Functional Analysis, Winter Semester 2025–26, HU Berlin

Chris Wendl

Contents

These are lecture notes for the course Functional Analysis (*Funktionalanalysis*) in the 2025–26 winter semester at Humboldt University, Berlin.

Since the notes were designed for use at a German university, I have made an effort to include the German translations (*geschrieben in dieser Schriftart*) of important terms wherever they are introduced, but I have not tried to do this comprehensively.

Disclaimer: These lecture notes were written quickly, so there will be typos, and possibly even some mathematical errors. If you notice any, please send me an e-mail and I will correct. Many thanks to Yutong Dai for pointing out many small errors that have since been corrected.

For more detailed treatments of many of the topics in these notes, I mainly recommend the books by Bühler-Salamon [**BS18**], Reed-Simon [**RS80**] (especially for general Banach/Hilbert space and spectral theory), and Lieb-Loss [**LL01**] (for L^p theory, distributions and Sobolev spaces).

Most recent update: **April 5, 2026**

Part 1: Fundamentals

In this first portion of the course, the aim is to give a general idea of what functional analysis is and what it's good for, and then to develop some basic building blocks in preparation for the deeper results to come. Most theorems in Part 1 are relatively easy; the deepest is probably the Riesz representation theorem, which gives a natural isomorphism between every Hilbert space and its own dual space.

1. Banach spaces

1.1. What is functional analysis? One can distinguish between *linear* and *nonlinear* functional analysis; the first is necessarily a prerequisite for the second. On a first pass, I would describe linear functional analysis as an attempt to extend the familiar principles of linear algebra as far as possible into infinite-dimensional vector spaces. Nonlinear functional analysis, in turn, tries to extend differential geometry into the realm of infinite-dimensional manifolds. Both descriptions are overly abstract and incomplete, because both, in particular, fail to explain why the word "functional" appears in the name of the subject. There is a simple reason: the actual examples of infinite-dimensional spaces we consider in this subject are almost always spaces of *functions*, and the main motivation for doing so is to prove theorems about (ordinary or partial) differential equations, i.e. ODEs or PDEs.

You've likely seen at least one example of this in your previous analysis courses: in the basic theory of ODEs, the local existence and uniqueness of solutions to initial value problems (i.e. the Picard-Lindelöf theorem) is proved by constructing a Banach space of functions, and arguing that one of those functions solves the problem if and only if it is a fixed point of a certain transformation on the Banach space. In this way, the existence and uniqueness of solutions can be deduced from a general result about abstract Banach spaces, namely the contraction mapping principle, also known as the Banach fixed point theorem.

There are many more ways to deduce highly non-obvious results about differential equations from general theorems about Banach spaces, and we will see a few more examples in this course, starting already in §1.4 below. If you continue to an actual course on PDEs next semester, it will likely make heavy use of the results introduced in this one. That said, this is explicitly a course on *linear* functional analysis, and the differential equations to which we see the methods applied will therefore be linear.

1.2. Bounded linear operators. Before anything more can be said, we need a few definitions. We assume the reader is already familiar with the notions of normed vector spaces (*normierte Vektorräume*), inner products (*Skalarprodukte*), linear maps (*lineare Abbildungen*), metric spaces (*metrische Räume*), continuity (*Stetigkeit*), Cauchy sequences (*Cauchyreihen*) and completeness (*Vollständigkeit*). You have probably also seen Banach spaces before, but since it is one of the most central notions in functional analysis, we'll repeat the definition here.

Definition 1.1. A **Banach space** (*Banachraum*) is a normed vector space $(X, \|\cdot\|)$ that is complete, i.e. in which every Cauchy sequence converges.

EXAMPLE 1.2. $\mathbb{R}^n$ with the standard Euclidean norm is a Banach space, and so is $\mathbb{C}^n$ with the norm derived from the standard Hermitian inner product (which is equivalent to the Euclidean $\mathbb{R}^{2n}$). In fact, it is not hard to show that all *finite-dimensional* normed vector spaces are complete, and are therefore Banach spaces.

EXAMPLE 1.3. We equip the infinite-dimensional vector space

$$C^0([0,1]) := \{\text{continuous functions} f : [0,1] \to \mathbb{R}\}$$

with the so-called **sup-norm**, also known as the C^0**-norm**,

$$\|f\|_{C^0} := \sup_{t\in[0,1]} |f(t)|.$$

Note that since $[0,1]$ is compact, a standard theorem from first-year analysis guarantees that $|f(t)|$ actually attains a maximum, so long as $f : [0,1] \to \mathbb{R}$ is continuous, thus the C^0-norm on $C^0([0,1])$ can equivalently be written as $\max_{t\in[0,1]} |f(t)|$. A sequence $f_n \to f$ of functions in $C^0([0,1])$ converges in the C^0-norm if and only if it converges *uniformly*. By another standard result in first-year analysis, uniformly Cauchy sequences of continuous functions also converge uniformly to continuous functions, thus $C^0([0,1])$ with the C^0-norm is a Banach space.

An easy example of an *incomplete* normed vector space is obtained by restricting the C^0-norm to the subspace

$$C^\infty([0,1]) \subset C^0([0,1])$$

consisting of smooth (i.e. infinitely differentiable) functions. A uniformly Cauchy sequence of smooth functions can easily converge to a function that is continuous but not smooth, thus producing a Cauchy sequence with no limit in $C^\infty([0,1])$.

Many "obvious" facts from finite-dimensional linear algebra become false in general if one allows infinite-dimensional vector spaces. One such fact is that linear subspaces $V \subset X$ are also *closed* subsets. In infinite dimensions, the subspace $C^\infty([0,1]) \subset C^0([0,1])$ provides an easy counterexample; in fact, by the famous theorem of Weierstrass that continuous functions can be approximated uniformly by polynomials, $C^\infty([0,1])$ is *dense* in $C^0([0,1])$, meaning that the closure of $C^\infty([0,1])$ in $C^0([0,1])$ with respect to the C^0-norm is $C^0([0,1])$ itself.

Another important fact that you are unlikely to have questioned before is that linear maps between finite-dimensional vector spaces are also *continuous* maps. In infinite dimensions, this is catastrophically false, though I will need you to be patient and wait for §2.7 before I can show you an actual counterexample. A related fact we will be able to see more explicitly is that in an infinite-dimensional normed vector space $(X, \|\cdot\|)$, the unit sphere

$$\{x \in X \mid \|x\| = 1\} \subset X$$

is not a compact subset. The next definition and subsequent theorem provide a useful connection between these two observations.

DEFINITION 1.4. A linear map $A : X \to Y$ between normed vector spaces is called **bounded** (*beschränkt*) if there exists a constant $c > 0$ such that $\|Ax\| \leqslant c\|x\|$ holds for all $x \in X$, i.e. the supremum

$$\|A\| := \sup_{x\in X\setminus\{0\}} \frac{\|Ax\|}{\|x\|} \tag{1.1}$$

is finite.

REMARK 1.5. In the general theory of (not necessarily linear) maps $X \to Y$ between metric spaces, a map is called bounded if and only if its image is a bounded subset of Y. Nontrivial linear maps *never* have this property, thus the word "bounded" carries a slightly different meaning in Definition 1.4.

Observe that for a linear map $A : X \to Y$, rescaling allows us to write $\|A\|$ equivalently as

$$\|A\| = \sup_{\|x\|=1} \|Ax\|.$$

If the unit sphere in X were compact and both $A : X \to Y$ and the norm $\|\cdot\| : Y \to \mathbb{R}$ were assumed continuous, it would follow immediately that the maximum of $\|Ax\|$ is attained and $\|A\|$ is thus finite. In fact, the norm $\|\cdot\| : Y \to \mathbb{R}$ is always a continuous function; this follows from the so-called "reverse" triangle inequality,

$$\|x - y\| \geqslant \big|\|x\| - \|y\|\big|, \tag{1.2}$$

which can be derived from the usual triangle inequality by writing

$$\|x\| - \|y\| = \|(x - y) + y\| - \|y\| \leqslant \|x - y\| + \|y\| - \|y\| = \|x - y\|,$$

and similarly $\|y\| - \|x\| \leqslant \|y - x\|$. In general, however, the unit sphere is not compact, and $A : X \to Y$ need not be continuous. We nonetheless have:

THEOREM 1.6. *A linear map $A : X \to Y$ between normed vector spaces is continuous if and only if it is bounded.*

PROOF. If A is bounded, then for any sequence $x_n \to x$ in X, one has

$$\|Ax - Ax_n\| = \|A(x - x_n)\| \leqslant \|A\| \cdot \|x - x_n\| \to 0,$$

implying $Ax_n \to Ax$, thus A is continuous. Conversely, if A is not bounded, then there exists a sequence $x_n \neq 0 \in X$ such that

$$\frac{\|Ax_n\|}{\|x_n\|} \to \infty.$$

The sequence $y_n := x_n / \|Ax_n\|$ then satisfies $y_n \to 0$, but $\|Ay_n\| = 1$ for all n, implying that Ay_n cannot converge to $0 = A(0)$, and A is therefore discontinuous at 0. □

