compact, there exists a finite subset $\{y_1, \dots, y_n\} \subset S$ such that $S \subset \bigcup_{i=1}^n (y_i + U)$. For each $i \in \{1, \dots, n\}$ we set

$$A(y_i) := \{x \in K : F(x) \notin y_i + U\}.$$

Since $y_i + U$ is open, the set $A(y_i)$ is closed. Note that

$$\bigcap_{i=1}^{n} A(y_i) = \Big\{ x \in K : F(x) \notin \bigcup_{i=1}^{n} (y_i + U) \Big\} = \emptyset.$$

Applying the contraposition of Lemma 4.4 to the set $B := \{y_1, \dots, y_n\}$, we deduce there exists a subset $J \subset \{1, \dots, n\}$ and $x_U \in \operatorname{co}\{y_i : i \in J\}$ with the property $x_U \notin \bigcup_{i \in J} A(y_i)$. Since $x_U \in K$ by convexity, this implies $F(x_U) \in y_i + U$ for every $i \in J$. In turn, there exists $u_i \in U$ with $F(x_U) = y_i + u_i$. Write $x_U = \sum_{i \in J} \lambda_i y_i$, where $\sum_{i \in J} \lambda_i = 1$ and all coefficients are nonnegative. By convexity of U,

$$F(x_U) = \sum_{i \in J} \lambda_i \, F(x_U) = \sum_{i \in J} \lambda_i (y_i + u_i) = x_U + \sum_{j \in J} \lambda_i u_i \in x_U + U.$$

Hence, for fixed U there exists $x_U \in K$ such that $F(x_U) \in x_U + U$. Thanks to the compactness of the reference set S, we may and will therefore assume that net $\{F(x_U) : U$ balanced, convex, open neighborhood of the origin$\}$[36] converges to some x_0 in K. But then x_U converges to x_0. Indeed, let V be a convex, balanced neighborhood of the origin. Then eventually (along a subnet) $x_U \in F(x_U) + U \subset x_0 + V/2 + V/2 \subset x_0 + V$. By continuity of F, it follows that $F(x_0) = x_0$. □

Remark 4.6 (Banach space version of Theorem 4.5). The Schauder–Tychonoff fixed point theorem is usually stated in Banach spaces. For the proof presented here, the only simplification would be that one can replace U with balls of radius $1/m$ and directly construct a sequence instead of a net. However, in Banach spaces there is a simpler proof which avoids the application of Lemma 4.4 and uses Brouwer's fixed point theorem in a more direct way. We present this proof in the appendix and we highlight analogies with the proof on LCTVS discussed here. ■

Next, we present an interesting application of the Schauder fixed point theorem.

Theorem 4.7 (Peano)**.** *Let $(t_0, x_0) \in \mathbf{R} \times \mathbf{R}^n$ and consider the ODE $x'(t) = f(t, x(t))$ and $x(t_0) = x_0$. If the function $f\colon [-a + t_0, a + t_0] \times \overline{B_R(y_0)} \to \mathbf{R}^n$ is continuous for some $a, R > 0$, then there exists $\delta > 0$ such that the above initial value problem has a solution $x\colon [-\delta + t_0, \delta + t_0] \to \mathbf{R}^n$.*

Proof. Exercise 12.3. □

In general, applying the Schauder–Tychonoff fixed point theorem can be quite tricky since one needs to find the set K for which F maps K to K. The following consequence avoids this problem on Banach spaces.

Theorem 4.8 (Schaefer's fixed point theorem)**.** *Let X be a Banach space and let $F\colon X \to X$ be continuous such that $\overline{F(B)}$ is compact for every bounded set $B \subset X$. Assume further that the set $\{x \in X : x = \lambda F(x)$ for some $\lambda \in [0, 1]\}$ is bounded in X. Then F has a fixed point.*

Proof. Exercise 12.2. □

Lecture 12.

[36] The family of balanced, convex, open neighborhoods of the origin can be turned into a directed set with respect to set inclusion. Note that this family is stable under intersections and therefore the net is well defined. If you have not seen nets, you should imagine that they are natural generalization of sequences with uncountably many indices.

4.3. **Brouwer's fixed point theorem.** Here we prove Brouwer's fixed point theorem in the general version we used to show the Schauder–Tychonoff fixed point theorem. There are several proofs, which use quite difficult approaches. We use here an analytical one which only requires a change of variables for the occurring integrals. Let us first state the theorem.

Theorem 4.9 (Brouwer)**.** *Let $K \subset \mathbf{R}^n$, where $n \in \mathbf{N}$, be convex, compact, and nonempty. Moreover, let $f\colon K \to K$ be continuous. Then f has a fixed point in K.*

Remark 4.10 (From Euclidean space to finite dimensions)*.* In the proof of Lemma 4.4 we used the above theorem in a finite-dimensional subspace of a Hausdorff TVS. By Proposition 1.33, such spaces are linearly homeomorphic to $\mathbf{R}^n$ for some $n \in \mathbf{N}$. All these identifications preserve convexity and compactness, so that the restriction to maps defined on subsets of $\mathbf{R}^n$ imposes no restriction. ■

We will prove Theorem 4.9 in several stages. First we reduce the analysis to the case when K is the closed unit ball $\overline{B}_1(0)$.

Lemma 4.11 (Reduction lemma)**.** *Let us assume that every continuous function $f\colon \overline{B}_1(0) \to \overline{B}_1(0)$ has a fixed point. Then the statement from Theorem 4.9 holds.*

Proof. First note the assumption implies every continuous map $f\colon \overline{B}_R(0) \to \overline{B}_R(0)$ has a fixed point, where $R > 0$ is arbitrary. Indeed, given such an f, consider the map $f_R\colon \overline{B}_1(0) \to \overline{B}_1(0)$ defined by $f_R(x) := f(Rx)/R$, which has a fixed point $x_0 \in \overline{B}_1(0)$ by assumption. Then Rx_0 is clearly a fixed point of f.

Now consider a convex, compact, and nonempty set $K \subset \mathbf{R}^n$ and a continuous function $f\colon K \to K$. Take a finite or countable dense subset $\{a_i : i \in I\} \subset K$. Given any $i \in I$, define the continuous cutoffs $\varphi_i\colon \mathbf{R}^n \setminus K \to [0,1]$ by

$$\varphi_i(x) = \max\Big\{2 - \frac{|x - a_i|}{\mathsf{d}(x,K)}, 0\Big\}.$$

We then define $\widetilde{f}\colon \mathbf{R}^n \to \mathbf{R}^n$ by

$$\widetilde{f}(x) = \begin{cases} f(x) & \text{if } x \in K,\\ \Big[\sum_{i\in I} 2^{-i}\varphi_i(x)\Big]^{-1} \sum_{i\in I} 2^{-i}\varphi_i(x)\, f(a_i) & \text{otherwise} \end{cases}$$

Clearly, $\widetilde{f}$ extends f. Moreover, $\widetilde{f}$ is continuous on the open set $\mathbf{R}^n \setminus K$. We claim it is also continuous at every given $x \in K$. To this aim, it suffices to consider a sequence $(x_n)_{n\in\mathbf{N}}$ in $\mathbf{R}^n \setminus K$ converging to x. Given $\varepsilon > 0$, the continuity of f on K implies there exists $\delta > 0$ such that $|f(x) - f(a_i)| < \varepsilon$ for all $i \in I$ with $|x - a_i| < \delta$. Clearly $\varphi_i(x_n) = 0$ whenever $|x_n - a_i| \geq 2\,\mathsf{d}(x_n, K)$. Since last term tends to 0 as $n \to \infty$, for $n \in \mathbf{N}$ large enough we only need to consider those $i \in I$ such that $|x - a_i| < \delta$. For those, an elementary estimate shows $|\widetilde{f}(x_n) - \widetilde{f}(x)| \leq \varepsilon$, which proves continuity. In total, this shows $\widetilde{f}$ is a continuous extension of f. Finally, we note that since K is closed and convex, by the definition of the extension we get

$$\widetilde{f}(\mathbf{R}^n) \subset \overline{\mathrm{co}(f(K))} \subset K.$$

Consider now a closed ball $\overline{B}_R(0)$ that contains K. Then by the first part of the proof, $\widetilde{f}$ has a fixed point in $\overline{B}_R(0)$. But we know this fixed point has to belong to K. Hence f has a fixed point as well. □

Remark 4.12 (Alternative proof of Lemma 4.11)*.* One could alternatively try to find a homeomorphism $g\colon \overline{B}_1(0) \to K$. Indeed, assume such a map exists. Then $F := g^{-1} \circ f \circ g\colon \overline{B}_1(0) \to \overline{B}_1(0)$ is continuous and therefore there exists a fixed point x_0 of F by assumption. This easily implies $g(x_0) \in K$ is a fixed point for f. Note when $K = B_R(0)$ we have $g(x) = Rx$, as shown during the above proof.

To construct the map g, we need an extra assumption on K. This is because K can be a lower dimensional object in general. For instance, we cannot hope to construct a homeomorphism from the circle (which is intrinsically convex and compact) to the unit disk. Therefore, let us assume there exist $x_0 \in \mathbf{R}^n$ and $\delta > 0$ such that $B_\delta(x_0) \subset K$. Up to translation, we can may and will assume $x_0 = 0$. Then, set $p(x) := \inf\{t \geq 0 : tx \notin K\}$ and $p(0) := 0$ (note the similarity to the Minkowski functional). Since K is convex, it is not hard to show p is a norm and K is the unit ball in this norm. Now, define $h\colon K \to \overline{B}_1(0)$ by

$$h(x) := \begin{cases} \dfrac{p(x)}{|x|}\, x & \text{if } x \neq 0, \\ 0 & \text{otherwise,} \end{cases}$$

where $|\cdot|$ is the Euclidean norm. This map can have problems of continuity at the origin per se. However, since all norms in $\mathbf{R}^n$ are equivalent, h is indeed continuous at 0. (The continuity at other points follows by continuity of the involved norms.) Consequently, we also know h is a homeomorphism and we define $g := h^{-1}$. ■

Having reduced the proof to the unit ball, we now formulate an equivalent statement. We will only prove that this theorem implies Brouwer's fixed point theorem. The reverse implication will be part of the exercises.

Lemma 4.13 (Sufficient condition for Lemma 4.11)**.** *Assume there exists no continuous map $R\colon \overline{B}_1(0) \to \partial B_1(0)$ such that $R(x) = x$ for every $x \in \partial B_1(0)$. Then every continuous map $f\colon \overline{B}_1(0) \to \overline{B}_1(0)$ has a fixed point.*

Proof. Assume by contradiction there exists $f\colon \overline{B}_1(0) \to \overline{B}_1(0)$ such that $f(x) \neq x$ for every $x \in \overline{B}_1(0)$. Set $g(x) := x - f(x)$ and $h_x(t) := |x + tg(x)|^2 - 1$ for $x \in \overline{B}_1(0)$ and $t \in \mathbf{R}$. Note $h_x(0) \leq 0$ and $\lim_{t\to\infty} h_x(t) = \infty$ since $g(x) \neq 0$. Hence there exists $t_x \geq 0$ such that $h_x(t_x) = 0$. Since $t \mapsto h_x(t)$ is a second order polynomial, the number t_x can be calculated explicitly in terms of x and $g(x)$ (in particular it is unique) and one can show the dependency of t_x on $x \in \overline{B}_1(0)$ is continuous. Define then the continuous assignment $R(x) := x + t_x g(x)$, where $x \in \overline{B}_1(0)$. By construction, $|R| = 1$ everywhere, so that $R\colon \overline{B}_1(0) \to \partial B_1(0)$. Moreover, for any $x \in \partial B_1(0)$ we have $h_x(0) = 0$, so that uniqueness of t_x implies $t_x = 0$. Hence $R(x) = x$ on $\partial B_1(0)$, which gives a contradiction. □

A map R as in the previous lemma is often called a retraction. We are thus left to show that there exists no retraction from $\overline{B}_1(0)$ to its boundary. We first prove there exists no C^1-retraction.

Lemma 4.14 (C^1-retraction theorem)**.** *There exists no function $f\colon \overline{B}_1(0) \to \partial B_1(0)$ that is continuously differentiable on a neighborhood of $\overline{B}_1(0)$ which obeys $f(x) = x$ for every $x \in \partial B_1(0)$.*

Proof. We argue again by contradiction and assume such a map f exists. For $t \in [0,1]$ set $f_t = (1-t)\,\mathrm{Id} + t\,f$. By compactness of $\overline{B}_1(0)$, the function $f - \mathrm{Id}$ is Lipschitz continuous on $\overline{B}_1(0)$ with Lipschitz constant $c \geq 1$. As a consequence, given any $x, y \in \overline{B}_1(0)$ this shows

$$|f_t(x) - f_t(y)| \geq |x - y| - t\,|(f - \mathrm{Id})(x) - (f - \mathrm{Id})(y)| \geq (1 - ct)|x - y|.$$

In particular, for every $t \in [0, 1/c)$ the map f_t is injective and the inverse function is Lipschitz continuous as well. Moreover, the assignment $(t, x) \mapsto \det \mathrm{D}_x f_t(x)$ is continuous and equal to one for $t = 0$. Hence there exists $\varepsilon > 0$ such that for all $(t, x) \in [0, \varepsilon] \times \overline{B}_1(0)$, we have $\det \mathrm{D}_x f_t(x) > 0$[37]. By the inverse function

[37]A priori, for every $x \in \overline{B}_1(0)$ there exists a neighborhood of the form $[0, \varepsilon_x] \times B_\delta(x)$ such that the statement holds on this set. The sets $B_\delta(x)$ form an open cover of $\overline{B}_1(0)$. By compactness, we find a finite subcover and thence a common $\varepsilon > 0$ that works for all $x \in \overline{B}_1(0)$.

theorem, we conclude for all $t \in [0,\varepsilon]$ the image $f_t(B_1(0))$ is open and moreover, by convexity, it is a subset of $\overline{B}_1(0)$. Let us prove $f_t(B_1(0)) = B_1(0)$ for sufficiently small $t > 0$. By the arguments hitherto, we know $f_t(B_1(0)) \subset B_1(0)$. Note that $f_t(\partial B_1(0)) = \partial B_1(0)$, which implies

$$B_1(0) \setminus f_t(B_1(0)) = B_1(0) \setminus f_t(\overline{B}_1(0)).$$

By continuity of f_t, the set $f_t(\overline{B}_1(0))$ is compact and in particular closed. We conclude $B_1(0) \setminus f_t(B_1(0))$ is also open. Therefore, this enables us to decompose $B_1(0) = f_t(B_1(0)) \cup (B_1(0) \setminus f_t(B_1(0)))$ which is the union of two open and disjoint sets. Since $B_1(0)$ is connected, we deduce $B_1(0) \setminus f_t(B_1(0))$ has to be empty; and indeed, $f_t(B_1(0)) = B_1(0)$.

Lastly, given any $t \in [0,1]$ we consider the assignment

$$v(t) := \int_{B_1(0)} \det \mathrm{D}_x f_t(x)\,\mathrm{d}x.$$

The change of variables formula implies $v(t) = \mathcal{L}^n[f_t(B_1(0))] = \mathcal{L}^n[B_1(0)]$ for all $t \in [0, \min\{\varepsilon, 1/c\})$. Since $v(t)$ is a polynomial, we know $v(t)$ equals a positive constant on $[0,1]$. However, for $t = 1$ we have $f_t = f$ and since $|f(x)|^2 = 1$ for all $x \in B_1(0)$, we know $\mathrm{D}f(x)\,f(x) = 0$. This means $\mathrm{D}f(x)$ has a nontrivial kernel. In particular, $\det \mathrm{D}f(x) = 0$ for all $x \in B_1(0)$, a contradiction. □

Finally, we remove the smoothness assumption by an approximation argument.