You will easily convince yourself that the set of all continuous linear maps between two given normed vector spaces defines a vector space in an obvious way, with vector addition and scalar multiplication defined pointwise. In this course, we'll denote this vector space by

$$\mathscr{L}(X, Y) := \{A : X \to Y \mid A \text{ continuous and linear}\}.$$

In functional analysis, linear maps $A : X \to Y$ are also often called **operators**, and elements of $\mathscr{L}(X, Y)$ are thus called **continuous/bounded linear operators** (*stetige/beschränkte lineare Operatoren*). We use the so-called **operator norm** (*Operatornorm*) (1.1) to make $\mathscr{L}(X, Y)$ into a normed vector space. To see that the operator norm satisfies the triangle inequality, suppose $A, B \in \mathscr{L}(X, Y)$, so the inequalities $\|Ax\| \leqslant \|A\| \cdot \|x\|$ and $\|Bx\| \leqslant \|B\| \cdot \|x\|$ hold for all $x \in X$. Then using the triangle inequality for the norm on Y, we have

$$\|(A + B)x\| = \|Ax + Bx\| \leqslant \|Ax\| + \|Bx\| \leqslant \|A\| \cdot \|x\| + \|B\| \cdot \|x\| = (\|A\| + \|B\|)\,\|x\|$$

for all $x \in X$, proving $\|A + B\| \leqslant \|A\| + \|B\|$.

REMARK 1.7. There is no obvious way to introduce a norm on the space of *all* linear maps $X \to Y$, but we can do this if continuity is required, because Theorem 1.6 then makes the operator norm finite.

THEOREM 1.8. *If Y is complete, then so is $\mathscr{L}(X, Y)$.*

Proof. Assume $A_n \in \mathscr{L}(X,Y)$ is a Cauchy sequence. Then for every $x \in X$,

$$\|A_n x - A_m x\| = \|(A_n - A_m)x\| \leqslant \|A_n - A_m\| \cdot \|x\|$$

is small whenever m and n are large, proving that $A_n x$ is a Cauchy sequence in Y, and therefore converges. Define a map $A : X \to Y$ by

$$Ax := \lim_{n\to\infty} A_n x.$$

One easily checks that A is linear. It remains to show that A is bounded and $A_n \to A$ in the operator norm.

We claim: $\|A_n\|$ is a Cauchy sequence in $\mathbb{R}$. Indeed, using the reverse triangle inequality (1.2), we have

$$\|A_n - A_m\| \geqslant \big|\|A_n\| - \|A_m\|\big|,$$

proving that the latter is small for m, n large. It follows that $\|A_n\|$ is a convergent sequence of real numbers, and we can therefore define

$$C := \lim_{n\to\infty} \|A_n\| \geqslant 0.$$

Now for any given $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that

$$n \geqslant N \quad \Rightarrow \quad \|A_n\| \leqslant C + \epsilon,$$

implying that for all $x \in X$,

$$n \geqslant N \quad \Rightarrow \quad \|A_n x\| \leqslant (C + \epsilon)\|x\|.$$

Since $A_n x \to Ax$ and the norm is a continuous function, it follows that $\|Ax\| \leqslant (C + \epsilon)\|x\|$ for all x, hence A is bounded, i.e. $A \in \mathscr{L}(X,Y)$.

To see that $A_n \to A$ in the operator norm, suppose $\epsilon > 0$ is given, and fix $N \in \mathbb{N}$ large enough so that

$$m, n \geqslant N \quad \Rightarrow \quad \|A_m - A_n\| \leqslant \epsilon,$$

hence for all $x \in X$,

$$m, n \geqslant N \quad \Rightarrow \quad \|A_m x - A_n x\| \leqslant \epsilon\|x\|.$$

Fixing n but letting $m \to \infty$, we have $A_m x \to Ax$ and thus deduce from this inequality

$$n \geqslant N \quad \Rightarrow \quad \|Ax - A_n x\| \leqslant \epsilon\|x\|$$

for all $x \in X$. This can be interpreted as the statement that $\|A - A_n\| \leqslant \epsilon$ for all $n \geqslant N$, and thus proves $A_n \to A$. □

Corollary 1.9. *For any Banach space X, the space*

$$\mathscr{L}(X) := \mathscr{L}(X,X)$$

of bounded linear operators from X to itself, equipped with the operator norm, is also a Banach space. □

1.3. Convergent series and invertible operators. Recall that a series $\sum_{n=1}^{\infty} x_n$ in a normed vector space is said to **converge absolutely** if

$$\sum_n \|x_n\| < \infty.$$

An important theorem from first-year analysis states that if X is a Banach space, then every absolutely convergent series in X is also convergent, meaning its sequence of partial sums converges with respect to the norm of X.

Remark 1.10. Depending what kind of analysis courses you've had before, you may or may not have seen the theorem about absolute convergence I just mentioned in the general form we're about to need. What you have definitely seen is a theorem stating that absolutely convergent series of *real numbers* are also convergent. If you go back and look at the proof of that theorem, you'll find that the most important property of the real numbers it uses is that Cauchy sequences converge, thus it generalizes easily to prove the same result for series in arbitrary Banach spaces.

Lemma 1.11. *Assume X is a Banach space and let $\mathbb{1} \in \mathscr{L}(X)$ denote the identity operator from X to itself. Then for every $A \in \mathscr{L}(X)$ with $\|A\| < 1$, the map $\mathbb{1} + A \in \mathscr{L}(X)$ has a bounded inverse $(\mathbb{1} + A)^{-1} \in \mathscr{L}(X)$.*

Proof. Let $B := \mathbb{1} - A + A^2 - A^3 + \ldots$. One easily checks that $\|A^k\| \leqslant \|A\|^k$ for all $k \in \mathbb{N}$, thus

$$\sum_{n=0}^{\infty} \|(-1)^n A^n\| \leqslant \sum_{n=0}^{\infty} \|A\|^n < \infty$$

since $\|A\| < 1$, and the series defining B therefore converges absolutely, making B a well-defined element of $\mathscr{L}(X)$. We leave it as an exercise to check

$$B(\mathbb{1} + A) = (\mathbb{1} + A)B = \mathbb{1}.$$

□

Corollary 1.12. *If X, Y are Banach spaces and $A \in \mathscr{L}(X, Y)$ has a bounded inverse $A^{-1} \in \mathscr{L}(Y, X)$, then so does $A + B$ for all $B \in \mathscr{L}(X, Y)$ with $\|B\|$ sufficiently small.*

Proof. Write $A + B = A\left(\mathbb{1} + A^{-1}B\right)$, then

$$(A + B)^{-1} = (\mathbb{1} + A^{-1}B)^{-1}A^{-1}$$

makes sense whenever $\|B\| < 1/\|A^{-1}\|$, so that $\|A^{-1}B\| \leqslant \|A^{-1}\| \cdot \|B\| < 1$. □

1.4. Application: A boundary value problem. The small amount of general theory developed so far already has nontrivial applications to the study of differential equations. Here is a demonstration.

Given continuous functions $P, f : [0, 1] \to \mathbb{R}$, consider the following second-order boundary value problem for functions $x : [0, 1] \to \mathbb{R}$,

$$\begin{aligned} \ddot{x}(t) + P(t)x(t) &= f(t), \\ x(0) = x(1) &= 0, \end{aligned} \tag{1.3}$$

where $\dot{x}$ and $\ddot{x}$ denote the first and second derivatives of x respectively. For a functional-analytic approach to this problem, we can define a normed vector space

$$X := \left\{ x : [0, 1] \to \mathbb{R} \;\middle|\; x \text{ is of class } C^2 \text{ and } x(0) = x(1) = 0 \right\},$$

with the so-called C^2-norm

$$\|x\|_{C^2} := \|x\|_{C^0} + \|\dot{x}\|_{C^0} + \|\ddot{x}\|_{C^0}.$$

Let $Y := C^0([0, 1])$, with the C^0-norm, and for a given continuous function $P : [0, 1] \to \mathbb{R}$, define the linear map

$$T_P : X \to Y : x \mapsto \ddot{x} + Px.$$

As observed in §1.2, Y is a Banach space.

Exercise 1.13. Show that X with the C^2-norm is also a Banach space, and $T_P : X \to Y$ is a bounded linear operator for every $P \in C^0([0, 1])$.