Theorem 4.15 (Retraction theorem)**.** *There exists no continuous function $f\colon \overline{B}_1(0) \to \partial B_1(0)$ such that $f(x) = x$ for every $x \in \partial B_1(0)$.*

Proof. Again we argue by contradiction and construct a C^1-retraction, which is absurd in view of the previous lemma. First extend the given retraction to $\mathbf{R}^n$ by setting $f(x) := x$ whenever $|x| > 1$. This extension is continuous. By a convolution argument, we find a sequence $(f_k)_{k\in\mathbf{N}}$ of functions in $C^\infty(\mathbf{R}^n, \mathbf{R}^n)$ such that $f_k \to f$ locally uniformly on $\mathbf{R}^n$ as $k \to \infty$. Let $h \in C_c^\infty((-1,1);[0,1])$ such that $h(0) = 1$ and define $h_k\colon \mathbf{R}^n \to \mathbf{R}$ by $h_k(x) := h(k|x|^2 - k)$. Then we have $h_k(x) = 1$ for every $x \in \partial B_1(0)$, while $h_k(x) \to 0$ as $k \to \infty$ whenever $|x| \neq 1$. Define then

$$g_k(x) := h_k(x)\,x + (1 - h_k(x))\,f_k(x).$$

Note that $g_k \in C^\infty(\mathbf{R}^n, \mathbf{R}^n)$. Moreover, the above construction yields $g_k(x) = x$ for $x \in \partial B_1(0)$ since $h_k = 1$ on the unit sphere. The idea now is to normalize g_k by considering $g_k/|g_k|$. To this end, we show that for every sufficiently large $k \in \mathbf{N}$, we have $|g_k| \geq c$ on $\overline{B}_1(0)$. Clearly this condition allows us to normalize g_k also in a (k-dependent) neighborhood of $\overline{B}_1(0)$ and then we are done.

Assume by contradiction there is a subsequence $(k_j)_{j\in\mathbf{N}}$ with $\min_{\overline{B}_1(0)} |g_{k_j}| \to 0$ as $j \to \infty$. There exists a sequence $(x_j)_{j\in\mathbf{N}}$ in $\overline{B}_1(0)$ with $g_{k_j}(x_j) \to 0$ as $j \to \infty$ and — up to a further subsequence — $h_{k_j}(x_j) \to t \in [0,1]$ and $x_j \to x_0 \in \overline{B}_1(0)$. By locally uniform convergence of $(f_k)_{k\in\mathbf{N}}$ to f, we get $0 = tx_0 + (1-t)f(x_0)$. Note that $|x_0| = 1$ is impossible since otherwise $f(x_0) = x_0$ and therefore $x_0 = 0$, which is absurd. But if $|x_0| < 1$, then for $j \in \mathbf{N}$ large enough we have $h_{k_j}(x_j) = 0$, which implies $t = 0$ and then $f(x_0) = 0$, which gives again a contradiction. Hence for $k \in \mathbf{N}$ large enough, the map $g_k/|g_k|$ is a well-defined retraction that is regular in a neighborhood of $\overline{B}_1(0)$. This contradicts the previous lemma. □

5. Gradient flows in Hilbert spaces

Recall that the gradient flow in $\mathbf{R}^n$, where $n \in \mathbf{N}$, of a (say) smooth potential $V\colon \mathbf{R}^n \to \mathbf{R}$ starting at $o \in \mathbf{R}^n$ is a continuous curve $x\colon \mathbf{R}_+ \to \mathbf{R}^n$ which is differentiable in $(0,\infty)$ and which obeys $x_t' = -\nabla V(x_t)$ for every $t > 0$ and $x_0 = o$. The following brief survey provides an introduction into gradient flows on general

Hilbert spaces. It was introduced by Brézis [2]. Although you will find all relevant results that are used to date in this book, overall it is rather outdated; for a modern account on gradient flows in general metric spaces, we refer to the book of Ambrosio–Gigli–Savaré [1].

Gradient flows have many simple properties (good solution theory, quantitative estimates, equilibrium points, etc.) that make them interesting. They arise in numerous geometries and are a ubiquitous tool in modern analysis and its applications to probability theory, geometry, machine learning, artificial intelligence, etc. One particularly rich area where they have become central is optimal transport. This trend started from the papers of Jordan–Kinderlehrer–Otto and Otto. As a toy example, consider a parabolic PDE of the form $\partial_t u_t = \mathsf{L} u_t$ on $\mathbf{R}^n$, where L is usually an elliptic second-order differential operator acting spatially (e.g. the Laplacian). By endowing the space $\mathcal{P}(\mathbf{R}^n)$ of probability measures over $\mathbf{R}^n$ with a certain optimal transport geometry, [7,8] were able to show this PDE can be interpreted as a gradient flow $\partial_t \mu_t = -\nabla E(\mu_t)$ on $\mathcal{P}(\mathbf{R}^n)$ (understood appropriately), where E is an entropy functional. The correspondence is given by $\mu_t = u_t\,\mathcal{L}^n$. The point of this identification is that one can trade a PDE on a finite-dimensional space for an ODE (which are often simpler to study than PDEs) on an infinite-dimensional space. This is no Hilbert space theory, but working on Hilbert spaces is simpler and conveys many of the key ideas used in this metric context as well.

In the first part, we will clarify the meaning of the "gradient" of a functional in a Hilbert space. This happens by means of convex analysis, for basics of which we refer to Rockafellar [9]. In the second part, we outline the general theory of existence, uniqueness, and fundamental properties.

5.1. **Convexity and subdifferentials.** Let H be a real Hilbert space. Let $\langle\cdot,\cdot\rangle$ denote the inherent scalar product and $\|\cdot\| := \sqrt{\langle\cdot,\cdot\rangle}$ the induced norm. Let $E\colon H \to \mathbf{R}_+ \cup \{\infty\}$ be an "energy" functional which we assume to be

- **convex**, i.e. $E((1-t)x + ty) \le (1-t)E(x) + tE(y)$ for every $x, y \in H$ and every $t \in [0,1]$ and
- **lower semicontinuous**, i.e. whenever $(x_n)_{n\in\mathbf{N}}$ converges to $x \in H$,
$$E(x) \le \liminf_{n\to\infty} E(x_n).$$
For more on lower semicontinuity, we refer to Definition B.7 et seq.

Let $\mathcal{D}(E) := \{x \in H : E(x) \in \mathbf{R}\}$ denote the convex **domain** of E. To create a nonpathological theory, we will assume $\mathcal{D}(E)$ is nonempty.

Remark 5.1 (Euclidean Dirichlet energy). The basic example we will be interested in is the following. On $H := L^2(\mathbf{R}^n, \mathcal{L}^n)$ we consider the functional $E\colon L^2(\mathbf{R}^n,\mathcal{L}^n) \to \mathbf{R}_+ \cup \{\infty\}$ defined through

$$E(u) := \begin{cases} \dfrac{1}{2}\displaystyle\int_{\mathbf{R}^n} |\nabla u|^2 \,\mathrm{d}\mathcal{L}^n & \text{if } u \in W^{1,2}(\mathbf{R}^n),\\ \infty & \text{otherwise.}\end{cases}$$

It is clearly convex. Lower semicontinuity with respect to L^2-convergence is straightforward and left as an exercise. You may want to use Corollary B.8. ■

The following is a general notion from convex analysis generalizing the concept of a "differential" for a convex function, even though a convex function is in general not differentiable everywhere[38].

[38]However, convex functions on $\mathbf{R}^n$ enjoy good regularity properties. They are in fact twice differentiable $\mathcal{L}^n$-a.e. and locally Lipschitz continuous on the interior of their domain. On $\mathbf{R}$, the latter fact can easily be derived from the monotonicity of difference quotients implied by convexity. An analogous principle holds in $\mathbf{R}^n$; compare with Proposition 5.4.

Definition 5.2 (Subdifferential)**.** *The* **subdifferential** *of E at a point $x \in \mathcal{D}(E)$ is the set $\partial^- E(x)$ of all $x^* \in H$ such that for every $y \in H$,*

$$\langle x^*, y - x\rangle \leq E(y) - E(x).$$

We also set $\partial^- E(x) := \emptyset$ provided $x \notin \mathcal{D}(E)$.

We write $\mathcal{D}(\partial^- E)$ for the set of all $x \in H$ such that $\partial^- E(x) \neq \emptyset$. In particular, by the above convention we have $\mathcal{D}(\partial^- E) \subset \mathcal{D}(E)$.

By definition, given any $x \in \mathcal{D}(E)$ we have $0 \in \partial^- E(x)$ if and only if E attains its minimum at x.

In more pictorial words, the subdifferential of E at $x \in \mathcal{D}(E)$ is the set of all slopes of tangents that touch the graph of E from below at x.

Example 5.3 (Absolute value). On the elementary Hilbert space $H := \mathbf{R}$ consider the convex and continuous function $E\colon \mathbf{R} \to \mathbf{R}_+$ given by $E(x) := |x|$. Check as an exercise that its subdifferentials have the following form for every $x \in \mathbf{R}$:

$$\partial^- E(x) = \begin{cases} \{1\} & \text{if } x > 0,\\ [-1,1] & \text{if } x = 0,\\ \{-1\} & \text{otherwise.}\end{cases} \qquad \blacksquare$$

Proposition 5.4 (Properties of the subdifferential)**.** *The following hold.*

(i) **Monotonicity.** *The multivalued map $\partial^- E\colon H \to 2^H$ constitutes a monotone operator*[39]*. That is, for every $x, y \in H$, every $x^* \in \partial^- E(x)$, and every $y^* \in \partial^- E(y)$,*

$$\langle y - x, y^* - x^*\rangle \geq 0.$$

(ii) **Strong-weak closure.** *The graph of $\partial^- E$ is strongly-weakly closed in H^2. That is, assume $(x_n)_{n\in\mathbf{N}}$ is a sequence in H converging to $x \in H$. Let $(x_n^*)_{n\in\mathbf{N}}$ be a sequence of elements $x_n^* \in \partial^- E(x_n)$ which converges weakly in H to $x^* \in H$. Then $x^* \in \partial^- E(x)$.*

Proof. Exercise 13.4. □

Lecture 13.

5.2. **Existence, uniqueness, and properties of gradient flows.** After having clarified how the quantity "∇E" of the targeted gradient flow equation should be understood in general Hilbert spaces, we now turn to the precise definition of the gradient flow equation itself.

As a first step, we clarify the targeted regularity in time. For more details on absolutely continuous H-valued functions, we refer to §E.

Definition 5.5 (Local absolute continuity)**.** *We will call a curve $x\colon (0,\infty) \to H$* **locally absolutely continuous** *if its restriction to every compact subset of $(0,\infty)$ is 1-absolutely continuous according to Definition E.5. Equivalently, for every compact interval $I \subset (0,\infty)$, there is $f_I \in L^1(I, \mathcal{L}^1)$ such that for every $s, t \in I$ with $s < t$,*

$$\big\|x(t) - x(s)\big\| \leq \int_s^t f_I(r)\,\mathrm{d}r.$$

Remark 5.6 (Fundamental theorem of calculus). Combining Proposition E.6 and Theorem E.7, for every locally absolutely continuous function $x\colon (0,\infty) \to H$ and $\mathcal{L}^1$-a.e. $t > 0$, the following derivative exists:

$$x'(t) := \lim_{h\to 0} \frac{x(t+h) - x(t)}{h}.$$

[39]Historically, this fact embeds the theory of gradient flows for convex and lower semicontinuous functionals into the theory of gradient flows for maximal monotone operators by Brézis [2].

It belongs to $L^1(I;\mathcal{L}^1)$ for every compact interval $I \subset (0,\infty)$.

Moreover, for every $s,t \in [0,1]$ with $s < t$, we have

$$x(t) - x(s) = \int_s^t x'(r)\,\mathrm{d}r$$

in the sense of Bochner integration, cf. §E. Moreover, the map x' in this formula is uniquely determined up to modifications on $\mathcal{L}^1$-null sets. ■

Definition 5.7 (Gradient flow trajectory)**.** *A* **gradient flow trajectory** *of E is a continuous curve $x\colon \mathbf{R}_+ \to H$ which is locally absolutely continuous on $(0,\infty)$, obeys $x(t) \in \mathcal{D}(\partial^- E)$ for every $t > 0$, and satisfies the following for $\mathcal{L}^1$-a.e. $t > 0$:*

$$-x'(t) \in \partial^- E(x(t)). \tag{5.1}$$

For now, the gradient flow equation on Hilbert spaces holds only for a.e. time and is defined by a differential inclusion rather than a genuine identity. Both properties can be improved, as stated in Theorem 5.12 below.

Typically, solving the gradient flow equation with gradient flow trajectory is coupled with fixing an initial condition $o \in H$, i.e. requiring $x(0) = o$.

5.2.1. *Existence.* The goal of this part is Theorem 5.11. It states general existence of gradient flows in Hilbert spaces.

To this aim, we prepare some material. Existence will be based on a numerical scheme called *minimizing movement scheme*, which was introduced by the Italian mathematician De Giorgi. Given any $x \in H$ and a step size $\tau > 0$, we define the convex and lower semicontinuous functional $F_{x,\tau}\colon H \to \mathbf{R}_+ \cup \{\infty\}$ by

$$F_{x,\tau} := E + \frac{\|\cdot - x\|^2}{2\tau}.$$

We will then generate a sequence $(x^\tau_{(k)})_{k\in\mathbf{N}_0}$ as follows. Define $x^\tau_{(0)} := o$, the given initial point. Inductively, given $x^\tau_{(k)}$ for $k \in \mathbf{N}_0$, we choose a point

$$x^\tau_{(k+1)} \in \operatorname{argmin}\big\{F_{x^\tau_{(k)},\tau}(y) : y \in H\big\}. \tag{5.2}$$

We interpret $x^\tau_{(k)}$ as the point interpolated by a piecewise affine curve $x^\tau\colon \mathbf{R}_+ \to H$ at $t = k\tau$. Observe that formally — and rigorously if e.g. $H = \mathbf{R}^n$ for some $n \in \mathbf{N}$ and E is continuously differentiable —, (5.2) implies

$$0 = \nabla F_{x^\tau_{(k)},\tau}(x^\tau_{(k+1)}) = \nabla E(x^\tau_{(k+1)}) + \frac{x^\tau_{(k+1)} - x^\tau_{(k)}}{\tau} \tag{5.3}$$

for every $k \in \mathbf{N}$. This is why we hope and expect x^τ converges to a gradient flow trajectory of E as $\tau \to 0$ in a sense yet to be specified.

The subsequent lemma makes the second identity of (5.3) rigorous in general Hilbert spaces (in terms of subdifferentials).

Lemma 5.8 (Correspondence of subdifferentials)**.** *For every $x \in H$ and every $\tau > 0$, we have $\mathcal{D}(\partial^- F_{x,\tau}) = \mathcal{D}(\partial^- E)$, and for every $y \in H$,*

$$\partial^- F_{x,\tau}(y) = \partial^- E(y) + \frac{y - x}{\tau}.$$

Proof. Exercise 14.1. □

Now we show the scheme (5.2) is well-defined.

Proposition 5.9 (Existence of minimizers)**.** *Given any $x \in H$ and any $\tau > 0$, there exists a unique minimizer $x_\tau \in H$ of the functional $F_{x,\tau}$.*

Moreover, we have $x_\tau \in \mathcal{D}(\partial^- E)$ and

$$-\frac{x_\tau - x}{\tau} \in \partial^- E(x_\tau).$$

Proof. To show uniqueness, assume $y, y' \in H$ are two different minimizers of $F_{x,\tau}$. Since $\mathcal{D}(E)$ is nonempty, we necessarily have $y, y' \in \mathcal{D}(E)$. Moreover, note that the squared norm on a Hilbert space is strictly convex. Setting $y'' := y/2 + y'/2$, together with convexity of E these observations imply

$$\begin{aligned} F_{x,\tau}(y'') &< \frac{1}{2}\Big[E(y) + \frac{\|y-x\|^2}{2\tau}\Big] + \frac{1}{2}\Big[E(y') + \frac{\|y'-x\|^2}{2\tau}\Big] \\ &= \min\big\{F_{x,\tau}(z) : z \in H\big\}, \end{aligned}$$

which is of course absurd.

Existence is a standard application of the so-called direct method of calculus of variations. Since $\mathcal{D}(E)$ is nonempty and $F_{x,\tau}$ is positive, the minimum of $F_{x,\tau}$ is a real number. Hence, there exists a sequence $(x_n)_{n\in\mathbf{N}}$ with the property that $\lim_{n\to\infty} F_{x,\tau}(x_n) = \min\{F(z) : z \in H\}$. We claim the existence of $R > 0$ such that $\{x_n : n \in \mathbf{N}\} \subset \overline{B}_R(0)$. Indeed, along a suitable subsequence, $(\|x_{n_j}\|)_{j\in\mathbf{N}}$ would otherwise diverge to infinity. The triangle inequality yields

$$\liminf_{j\to\infty} \|x_{n_j} - x\| \geq \liminf_{j\to\infty} \|x_{n_j}\| - \|x\| = \infty.$$

By definition of $F_{x,\tau}$ and nonnegativity of E, we get $\lim_{j\to\infty} F_{x,\tau}(x_{n_j}) = \infty$, in contradiction to the finiteness of this limit. Thus, by the Banach–Alaoglu–Bourbaki Theorem C.1, there exists $m \in \overline{B}_R(0)$ such that $(x_n)_{n\in\mathbf{N}}$ converges weakly to m, up to a subsequence we will not relabel. We finally claim m is a minimizer of F. Indeed, recall from Corollary B.8 that $F_{x,\tau}$ is sequentially lower semicontinuous with respect to the weak topology. The same applies to the norm (and hence its square) by basic properties of weak convergence. Hence,

$$\min\{F(z) : z \in H\} \leq F_{x,\tau}(m) \leq \liminf_{n\to\infty} F_{x,\tau}(x_n) = \min\{F(z) : z \in H\}.$$

This forces equality to hold throughout.

Finally, since m minimizes $F_{x,\tau}$, we have $m \in \mathcal{D}(\partial^- F_{x,\tau})$ and $0 \in \partial^-(F_{x,\tau})$. By Lemma 5.8, this translates into $m \in \mathcal{D}(\partial^- E)$ and

$$0 \in \partial^- E(m) + \frac{m-x}{\tau}. \qquad \square$$

Remark 5.10 (Minimal norms). With a similar procedure by minimizing the strictly convex norm $\|\cdot\|$ combined with Proposition 5.4, one can prove the following. If $x \in \mathcal{D}(\partial^- E)$, there exists a unique element $x^* \in \partial^- E(x)$ whose norm is minimal among all elements in $\partial^- E(x)$. ■

Theorem 5.11 (Existence of gradient flow trajectories)**.** *Let $o \in \overline{\mathcal{D}(E)}$ be given. Then there exists a gradient flow trajectory of E starting at o.*

Proof. We will present the proof for $o \in \mathcal{D}(E)$. The more general case $o \in \overline{\mathcal{D}(E)}$ will be addressed in Exercise 14.3.

Given any $\tau > 0$, let $(x^\tau_{(k)})_{k\in\mathbf{N}}$ be the minimizing movement sequence with step size τ constructed by (5.2) with initial datum $x^\tau_{(0)} := o$. Here, the (unique) existence of a minimizer from (5.2) is provided by Proposition 5.9. We define $x^\tau : \mathbf{R}_+ \to H$ by $x^\tau(k\tau) := x^\tau_{(k)}$ with affine interpolation on $(k\tau, (k+1)\tau)$, where $k \in \mathbf{N}_0$. More precisely, it is explicitly given by the formula

$$x(t) = x^\tau_{(\lfloor t/\tau\rfloor)} + \Big[\frac{x^\tau_{(\lfloor t/\tau\rfloor+1)} + x^\tau_{(\lfloor t/\tau\rfloor)}}{\tau}\Big]\Big[t - \tau\Big\lfloor\frac{t}{\tau}\Big\rfloor\Big].$$

The affine interpolation ensures that for every such k and every $t \in (k\tau, (k+1)\tau)$,

$$(x^\tau)'(t) = \frac{x^\tau_{(k+1)} - x^\tau_{(k)}}{\tau}. \tag{5.4}$$

By the minimizing property from (5.2),

$$E(x^\tau_{(k+1)}) + \frac{\big\Vert x^\tau_{(k+1)} - x^\tau_{(k)}\big\Vert^2}{2\tau} \leq E(x^\tau_{(k)}).$$

In particular, by a telescopic sum and since $o \in \mathcal{D}(E)$, we observe

$$\frac{1}{2}\int_0^\infty \big\Vert (x^\tau)'(t)\big\Vert^2 \,\mathrm{d}t = \sum_{k\in\mathbf{N}_0} \frac{\big\Vert x^\tau_{(k+1)} - x^\tau_{(k)}\big\Vert^2}{2\tau} \leq E(o) < \infty. \tag{5.5}$$

Now we claim there exists a continuous curve $x\colon \mathbf{R}_+ \to H$ which is the uniform limit of x^τ as $\tau \to 0$. Since the space of H-valued continuous maps on $\mathbf{R}_+$ endowed with the supremum norm is complete, it suffices to show

$$\lim_{\substack{\tau\to 0,\\ \eta\to 0}} \sup_{t\in\mathbf{R}_+} \big\Vert x^\tau(t) - x^\eta(t)\big\Vert = 0. \tag{5.6}$$

Let $\tau, \eta > 0$ be fixed. Let $t \in ((k-1)\tau, k\tau) \cap ((k'-1)\eta, k'\eta)$, where $k, k' \in \mathbf{N}$. Then

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\frac{\big\Vert x^\tau(t) - x^\eta(t)\big\Vert^2}{2} &= \big\langle (x^\tau)'(t) - (x^\eta)'(t), x^\tau(t) - x^\eta(t)\big\rangle\\ &= \big\langle (x^\tau)'(t) - (x^\eta)'(t), x^\tau(k\tau) - x^\eta(k'\eta)\big\rangle\\ &\qquad + \big\langle (x^\tau)'(t) - (x^\eta)'(t), (x^\tau(t) - x^\tau(k\tau))\big\rangle\\ &\qquad - \big\langle (x^\tau)'(t) - (x^\eta)'(t), (x^\eta(t) - x^\eta(k'\eta))\big\rangle.\end{aligned}$$

Here we used the standard differentiation formula for the squared norm on Hilbert spaces and a simple zero addition. Combining (5.4) with the second part of Proposition 5.9 and Proposition 5.4, we infer

$$\big\langle (x^\tau)'(t) - (x^\eta)'(t), x^\tau(k\tau) - x^\eta(k'\eta)\big\rangle \leq 0.$$

For the remaining terms, the Cauchy–Schwarz inequality and (5.4) again yield

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\frac{\big\Vert x^\tau(t) - x^\eta(t)\big\Vert^2}{2} &\leq \big[\big\Vert (x^\tau)'(t)\big\Vert + \big\Vert (x^\eta)'(t)\big\Vert\big]\,\big[\tau\,\big\Vert (x^\tau)'(t)\big\Vert + \eta\,\big\Vert (x^\eta)'(t)\big\Vert\big]\\ &= \tau\,\big\Vert (x^\tau)'(t)\big\Vert^2 + \eta\,\big\Vert (x^\eta)'(t)\big\Vert^2\\ &\qquad + (\tau+\eta)\,\big\Vert (x^\tau)'(t)\big\Vert\,\big\Vert (x^\eta)'(t)\big\Vert\\ &\leq \big\Vert (x^\tau)'(t)\big\Vert^2\Big[\tau + \frac{\tau+\eta}{2}\Big] + \big\Vert (x^\eta)'(t)\big\Vert^2\Big[\eta + \frac{\tau+\eta}{2}\Big].\end{aligned}$$

In the last step, we used Young's inequality. Integrating the preceding estimate over $[0, T]$, where $T > 0$ is fixed, and employing (5.5) leads to

$$\frac{\big\Vert x^\tau(T) - x^\eta(T)\big\Vert^2}{2} \leq 2\,E(o)\,(\tau+\eta).$$

Since the right-hand side does not depend on T, this shows (5.6).

Next, we upgrade the regularity of the curve $x\colon \mathbf{R}_+ \to H$ thus obtained. Since the estimate (5.6) is independent of τ, the set $\{(x^\tau)' : \tau > 0\}$ is norm bounded in the Hilbert space $L^2(\mathbf{R}_+; H)$, cf. Definition E.3. By the Banach–Alaoglu–Bourbaki Theorem C.1, there are $v \in L^2(\mathbf{R}_+; H)$ and a sequence $(\tau_n)_{n\in\mathbf{N}}$ in $(0,\infty)$ decreasing to zero such that $(x^{\tau_n})' \rightharpoonup v$ in $L^2(\mathbf{R}_+; H)$ as $n \to \infty$. Let $s, t \in (0, \infty)$ with $s < t$. By pairing the integral in question against an arbitrary vector from H, it is not difficult to prove that, with respect to the weak topology of H,

$$\lim_{n\to\infty} \int_s^t (x^{\tau_n})'(r)\,\mathrm{d}r = \int_s^t v(r)\,\mathrm{d}r.$$

On the other hand, by the above paragraph and the construction of our piecewise affine interpolants we know

$$\lim_{n\to\infty}\int_s^t (x^{\tau_n})'(r)\,\mathrm{d}r = \lim_{n\to\infty}\big[x^{\tau_n}(t) - x^{\tau_n}(s)\big] = x(t)-x(s).$$

Uniqueness of weak limits implies

$$x(t)-x(s) = \int_s^t v(r)\,\mathrm{d}r.$$

Using (E.2), this implies x is locally absolutely continuous on $(0,\infty)$ — in fact, it belongs to $\mathrm{AC}^2(\mathbf{R}_+,H)$ according to Definition E.5. In particular, by Remark 5.6 the map x is differentiable $\mathcal{L}^1$-a.e. with $x'(t)=v(t)$ for $\mathcal{L}^1$-a.e. $t>0$.

It remains to prove x obeys the differential inclusion (5.1). We claim for every $y\in H$ and every $t_0,t_1\in\mathbf{R}_0$ with $t_0<t_1$, we have

$$\int_{t_0}^{t_1} E(x(t))\,\mathrm{d}t + \int_{t_0}^{t_1}\big\langle x'(t), x(t)-y\big\rangle\,\mathrm{d}t \leq E(y)\,(t_1-t_0).$$

Note that this is an integrated version of the claim. First, since $E\circ x$ is a piecewise convex and nonnegative function and as $x\in\mathrm{AC}^2(\mathbf{R}_+,H)$, both integrals on the left-hand side are well-defined and finite. Without restriction, we may and will assume $y\in\mathcal{D}(E)$. Using lower semicontinuity and Fatou's lemma,

$$\begin{aligned}
&\int_{t_0}^{t_1} E(x(t))\,\mathrm{d}t + \int_{t_0}^{t_1}\big\langle x'(t), x(t)-y\big\rangle\,\mathrm{d}t\\
&\qquad\leq \liminf_{\tau\to 0}\Big[\int_{t_0}^{t_1} E(x^\tau(t))\,\mathrm{d}t + \int_{t_0}^{t_1}\big\langle (x^\tau)'(t), x^\tau(t)-y\big\rangle\,\mathrm{d}t\Big]\\
&\qquad\leq \liminf_{\tau\to 0}\Big[\int_{t_0}^{t_1} E(x^\tau(\lfloor\tau^{-1}t\rfloor+1)\,\mathrm{d}t\\
&\qquad\qquad + \int_{t_0}^{t_1}\big\langle (x^\tau)'(t), x^\tau(\lfloor\tau^{-1}t\rfloor+1)-y\big\rangle\,\mathrm{d}t\Big]\\
&\qquad\leq \int_{t_0}^{t_1} E(y)\,\mathrm{d}t\\
&\qquad= E(y)\,(t_1-t_0).
\end{aligned}$$

This shows the claim. The desired differential inclusion now follows by differentiating this integral inequality at every $t>0$ which is a Lebesgue point of both $E\circ x$ and x'. This set is notably independent of y. □

Lecture 14.

5.2.2. *Uniqueness, fundamental properties, and infinitesimal generator.*

Theorem 5.12 (Uniqueness and properties of gradient flow trajectories)**.** *Let $x\colon\mathbf{R}_+\to H$ be a gradient flow trajectory of E. Then the following properties hold.*

(i) **Contraction.** *Given any other gradient flow trajectory $y\colon\mathbf{R}_+\to H$, every $t\in\mathbf{R}_+$ satisfies the inequality*

$$\big\|x(t)-y(t)\big\| \leq \big\|x(0)-y(0)\big\|.$$

In particular, gradient flow trajectories with fixed initial points are unique.

(ii) **Energy dissipation.** *The assignment $t\mapsto E(x(t))$ is nondecreasing on $\mathbf{R}_+$ and locally Lipschitz continuous on $(0,\infty)$.*

(iii) **A priori estimate.** *For every* $z \in H$ *and every* $t > 0$,

$$E(x_t) \leq E(z) + \frac{\|x(0) - z\|^2}{2t}.$$

(iv) **Laplacian.** *For* every $t > 0$, *the right derivative*

$$x'^{+}(t) := \lim_{h \to 0+} \frac{x(t+h) - x(t)}{h}$$

exists in H. *It is equal to minus the unique element of minimal norm in* $\partial^- E(x_t)$, *cf. Remark 5.10. The same holds at zero if* $x(0) \in \mathcal{D}(\partial^- E)$.

Proof. Exercises 14.2 and 14.4. □

The last item from the previous theorem motivates the following definition.

Definition 5.13 (Infinitesimal generator)**.** *We define the domain of the* **infinitesimal generator** L *of* E *as* $\mathcal{D}(\mathsf{L}) := \mathcal{D}(\partial^- E)$.

Given any $x \in \mathcal{D}(\mathsf{L})$, *we define* $\mathsf{L}x \in H$ *as minus the unique element of minimal norm in* $\partial^- E(x)$, *cf. Remark 5.10.*

In particular, given a gradient flow trajectory $x\colon \mathbf{R}_+ \to H$ of E, the last claim of Theorem 5.12 implies for every $t > 0$, we have

$$x'^{+}(t) = \mathsf{L}x(t). \tag{5.7}$$

This looks much like the classical Euclidean heat equation. This correspondence is not coincidental; making it rigorous is the objective of this lecture. This is a very functional analytic approach to the heat equation, which extends to many other settings without essential changes. The reader interested in the heat equation from the PDE point of view is invited to consult Evans' book [6][40].

5.3. **Euclidean heat equation as gradient flow of the Dirichlet energy.** We will now devote our entire attention to the setting of Remark 5.1. That is, we take $H := L^2(\mathbf{R}^n, \mathcal{L}^n)$, where $n \in \mathbf{N}$, and we consider the Dirichlet energy E defined there. We denote the standard Euclidean scalar product on $\mathbf{R}^n$ by $\cdot$.