Looking at the special case $P = 0$, we claim that the operator $T_0 : X \to Y : x \mapsto \ddot{x}$ has a bounded inverse. Indeed, given any function $f \in Y = C^0([0,1])$, one can integrate twice to find a unique C^2-function $y : [0,1] \to \mathbb{R}$ such that $\ddot{y} = f$ and $y(0) = \dot{y}(0) = 0$. The unique C^2-function $x : [0,1] \to \mathbb{R}$ satisfying $\ddot{x} = f$ and $x(0) = x(1) = 0$ is then given by

$$x(t) = y(t) - y(1)t.$$

To see that the inverse $f \mapsto x$ obtained in this way is bounded, we can first use the fundamental theorem of calculus to write

$$\dot{y}(t) = \int_0^t f(s)\,ds, \qquad y(t) = \int_0^t \dot{y}(s),$$

deducing from the first relation the estimate

$$\|\dot{y}\|_{C^0} = \max_{t\in[0,1]} |\dot{y}(t)| = \max_{t\in[0,1]} \left| \int_0^t f(s)\,ds \right| \leqslant \max_{t\in[0,1]} \int_0^t |f(s)|\,ds \leqslant \int_0^1 |f(s)|\,ds \leqslant \int_0^1 \|f\|_{C^0}\,ds = \|f\|_{C^0},$$

and subsequently, by the same trick, the estimate

$$\|y\|_{C^0} \leqslant \|\dot{y}\|_{C^0} \leqslant \|f\|_{C^0}.$$

These imply

$$|x(t)| \leqslant |y(t)| + |y(1)|t \leqslant 2\|y\|_{C^0} \leqslant 2\|f\|_{C^0}$$

for all $t \in [0,1]$, hence $\|x\|_{C^0} \leqslant 2\|f\|_{C^0}$, and similarly

$$|\dot{x}(t)| \leqslant |\dot{y}(t)| + |y(1)| \leqslant \|\dot{y}\|_{C^0} + \|y\|_{C^0} \leqslant 2\|f\|_{C^0}$$

for all $t \in [0,1]$, hence $\|\dot{x}\|_{C^0} \leqslant 2\|f\|_{C^0}$. We conclude

$$\|x\|_{C^2} = \|x\|_{C^0} + \|\dot{x}\|_{C^0} + \|\ddot{x}\|_{C^0} \leqslant 2\|f\|_{C^0} + 2\|f\|_{C^0} + \|f\|_{C^0} = 5\|f\|_{C^0},$$

proving that the linear map $Y \to X : f \mapsto x$ is bounded.

Now, observe that for all $x \in X$ and for all $P \in C^0([0,1])$,

$$\|(T_P - T_0)x\|_{C^0} = \|Px\|_{C^0} \leqslant \|P\|_{C^0}\|x\|_{C^0} \leqslant \|P\|_{C^0}\|x\|_{C^2},$$

so the operator norm for $T_P - T_0 : X \to Y$ satisfies

$$\|T_P - T_0\| \leqslant \|P\|_{C^0}.$$

Plugging this information into Corollary 1.12, we conclude:

THEOREM 1.14. *There exists a constant $c > 0$ such that for every pair of continuous functions $P, f : [0,1] \to \mathbb{R}$ satisfying $\|P\|_{C^0} < c$, the boundary value problem* (1.3) *has a unique solution.* □

With a little more effort, one can also write down precise bounds on the constant $c > 0$ in this result and an explicit formula for the solution $x(t)$ as an absolutely and uniformly convergent series of functions.

1.5. Some examples of Banach spaces. We conclude this preliminary introduction to Banach spaces with some more examples that will be mentioned again many times in the rest of the course.

CONVENTION 1.15. Unless otherwise noted, functions throughout this course will be assumed to take values in a fixed finite-dimensional inner product space $(V, \langle\ ,\ \rangle)$ over the field $\mathbb{K}$, which is either $\mathbb{R}$ or $\mathbb{C}$, and the norm on V determined by this inner product will be denoted by $|\cdot|$. It will not usually be necessary to specify a concrete choice of the space V, but the most popular choices in applications are $\mathbb{R}$ and $\mathbb{C}$ with their standard inner products, followed by $\mathbb{R}^m$ and $\mathbb{C}^m$ for arbitrary $m \in \mathbb{N}$.

CONVENTION 1.16. We will make frequent use of multi-index notation for partial derivatives of arbitrary order. A **multi-index** of **order** $k \geqslant 0$ in $n \geqslant 0$ variables is by definition an n-tuple of nonnegative integers

$$\alpha = (\alpha_1, \dots, \alpha_n) \qquad \text{with} \qquad |\alpha| := \sum_{j=1}^{n} \alpha_j = k.$$

Every such multi-index determines a differential operator of order k for functions of n real variables $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, defined by

$$\partial^\alpha := \partial_1^{\alpha_1} \dots \partial_n^{\alpha_n} = \left(\frac{\partial}{\partial x_1}\right)^{\alpha_1} \dots \left(\frac{\partial}{\partial x_n}\right)^{\alpha_n},$$

as well as a real-valued polynomial function on $\mathbb{R}^n$, written

$$x^\alpha := x_1^{\alpha_1} \dots x_n^{\alpha_n}.$$

EXAMPLE 1.17. Given an open subset $\Omega \subset \mathbb{R}^n$ and an integer $m \geqslant 0$, the C^m**-norm** can be defined for C^m-functions on Ω by

$$\|f\|_{C^m} := \sum_{|\alpha| \leqslant m} \sup_{x \in \Omega} |\partial^\alpha f(x)|,$$

where the sum ranges over the set of all multi-indices of order at most m. Since continuous functions on Ω need not be bounded, this norm can be infinite, but it is finite on the subspace

$$C_b^m(\Omega) := \left\{ f : \Omega \to V \;\middle|\; f \text{ is of class } C^m \text{ and } \|f\|_{C^m} < \infty \right\},$$

consisting of all bounded C^m-functions on Ω whose partial derivatives up to order m are also bounded. One then has convergence $f_k \to f$ of functions in $C_b^m(\Omega)$ if and only if the partial derivatives of f_k up to order m converge uniformly on Ω to the corresponding partial derivatives of f. We will sometimes denote this notion of convergence by

$$f_k \xrightarrow{C^m} f.$$

THEOREM 1.18. *The space $C_b^m(\Omega)$ with the C^m-norm is a Banach space.*

PROOF. If $f_k \in C_b^m(\Omega)$ is a Cauchy sequence in the C^m-norm, then $\partial^\alpha f_k$ is uniformly Cauchy for each multi-index with $|\alpha| \leqslant m$, so by first-year analysis, there exists a continuous function g_α on Ω such that $\partial^\alpha f_k \to g_\alpha$ uniformly. This is true in particular for $\alpha := 0 = (0, \dots, 0)$, so we can write $f := g_0$ and say $f_k \to f$ uniformly. We claim that f is of class C^m and satisfies $\partial^\alpha f = g_\alpha$ for each $|\alpha| \leqslant m$; this will imply $f_k \to f$ in the C^m-norm.

Consider first the case $|\alpha| = 1$: the task is then to show that for each $j = 1, \dots, n$, if $\partial_j f_k$ converges uniformly on Ω as $k \to \infty$ to a continuous function g_j, then the jth partial derivative of f exists and matches g_j. For any $x \in \Omega$, pick $\epsilon > 0$ small enough so that the ball of radius ϵ about x is contained in Ω, pick $h \in \mathbb{R}$ with $|h| < \epsilon$, and let $e_j \in \mathbb{R}^n$ denote the jth standard basis vector. Then letting $k \to \infty$ on both sides of the relation

$$f_k(x + he_j) = f_k(x) + h \int_0^1 \partial_j f_k(x + the_j)\, dt$$

produces the relation

$$f(x + he_j) = f(x) + h \int_0^1 g_j(x + the_j)\, dt,$$

where for the right hand side, we are using the uniform convergence $\partial_j f_k \to g_j$ to conclude the convergence of the sequence of integrals. Taking the limit as $h \to 0$, this proves that $\partial_j f(x)$ exists and equals $g_j(x)$.

Similar arguments prove the result for all $|\alpha| \leqslant m$ by induction on the order of the multi-index. $\square$

EXAMPLE 1.19. The following spaces should be at least slightly familiar to you from an introductory course on measure theory, and we will have a lot more to say about them in Part 2 of this course. Assume (X, μ) is a measure space (*Maßraum*), i.e. X is a set, and μ is a function assigning numbers $\mu(\mathcal{U}) \in [0, \infty]$ to suitable subsets $\mathcal{U} \subset X$ that can be interpreted as the "volumes" of these subsets, so that integrals of suitable functions on X can be defined. Recall that a function on X is then said to have a certain property **almost everywhere** (*fast überall*), abbreviated "a.e.", if its restriction to some subset $X' \subset X$ with $\mu(X \backslash X') = 0$ has that property. For any $p \geqslant 1$, we define the **L^p-norm** for measurable functions on X by

$$\|f\|_{L^p} := \|f\|_{L^p(X)} := \left(\int_X |f|^p \, d\mu \right)^{1/p} \in [0, \infty].$$

For reasons that will be discussed in Part 2, it is also natural to extend this definition to $p = \infty$ by defining the L^∞-norm of f to be its **essential supremum**,

$$\|f\|_{L^\infty} := \|f\|_{L^\infty(X)} := \operatorname{ess\,sup}_{x \in X} |f(x)| := \inf \left\{ c \geqslant 0 \;\middle|\; |f| \leqslant c \text{ almost everywhere} \right\} \in [0, \infty].$$

In other words, if $\|f\|_{L^\infty} < \infty$, then f satisfies the bound $|f(x)| \leqslant \|f\|_{L^\infty}$ for almost all $x \in X$, and $\|f\|_{L^\infty}$ is the smallest possible constant for which this is true. It is straightforward to verify that the L^∞-norm satisfies the triangle inequality, while for $1 \leqslant p < \infty$,

$$\|f + g\|_{L^p} \leqslant \|f\|_{L^p} + \|g\|_{L^p}$$

is a basic result of measure theory known as *Minkowski's inequality*. Another basic result of measure theory is that L^p-Cauchy sequences are L^p-convergent. Proofs of these results can be found e.g. in [**Sal16**, Chapter 4].