Definition 5.14 (Laplacian)**.** *We say a function* $u \in L^2(\mathbf{R}^n, \mathcal{L}^n)$ *belongs to the domain of the* **Laplacian**, *symbolically* $u \in \mathcal{D}(\Delta)$, *if* $u \in W^{1,2}(\mathbf{R}^n)$ *and there is a function* $g \in L^2(\mathbf{R}^n, \mathcal{L}^n)$ *such that for every* $v \in W^{1,2}(\mathbf{R}^n, \mathcal{L}^n)$,

$$\int_{\mathbf{R}^n} \nabla u \cdot \nabla v \, \mathrm{d}\mathcal{L}^n = -\int_{\mathbf{R}^n} g\, v \, \mathrm{d}\mathcal{L}^n.$$

In this case, g *is uniquely determined and we write* Δu *in place of it.*

Remark 5.15 (Symmetry)*.* The Laplacian defined above is a symmetric operator, in the sense that for every $u, v \in \mathcal{D}(\Delta)$,

$$\int_{\mathbf{R}^n} u\, \Delta v \, \mathrm{d}\mathcal{L}^n = \int_{\mathbf{R}^n} v\, \Delta u \, \mathrm{d}\mathcal{L}^n.$$

This is a straightforward consequence of its definition. ■

Proposition 5.16 (Laplacian vs. infinitesimal generator)**.** *We have the identity*

$$\mathcal{D}(\Delta) = \mathcal{D}(\partial^- E).$$

Moreover, if $u \in L^2(\mathbf{R}^n, \mathcal{L}^n)$ *belongs to either set, the subdifferential* $\partial^- E(u)$ *is single-valued and contains* $-\Delta u$ *as its only element.*

[40]We recommend this resource for complementary reading. As Evans points out in his introduction, the theory of PDEs should not be regarded as a subbranch of functional analysis, but as a domain in its own right.

Proof. As a preparation, we observe two simple facts given any $u, v \in W^{1,2}(\mathbf{R}^n)$. First, convexity of E implies convexity of the assignment $\varepsilon \mapsto E(u + \varepsilon\, v)$ on $\mathbf{R}$. And second, the bilinearity of the Euclidean scalar product easily implies

$$\begin{aligned}\lim_{\varepsilon\to 0} \frac{E(u+\varepsilon\, v) - E(u)}{\varepsilon} &= \lim_{\varepsilon \to 0} \frac{1}{2\varepsilon}\Big[\int_{\mathbf{R}^n} \big[|\nabla u|^2 + 2\varepsilon\, \nabla u\cdot\nabla v + \varepsilon^2\, |\nabla v|^2\big]\, \mathrm{d}\mathcal{L}^n \\ &\qquad - \int_{\mathbf{R}^n} |\nabla u|^2\, \mathrm{d}\mathcal{L}^n\Big] \\ &= \int_{\mathbf{R}^n} \nabla u \cdot \nabla v\, \mathrm{d}\mathcal{L}^n.\end{aligned} \tag{5.8}$$

We turn to the proof of the first claim. Assume $u \in \mathcal{D}(\Delta)$, so that $u \in W^{1,2}(\mathbf{R}^n)$ by definition. We claim for every $v \in W^{1,2}(\mathbf{R}^n)$,

$$-\int_{\mathbf{R}^n} v\, \Delta u\, \mathrm{d}\mathcal{L}^n \leq E(u+v) - E(u), \tag{5.9}$$

from which the claim follows by noting $\mathcal{D}(E) = W^{1,2}(\mathbf{R}^n)$. Given any $\varepsilon \in (0,1)$, convexity of E readily implies

$$E(u+\varepsilon\, v) = E\big[(1-\varepsilon)\, u + \varepsilon\,(u+v)\big] \leq (1-\varepsilon)\, E(u) + \varepsilon\, E(u+v).$$

Rearranging this inequality, dividing by ε, and using (5.8) yields

$$E(u+v) - E(v) \geq \lim_{\varepsilon\to 0+} \frac{E(u+\varepsilon v) - E(u)}{\varepsilon} = \int_{\mathbf{R}^n} \nabla u \cdot \nabla v\, \mathrm{d}\mathcal{L}^n.$$

This shows (5.9) and hence $-\Delta u \in \partial^- E(u)$ by Definition 5.14.

Conversely, assume $u \in \mathcal{D}(\partial^- E)$. Let $u^* \in \partial^- E(u)$. We claim that $u^* = -\Delta u$, which establishes the desired reverse inclusion and the single-valuedness of $\partial^- E(u)$ simultaneously. Given any $\varepsilon \in \mathbf{R}$ and any $v \in W^{1,2}(\mathbf{R}^n)$, we have

$$\varepsilon \int_{\mathbf{R}^n} v\, u^*\, \mathrm{d}\mathcal{L}^n \leq E(u+\varepsilon v) - E(u).$$

Dividing this inequality by $\varepsilon > 0$ and $\varepsilon < 0$, respectively, and using (5.8) twice,

$$\begin{aligned}\int_{\mathbf{R}^n} \nabla u\cdot\nabla v\, \mathrm{d}\mathcal{L}^n &= \lim_{\varepsilon\to 0+} \frac{E(u-\varepsilon\, v) - E(u)}{-\varepsilon} \\ &\leq \int_{\mathbf{R}^n} v\, u^*\, \mathrm{d}\mathcal{L}^n \\ &\leq \lim_{\varepsilon\to 0+} \frac{E(u+\varepsilon\, v) - E(u)}{\varepsilon} \\ &= \int_{\mathbf{R}^n} \nabla u\cdot\nabla v\, \mathrm{d}\mathcal{L}^n.\end{aligned}$$

This forces equality to hold throughout and shows $u^* = -\Delta u$, as desired. □

Since the L^2-closure of $W^{1,2}(\mathbf{R}^n)$ coincides with $L^2(\mathbf{R}^n, \mathcal{L}^n)$, by Theorem 5.11 every $u \in L^2(\mathbf{R}^n, \mathcal{L}^n)$ forms the starting point of a unique gradient flow trajectory denoted by $\mathsf{h}.u\colon \mathbf{R}_+ \to L^2(\mathbf{R}^n, \mathcal{L}^n)$. For every $t \in \mathbf{R}_+$, this procedure defines an operator $\mathsf{h}_t\colon L^2(\mathbf{R}^n, \mathcal{L}^n) \to L^2(\mathbf{R}^n, \mathcal{L}^n)$, where h_0 is simply the identity operator. It is a **semigroup** of operators in the sense of Hille–Yosida, i.e. it obeys $\mathsf{h}_{t+s} = \mathsf{h}_t \circ \mathsf{h}_s$ for every $s, t \in \mathbf{R}_+$. This elementary consequence of uniqueness of gradient flow trajectories stipulated in Theorem 5.12 is left as an exercise to the reader.

Definition 5.17 (Heat flow)**.** *The above family* $\mathsf{h}.$ *of operators is called* **heat flow***.*

The rest of these notes establishes some very basic properties.

Proposition 5.18 (Basic properties)**.** *The operator* h_t *is linear for every* $t \in \mathbf{R}_+$. *Also,* h_t *is a contraction for every* $t \in \mathbf{R}_+$. *That is, for every* $u \in L^2(\mathbf{R}^n, \mathcal{L}^n)$,

$$\big\|\mathsf{h}_t u\big\|_{L^2(\mathbf{R}^n,\mathcal{L}^n)} \leq \|u\|_{L^2(\mathbf{R}^n,\mathcal{L}^n)}.$$

In particular, the heat operator h_t *is bounded, hence continuous.*

Proof. The first claim is a direct consequence of Proposition 5.16 and the linearity of the Laplacian Δ, which is a trivial consequence of its definition.

The second point follows from the first item of Theorem 5.12 and the trivial observation that the heat flow starting at zero remains stationary. Moreover, linear operators between Banach spaces are bounded if and only if they are continuous. □

Proposition 5.19 (Commutation)**.** *For every* $u \in \mathcal{D}(\Delta)$ *and every* $t \in \mathbf{R}_+$, *we have the identity* $\mathsf{h}_t\Delta u = \Delta \mathsf{h}_t u$.

Proof. This follows from the identities

$$\Delta \mathsf{h}_t u = \lim_{\varepsilon\to 0+} \frac{\mathsf{h}_t(\mathsf{h}_\varepsilon u) - \mathsf{h}_t u}{\varepsilon} = \mathsf{h}_t\Big[\lim_{\varepsilon\to 0} \frac{\mathsf{h}_\varepsilon u - u}{\varepsilon}\Big] = \mathsf{h}_t u.$$

Here we used (5.7) in the first, Proposition 5.18 in the second, and the last clause from Theorem 5.12 in conjunction with Proposition 5.16 in the third identity. □

Corollary 5.20 (Symmetry)**.** *For every* $t \in \mathbf{R}_+$, *the heat operator* h_t *is symmetric, in the sense that for every* $u, v \in L^2(\mathbf{R}^n, \mathcal{L}^n)$,

$$\int_{\mathbf{R}^n} u\,\mathsf{h}_t v \,\mathrm{d}\mathcal{L}^n = \int_{\mathbf{R}^n} v\,\mathsf{h}_t u \,\mathrm{d}\mathcal{L}^n.$$

Proof. We are employing a common trick by interpolation and differentiation. Given any $t > 0$, define the function $F\colon [0,t] \to \mathbf{R}$ by

$$F(s) := \int_{\mathbf{R}^n} \mathsf{h}_s u\,\mathsf{h}_{t-s}\,v\,\mathrm{d}\mathcal{L}^n.$$

Since gradient flow trajectories are continuous on $\mathbf{R}_+$ with values in $L^2(\mathbf{R}^n, \mathcal{L}^n)$ and since the scalar product is continuous, F is easily seen to be continuous. Moreover, it is easily seen to be continuously differentiable on $(0,t)$ with derivative

$$F'(s) = \int_{\mathbf{R}^n} \Delta \mathsf{h}_s u\,\mathsf{h}_{t-s} v\,\mathrm{d}\mathcal{L}^n - \int_{\mathbf{R}^n} \mathsf{h}_s u\,\mathsf{h}_{t-s} v\,\mathrm{d}\mathcal{L}^n = 0.$$

The last identity follows from Remark 5.15. This, together with continuity on all of $[0,t]$, forces F to be constant; in particular, we have $F(0) = F(t)$, as desired. □

Theorem 5.21 (Heat flow characterization of the Laplacian)**.** *Let* $u \in L^2(\mathbf{R}^n, \mathcal{L}^n)$. *Then* $u \in \mathcal{D}(\Delta)$ *if and only if the limit*

$$\lim_{t\to 0+} \frac{\mathsf{h}_t u - u}{t} \tag{5.10}$$

exists in the strong topology of $L^2(\mathbf{R}^n, \mathcal{L}^n)$*; in this case,* (5.10) *equals* Δu.

Proof. If $u \in \mathcal{D}(\Delta)$, existence of the limit (5.10) follows from the last statement of Theorem 5.12. Moreover, its equality to Δu follows from Proposition 5.16.

Conversely, suppose the limit (5.10) — which we call g — exists in $L^2(\mathbf{R}^n, \mathcal{L}^n)$. We first claim $u \in W^{1,2}(\mathbf{R}^n)$. Given any $\varepsilon > 0$, Corollary 5.20 yields

$$\begin{aligned}\int_{\mathbf{R}^n} \mathsf{h}_\varepsilon u\, g\,\mathrm{d}\mathcal{L}^n &= \lim_{t\to 0+}\int_{\mathbf{R}^n} \mathsf{h}_\varepsilon u\,\frac{\mathsf{h}_t u - u}{t}\,\mathrm{d}\mathcal{L}^n\\ &= \lim_{t\to 0+}\int_{\mathbf{R}^n} u\,\frac{\mathsf{h}_t(\mathsf{h}_\varepsilon u) - \mathsf{h}_\varepsilon u}{t}\,\mathrm{d}\mathcal{L}^n.\end{aligned}$$

By Proposition 5.19, Corollary 5.20 again, and Definition 5.14,

$$\begin{aligned}\int_{\mathbf{R}^n} \mathsf{h}_\varepsilon u\, g\, \mathrm{d}\mathcal{L}^n &= \int_{\mathbf{R}^n} u\, \Delta \mathsf{h}_\varepsilon u\, \mathrm{d}\mathcal{L}^n\\ &= \int_{\mathbf{R}^n} \mathsf{h}_{\varepsilon/2} u\, \Delta \mathsf{h}_{\varepsilon/2} u\, \mathrm{d}\mathcal{L}^n\\ &= -\int_{\mathbf{R}^n} \big|\nabla \mathsf{h}_{\varepsilon/2} u\big|^2\, \mathrm{d}\mathcal{L}^n.\end{aligned}$$

Lower semicontinuity of E then implies

$$\begin{aligned}E(u) &\leq \liminf_{\varepsilon\to 0+} E(\mathsf{h}_\varepsilon u)\\ &= \frac{1}{2} \liminf_{\varepsilon\to 0+} \int_{\mathbf{R}^n} \big|\nabla \mathsf{h}_\varepsilon u\big|^2\, \mathrm{d}\mathcal{L}^n\\ &= -\frac{1}{2} \limsup_{\varepsilon\to 0+} \int_{\mathbf{R}^n} \mathsf{h}_\varepsilon u\, g\, \mathrm{d}\mathcal{L}^n\\ &< \infty.\end{aligned}$$

Since the domain of E coincides with $W^{1,2}(\mathbf{R}^n)$, the claim is proven.

It remains to show $u \in \mathcal{D}(\Delta)$. Recall by Theorem 5.12 that $E(\mathsf{h}_\varepsilon u) \leq E(u)$ for every $\varepsilon \in \mathbf{R}_+$. Consequently, the family $\{\mathsf{h}_\varepsilon u : \varepsilon \in \mathbf{R}_+\}$ is bounded in $W^{1,2}(\mathbf{R}^n)$, hence weakly precompact. However, since $\mathsf{h}_\varepsilon u \to u$ strongly in $L^2(\mathbf{R}^n, \mathcal{L}^n)$, this forces $\mathsf{h}_\varepsilon u \rightharpoonup u$ weakly in $L^2(\mathbf{R}^n, \mathcal{L}^n)$. Thus, applying the version of Lebesgue's differentiation theorem for continuous functions, every $v \in W^{1,2}(\mathbf{R}^n)$ obeys

$$\begin{aligned}\int_{\mathbf{R}^n} g\, v\, \mathrm{d}\mathcal{L}^n &= \lim_{t\to 0+} \int_{\mathbf{R}^n} \frac{\mathsf{h}_t u - u}{t}\, v\, \mathrm{d}\mathcal{L}^n\\ &= \lim_{t\to 0+} \frac{1}{t} \int_0^t \int_{\mathbf{R}^n} \Delta \mathsf{h}_s u\, v\, \mathrm{d}\mathcal{L}^n\, \mathrm{d}s\\ &= -\lim_{t\to 0+} \frac{1}{t} \int_0^t \int_{\mathbf{R}^n} \nabla \mathsf{h}_s u \cdot \nabla v\, \mathrm{d}\mathcal{L}^n\, \mathrm{d}s\\ &= -\int_{\mathbf{R}^n} \nabla u \cdot \nabla v\, \mathrm{d}\mathcal{L}^n.\end{aligned}$$

This is the desired identity. □

Appendices[41]

Appendix A. Weak topologies induced by families of functions

The following presentation is loosely based on [3, §§3.1–3.4], to which we refer for a deeper discussion with more advanced results.

Definition A.1 (Weak topology)**.** *Let X be a set, (Y, ρ) a topological space and $\mathcal{F} := \{f_i : i \in I\}$ constitute a collection of maps $f_i \colon X \to Y$. We define the* **weak topology on X induced by $\mathcal{F}$** *as the coarsest topology $\tau_{\mathcal{F}}$ on X such that $f_i \colon (X, \tau_{\mathcal{F}}) \to (Y, \rho)$ is continuous for every $i \in I$.*

Standard arguments from topology ensure $\tau_{\mathcal{F}}$ exists and the above definition is meaningful. Here are some known facts about $\tau_{\mathcal{F}}$; you may want to try to prove them by yourself using the definitions.

- If τ is another topology on X such that $f_i \colon (X, \tau) \to (Y, \rho)$ is continuous for every $i \in I$, then $\tau_{\mathcal{F}} \subset \tau$.

[41] The content of these appendices is not examinable. It is only some extra material which we hope is useful for you to understand weak topologies and their relevance for this course better, but also their overall importance in functional analysis.

- For every $V \in \rho$ and every $i \in I$, we have $f_i^{-1}(V) \in \tau_{\mathcal{F}}$.
- A basis of the topology $\tau_{\mathcal{F}}$ is given by sets of the form $\bigcap_{i \in I_0} f_i^{-1}(V_i)$, where $I_0 \subset I$ is finite and $V_i \in \rho$ for every $i \in I_0$. Analogously, a neighborhood basis of $x \in X$ is given by sets of the form $\bigcap_{i \in I_0} f_i^{-1}(V_i)$, where I_0 is a finite subset of I and $f^i(x) \in V_i \in \rho$ for every $i \in I_0$.
- A sequence $(x_n)_{n\in\mathbf{N}}$ in X converges to x with respect to $\tau_{\mathcal{F}}$ **if and only if** for every $i \in I$, the sequence $(f_i(x_n))_{n\in\mathbf{N}}$ in the target space Y converges to $f_i(x)$ with respect to ρ.

For simplicity, in the above we took Y to be fixed and independent of i. This is not necessary, and everything generalizes to the case of varying $\{Y_i : i \in I\}$; a canonical example of this situation is the following.

Example A.2 (Product topology). Given an arbitrary family of topological spaces (X_i, τ_i), recall that in the initial compendium we defined the product topology τ on the product space $X := \prod_{i\in I} X_i$. Then τ is nothing but the weak topology induced by the family $\{\pi_i : i \in I\}$ of projection maps $\pi_i(x) := x_i$. Given $I_0 \subset I$ finite and $V_i \in \tau_i$ for $i \in I_0$, we see that

$$\bigcap_{i\in I_0} \pi_i^{-1}(V_i) = \{x \in X : \pi_i(x) \in V_i \text{ for every } i \in I_0\}$$

defines an open set; the collection of all such sets constitutes a basis of τ (and a neighborhood basis of $\bar{x}$, respectively, if we additionally require V_i to contain $\bar{x}_i$).

A sequence $(x^n)_{n\in\mathbf{N}}$ in X converges to $x \in X$ with respect to τ if and only if, for every $i \in I$, $(x_i^n)_{n\in\mathbf{N}}$ converges to x_i with respect to τ_i. ■

In all the next examples, for simplicity we restrict ourselves to the case where the target space Y does not depend on i and is given by $\mathbf{R}$ with its Euclidean topology (that we will not specify notationally).

Example A.3 (Metric spaces). Let (X, d) be a metric space. Given any $x \in X$, define a function $f_x \colon (X, \mathsf{d}) \to \mathbf{R}$ by $f_x(y) := \mathsf{d}(x, y)$. Then the topology τ^{d} induced by d corresponds to the weak topology induced by the family of functions $\{f_x : x \in X\}$.

One can also look at it differently, by rather considering d as acting "globally" on the product space X^2. Then τ^{d} is equivalent to the coarsest topology τ on X such that $\mathsf{d} \colon (X^2, \tau^2) \to \mathbf{R}$ is a continuous map. ■

Example A.4 (LCTVS). Let X be a LCTVS with topology τ induced by the family of seminorms $\{p_i : i \in I\}$. Then τ is nothing but the weak topology on X induced the family of maps $\{p_i(x - \cdot) : i \in I, x \in X\}$.

Having presented these examples which connect the use of weak topologies to all the relevant examples we have seen in the first lectures, we now discuss what is truly usually referred to as "weak topologies".

Given a LCTVS (X, τ), let $X' := \mathcal{L}(X, \mathbf{R})$ denote its topological dual, namely the collection of all linear, continuous maps $x' : X \to \mathbf{R}$. Note the definition of X' actually depends on τ.

Definition A.5 (Weak topology)**.** *The* **weak topology** τ_w *of* (X, τ) *is the weak topology induced by the family* $\{x' : x' \in X'\}$ *— compare with Definition 1.38.*

Remark A.6 (Basic properties). We can now rephrase several facts about weak topologies and weak convergence.

- By definition, we always have $\tau_w \subset \tau$. (This justifies the terminology "weak topology", as opposed to the strong topology τ.)
- A sequence $(x_n)_{n\in\mathbf{N}}$ in X **converges weakly** to $x \in X$ if and only if it converges with respect to τ_w. That is, $x'(x^n) \to x'(x)$ as $n \to \infty$ as real numbers for every $x' \in X'$.

- Since $\tau_w \subset \tau$, if a sequence $(x_n)_{n\in\mathbf{N}}$ in X converges to $x \in X$ with respect to τ, it also converges weakly to x, namely with respect to τ_w. The converse in general is not true. (This justifies the terminology "weak convergence", as opposed to the strong convergence induced by τ.)
- Given any $x \in X$, a basis of neighborhoods of x are sets of the form
$$W_{\varepsilon,J_0}(x) := \{y \in X : |x'(x) - x'(y)| < \varepsilon \text{ for every } x' \in J_0\}$$
where $\varepsilon > 0$ and $J_0 \subset X'$ is finite; compare with Definition 1.8. ■

We recall (X, τ_w) is again a LCTVS by Corollary 1.39.

Consider now the dual space X', which has a natural vector space structure. For any $x \in X$, we can define the **evaluation map** φ_x on X' by $\varphi_x(x') := x'(x)$. Observe that φ_x is a linear map from X' to $\mathbf{R}$.

Definition A.7 (Weak* topology)**.** *The* **weak* topology** τ_{w^*} *of* X' *is defined as the weak topology induced by the family* $\{\varphi_x : x \in X\}$ *— compare with Definition 1.40.*

It can be shown (X', τ_{w^*}) is again a LCTVS, cf. the proof of Corollary 1.39 and the discussion around Definition 1.40. Thus far, we did not have any other candidate topology on X'; but the structure of X' inherited from being the dual of a LCTVS naturally induces one. In particular, the weak* topology τ_{w^*} can be regarded as the topology induced on X' by the action of its **predual** X.

Example A.8 (Banach spaces). Suppose now X is a Banach space, with strong topology τ_X induced by a norm $\|\cdot\|_X$. In this case, we know X' also has a normed structure: linear operators are continuous if and only if they are bounded, yielding $\|\cdot\|_{X'}$ as defined by
$$\|x'\|_{X'} := \sup_{x\in X\setminus\{0\}} \frac{|x'(x)|}{\|x\|_X} = \sup_{\substack{x\in X,\\ \|x\|_X=1}} |x'(x)|,$$
which turns $(X', \|\cdot\|_{X'})$ into a Banach space. We can then iterate this procedure and define X'' as the topological dual of X', which will also be a Banach space with norm $\|\cdot\|_{X''}$, and so on. In this situation, several topologies are available.

- We can endow X with either the strong topology τ_X or the weak topology $\tau_{w,X}$, which satisfy $\tau_{w,X} \subset \tau_X$.
- We can endow X' with either the strong topology $\tau_{X'}$ induced by $\|\cdot\|_{X'}$, the weak topology $\tau_{w,X'}$ induced by its dual X'', or the weak* topology τ_{w^*} induced by the evaluations maps $\{\varphi_x : x \in X\}$.

We claim that
$$\tau_{w^*} \subset \tau_{w,X'} \subset \tau_{X'}. \tag{A.1}$$
In other words, the weak* topology is weaker than the weak topology, which in turn is weaker than the strong one. A similar statement holds when comparing notions of convergence. One can produce examples of spaces X and X' where all the inclusions appearing in (A.1) are strict.[42]

To prove (A.1), first observe that the inclusion $\tau_{w,X'} \subset \tau_{X'}$ follows from the properties of the weak topology. To prove $\tau_{w^*} \subset \tau_{w,X'}$, note that for any $x \in X$, the evaluation map φ_x is a bounded linear function on X', since the relation
$$|\varphi_x(x')| = |x'(x)| \leq \|x'\|_{X'}\,\|x\|_X$$
valid for every $x' \in X'$, implies
$$\|\varphi_x\|_{X''} \leq \|x\|_X. \tag{A.2}$$

[42]For the interested reader, a basic example is given by $X = L^1$ and $X' = L^\infty$. Time permitting, we will discuss this deeper in a later version of these notes.

In particular, $\{\varphi_x : x \in X\} \subset X''$, which implies φ_x is continuous under $\tau_{w,X'}$ for every $x \in X$. Since by definition, τ_{w^*} is the coarsest topology with this property, we deduce $\tau_{w^*} \subset \tau_{w,X'}$. ∎

Appendix B. Weak topologies and infinite dimensional Banach spaces

We now aim for a better understanding of weak topologies τ_w on Banach spaces X, their key properties and their relations to the aforementioned τ_{w^*} and $\tau_{w,X'}$. Throughout this section, we will always tacitly assume $(X, \|\cdot\|_X)$ is a Banach space. We start with some basic facts about X, X', and X''.

Lemma B.1 (Isometry). *Let X be a Banach space and consider the map $J : X \to X''$ given by $J(x) := \varphi_x$, where φ_x is the evaluation map defined before Definition A.7. Then J is an injective linear isometry from X to a closed subspace of X''; in particular, every $x \in X$ satisfies*

$$\|\varphi_x\|_{X''} = \sup_{\|x'\|_{X'}=1} |x'(x)| = \|x\|_X. \tag{B.1}$$

Proof. It is easy to check J is linear. We already showed in (A.2) that $\|\varphi_x\|_{X''} \leq \|x\|_X$ for every $x \in X$, so it remains to prove the reverse estimate. Fix $\bar{x} \in X \setminus \{0\}$ and define a linear functional $\bar{x}'$ on the line $\mathbf{R}\bar{x}$ by $\bar{x}'(\lambda\bar{x}) = \lambda\|\bar{x}\|_X$. By the analytic version of the Hahn–Banach Theorem 1.35, we can extend $\bar{x}'$ to a nonrelabeled linear functional on all of X. For every $x \in X$, it satisfies

$$|x'(x)| \leq \|x\|_X.$$

This implies $\|\bar{x}'\|_{X'} \leq 1$ and $\bar{x}'(\bar{x}) = \|\bar{x}\|_X$, so that $\|\bar{x}'\|_{X'} = 1$. But then

$$\varphi_{\bar{x}}(\bar{x}') = \bar{x}'(\bar{x}) = \|\bar{x}\|_X = \|\bar{x}\|_X \|\bar{x}'\|_{X'};$$

by the definition of the X''-norm, this implies $\|\varphi_{\bar{x}}\|_{X''} \geq \|\bar{x}\|_X$. □

The map J defined in Lemma B.1 is sometimes refereed to as the **canonical injection** of X into X''.

Corollary B.2 (Invariance of infinite dimensionality). *Let X be a Banach space. Then X is infinite dimensional if and only if X' is infinite dimensional.*

Proof. If X is finite-dimensional, then by Remark 1.34 it is isomorphic to $\mathbf{R}^d$ with equivalent norm. By basic linear algebra, its dual X' is isomorphic to $(\mathbf{R}^d)' = \mathbf{R}^d$ and thus finite-dimensional.

Now assume X' is finite-dimensional, then so is X'' by the above argument. By Lemma B.1, X'' contains an isometric copy $J(X)$ of X, which implies X must be finite-dimensional as well. □

Similarly to the situation presented in §2, in practical applications the notion of convergence of sequences in weak topologies is the most useful to use. This is because weak topologies often have a lot of compact (or sequentially compact) sets, cf. the Banach–Alaoglu–Bourbaki Theorem C.1. At the same time, one has to be careful about it: **sequentially closed sets need not be closed**.

Therefore we aim to better understand properties encoded by weak convergence.

We start with a basic yet useful fact: if a sequence $(x_n)_{n\in\mathbf{N}}$ in X converges weakly to $x \in X$ — namely with respect to τ_w —, it is bounded with respect to $\|\cdot\|_X$. In the following, we will sometimes denote weak convergence by "$x_n \rightharpoonup x$" (contrary to strong convergence with respect to $\|\cdot\|_X$, denoted by "$x_n \to x$".

Lemma B.3 (Boundedness). *Let $(x_n)_{n\in\mathbf{N}}$ be a sequence in X and $x \in X$ such that $x_n \rightharpoonup x$ as $n \to \infty$. Then $\sup_{n\in\mathbf{N}} \|x_n\|_X < \infty$ and*

$$\|x\|_X \leq \liminf_{n\to\infty} \|x_n\|_X. \tag{B.2}$$

Property (B.2) is often refereed to as the lower semicontinuity of the norm in question in the weak(*) topology[43].

The proof is based on the Banach–Steinhaus theorem. For its proof (on the Baire category theorem), we refer to [3, Thm. 2.2].

Theorem B.4 (Banach–Steinhaus theorem viz. uniform boundedness principle)**.** *Let E and F be Banach spaces and let $\{T_i : i \in I\}$ be a family of continuous linear operators from E into F. Assume that for every $x \in E$,*

$$\sup_{i\in I} \|T_i x\|_F < \infty.$$

Then the family $\{T_i : i \in I\}$ is bounded in the operator norm; in other words, there exists a constant $C > 0$ such that for every $i \in I$ and every $x \in E$,

$$\|T_i x\|_F \leq C\, \|x\|_E.$$

This theorem is arguably one of the most surprising yet powerful elements of functional analysis. It derives *uniform* boundedness out of *pointwise* boundedness, which is per se strictly weaker.

Proof of Lemma B.3. Since $x_n \rightharpoonup x$ as $n \to \infty$, for any $x' \in X'$ we have $\varphi_{x_n}(x') = x'(x_n) \to x'(x)$ as $n \to \infty$. Since any convergent sequence in $\mathbf{R}$ is bounded, we deduce that, for any fixed $x' \in X'$, we have $\sup_{n\in\mathbf{N}} |\varphi_{x_n}(x')| < \infty$. By the uniform boundedness principle (applied with $E = X'$, $F = \mathbf{R}$, and $\{T_n : n \in \mathbf{N}\} = \{\varphi_{x_n} : n \in \mathbf{N}\} \subset X''$) we deduce

$$\sup_{n\in\mathbf{N}} \|x_n\|_X = \sup_{n\in\mathbf{N}} \|\varphi_{x_n}\|_{X''} < \infty$$

where the first equality comes from Lemma B.1.

Since $x_n \rightharpoonup x$ as $n \to \infty$, the same holds for any subsequence we can extract. In particular, we may extract $(x_{n_k})_{k\in\mathbf{N}}$ with $\lim_{k\to\infty} \|x_{n_k}\|_X = \liminf_{n\to\infty} \|x_n\|_X$. It follows that, given any $x' \in X'$ with $\|x'\|_{X'} = 1$,

$$|x'(x)| = \lim_{k\to\infty} |x'(x_{n_k})| \leq \lim_{k\to\infty} \|x'\|_{X'}\, \|x_{n_k}\|_X = \liminf_{n\to\infty} \|x_n\|_X.$$

Taking the supremum over $x' \in X'$ with $\|x'\|_{X'} = 1$ in this inequality and applying (B.1), one gets (B.2). □

Remark B.5 (Boundedness in the weak* topology). The same argument gives the following fact: given a sequence $(x'_n)_{n\in\mathbf{N}}$ in X' which converges weakly* to $x' \in X'$, we have $\sup_{n\in\mathbf{N}} \|x'_n\|_{X'} < \infty$ and $\|x\|_{X'} \leq \liminf_{n\to\infty} \|x'_n\|_{X'}$. ■

The next result gives simple conditions to verify a set $E \subset X$ is weakly closed (thus also weakly sequentially closed).