For every $p \in [1, \infty]$, $\|f\|_{L^p} = 0$ if and only if the function f vanishes almost everywhere, so in order to produce a space on which the norm is positive for all nontrivial elements, we need to divide out a subspace, producing the quotient space

$$L^p(X) = L^p(X, \mu) := \left\{ f : \Omega \to V \;\middle|\; f \text{ is measurable and } \|f\|_{L^p} < \infty \right\} \Big/ \left\{ f \;\middle|\; f \equiv 0 \text{ a.e.} \right\}.$$

This, equipped with the L^p-norm, is a Banach space. One often abuses notation by pretending that elements $f \in L^p(X)$ are actual functions, but it's important to keep in mind that they are in reality only *equivalence classes* of functions, where two functions are equivalent if and only if they are equal almost everywhere. In particular, since a single point in a reasonable measure space is typically a set of measure zero, elements of $L^p(X)$ do not generally have well-defined *values* at individual points; one can instead expect them to have well-defined *averages* over measurable subsets.

If you haven't seen much PDE theory before, you might expect a space like $C^k_b(\Omega)$ to be useful for the study of kth-order PDEs on a region $\Omega \subset \mathbb{R}^n$. This is true to a limited extent—for instance, the Banach space used in §1.4 was a variation on $C^2_b(\Omega)$—but in practice, $C^k_b(\Omega)$ does not have many of the most useful properties that would make the tools of functional analysis applicable. We will see that the L^p-spaces are typically better for this purpose, in spite of the obvious drawback that functions of class L^p cannot even be expected to be continuous, much less differentiable. The desire to discuss derivatives in the context of L^p-spaces leads naturally to the *Sobolev spaces*, which play an enormous role in the modern theory of PDEs, and we will touch upon this in Part 2.

2. Topological vector spaces

As a reminder, functions in this course will typically be assumed to take values in an unspecified finite-dimensional inner product space $(V, \langle\ ,\ \rangle)$ over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$; cf. Convention 1.15.

2.1. Some non-examples of Banach spaces. The following examples are meant to illustrate why it is sometimes useful to consider infinite-dimensional vector spaces that satisfy some conditions strictly more general than Banach spaces.

EXAMPLE 2.1. Recall that a function on an open domain in $\mathbb{R}^n$ is called **smooth** (*glatt*) or **of class** C^∞ if it has well-defined derivatives of all orders. Given an open set $\Omega \subset \mathbb{R}^n$, we define

$$C_b^\infty(\Omega) := \bigcap_{m=0}^{\infty} C_b^m(\Omega) = \left\{ f : \Omega \to V \;\middle|\; f \text{ is smooth and } \partial^\alpha f \text{ is bounded for all multi-indices } \alpha \right\}.$$

We define the convergence of a sequence $f_k \to f$ in $C_b^\infty(\Omega)$ to mean that $\partial^\alpha f_k \to \partial^\alpha f$ uniformly on Ω for all multi-indices α. Equivalently, this means $\|f - f_k\|_{C^m} \to 0$ as $k \to \infty$ for every $m \geqslant 0$, i.e. f_k converges to f with respect to all of the Banach space norms $\|\cdot\|_{C^m}$. This notion of convergence is often called C^∞**-convergence**, and we will sometimes denote it by

$$f_k \xrightarrow{C^\infty} f.$$

EXAMPLE 2.2. For $0 \leqslant m \leqslant \infty$, let

$$C^m_{\mathrm{loc}}(\Omega) = C^m(\Omega) := \left\{ f : \Omega \to V \;\middle|\; f \text{ is of class } C^m \right\},$$

so in contrast to Examples 1.17 and 2.1, we are no longer requiring anything to be bounded; note however that for any given compact subset $K \subset \Omega$, the restriction to K of every derivative of a function $f \in C^m(\Omega)$ will automatically be bounded. With this in mind, we define C^m_{loc}-**convergence** of a sequence

$$f_k \xrightarrow{C^m_{\mathrm{loc}}} f$$

to mean that $\partial^\alpha f_k$ converges uniformly to $\partial^\alpha f$ on all compact subsets for every multi-index α of order $|\alpha| \leqslant m$. Equivalently, this means that for every compact subset $K \subset \Omega$ and every $j \leqslant m$, $f_k|_K$ converges to $f|_K$ in the C^j-norm for functions on K; the latter is defined by taking the supremum in Example 1.17 over all $x \in K$ rather than $x \in \Omega$, and we can write

$$f_k \xrightarrow{C^j(K)} f \quad \text{for all } j \leqslant m \text{ and } K \subset \Omega \text{ compact.}$$

In both of these examples, the notion of convergence is natural for the class of functions under consideration, and there is an infinite collection of norms that are obviously relevant, but the convergence cannot be characterized via any one of those norms. There is a worse problem in Example 2.2: for any *finite* collection of norms of the form $\|\cdot\|_{C^j(K)}$ with $j \leqslant m$ and $K \subset \Omega$ compact, one can easily find nontrivial functions in $C^m(\Omega)$ for which all of those norms vanish.

2.2. Topologies. In order to place the examples of §2.1 into the appropriate framework, we need to go beyond normed vector spaces and consider more general notions of convergence. Some later examples will even require convergence to be defined in a way that cannot be described in terms of metric spaces. The correct language to use for this is that of *topological* spaces, to which we now give a quick introduction.

As motivation for the definitions to follow, let us recall two basic facts about metric spaces. First, in a metric space M, one has convergence $x_n \to x$ if and only if for all neighborhoods $\mathcal{U} \subset M$ of x, $x_n \in \mathcal{U}$ for all n sufficiently large. Second, a map $f : M \to M'$ between two metric spaces is continuous if and only if for every open subset $\mathcal{U} \subset M'$, the preimage $f^{-1}(\mathcal{U})$ is likewise an open subset of M. These statements are usually not considered to be the actual *definitions* of convergence or continuity in metric spaces—those definitions typically involve ϵ's and/or δ's. But they are equivalent to those definitions, and this leads to the following insight: one doesn't actually need a *metric* in order to define what convergence and continuity mean, one only really needs to know what it means for a set to be *open*.

DEFINITION 2.3. A **topology** (*Topologie*) on a set X is a collection $\mathcal{T}$ of subsets of X (called the "**open**" sets) satisfying the following conditions:

(i) $\varnothing \in \mathcal{T}$ and $X \in \mathcal{T}$;
(ii) Arbitrary unions of sets from $\mathcal{T}$ are also in $\mathcal{T}$;
(iii) If $A \in \mathcal{T}$ and $B \in \mathcal{T}$, then $A \cap B \in \mathcal{T}$.

A **topological space** (*topologischer Raum*) $(X, \mathcal{T})$ is a set X equipped with a topology $\mathcal{T}$.

Continuity (*Stetigkeit*) of a map $f : X \to Y$ between two topological spaces is now defined exactly as described above: it means that for every open set $\mathcal{U} \subset Y$, $f^{-1}(\mathcal{U}) \subset X$ is also open. **Convergence** (*Konvergenz*) is similarly defined in terms of neighborhoods, where a set $\mathcal{U} \subset X$ is called a **neighborhood** (*Umgebung*) of a point $x \in X$ if and only if it contains an open subset containing x. Axiomatizing the notion of "openness" in this way makes it unnecessary to mention any ϵ's or δ's.

DEFINITION 2.4. A map $f : X \to Y$ between two topological spaces is called a **homeomorphism** (*Homöomorphismus*) if it is continuous and bijective and its inverse $f^{-1} : Y \to X$ is also continuous. In particular, the latter means that for every $\mathcal{U} \subset X$ open, the image $f(\mathcal{U}) \subset Y$ is likewise open.

EXAMPLE 2.5. Metric spaces (and in particular normed vector spaces) are naturally also topological spaces, with the usual notion of open sets. A topology is called **metrizable** (*metrisierbar*) whenever it can be defined in this way via a metric.

REMARK 2.6. Note that there are always multiple distinct choices of metric that will define the same metrizable topology; see e.g. Exercise 2.12 below. In many Banach spaces of interest in applications, the definition of the norm depends on noncanonical choices, but the induced topology does not.

EXAMPLE 2.7. If X and Y are topological spaces, their Cartesian product $X \times Y$ inherits a natural **product topology**, defined as the smallest topology containing all sets of the form $\mathcal{U} \times \mathcal{V} \subset X \times Y$ with $\mathcal{U} \subset X$ and $\mathcal{V} \subset Y$ both open. (Note that since arbitrary unions and finite intersections of such product sets must be allowed to be open in $X \times Y$, the product topology contains much more than just the products of open sets.)