Lemma B.6 (Closedness)**.** *Let $E \subset X$ be convex. Then E is strongly closed if and only if it is weakly closed.*

Proof. Since $\tau_w \subset \tau_X$, weakly closed sets are always strongly closed. Thus we only need to show that a convex, strongly closed set E is also weakly closed. Let E be convex and strongly closed, $x \in E^c$. We apply the geometric version of the Hahn–Banach Theorem 1.36 to $A = \{x\}$ and $B = E$ to find $x' \in X'$ such that $x'(x) < \alpha < \beta < x'(y)$ for all $y \in E$. Observing that the set $U_x = \{z \in X : x'(z) < \alpha\}$ is open in the weak topology and that $x \in U_x \subset E^c$, we conclude E^c is open in τ_w and thus E is weakly closed. □

[43]It is also naturally related to the *Fatou property* that many function spaces and spaces of distributions have, which — as the name suggests — is linked to Fatou's lemma in measure theory, which asserts lower semicontinuity of the Lebesgue integral.

The simplest example of a weakly closed set in X is the closed ball $\overline{B}_1(0)$, by virtue of Lemma B.6. Similarly, the set $\overline{B}_r(x_0)$ is weakly closed for every $r > 0$ and every $x_0 \in X$.

There is another natural class of weakly closed sets, which are of fundamental importance in the **direct method of calculus of variations**. To introduce them, we need to define a class of functionals first.

Definition B.7 (Lower semicontinuous functionals)**.** *Let (Y, τ_Y) designate a topological space and let $F\colon Y \to \mathbf{R}$.*

a. *We call F* **lower semicontinuous** *if its* **sublevel sets** *are closed. Namely, given any $\lambda \in \mathbf{R}$, $E_\lambda := \{y \in Y : F(y) \leq \lambda\}$ is closed with respect to τ_Y.*
b. *We call F* **sequentially lower semi-continuous** *if for every sequence $(y_n)_{n\in\mathbf{N}}$ convergent to $y \in Y$ with respect to τ_Y, $F(y) \leq \liminf_{n\to\infty} F(y_n)$.*

It can be shown F is sequentially lower semicontinuous if and only if its level sets are sequentially closed. As a consequence, lower semicontinuous functionals are always sequentially lower semicontinuous, but the converse might not be true. If the topology τ_Y is metrizable, then both notions of lower semicontinuity coincide.

Recall $F\colon X \to \mathbf{R}$ is convex if for every $x, y \in X$ and every $\lambda \in [0, 1]$,

$$F((1-\lambda)x + \lambda y) \leq (1-\lambda)F(x) + \lambda F(y).$$

By Lemma B.6, we can characterize weakly lower semicontinuous functionals on X as soon as we additionally impose the geometric constraint of convexity.

Corollary B.8 (Characterizations of lower semicontinuity)**.** *Let $F\colon X \to \mathbf{R}$ be a convex functional. Then the following are equivalent.*

a. *F is strongly lower semicontinuous, i.e. lower semicontinuous with respect to $\|\cdot\|_X$ or equivalently the strong topology τ_X.*
b. *F is weakly lower semicontinuous, i.e. lower semicontinuous with respect to the weak topology τ_w.*
c. *F is weakly sequentially lower semicontinuous with respect to the weak topology τ_w.*

Proof. (i) $\Longrightarrow$ (ii). Fix $\lambda \in \mathbf{R}$ and consider the level set E_λ. Since F is convex, E_λ is convex. Since F is strongly lower semicontinuous, E_λ is strongly closed. By Lemma B.6, we deduce E_λ is weakly closed, which shows (ii).

(ii) $\Longrightarrow$ (iii). This follows from the aforementioned facts.

(iii) $\Longrightarrow$ (i). It suffices to show E_λ is sequentially closed in the strong topology. Let $(x_n)_{n\in\mathbf{N}}$ be a sequence in E_λ such that $\|x_n - x\|_X \to 0$ as $n \to \infty$. This forces $x_n \rightharpoonup x$ as $n \to \infty$. Since F is weakly sequentially lower semicontinuous, we deduce $F(x) \leq \liminf_{n\to\infty} F(x_n) \leq \lambda$. This gives $x \in E_\lambda$, which terminates the proof. □

We now turn our attention to open sets in weak topologies. They suffer the pathology of being naturally unbounded.

Lemma B.9 (Unboundedness)**.** *Let X be an infinite-dimensional Banach space and let τ_w denote its weak topology. Then every open set $U \in \tau_w$ is unbounded, both with respect to the weak topology τ_w and the strong topology τ_X.*

Proof. As $\tau_w \subset \tau_X$, it suffices to show unboundedness in τ_w. Up to a translation, we may and will assume $0 \in U$. Since $0 \in U \in \tau_w$, there must exist $x_1', \dots x_n' \in X'$ and $\varepsilon > 0$ such that

$$B_{\varepsilon,\{x_1',\dots,x_n'\}}(0) := \{x \in X : |x_i'(x)| < \varepsilon \text{ for every } i \in \{1,\dots,n\}\} \subset U.$$

We claim there exists $\bar{x} \in X \setminus \{0\}$ such that $x_i'(\bar{x}) = 0$ for every $i \in \{1,\dots,n\}$. Suppose to the contrary this claim is false. The continuous linear map $\Phi\colon X \to \mathbf{R}^n$

given by $\Phi(x) := (x_1'(x), \dots x_n'(x))$ would then be injective. Denoting by $\Phi(X) \subset \mathbf{R}^n$ its range, applying Proposition 1.33 one then deduces $\Phi\colon X \to \Phi(X)$ is a linear isomorphism, contradicting the infinite-dimensionality of X.

Next, given such a point $\bar{x}$, arguing as in Lemma B.1, one can construct $\bar{x}' \in X'$ such that $\bar{x}'(\bar{x}) = 1$. Let us define

$$V := \{x \in X : |x_i'(x)| < \varepsilon \text{ for every } i \in \{1, \dots, n\},\ |\bar{x}'(x)| < 1\}.$$

Clearly, V is open in τ_w. We claim there exists no $s >$ such that $U \subset sV$, yielding unboundedness. Indeed, by construction $\lambda\bar{x} \in B_{\varepsilon,\{x_1',\dots,x_n'\}}(0) \subset U$ for every $\lambda \in \mathbf{R}$, but $\lambda\bar{x} \in sV$ if and only if $|\lambda| < s$. □

As a byproduct of Lemma B.9, if a set E is strongly bounded, then it cannot be weakly open! The standard example is the open unit ball $B_1(0)$.

This contrast between boundedness of convergent sequences and unboundedness of open sets is the cause of several pathologies concerning weak topologies. In particular, it implies their lack of metrizability.

Lemma B.10 (Weak metrizability)**.** *Let $(X, \|\cdot\|_X)$ be a Banach space. Then the weak topology τ_w is metrizable if and only if X is finite-dimensional.*

Proof. If X is finite-dimensional, it is isomorphic to $\mathbf{R}^d$ and its weak topology τ_w coincides with the strong topology induced by $\|\cdot\|_X$, which is equivalent to the Euclidean one.

Conversely, assume X is infinite-dimensional and assume by contradiction there exists a metric d inducing τ_w. Consider the sets $U_n := B_{1/n}(0)$, where $n \in \mathbf{N}$. By construction, $U_n \in \tau_w$ and hence Lemma B.9 implies it is unbounded in both the weak and strong topologies. In turn, by Exercise 3.1, for every $n \in \mathbf{N}$ we can find $x_n \in U_n$ such that $\|x_n\|_X \geq n$. On the other hand, the sequence $(x_n)_{n\in\mathbf{N}}$ constructed in this way must converge weakly to 0 since $\mathsf{d}(x_n, 0) < 1/n$. But by Lemma B.3, the sequence $(x_n)_{n\in\mathbf{N}}$ is bounded in $(X, \|\cdot\|_X)$, which contradicts the property $\|x_n\|_X \geq n$ coming from the construction. □

The same kind of issue applies to the weak* topology.

Lemma B.11 (Weak* metrizability)**.** *Let X be a Banach space with dual space X'. Then the weak* topology τ_{w^*} is metrizable if and only if X is finite-dimensional.*

Proof. If X is finite-dimensional, the same argument as in Lemma B.10 applies.

Conversely, assume τ_{w^*} is metrizable. By Exercise 5.1 X' admits an at most countable algebraic base. We claim this implies X' is finite-dimensional, from which the conclusion will follow by Corollary B.2.

To see the claim, consider an at most countable algebraic base $\{y_n : n \in \mathbf{N}\}$ of X', and define the increasing subspaces $Y_n := \mathbf{R}y_1 + \cdots + \mathbf{R}y_n$. Since Y_n is finite-dimensional by construction, Proposition 1.33 ensures it is closed in X. Moreover, since $\{y_n : n \in \mathbf{N}\}$ is an algebraic base, we have

$$X' = \bigcup_{n\in\mathbf{N}} Y_n. \tag{B.3}$$

Now there are two options, namely

- the set $\{y_n : n \in \mathbf{N}\}$ is finite, so $X' = Y_{n_0}$ for some $n_0 \in \mathbf{N}$, which shows finite-dimensionality of X', or
- the sequence is countably infinite, in which case Y_n is a proper subspace of X' for every $n \in \mathbf{N}$.

In the second, since proper linear subspaces in a TVS always have empty interior (prove this as an exercise), by (B.3) it would follow X' can be written as a countable union of closed sets with empty interior. By the Baire category theorem, this would imply X' has empty interior, which is a contradiction. □

Appendix C. Weak compactness, separability, and reflexivity

This section contains some of the most classical results at the heart of functional analysis. The first one we need to mention is the following.

Theorem C.1 (Banach–Alaoglu–Bourbaki theorem)**.** *Let X be a Banach space, X' its dual. Then the closed unit ball $\{x' \in X' : \|x'\|_{X'} \leq 1\}$ is compact with respect to the weak* topology.*

See [3, Thm. 3.16] for a proof. The compactness of the unit ball is the most essential property of the weak* topology, and the main reason for its introduction. Compare it to Remark 1.34: if X is infinite-dimensional, the same can never be true in the strong topology induced by $\|\cdot\|_X$!

Theorem C.1 tells us there are many compact sets in (X', τ_{w^*}). This is useful in optimization problems: if K is a compact set in (X', τ_{w^*}) and $F\colon (X', \tau_{w^*}) \to \mathbf{R}$ is a continuous function, then there exists $\bar{x} \in K$ with

$$\inf_{x' \in K} F(x') = \min_{x' \in K} F(x') = F(\bar{x}).$$

Indeed, the set $F(K)$ is compact in $\mathbf{R}$, thus of the form $[\min_K F(x), \max_K F(x)]$. Compactness therefore yields the existence of minimizers in optimization problems.

There are some adjacent remarks about Theorem C.1:

- Our original space was X, on which we defined the weak topology (X, τ_w). So we would like to obtain compactness results in (X, τ_w).
- In applications, one would often like to construct approximate minimizers in an algorithmic way. This often results in an approximation sequence $(x_n)_{n\in\mathbf{N}}$. While compactness is useful, we would like to understand *sequential compactness* as well, as to guarantee the sequence $(x_n)_{n\in\mathbf{N}}$ in fact converges to a minimizer.
- By Lemma B.11, the weak* is not metrizable. Therefore, compact sets and sequentially compact sets might not coincide, possibly heavily limiting the effect of Theorem C.1. A similar issue applies for the weak topology, in light of Lemma B.10.

To overcome these issues, we need to introduce some concepts.

Definition C.2 (Reflexivity)**.** *A Banach space X is* **reflexive** *if the canonical injection $J\colon X \to X''$ from Lemma B.1 is surjective.*

Remark C.3 (Weak vs. weak* convergence). Reflexivity allows to link weak convergence in X to weak* convergence in X''. A sequence $(x_n)_{n\in\mathbf{N}}$ converges weakly to $x \in X$ if and only if $(Jx_n)_{n\in\mathbf{N}}$ converges weakly* to Jx, since for every $x' \in X'$,

$$\lim_{n\to\infty} (Jx_n)(x') = \lim_{n\to\infty} x'(x_n) = x'(x) = (Jx)(x').$$

Similarly, the weak topology on X and the weak* topology on X'' coincide. One can then apply Theorem C.1 (with X'' in place of X) we deduce that, if X is reflexive, the closed unit ball $\{x \in X : \|x\|_X \leq 1\}$ is compact in the weak topology! ■

In the following, we will not rely on Theorem C.1 or its consequences due to Remark C.3. Instead, the main goal of this appendix is to provide a sufficiently self-contained proof of the following fundamental result.

Theorem C.4 (Boundedness implies sequential precompactness)**.** *Let X be a reflexive Banach space. Let $(x_n)_{n\in\mathbf{N}}$ be a bounded sequence in X. Then there exists* *a subsequence* $(x_{n_k})_{k\in\mathbf{N}}$ *and a point* $x \in X$ *such that the subsequence converges to* x *in the weak topology of* X*.*

In other words, if X *is reflexive, then bounded sequences in* X *are weakly sequentially precompact.*

Before delving into the proof, let us briefly mention to other fundamental results from Functional Analysis which are closed related to Theorem C.4, although we will not use them.

The **Eberlein–Smulian theorem** states that, on an arbitrary Banach space X, weakly compact and sequentially weakly compact sets coincide.

The **Kakutani theorem** states that the unit ball B_X in X is weakly compact if and only if X is reflexive. We see in particular that the reflexivity assumption in Theorem C.4 is actually not just a sufficient condition, but also a necessary one, in order to guarantee the existence of a weak limit (after extraction of subsequence, as usual). These theorems also nicely address the third issue from the above list.

To prove Theorem C.4, we need some preparations. We start by collecting some basic properties of reflexive spaces.

Lemma C.5 (Properties of reflexive spaces)**.** *The following hold.*

(i) *If* $(X, \|\cdot\|_X)$ *is reflexive and* M *is a closed linear subspace of* X*, then* $(M, \|\cdot\|_X)$ *is reflexive.*

(ii) X *is reflexive if and only if its dual* X' *is reflexive.*

Proof. For detailed proofs, we refer to [3, Prop. 3.20, Cor. 3.21]. Here are the main ideas in order to get a simple self-contained proof.

(i) Given any $x' \in X'$, we define an element of $m' \in M'$ by considering its restriction $x'\big|_M$ to M. Given $m'' \in M''$, by duality we can define $\tilde{m}'' \in X''$ by $\tilde{m}''(x') = m''(x'\big|_M)$. By reflexivity of X, this implies the existence of $\bar{x} \in X$ such that $J_X\bar{x} = \tilde{m}''$. To complete the proof, it remains to show $\bar{x} \in M$ and $J_M\bar{x} = m''$.

For the first claim, if by contradiction $\bar{x} \notin M$, by the analytic Hahn–Banach theorem and some additional technical work, one can construct a linear functional $x' \in X$ such that $x' \equiv 0$ on M and $x'(\bar{x}) = 1$. But then by construction

$$1 = x'(x) = (J_X\bar{x})(x') = \tilde{m}''(x') = m''(x'\big|_M) = m''(0) = 0,$$

which is a contradiction.

For the second claim, again by the analytic Hahn–Banach theorem, any $m' \in M'$ admits an extension $x' \in X'$ such that $x'\big|_M = m'$, so that

$$(J_M\bar{x})(m') = m'(\bar{x}) = x'(\bar{x}) = (J_X\bar{x})(x') = \tilde{m}''(x') = m''(x'\big|_M) = m''(m');$$

as the identity holds for every $m' \in M'$, we conclude $J_M\bar{x} = m''$.

(ii) Let us show that, if X is reflexive, so is X'. Given $\varphi \in X'''$, we need to find $x' \in X'$ such that $\varphi = J_{X'}x'$, i.e. $\varphi(x'') = x''(x')$. By assumption, any $x'' \in X''$ is of the form J_Xx for some $x \in X$, so this is equivalent to constructing x' with the following property for every $x \in X$:

$$x'(x) = (J_Xx)(x') = \varphi(J_Xx).$$

Since φ and J_X are continuous, we can *define* $x' \in X'$ by the relation $x' = \varphi \circ J_X$, which concludes the proof.

By the above, if X' is reflexive, so is X''; but then by part (i) so is X, since we can identify it with $J_X(X)$ which is a closed linear subspace of X''. □

Definition C.6 (Separability)**.** *A Banach space* X *is* **separable** *if there exists a countable* $D \subset X$ *which is dense in* X *with respect to* $\|\cdot\|_X$*.*

We recall the not entirely obvious fact that if X is separable and $Y \subset X$, then Y is also separable [3, Prop. 3.25].

Lemma C.7 (Characterization of separability)**.** *X is separable if and only if it admits a countable* **linearly dense** *subset, namely there exists $D \subset X$ whose linear span is dense in X.*

Proof. Clearly if D is dense, it is also linearly dense.

Conversely, if D is countable and linearly dense, then

$$E := \Big\{ \sum_{i=1}^{n} \lambda_i \, y_i : n \in \mathbf{N}, \, y_i \in D, \, \lambda_i \in \mathbf{Q} \Big\}$$

is still countable (since we restricted to rational coefficients). Since elements in $\operatorname{span} D$ can be approximated arbitrarily well by E, we get $X = \overline{\operatorname{span} D} \subset \overline{E}$. □

The next lemma provides an easy-to-check condition for a set to be linearly dense.

Lemma C.8 (Characterization of linear density)**.** *A set $D \subset X$ is linearly dense in X if and only if for every $x' \in X$, the following holds. If x' vanishes on all of D — i.e. $x'(y) = 0$ for every $y \in D$ — then $x' = 0$.*

Proof. One implication is trivial. Suppose D is linearly dense. Since any $x' \in X'$ is linear and continuous, x' is entirely determined by its values in D. In particular if $x'(y) = 0$ for every $y \in D$, it must be 0 everywhere.

Conversely, suppose D is not linearly dense, so that $Y := \overline{\operatorname{span} D}$ is a closed, proper linear subspace of X. Then there exists a point $x \in X \setminus Y$. By the analytic Hahn–Banach theorem we can construct $x' \in X'$ such that $x'(y) = 0$ for all $y \in Y$ (in particular for $y \in D$), but $x'(x) = 1$, a contradiction. □

Proposition C.9 (A sufficient condition for separability)**.** *Let X be a Banach space such that X' is separable. Then X is separable.*

The converse statement is not true. The space $L^1(\mathbf{R}^d, \mathcal{L}^d)$ is separable, but its dual $L^\infty(\mathbf{R}^d, \mathcal{L}^d)$ is not.

Proof of Proposition C.9. Let $E = \{x'_n : n \in \mathbf{N}\}$ be a dense subset in X'. By definition of $\|\cdot\|_{X'}$, for each x'_n, there exists $x_n \in X$ such that $\|x_n\|_X = 1$ and $x'_n(x_n) > \|x'_n\|_{X'}/2$. We claim $D := \{x_n : n \in \mathbf{N}\}$ is linearly dense in X, which gives the conclusion by Lemma C.7. By Lemma C.8, it suffices to show that, given any $x' \in X$ such that $x'(x_n) = 0$ for all x_n, we have $x' \equiv 0$. Given such an x', by separability we can find a sequence $(x'_m)_{m \in \mathbf{N}}$ in E such that $\|x'_m - x'\|_{X'} \to 0$ as $m \to \infty$. But then by construction

$$\|x' - x'_m\|_{X'} \geq |x'(x_m) - x'_m(x_m)| = |x'_m(x_m)| \geq \frac{\|x'_m\|_{X'}}{2}$$

so that sending $m \to \infty$ we find $0 \geq \|x'\|_{X'}/2$, yielding $x' = 0$. □

Theorem C.10 (Reflexivity and separability combined)**.** *Let X be a Banach space. Then X is reflexive and separable if and only if X' is reflexive and separable.*

Proof. By Lemma C.5 and Proposition C.9, if X is separable and reflexive, so if X'. By the first implication, if X' is separable and reflexive, then so is X''; but if X' is reflexive, so is X, meaning $J\colon X \to X''$ is an isometry. Therefore X is separable since it is isometric to the separable space X''. □

Our main interest in separability is due to the following result.

Theorem C.11 (Topological consequences of separability)**.** *Let X be a Banach space with dual X'. Then the following hold.*

(i) *The closed unit ball $B_{X'} := \{x' \in X' : \|x'\|_{X'} \leq 1\}$, endowed with the weak topology τ_{w^*}, is metrizable if and only if X is separable.*

(ii) *If X is separable, then $B_{X'}$ is sequentially compact w.r.t. τ_{w^*}.*

Remark C.12 (About Theorem C.11). We invite the reader to compare item (i) to Lemma B.11. Even though (X', τ_{w^*}) is not metrizable, the restriction of the topology to the unit ball (in fact, to any ball of arbitrary but finite radius) is metrizable! Additionally, we know by Lemma B.3 that, if $(x'_n)_{n\in\mathbf{N}}$ converges weakly* to $x' \in X'$, there exists $R > 0$ large enough such that $(x'_n)_{n\in\mathbf{N}} \subset \{x' \in X' : \|x'\| \leq R\}$, a set on which the topology is metrizable. Thus, if X is separable, the weak* topology is (roughly speaking) "locally metrizable". ■

Proof of Theorem C.11. We start by showing separately the two implications from (i), following [3, Thm 3.28].

Assume first X is reflexive and let $D = \{x_n : n \in \mathbf{N}\}$ be a countable dense set in its closed unit ball $B_X := \{x \in X : \|x\|_X \leq 1\}$. Define a metric d on $B_{X'}$ by

$$\mathsf{d}(x', y') := \sum_{n\in\mathbf{N}} 2^{-n}\, |(x' - y')(x_n)|.$$

The d is well-defined since $\|x' - y'\|_{X'} \leq 2$ for every $x', y' \in B_{X'}$ and $\|x_n\|_X \leq 1$ for every $n \in \mathbf{N}$. Moreover, d is translation-invariant by definition.

We claim this metric induces the weak* topology on $B_{X'}$. We do not give a full proof here; let us only show a simpler fact, namely that if $\mathsf{d}(x'_m, 0) \to 0$ as $m \to \infty$, then $(x'_m)_{m\in\mathbf{N}}$ converges to 0 in the weak* topology. Indeed, by definition of d, one readily checks $\lim_{n\to\infty} \mathsf{d}(x'_m, 0) = 0$ as $m \to \infty$ if and only if $\lim_{m\to\infty} x'_m(x_n) = 0$ for every fixed $n \in \mathbf{N}$. Now let $x \in X \setminus \{0\}$ and set $\bar{x} := x/\|x\|_X$. Given any $\varepsilon > 0$, there exists $n \in \mathbf{N}$ such that $\|x_n - \bar{x}\|_X \leq \varepsilon/\|x\|_X$. We deduce

$$\begin{aligned} |x'_m(x)| &= \|x\|_X\, |x'(\bar{x})| \\ &\leq \|x\|_X \left[|x'_m(x_n)| + |x'_m(\bar{x} - x_n)|\right] \\ &\leq \|x\|_X \left[|x'_m(x_n)| + \|x'_m\|_{X'} \|\bar{x} - x_n\|_X\right] \\ &\leq \|x\|_X\, |x'_m(x_n)| + \varepsilon. \end{aligned}$$

Sending $m \to \infty$, we find

$$\limsup_{m\to\infty} |x'_m(x)| \leq \varepsilon.$$

By the arbitrariness of ε, as the argument holds for any $x \in X$, we conclude that $|x'_m(x)| \to 0$ as $m \to \infty$, namely $(x'_m)_{m\in\mathbf{N}}$ converges to 0 in the weak* topology.

We turn to the converse implication from (i). Suppose $B_{X'}$ is metrizable with distance d. Given any $n \in \mathbf{N}$, consider the set $U_n := \{x' \in B_{X'} : \mathsf{d}(x', 0) < 1/n\}$. Since d induces τ_{w^*}, there exists a sequence $\{E_n : n \in \mathbf{N}\}$ of finite subsets of X — say $E_n = \{x_1^n, \dots, x_{N_n}^n\}$ — and a sequence $(\varepsilon_n)_{n\in\mathbf{N}}$ of positive real numbers with

$$V_n := \{x' \in B_{X'} : |x'(x_i^n)| < \varepsilon_n \text{ for every } i \in \{1, \dots, N_n\}\} \subset U_n.$$

Define the countable set $D := \bigcup_{n\in\mathbf{N}} E_n$.

We claim it is linearly dense, which will terminate the proof. Indeed, let $x' \in X'$ be such that $x'(y) = 0$ for every $y \in D$. Either $\|x'\|_{X'} = 0$, or $\tilde{x}' := x'/\|x'\|_{X'} \in B_{X'}$ and $\tilde{x}'(y) = 0$ for all $y \in D$; but then by construction $\tilde{x}' \in U_n$ for all $n \in \mathbf{N}$, i.e. $\mathsf{d}(\tilde{x}', 0) = 0$, implying $\tilde{x}' = 0$, contradiction. We deduce $x' = 0$, so that D is linearly dense by Lemma C.8.

(ii) In view of the arguments from (i) — in particular, the construction of D and the choice of the metric d —, it is enough to show that, given a sequence $(x'_m)_{m\in\mathbf{N}}$ in $B_{X'}$, there exist a subsequence $(x'_{m_k})_{k\in\mathbf{N}}$ and $x' \in B_{X'}$ with the property $\lim_{k\to\infty} x'_{m_k}(x_n) = x'(x_n)$ for every $n \in \mathbf{N}$. This is a classical and general procedure called *Cantor's diagonal argument*. To construct the subsequence, we will actually inductively define a countable family of subsequences. To make the notation more manageable, it is convenient to denote subsequences by $(x'_{f(n)})_{n\in\mathbf{N}}$, or simply $(f(n))_{n\in\mathbf{N}}$, for suitable increasing functions $f\colon \mathbf{N}\to\mathbf{N}$.

We start by looking at $(x'_m(x_1))_{m\in\mathbf{N}}$. Because $|x'_m(x_1)|\le \|x'_m\|_{X'}\|x_1\|_X\le 1$ for every $m\in\mathbf{N}$, the sequence is bounded in $\mathbf{R}$. By the Bolzano–Weierstraß theorem, there exist a subsequence f^1 and $c_1\in\mathbf{R}$ such that $\lim_{m\to\infty} x'_{f^1(m)}(x_1)=c_1$. We can now run the same argument by looking at $(x'_{f^1(m)}(x_2))_{m\in\mathbf{N}}$. Then we find a subsequence f^2 of f^1 and $c_2\in\mathbf{R}$ such that $\lim_{m\to\infty} x'_{f^2(m)}(x_2)=c_2$. By an inductive procedure, starting from f^j, where $j\in\mathbf{N}$, we find a subsequence f^{j+1} and $c_{j+1}\in\mathbf{R}$ with the property $\lim_{m\to\infty} x'_{f^{j+1}(m)}(x_{j+1})=c_{j+1}$. This defines a family $\{f^j : j\in\mathbf{N}\}$ of increasing maps from $\mathbf{N}$ to $\mathbf{N}$, where f^{j+1} is a subsequence of f^j for every $j\in\mathbf{N}$. Finally, we define the sequence

$$F(m) := f^m(m)$$

which is the *diagonalization step*. By construction, F is eventually a subsequence of f^j for every $j\in\mathbf{N}$, therefore the limit of $(x'_{F(m)}(x_j))_{m\in\mathbf{N}}$ exists and coincides with that of $(x'_{f^j(m)}(x_j))_{m\in\mathbf{N}}$. In other words, every $n\in\mathbf{N}$ satisfies

$$\lim_{m\to\infty} x'_{F(m)}(x_n)=c_n.$$

We now define a linear operator A on $\operatorname{span} D$ by

$$A\Big[\sum_{n=1}^{k}\lambda_n x_n\Big] := \sum_{n=1}^{k}\lambda_n c_n. \tag{C.1}$$

Since the functionals x'_m belong to $B_{X'}$ as hypothesized, they are uniformly Lipschitz continuous with constant 1. This property is inherited by A: for every $y\in\operatorname{span} D$,

$$\|A(y)\|_X=\lim_{m\to\infty}\|x'_{F(m)}(y)\|_X\le \limsup_{m\to\infty}\|x'_{F(m)}\|_{X'}\,\|y\|_X\le\|y\|_X.$$

By a classical extension theorem, A admits a unique linear and Lipschitz continuous extension to $\overline{\operatorname{Span} D}=X$ that we will denote by x'. By the same argument, we have $\|x'\|_X\le 1$, implying $x'\in B_{X'}$. By (C.1), all in all we have constructed a subsequence $(x'_{F(m)})_{m\in\mathbf{N}}$ and $x'\in B_{X'}$ such that for every $n\in\mathbf{N}$,

$$\lim_{m\to\infty} x'_{F(m)}(x_n)=x'(x_n).$$

This implies $(x'_{F(m)})_{m\in\mathbf{N}}$ converges to x' in the weak* topology, as desired. □

Remark C.13 (Bounded sets). The second part of Theorem C.11 is stated for $B_{X'}$, but one can extend it as follows. Given a bounded set $E\subset X'$ and a sequence $(x'_n)_{n\in\mathbf{N}}$ in E, there exists a subsequence $(x'_{n_k})_{k\in\mathbf{N}}$ which converges weakly* to some point $x'\in X$. In other words, if X is separable, then **bounded sets in x' are sequentially precompact with respect to τ_{w^*}**. This is because, by dilations, one can immediately show $RB_{X'}=\{x'\in X' : \|x'\|_{X'}\le R\}$ is sequentially compact with respect to τ_{w^*} for every $R>0$. ■

Proof of Theorem C.4. Let $(x_n)_{n\in\mathbf{N}}$ be a bounded sequence in X and let Y denote the closure of the linear span of that sequence. By construction, Y is a closed linear subspace of X and $(Y,\|\cdot\|_X)$ is a separable Banach space. By Lemma C.5, Y is reflexive. By Theorem C.10, Y' and Y'' are separable and reflexive. Let J_Y be the canonical injection from Y to Y''. Then the sequence $(J_Y x_n)_{n\in\mathbf{N}}$ constitutes a bounded sequence in Y''. By Theorem C.11, we can extract a subsequence $(J_Y x_{n_k})_{k\in\mathbf{N}}$ which converges weakly* in Y'', thus by reflexivity to some $J_Y y$, where $y\in Y$. This means that, given any $y'\in Y'$,

$$\lim_{k\to\infty} y'(x_{n_k})=\lim_{k\to\infty} J_Y x_{n_k}(y')=J_Y y(y')=y'(y).$$

In other words, the sequence $(x_{n_k})_{k\in\mathbf{N}}$ converges weakly in Y to $y \in Y$. Since any $x' \in X'$ defines an element of Y' by its restriction $y' := x'\big|_Y$, we conclude $(x_{n_k})_{k\in\mathbf{N}}$ converges to y in the weak topology of X as well. □

Thus far, we have considered the abstract property of reflexivity, but given a Banach space X it might be quite hard to say whether it is reflexive or not. Many standard classes however have been extensively studied, cf. [3].