The proof of the following is a straightforward exercise:

THEOREM 2.8. *Continuous maps $f : X \to Y$ between topological spaces are also* ***sequentially continuous*** *(folgenstetig), meaning that for every convergent sequence $x_n \to x$ in X, the sequence $f(x_n)$ converges in Y to $f(x) \in Y$.* □

You may recall from first-year analysis that in metric spaces, the converse of Theorem 2.8 also holds: sequentially continuous maps between metric spaces are also continuous. In topological spaces, this is not true in general, though writing down an actual counterexample takes some effort.

We now come to the most important definition of this section.

DEFINITION 2.9. A **topological vector space** (*topologischer Vektorraum*) over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ is a vector space X endowed with a topology such that the maps

$$X \times X \to X : (x, y) \mapsto x + y, \qquad \text{and} \qquad \mathbb{K} \times X \to X : (\lambda, x) \mapsto \lambda x$$

are both continuous.

EXAMPLE 2.10. Normed vector spaces are also topological vector spaces, with the topology defined via the usual notion of open sets.

DEFINITION 2.11. Two norms $\|\cdot\|_0$ and $\|\cdot\|_1$ on a vector space X are called **equivalent** if there exists a constant $c > 0$ such that

$$\frac{1}{c}\|f\|_0 \leqslant \|f\|_1 \leqslant c\|f\|_0 \qquad \text{for all } f \in X.$$

EXERCISE 2.12. Prove that two norms on a vector space are equivalent if and only if they define the same topology.

EXAMPLE 2.13. It is not so difficult to show that the C^m-norm defined in Example 1.17 is equivalent to the norm

$$\|f\| := \max_{|\alpha| \leqslant m} \|\partial^\alpha f\|_{C^0}.$$

The following theorem gives another important class of examples of topological vector spaces, though not the most useful for functional analysis, which is why I will not feel guilty about omitting the proof. The proof amounts to showing that on a finite-dimensional topological vector space, the topology can *always* be defined in terms of a norm, and any two norms are equivalent. We will see that neither is generally true in infinite-dimensional spaces.

THEOREM 2.14. *Every finite-dimensional vector space admits a unique topology for which it is a topological vector space.* □

2.3. Locally convex spaces. In §2.1, we saw two examples of function spaces that each had a natural notion of convergence definable via an infinite family of norms. Strictly speaking, in the case of $C^m(\Omega)$ with C^m_{loc}-convergence (Example 2.2), the norms $\|\cdot\|_{C^k(K)}$ do not actually define norms on $C^m(\Omega)$, because they can vanish on functions that are nontrivial. We therefore need a slight extension of the usual concept of a norm before we can formalize such notions of convergence on function spaces.

DEFINITION 2.15. A **seminorm** on a vector space X over $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ is a function $\|\cdot\| : X \to [0, \infty)$ satisfying the following two conditions:

(i) $\|\lambda x\| = |\lambda| \cdot \|x\|$ for all $\lambda \in \mathbb{K}$ and $x \in X$;
(ii) $\|x + y\| \leqslant \|x\| + \|y\|$ for all $x, y \in X$.

What's missing from this definition is the condition that $\|x\| > 0$ whenever $x \neq 0$; if that condition is added, a seminorm becomes a norm.

DEFINITION 2.16. A **locally convex space** (*lokal konvexer Raum*) is a vector space X endowed with the smallest topology containing all sets of the form

$$B_R^\alpha(x_0) := \left\{ x \in X \;\middle|\; \|x - x_0\|_\alpha < R \right\}, \qquad \text{where} \qquad x_0 \in X,\; R > 0,\; \alpha \in I$$

for some family of seminorms $\{\|\cdot\|_\alpha\}_{\alpha \in I}$, indexed by a set I, satisfying the condition

$$\|x\|_\alpha = 0 \text{ for all } \alpha \in I \qquad \Rightarrow \qquad x = 0.$$

REMARK 2.17. We will see below that locally convex spaces are also topological vector spaces. It is easy to check that if $\{\|\cdot\|_\alpha\}_{\alpha \in I}$ is a family of seminorms defining the topology of a locally convex space X, then each of these seminorms is also a continuous function

$$\|\cdot\|_\alpha : X \to [0, \infty).$$

Indeed, given any $x_0 \in X$ and $\epsilon > 0$, one easily finds a neighborhood $\mathcal{U} \subset X$ of x_0 such that

$$x \in \mathcal{U} \quad \Rightarrow \quad \big|\|x\|_\alpha - \|x_0\|_\alpha\big| < \epsilon.$$

The ball $\mathcal{U} := B^\alpha_\epsilon(x_0) \subset X$ does the trick, since by the reverse triangle inequality, $x \in B^\alpha_\epsilon(x_0)$ implies

$$\epsilon > \|x - x_0\|_\alpha \geqslant |\|x\|_\alpha - \|x_0\|_\alpha| .$$

REMARK 2.18. The term *locally convex* arises from the observation that in a locally convex space X, the topology of X is generated by convex sets, namely the balls $B^\alpha_R(x_0) \subset X$. Indeed, one has the convexity condition

$$x, y \in B^\alpha_R(x_0) \qquad \Rightarrow \qquad tx + (1-t)y \in B^\alpha_R(x_0) \text{ for all } t \in [0,1],$$

since the two axioms satisfied by a seminorm imply

$$\begin{aligned}\|[tx + (1-t)y] - x_0\|_\alpha &= \|t(x - x_0) + (1-t)(y - x_0)\|_\alpha \leqslant \|t(x - x_0)\|_\alpha + \|(1-t)(y - x_0)\|_\alpha \\ &= t\|x - x_0\|_\alpha + (1-t)\|y - x_0\|_\alpha \leqslant tR + (1-t)R = R.\end{aligned}$$

PROPOSITION 2.19. *Suppose X is a locally convex space with a topology defined via the family of seminorms $\{\|\cdot\|_\alpha\}_{\alpha\in I}$. Then:*

(1) X is a topological vector space;
(2) A sequence $x_n \in X$ coverges to $x \in X$ if and only if for every $\alpha \in I$, $\|x - x_n\|_\alpha \to 0$ as $n \to \infty$.
(3) A set $\mathcal{U} \subset X$ is open if and only if for every $x_0 \in \mathcal{U}$, there exists a finite subset $I_0 \subset I$ and numbers $\epsilon_\alpha > 0$ for $\alpha \in I_0$ such that the set

$$\left\{x \in X \;\middle|\; \|x - x_0\|_\alpha < \epsilon_\alpha \text{ for all } \alpha \in I_0\right\}$$

is contained in $\mathcal{U}$.

PROOF. We first prove (3). In general, the smallest topology containing any given collection of sets consists of all unions of finite intersections of sets from that collection. Thus an open set $\mathcal{U} \subset X$ is a union of finite intersections of balls defined in terms of the given seminorms; it follows that if $x_0 \in \mathcal{U}$, then some finite intersection of balls contained in $\mathcal{U}$ contains x_0, i.e.

$$x_0 \in \bigcap_{\alpha\in I_0} B^\alpha_{R_\alpha}(x_\alpha) \subset \mathcal{U}$$

for some finite subset $I_0 \subset I$ and some $R_\alpha > 0$ and $x_\alpha \in X$ defined for $\alpha \in I_0$. Using the triangle inequality, one shows that if $\epsilon_\alpha > 0$ is chosen sufficiently small for each $\alpha \in I_0$, then $B^\alpha_{\epsilon_\alpha}(x_0) \subset B^\alpha_{R_\alpha}(x_\alpha)$.

Next, we prove (2). If $x_n \to x$ in X, then since $B^\alpha_\epsilon(x) \subset X$ is a neighborhood of x for each $\epsilon > 0$ and $\alpha \in I$, we have $x_n \in B^\alpha_\epsilon(x)$ for all n sufficiently large, implying $\|x - x_n\|_\alpha \to 0$ as $n \to \infty$. Conversely, suppose $\|x - x_n\|_\alpha \to 0$ for every $\alpha \in I$. Given a neighborhood $\mathcal{U} \subset X$ of x, we apply statement (2) to find a finite collection of balls $\{B^\alpha_{\epsilon_\alpha}(x)\}_{\alpha\in I_0}$ such that

$$x \in \bigcap_{\alpha\in I_0} B^\alpha_{\epsilon_\alpha}(x) \subset \mathcal{U}.$$

Since this is only a finite collection, we then have $x_n \in B^\alpha_{\epsilon_\alpha}(x)$ for all $\alpha \in I_0$ whenever n is sufficiently large, implying $x_n \in \mathcal{U}$, and hence $x_n \to x$.

The proof of (1) is contained in Exercise 2.20 below. □

EXERCISE 2.20. For a locally convex space X, prove:

(a) A set $\mathcal{U} \subset X$ is open if and only if for every $x_0 \in \mathcal{U}$, there exists a continuous seminorm $\|\cdot\| : X \to [0,\infty)$ such that $B_1(x_0) := \{x \in X \mid \|x - x_0\| < 1\}$ is contained in $\mathcal{U}$.
Hint: Every finite positive linear combination of continuous seminorms is a continuous seminorm.
(b) X is also a topological vector space.