- Every Hilbert space H is reflexive. Indeed, in this case H is even isomorphic to its dual H' by the Riesz–Fréchet Theorem.
- Let $(\mathsf{M}, \mathfrak{m})$ be a measure space with $\mathfrak{m}$ σ-finite. Let $L^p(\mathsf{M}, \mathfrak{m})$ denote the associated Lebesgue space, where $p \in [1, \infty]$. Then $L^p(\mathsf{M}, \mathfrak{m})$ is reflexive provided $p \in (1, \infty)$. In this situation, one can identify its dual space with $L^q(\mathsf{M}, \mathfrak{m})$, where $1/p + 1/q = 1$. Applying this identification twice yields $L^p(\mathsf{M}, \mathfrak{m})'' = L^p(\mathsf{M}, \mathfrak{m})$. In the extremal case, $L^1(\mathsf{M}, \mathfrak{m})$ can still be identified with $L^\infty(\mathsf{M}, \mathfrak{m})$, but the converse is not true. Neither space is reflexive in general. (And $L^\infty(\mathsf{M}, \mathfrak{m})$ is not even separable in general.)
- Let $\Omega \subset \mathbf{R}^d$ be an open set, endowed with the Lebesgue measure. Denote the corresponding Lebesgue spaces by $L^p(\Omega, \mathcal{L}^d)$, where $p \in [1, \infty]$. Then for every $p \in [1, \infty)$, $L^p(\Omega, \mathcal{L}^d)$ is separable and $C_c^\infty(\Omega)$ is dense in it (by convolution).
- Let $\Omega \subset \mathbf{R}^d$ be an open and bounded set. Let $C(\overline{\Omega})$ denote the Banach space of continuous functions defined on its closure with the supremum norm. Then $C(\overline{\Omega})$ is separable by the Stone–Weierstrass theorem. Its dual can be identified with the space of signed Radon measures on $\overline{\Omega}$. Neither this set nor $C(\overline{\Omega})$ are reflexive.

Applying the results from this section, one can then deduce the following.

- On Hilbert spaces, closed bounded balls are weakly compact and sequentially compact. The same applies for Lebesgue spaces with exponents in $(1, \infty)$.
- $L^\infty(\mathsf{M}, \mathfrak{m})$ is the dual of the separable space $L^1(\mathsf{M}, \mathfrak{m})$, thus closed bounded balls in $L^\infty(\mathsf{M}, \mathfrak{m})$ are sequentially compact in the weak* topology.
- If Ω is bounded, the space of signed Radon measures on $\overline{\Omega}$ is the dual of the separable space $C(\overline{\Omega})$, thus closed bounded balls in the former are sequentially compact in the weak* topology.
- On the other hand, in the above situations boundedness is not a sufficient condition for weak precompactness in $L^1(\mathsf{M}, \mathfrak{m})$ or $C(\overline{\Omega})$. A characterization of weak compactness in $L^1(\mathsf{M}, \mathfrak{m})$ is given by the celebrated Dunford–Pettis theorem, cf. [3, Thm. 4.30].

Appendix D. Schauder–Tychonoff in Banach spaces

As mentioned in Remark 4.6, let us show the proof of Theorem 4.5 when X is a Banach space. The statement goes as follows.

Theorem D.1 (Schauder–Tychonoff in Banach spaces)**.** *Let X be a Banach space and $K \subset X$ be closed, convex, and nonempty. Let $F\colon K \to K$ be continuous such that $\overline{F(K)}$ is compact. Then F has a fixed point in K.*

Proof. Let $\varepsilon > 0$. Since $S := \overline{F(K)}$ is compact, we know there exists a finite subset $\{y_1, \dots, y_n\} \subset S$ such that $S \subset \bigcup_{i=1}^n B_\varepsilon(y_i)$[44]. Define the functions $g_1, \dots, g_n$ as

$$g_i(x) := \begin{cases} \varepsilon - \|x - y_i\| & \text{if } \|x - y_i\| \le \varepsilon, \\ 0 & \text{otherwise.} \end{cases}$$

[44]The balls are replacing the sets $y_i + U$ from the proof of Theorem 4.5. Indeed, $B_\varepsilon(y_i) = y_i + B_\varepsilon(0)$ and $B_\varepsilon(0)$ is clearly a convex, balanced and open neighborood of the origin.

Each g_i is continuous, $g_i(x) \geq 0$ and, since the balls are a covering of S, their sum is always strictly positive on S[45].

Define $C := \operatorname{co}\{y_1, \ldots, y_n\}$ and $g\colon S \to C$ by[46]

$$g(x) := \Big[\sum_{j=1}^{n} g_j(x)\Big]^{-1} \sum_{i=1}^{n} g_i(x)\, y_i. \tag{D.1}$$

This function is continuous and, since for any $x \in S$ we can write

$$x = \Big[\sum_{j=1}^{n} g_j(x)\Big]^{-1} \sum_{i=1}^{n} g_i(x)\, x$$

we deduce $\|g(x) - x\| \leq \varepsilon$ from the definition of $g_1, \ldots, g_n$.

Consider now the function $G := g \circ F\big|_C \colon C \to C$. Since C is a compact and convex subset of a finite-dimensional space and G is continuous, we can apply Brouwer's fixed point theorem to find $c_0 \in C$ such that $c_0 = G(c_0) = g(F(c_0))$. Hence, by the inequality above, we know that

$$\|F(c_0) - c_0\| = \|F(c_0) - g(F(c_0))\| \leq \varepsilon.$$

Replacing ε by $1/m$, given any $m \in \mathbf{N}$ there exists c_m such that

$$\|F(c_m) - c_m\| \leq \frac{1}{m}.$$

Since $\{F(c_m) : m \in \mathbf{N}\}$ is a sequence in the compact space S, there is a subsequence $\{c_{m_k} : k \in \mathbf{N}\}$ such that $F(c_{m_k}) \to x_0$ as $k \to \infty$ for some $x_0 \in S$. Moreover,

$$\|c_{m_k} - x_0\| \leq \|F(c_{m_k}) - c_{m_k}\| + \|F(c_{m_k}) - x_0\| \leq \frac{1}{m_k} + \|F(c_{m_k}) - x_0\|.$$

Sending $m_k \to \infty$ we see $c_{m_k} \to x_0$ since $F(c_{m_k}) \to x_0$. By the continuity of F, we conclude $F(c_{m_k}) \to F(x_0)$ and therefore $F(x_0) = x_0$. □

Appendix E. Bochner integration on Banach spaces

In this part, we collect some properties of integration of Banach-space valued functions. For a detailed account, we refer to e.g. the book of Diestel–Uhl [5]. Much of the material to follow should be strongly reminiscent of integration theory from measure theory for $\mathbf{R} \cup \{-\infty, \infty\}$-valued functions, modulo some peculiarities from the fact that the functions we study here take values in Banach spaces.

Let X be a Banach space. As usual, we denote its norm by $\|\cdot\|$ and its dual space by X'.

A function $x\colon [0,1] \to X$[47] is called simple if it assumes only finitely many values in X. More precisely, there exist $k \in \mathbf{N}$, Borel subsets $E_1, \ldots, E_k \subset [0,1]$, and $v_1, \ldots, v_k \in X$ such that for every $t \in [0,1]$,

$$x(t) = \sum_{i=1}^{k} 1_{E_i}(t)\, v_i. \tag{E.1}$$

A map $x\colon [0,1] \to X$ is called

- **strongly measurable** if there is a sequence $(x_n)_{n\in\mathbf{N}}$ of simple functions such that $(\|x_n - x\|)_{n\in\mathbf{N}}$ converges to zero $\mathcal{L}^1$-a.e. and
- **weakly measurable** if for every $x' \in X'$, the map $f\colon [0,1] \to \mathbf{R}$ given by $f(t) := x'(x(t))$ is Borel measurable.

[45] This is the analog the lower bound on the sum of the distances in Lemma 4.4.

[46] Notice the analogy with the function F defined in Lemma 4.4.

[47] We choose the domain $[0,1]$ to simplify the presentation. It could also be taken to be $\mathbf{R}$ or any nontrivial subinterval thereof.

Linear combinations of strongly measurable functions are strongly measurable; the analog holds for weak measurability. Moreover, if the map $x\colon [0,1] \to X$ is strongly measurable, then its norm $\|x\|$ is a Borel measurable map from $[0,1]$ to $\mathbf{R}_+$.

A precise relation between strong and weak measurability is the following.

Theorem E.1 (Pettis)**.** *A map $x\colon [0,1] \to X$ is strongly measurable if and only if it is weakly measurable and almost separably valued, i.e. there exists a Borel set $N \subset [0,1]$ with $\mathcal{L}^1[N] = 0$ such that $x([0,1] \setminus N)$ is a separable subset of X.*

In particular, if X is separable, strong and weak measurability are equivalent.

Now we turn to integration of such functions, the so-called **Bochner integrals**. Given a simple function $x\colon [0,1] \to X$ of the form (E.1), we set

$$\int_0^1 x(t)\,\mathrm{d}t := \sum_{i=1}^k \mathcal{L}^1[E_i]\,v_i.$$

Note this integral is X-valued by definition. It does not depend on the particular way (E.1) is written. Given $x\colon [0,1] \to X$ strongly measurable, we then say x is Bochner integrable if there exists a sequence $(x_n)_{n\in\mathbf{N}}$ of simple functions such that $\int_0^1 \|x(t) - x_n(t)\|\,\mathrm{d}t \to 0$ as $n \to \infty$. In this case, $(\int_0^1 x_n(t)\,\mathrm{d}t)_{n\in\mathbf{N}}$ is a Cauchy sequence in X, hence the following quantity is well-defined:

$$\int_0^1 x(t)\,\mathrm{d}t := \lim_{n\to\infty} \int_0^1 x_n(t)\,\mathrm{d}t.$$

This integral is independent of the choice of the sequence $(x_n)_{n\in\mathbf{N}}$ with the above properties. Moreover, the following “triangle inequality” holds:

$$\Big\|\int_0^1 x(t)\,\mathrm{d}t\Big\| \leq \int_0^1 \big\|x(t)\big\|\,\mathrm{d}t. \tag{E.2}$$

Given any Bochner integrable map $x\colon [0,1] \to X$ and any Borel measurable set $B \subset [0,1]$, we also define

$$\int_B x(t)\,\mathrm{d}t := \int_0^1 1_B(t)\,x(t)\,\mathrm{d}t.$$

The following is a convenient characterization of Bochner integrability.

Theorem E.2 (Bochner)**.** *A strongly measurable function $x\colon [0,1] \to X$ is Bochner integrable if and only if*

$$\int_0^1 \big\|x(t)\big\|\,\mathrm{d}t < \infty.$$

Armed with an integration theory for X-valued functions, we can now introduce Lebesgue and Sobolev spaces. Moreover, it allows us to define absolute continuity.

Definition E.3 (Lebesgue spaces)**.** *Given any $p \in [1,\infty]$, the space $L^p([0,1];X)$ is the space of (equivalence classes up to $\mathcal{L}^1$-a.e. equality of) those strongly measurable maps $x\colon [0,1] \to X$ such that $\|x\|_{L^p([0,1];X)} < \infty$, where*

$$\|x\|_{L^p([0,1];X)} := \begin{cases} \Big[\displaystyle\int_0^1 \big\|x(t)\big\|^p\,\mathrm{d}t\Big]^{1/p} & \textit{provided } p < \infty,\\ \mathcal{L}^1\text{-}\operatorname*{esssup}_{t\in[0,1]} \big\|x(t)\big\| & \textit{otherwise.}\end{cases}$$

Definition E.4 (Sobolev spaces)**.** *Given any $p \in [1,\infty]$, the space $W^{1,p}([0,1];X)$ consists of those $x \in L^p([0,1];X)$ such that there exists an element $x' \in L^p([0,1];X)$ such that for every $\varphi \in C_c^\infty((0,1))$,*

$$\int_0^1 \varphi'(t)\,x(t)\,\mathrm{d}t = -\int_0^1 \varphi(t)\,x'(t)\,\mathrm{d}t. \tag{E.3}$$

For every $p \in [1, \infty]$, the Sobolev space $W^{1,p}([0,1]; X)$ becomes a Banach space with respect to the norm

$$\|x\|_{W^{1,p}([0,1];X)} := \big[\|x\|_{L^p([0,1];X)} + \|x'\|_{L^p([0,1];X)}\big]^{1/p}.$$

In quite general settings, there is a one-to-one correspondence between Sobolev functions and absolutely continuous functions in Banach spaces.

Definition E.5 (Absolutely continuous curves)**.** *Given any $p \in [1, \infty]$, the space $\mathrm{AC}^p([0,1]; X)$ consists of all maps $x\colon [0,1] \to X$ such that there is $f \in L^p([0,1]; \mathcal{L}^1)$ such that for every $s, t \in [0,1]$ with $s < t$,*

$$\big\|x(t) - x(s)\big\| \leq \int_s^t f(r)\,\mathrm{d}r.$$

In particular, $\mathrm{AC}^\infty([0,1]; X)$ is the set of Lipschitz continuous maps $x\colon [0,1] \to X$.

The following two general results, stated without proof, verify the one-to-one correspondence outlined above.

Proposition E.6 (Absolutely continuous representative)**.** *Given any $p \in [1, \infty]$ and any $x \in W^{1,p}([0,1]; X)$, there exists $\tilde{x} \in \mathrm{AC}^p([0,1]; X)$ with $x = \tilde{x}$ $\mathcal{L}^1$-a.e.*

Moreover, for every $s, t \in [0,1]$ with $s < t$, the representative $\tilde{x}$ satisfies

$$\tilde{x}(t) - \tilde{x}(s) = \int_s^t x'(r)\,\mathrm{d}r.$$

In general Banach spaces, absolute continuity does not imply a.e. differentiability. In other words, the fundamental theorem of calculus cannot be turned into a statement about the derivative of the function in question. This property is connected to the so-called Radon–Nikodým property of Banach spaces. A simple sufficient criterion is separability; therefore, the one-to-one correspondence holds e.g. on every Hilbert space, the main setting of §5.

Theorem E.7 (Sobolev representative and a.e. differentiability)**.** *Assume X is reflexive. Given any $p \in [1, \infty]$ and any $x \in \mathrm{AC}^p([0,1]; X)$, for $\mathcal{L}^1$-a.e. $t \in [0,1]$ the following limit exists in X:*

$$x'(t) := \lim_{h \to 0} \frac{x(t+h) - x(t)}{h}.$$

The function x' thus defined — with e.g. constant extension beyond the set of all $t \in [0,1]$ for which the above limit does not exist — belongs to $L^p([0,1]; X)$ and satisfies (E.3) for every $\varphi \in C_c^\infty((0,1))$; in particular, $x \in W^{1,p}([0,1]; X)$.

References

[1] L. Ambrosio, N. Gigli, G. Savaré, *Gradient flows in metric spaces and in the space of probability measures.* Second edition, Birkhäuser, Basel, 2008.

[2] H. Brezis, *Opérateurs maximaux monotones et semi-groupes de contractions dans les espaces de Hilbert.* North-Holland Publishing Co., Amsterdam, 1973.

[3] ———, *Functional analysis, Sobolev spaces and partial differential equations.* Springer, New York, 2011.

[4] K. Deimling, *Nonlinear functional analysis,* Springer, Berlin-Heidelberg, 1985.

[5] J. Diestel, J. J. Uhl, *Vector measures,* American Mathematical Society, Providence, R.I., 1977

[6] L. C. Evans, *Partial differential equations*, Providence, RI, 2010.

[7] R. Jordan, D. Kinderlehrer, F. Otto, *The variational formulation of the Fokker-Planck equation.* SIAM J. Math. Anal. **29** (1998), no. 1, 1–17.

[8] F. Otto. *The geometry of dissipative evolution equations: the porous medium equation.* Comm. Partial Differential Equations **26** (2001), no. 1-2, 101–174.

[9] R. T. Rockafellar, *Convex Analysis.* Princeton University Press, Princeton, 1970.

[10] W. Rudin, *Functional analysis,* McGraw-Hill, Singapore, 1991.

Institute of Mathematics, EPFL, 1015 Lausanne, Switzerland

Email address: mathias.braun@epfl.ch