REMARK 2.21. The condition in Definition 2.16 that every nontrivial $x \in X$ satisfies $\|x\|_\alpha > 0$ for some $\alpha \in I$ guarantees in light of Proposition 2.19(2) that no sequence can converge simultaneously to two distinct points; in particular, to use the proper topological terminology, locally convex spaces are *Hausdorff* spaces. Perhaps it had not occurred to you to worry about the possibility that $x_n \to x$ and $x_n \to y$ might not imply $x = y$, but in fact, it is quite easy to cook up pathological examples of topological spaces in which this can happen, e.g. take any set X with more than one element and assign to it the *trivial* topology, consisting only of X and $\varnothing$. We fortunately have no need to worry about non-Hausdorff topological spaces in this course.

EXAMPLE 2.22. The space $C^\infty_b(\Omega)$ described in Example 2.1 is a locally convex space with topology defined via the countably infinite family of norms $\{\|\cdot\|_{C^m}\}_{m\geqslant 0}$.

Appying Proposition 2.19(3) to the space $C^\infty_b(\Omega)$, we deduce:

COROLLARY 2.23. *A set $\mathcal{U} \subset C^\infty_b(\Omega)$ is open if and only if for every $f_0 \in \mathcal{U}$, there exists $m \geqslant 0$ and $\epsilon > 0$ such that*

$$\|f - f_0\|_{C^m} < \epsilon \qquad \Rightarrow \qquad f \in \mathcal{U}.$$

□

EXAMPLE 2.24. The space $C^m_{\text{loc}}(\Omega)$ described in Example 2.2 is a locally convex space with topology defined via the uncountable family of seminorms

$$\left\{\|\cdot\|_{C^j(K)}\right\}_{0\leqslant j\leqslant m,\ K\subset\Omega \text{ compact}}$$

REMARK 2.25. If you read Definition 2.16 carefully, you will notice that we are considering the topology to be a part of the intrinsic structure of a locally convex space, but not the family of seminorms that determines it. Indeed, there can exist multiple distinct families of seminorms that define the same locally convex space. For instance, in Example 2.24, it is generally possible to present the domain $\Omega \subset \mathbb{R}^n$ as the union of a nested sequence of open subsets with compact closure,

$$\Omega_1 \subset \Omega_2 \subset \Omega_3 \subset \ldots \subset \bigcup_{j=1}^\infty \Omega_j = \Omega \qquad \text{with} \qquad K_j := \bar{\Omega}_j \text{ for all } j \in \mathbb{N}.$$

The countable family of seminorms

$$\left\{\|\cdot\|_{C^j(K_i)}\right\}_{0\leqslant j\leqslant m,\ i=1,2,3,\ldots}$$

then defines the same locally convex topology on $C^m_{\text{loc}}(\Omega)$.

THEOREM 2.26. *A locally convex space is metrizable if and only if its topology can be defined via a countable family of seminorms.*[1]

PARTIAL PROOF. We will sketch a proof that X is metrizable if it has a countable family of seminorms; the converse is interesting, but it is not something we'll actually need to know, so we shall omit its proof. Suppose X has a topology generated by the family of seminorms $\{\|\cdot\|_n\}_{n=1}^\infty$. We can then define a metric on X by

$$d(x, y) := \sum_{n=1}^\infty \frac{1}{2^n} \frac{\|x - y\|_n}{1 + \|x - y\|_n}.$$

We leave it as an exercise to check that d really is a metric and defines the same notion of open sets as the seminorms. □

[1] Here I am using the word **countable** to mean either *finite* or *countably infinite*.

EXERCISE 2.27. Here is an example of a topological vector space whose topology cannot be defined via a metric. Let $C^0_c(\mathbb{R}^n)$ denote the space of continuous functions $f : \mathbb{R}^n \to \mathbb{R}$ that vanish outside of compact subsets.[2] We endow $C^0_c(\mathbb{R}^n)$ with a locally convex topology defined via the family of seminorms $\{\|f\|_\varphi\}_{\varphi \in I}$ where I denotes the set of all continuous functions $\varphi : \mathbb{R}^n \to [0, \infty)$ and $\|f\|_\varphi := \|\varphi f\|_{C^0}$.

(a) Show that a sequence f_j converges to f in $C^0_c(\mathbb{R}^n)$ if and only if there exists a compact set $K \subset \mathbb{R}^n$ such that $f_j|_{\mathbb{R}^n \setminus K} \equiv 0$ for every $j \in \mathbb{N} \cup \{\infty\}$ and $f_j \to f$ uniformly on K.
(b) To show that $C^0_c(\mathbb{R}^n)$ is not metrizable, one can argue by contradiction and suppose there exists a metric d such that every neighborhood $\mathcal{U} \subset C^0_c(\mathbb{R}^n)$ of 0 contains an open set of the form $B_n := \{f \in C^0_c(\mathbb{R}^n) \mid d(0, f) < 1/n\}$ for $n \in \mathbb{N}$ sufficiently large. Show that in this situation, there must exist functions $\varphi_n \in I$ such that $A_n := \{f \in C^0_c(\mathbb{R}^n) \mid \|f\|_{\varphi_n} < 1\} \subset B_n$ for every n, then derive a contradiction by constructing a neighborhood $\mathcal{U}$ of 0 that does not contain A_n for any $n \in \mathbb{N}$.

EXAMPLE 2.28. We can use the seminorms $\|\cdot\|_\varphi$ described in Exercise 2.27 to write down an example of how *not* to put a topology on a vector space, i.e. the following is not a topological vector space. We observe that for functions $f \in C^0(\Omega)$ that do not have compact support, $\|f\|_\varphi$ is not always finite, but we can nonetheless define a topology on $C^0(\Omega)$ as the smallest topology containing the balls

$$B^\varphi_R(f_0) := \left\{f \in C^0(\Omega) \mid \|f - f_0\|_\varphi < R\right\}$$

for all $f_0 \in C^0(\Omega)$, all $R > 0$, and all continuous functions $\varphi : \Omega \to [0, \infty)$. Exercise 2.27 now shows that we have convergence $f_j \to f$ of continuous functions on Ω in this topology if and only if there is a compact subset $K \subset \Omega$ such that $f_j \equiv f$ on $\Omega \setminus K$ for all j and $f_j \to f$ uniformly on K. We claim: For this topology on $C^0(\Omega)$, the map

$$\mathbb{R} \times C^0(\Omega) \to C^0(\Omega) : (\lambda, f) \mapsto \lambda f$$

is not continuous. Indeed, suppose $f \in C^0(\Omega)$ does not have compact support. Then for any sequence $\lambda_j \in \mathbb{R}$ with $\lambda_j \neq 0$ for all j but $\lambda_j \to 0$, the sequence $\lambda_j f \in C^0(\Omega)$ does not converge to 0, because $\lambda_j f$ does not vanish on the complement of any compact subset of Ω.

2.4. Fréchet spaces. The interesting examples we've seen of locally convex spaces also have a stronger property that makes them *almost* as nice as Banach spaces for some purposes, though not as easy to work with.

DEFINITION 2.29. A metrizable locally convex space is called a **Fréchet space** (*Fréchetraum*) if it is complete with respect to some metric that defines its topology.

I'm not really going to have anything to say about Fréchet spaces in this course, but let's take note of the ones that we've already seen:

EXAMPLE 2.30. The spaces $C^\infty_b(\Omega)$ and $C^m_{\text{loc}}(\Omega)$ from Examples 2.1 and 2.2 are Fréchet spaces. For $C^\infty_b(\Omega)$, one can deduce this from the fact that $C^m_b(\Omega)$ is a Banach space for each $m < \infty$. For $C^m_{\text{loc}}(\Omega)$, one can use the countable family of seminorms described in Remark 2.25 to define a suitable metric as in Theorem 2.26, and then deduce completeness from the completeness of the norms $\|\cdot\|_{C^j(K)}$ for each finite j and compact set $K \subset \Omega$.

[2]We say in this case that the functions $f \in C^0_c(\mathbb{R}^n)$ have **compact support** (*kompakten Träger*) in $\mathbb{R}^n$.

2.5. Continuous extensions. We've seen already in the example of $C^\infty([0,1]) \subset C^0([0,1])$ that in infinite-dimensional normed vector spaces, a proper linear subspace can also be a *dense* subset. In fact, this fact is frequently useful: we will see that most of the Banach spaces of functions that arise in applications contain spaces of smooth functions as dense subspaces, and when we want to prove that those spaces have certain properties, proving it for the smooth functions is often easier. Such proofs can then be extended to the entire Banach space due to results such as the following:

LEMMA 2.31. *Suppose X is a normed vector space containing a dense subspace $X_0 \subset X$, and Y is a Banach space. Then every bounded linear operator $A : X_0 \to Y$ has a unique extension to a bounded linear operator $X \to Y$.*

PROOF. Given $x \in X$, density provides a sequence $x_n \in X_0$ with $x_n \to x$. Since $\|Ax_n - Ax_m\| \leqslant \|A\| \cdot \|x_n - x_m\|$, Ax_n is a Cauchy sequence in Y, and thus converges, allowing us to define

$$Ax := \lim_{n\to\infty} Ax_n.$$

This definition of Ax is independent of the choice of sequence x_n; indeed, if $x'_n \in X_0$ is another sequence converging to x, then

$$\|x_n - x'_n\| = \|(x_n - x) + (x - x'_n)\| \leqslant \|x_n - x\| + \|x - x'_n\| \to 0,$$

implying $\|Ax_n - Ax'_n\| \leqslant \|A\| \cdot \|x_n - x'_n\| \to 0$ and thus $\lim_n Ax_n = \lim_n Ax'_n$. The map $A : X \to Y$ defined in this way is clearly the *only* continuous map that extends $A : X_0 \to Y$; we leave it as an exercise to check that it is a bounded linear operator. □

2.6. Dual spaces. The notation $\mathscr{L}(X,Y)$ has been used so far to mean the space of bounded linear operators between two normed vector spaces. In the more general context of topological vector spaces, we do not have any notion of boundedness, but continuity makes sense, and it is thus natural to define

$$\mathscr{L}(X,Y) := \{\text{continuous linear maps } X \to Y\}.$$

DEFINITION 2.32. The **dual space** (*Dualraum*) of a topological vector space X over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ is defined as

$$X^* := \mathscr{L}(X, \mathbb{K}),$$

i.e. it is the space of continuous linear **functionals** (*Funktionale*) $\Lambda : X \to \mathbb{K}$.

Note that in infinite dimensions, the functional-analytic notion of a dual space is not the same as the *algebraic* dual space: the latter contains *all* linear functionals $X \to \mathbb{K}$, without regard for continuity. In fact, there is no guarantee in general that X^* is nontrivial. This is quite different from the situation in finite dimensions, where there are finite algorithms for writing down linear functionals $X \to \mathbb{K}$, which are automatically continuous. Having tried to convince you that the existence of continuous linear functionals is generally non-obvious, here are a few pieces of good news:

(1) Later in this course, we will prove the Hahn-Banach theorem, which implies that every locally convex space X has a nontrivial dual space. In fact, given any $x \neq 0 \in X$, the Hahn-Banach theorem enables the construction of a continuous linear functional $\Lambda \in X^*$ with $\Lambda(x) \neq 0$.
(2) In many of the most popular Banach spaces, there exist explicit descriptions of their dual spaces that imply $X^* \neq \{0\}$ without appealing to the Hahn-Banach theorem.[3] This is true in particular for all Hilbert spaces, and for the L^p-spaces with $1 \leqslant p < \infty$.

[3] The proof of the Hahn-Banach theorem requires the axiom of choice; some people find this unnerving and therefore prefer to avoid it.

(3) If X is a normed vector space, its dual X^* with the operator norm is automatically a Banach space. This follows from Theorem 1.8.

EXAMPLE 2.33. Assume the vector space V in which our functions take their values is the ground field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$. Then for any open set $\Omega \subset \mathbb{R}^n$, any finite measure μ on Ω determines a bounded linear functional

$$\Lambda_\mu : C_b^0(\Omega) \to \mathbb{K} : f \mapsto \int_\Omega f \, d\mu.$$

Since $|\Lambda_\mu(f)| \leqslant \mu(\Omega) \cdot \|f\|_{C^0}$, we have $\|\Lambda_\mu\| \leqslant \mu(\Omega)$.

EXAMPLE 2.34. For $1 \leqslant p, q \leqslant \infty$ with $\frac{1}{p} + \frac{1}{q} = 1$, any function $g \in L^q(X)$ on a measure space X determines a bounded linear functional

$$\Lambda_g : L^p(X) \to \mathbb{K} : f \mapsto \int_X \langle g, f \rangle \, d\mu.$$

The boundedness of Λ_g follows from *Hölder's inequality*, which gives

$$|\Lambda_g(f)| \leqslant \|g\|_{L^q} \cdot \|f\|_{L^p}, \qquad \text{hence} \qquad \|\Lambda_g\| \leqslant \|g\|_{L^q}.$$

DEFINITION 2.35. The **transpose** (*Transponierte*) of a continuous linear map $A : X \to Y$ is the linear map $A^\top : Y^* \to X^*$ defined by

$$(A^\top \lambda)(x) = \lambda(Ax) \qquad \text{for} \qquad \lambda \in Y^*, \ x \in X.$$

One easily checks that if X and Y are normed vector spaces, then $A^\top$ is also a bounded linear map and $\|A^\top\| \leqslant \|A\|$.

In the following result, we let

$$X^{**} := (X^*)^*$$

denote the dual of the dual space of X, also known as its **double dual**.

PROPOSITION 2.36. *For any normed vector space X, there is a canonical bounded linear map*

$$J : X \to X^{**}$$

given by $(Jx)(\lambda) := \lambda(x)$ for $x \in X$ and $\lambda \in X^*$.

PROOF. Clearly $\|(Jx)(\lambda)\| = |\lambda(x)| \leqslant \|\lambda\| \cdot \|x\|$, thus $\|Jx\| \leqslant \|x\|$, implying that J is bounded with $\|J\| \leqslant 1$. □

We will later deduce from the Hahn-Banach theorem that J is actually an *injective* map—in fact it is an **isometry** (*Isometrie*), meaning it satisfies $\|Jx\| = \|x\|$ for all $x \in X$. It therefore defines a natural inclusion $X \hookrightarrow X^{**}$.

DEFINITION 2.37. A Banach space X is called **reflexive** (*reflexiv*) if the natural map $J : X \to X^{**}$ is an isomorphism.

We will see that the L^p-spaces are reflexive for $1 < p < \infty$, though not for $p = 1$ or $p = \infty$.

EXERCISE 2.38. Here is an example of a topological vector space that is not locally convex and has a trivial dual space. Define $L^p([0,1])$ as usual to be the space of equivalence classes (up to equality almost everywhere) of measurable functions $f : [0,1] \to \mathbb{R}$ with $\|f\|_{L^p} := \left(\int_0^1 |f(x)|^p \, dx\right)^{1/p} < \infty$, but instead of $p \geqslant 1$, assume $0 < p < 1$. In this case, Minkowski's inequality does not hold, so $\|\cdot\|_{L^p}$ is not a norm, but we shall regard $L^p([0,1])$ as a metric space with the metric defined by $d(f,g) := \|f - g\|_{L^p}^p$.

(a) Show that d is a metric on $L^p([0,1])$.
Hint: Show first that $(x+y)^p \leqslant x^p + y^p$ holds for all $x, y \geqslant 0$. The latter can be deduced from the relation $a^q + b^q \leqslant (a+b)^q$ for $a, b \geqslant 0$ and $q := 1/p > 1$, which follows in turn from $(1+x)^q \geqslant 1 + x^q$ for $x \geqslant 0$, which you can prove by differentiating with respect to x.

(b) Prove that $L^p([0,1])$ with the topology defined via d is a topological vector space.

(c) Prove that the space of bounded measurable real-valued functions is dense in $L^p([0,1])$.
Hint: Given $f \in L^p([0,1])$, define functions f_n that match f wherever $|f| \leqslant n$.

(d) Prove by induction on $N \in \mathbb{N}$ that if K is any convex subset of a vector space, then for any finite collections $x_1, \dots, x_N \in K$ and $\tau_1, \dots, \tau_N \in [0,1]$ with $\sum_{i=1}^N \tau_i = 1$, $\sum_{i=1}^N \tau_i x_i \in K$. (This is known as a *convex combination* of $x_1, \dots, x_N$.)

(e) Prove that for any given $\epsilon > 0$, every bounded measurable function $f : [0,1] \to \mathbb{R}$ can be written as $f = \frac{1}{N}\sum_{i=1}^N f_i$ for some finite collection of functions $f_i \in L^p([0,1])$ satisfying $d(f_i, 0) < \epsilon$ for all $i = 1, \dots, N$. Conclude that the only closed convex subset of $L^p([0,1])$ containing a neighborhood of 0 is $L^p([0,1])$ itself.
Hint: Define each f_i to have support in an interval of length $1/N$, then make N large.

(f) Prove that all continuous linear functionals $\Lambda : L^p([0,1]) \to \mathbb{R}$ are trivial.
Hint: What kind of subset is $\{f \in L^p([0,1]) \mid |\Lambda(f)| \leqslant 1\}$?

2.7. Zorn's lemma and bases. I owe you an answer to the following question: Do there exist *unbounded* linear operators between infinite-dimensional Banach spaces?

The answer is yes, and there is an easy way to see it if you believe that every vector space admits a basis. Let us pause here to clarify what we mean exactly by the word "basis"; an infinite-dimensional spaces, there are multiple inequivalent things this word could mean, and the following notion is for instance *not* the same as the more useful notion of a Hilbert space basis that we will see in the next section.

DEFINITION 2.39. A **Hamel basis** (*Hamelbasis*) of a vector space X is a maximal linearly-independent subset. More precisely, a collection of elements $\{e_\alpha \in X\}_{\alpha \in I}$ is a Hamel basis of X if and only if every $x \in X$ can be written as

$$x = \sum_{\alpha \in I} c_\alpha e_\alpha$$

for uniquely-determined coefficients $\{c_\alpha \in \mathbb{K}\}_{\alpha \in I}$, at most finitely-many of which are nonzero.

Note that Definition 2.39 is purely algebraic; it has no analytical or topological content, and since the expression $x = \sum_\alpha c_\alpha e_\alpha$ is assumed to have at most finitely-many nonzero terms, there is no issue of convergence to worry about.

If you believe that Hamel bases always exist, then for any pair of normed vector spaces X, Y with $\dim X = \infty$, one can define an unbounded (and therefore discontinuous) linear operator $A : X \to Y$ as follows. Choose a Hamel basis $\{e_\alpha \in X\}_{\alpha \in I}$ of X, an infinite sequence $\{\alpha_n \in I\}_{n \in \mathbb{N}}$ and an element $y \neq 0 \in Y$, and define A as the unique linear map satisfying

$$Ae_{\alpha_n} := n\|e_{\alpha_n}\| y, \qquad \text{for all } n \in \mathbb{N}, \qquad Ae_\alpha := 0 \quad \text{for all other } \alpha \in I.$$

We then have $\frac{\|Ae_{\alpha_n}\|}{\|e_{\alpha_n}\|} = n\|y\| \to \infty$ as $n \to \infty$, thus A is unbounded.

The good news is that, indeed, Hamel bases always exist:

THEOREM 2.40. *Every vector space admits a Hamel basis.*

The bad news is that in typical cases of infinite-dimensional Banach spaces, Theorem 2.40 depends in an essential way on the axiom of choice. One can cook up stupid algebraic examples, e.g. if we define $\mathbb{R}^\infty$ to be the space of sequences $(x_1, x_2, x_3, \ldots)$ in which at most finitely-many terms are nonzero, then $\mathbb{R}^\infty$ has an obvious Hamel basis. But $\mathbb{R}^\infty$ is not a Banach space in any obvious way, and in general, if you were hoping to see an explicit example of a Hamel basis on an infinite-dimensional Banach space (and therefore also an explicit example of an unbounded operator), you will be disappointed. They must exist if you accept the axiom of choice, but they will typically contain uncountably many elements, and there is no reasonable algorithm to construct them.

We will not actually have any use in this course for Hamel bases or unbounded operators of the type described above, but we may as well take this opportunity to discuss the axiom of choice, which will be needed later for the Hahn-Banach theorem. The axiom of choice is equivalent to a somewhat mysterious-looking result known as Zorn's lemma:

THEOREM 2.41 (Zorn's lemma). *Suppose* $(S, \prec)$ *is a nonempty partially-ordered set in which every totally ordered subset* $T \subset S$ *has an upper bound. Then the upper bound for* T *can always be chosen to be a maximal element of* S. □

Let's be clear on the terminology: the assumption that $T \subset S$ is **totally ordered** means that every pair of elements $x, y \in T$ satisfies $x \prec y$ or $y \prec x$. An **upper bound** for T is an element $m \in S$ such that $t \prec m$ for all $t \in T$, and m is a **maximal element** of S if there is no $x \in S$ satisfying $x \neq m$ and $m \prec x$.

The proof that Zorn's lemma is equivalent to the axiom of choice is an unenlightening excursion into set theory that I will not reproduce here, but instead refer you to [**BS18**, Appendix A]. I recommend reading through it exactly once in your life. What follows is a practical demonstration of how Zorn's lemma can be used.

PROOF OF THEOREM 2.40. Given a vector space X, let S denote the set of all linearly-independent subsets of X, endowed with the partial order

$$\mathcal{A} \prec \mathcal{B} \qquad \Leftrightarrow \qquad \mathcal{A} \subset \mathcal{B}.$$

If $S_0 \subset S$ is a totally-ordered subset, then

$$\mathcal{B}_\infty := \bigcup_{\mathcal{B} \in S_0} \mathcal{B} \in S$$

is an upper bound for S_0, thus establishing the hypothesis of Zorn's lemma; the maximal element provided by Zorn's lemma is then a Hamel basis for X. □

3. Hilbert spaces

Hilbert spaces are the nicest class of vector spaces one can work with in functional analysis.

DEFINITION 3.1. A **Hilbert space** (*Hilbertraum*) is a complete inner product space $(\mathcal{H}, \langle\ ,\ \rangle)$, i.e. a Banach space $(\mathcal{H}, \|\cdot\|)$ endowed with an inner product $\langle\ ,\ \rangle$ such that $\|x\| = \sqrt{\langle x, x\rangle}$ for all $x \in \mathcal{H}$.

CONVENTION 3.2. For a complex inner product space $(\mathcal{H}, \langle\ ,\ \rangle)$, our convention in these notes is that the inner product is complex *antilinear* in the first argument and complex linear in the second:

$$\langle v, \lambda w\rangle = \lambda\langle v, w\rangle = \langle \bar{\lambda} v, w\rangle, \qquad \text{for } v, w \in \mathcal{H},\ \lambda \in \mathbb{C}.$$

Some sources prefer to do it the other way around.

EXAMPLE 3.3. For any measure space (X, μ), the space $L^2(X, \mu)$ of square-integrable functions $f : X \to V$ valued in the finite-dimensional inner product space $(V, \langle\ ,\ \rangle)$ is a Hilbert space with inner product defined by integrating the inner products of the values of two functions:

$$\langle f, g \rangle_{L^2} := \int_X \langle f, g \rangle \, d\mu.$$

The integral is well defined due to Hölder's inequality (also known as the Cauchy-Schwarz inequality in this special case), which gives

$$|\langle f, g \rangle| \leqslant \|f\|_{L^2} \cdot \|g\|_{L^2}.$$

We will see in Exercise 3.38 that up to isomorphism, every Hilbert space is equivalent to Example 3.3 for a suitable choice of measure space.

3.1. Geometry of inner product spaces. We begin with some general properties of inner product spaces that do not require finite-dimensionality, but also do not require completeness. For this discussion, we fix an inner product space $(\mathcal{H}, \langle\ ,\ \rangle)$, which we will usually abbreviate simply as $\mathcal{H}$.

Two elements $v, w \in \mathcal{H}$ are called **orthogonal** if $\langle v, w \rangle = 0$. The **orthogonal complement** (*orthogonales Komplement*) of a subset $W \subset \mathcal{H}$ is the subset

$$W^\perp := \left\{ x \in \mathcal{H} \;\middle|\; \langle w, x \rangle = 0 \text{ for all } w \in W \right\}.$$

It is easy to check that if W is a linear subspace, then so is $W^\perp$. A subset $S \subset \mathcal{H}$ is called an **orthonormal set** if

$$\|v\| = 1 \text{ for all } v \in S, \qquad \text{and} \qquad \langle v, w \rangle = 0 \text{ for all } v, w \in S \text{ with } v \neq w.$$

We will often write orthonormal sets as collections of elements $\{e_\alpha \in \mathcal{H}\}_{\alpha \in I}$ indexed by an auxiliary set I, in which case the orthonormality condition becomes

$$\langle e_\alpha, e_\beta \rangle = \begin{cases} 1 & \text{for } \alpha = \beta, \\ 0 & \text{for } \alpha \neq \beta. \end{cases}$$

EXAMPLE 3.4. One easily checks that the space $L^2([0, 1])$ of complex-valued L^2-functions $f : [0, 1] \to \mathbb{C}$ with the L^2-inner product

$$\langle f, g \rangle_{L^2} = \int_0^1 \bar{f}(t) g(t) \, dt$$

has

$$\left\{ e_k(t) := e^{2\pi i k t} \right\}_{k \in \mathbb{Z}}$$

as an orthonormal set.

We will see in §3.5 that a *maximal* orthonormal set in a Hilbert space $\mathcal{H}$ can be used in much the same way as a finite-dimensional basis or a Hamel basis: it allows us to express any element uniquely as a linear combination of basis elements, with the important difference that when $\dim \mathcal{H} = \infty$, the linear combinations in question will not generally be finite sums, but instead convergent series. That discussion requires a bit of preparation, and it depends crucially on the assumption that Hilbert spaces are complete. For the moment, we will have to be content with our understanding of how this works in finite-dimensional inner product spaces.

LEMMA 3.5. *In any inner product space $\mathcal{H}$, if $x \in \mathcal{H}$ satisfies $x = v + w$ with $\langle v, w \rangle = 0$, then $\|x\|^2 = \|v\|^2 + \|w\|^2$.*

PROOF. This follows from the bilinearity of the inner product: $\|x\|^2 = \langle v + w, v + w \rangle = \langle v, v \rangle + \langle w, w \rangle = \|v\|^2 + \|w\|^2$. □

THEOREM 3.6 (the "Pythagorean theorem"). *Suppose $\mathcal{H}$ is an inner product space, and $W \subset \mathcal{H}$ is a finite-dimensional subspace with an orthonormal basis $e_1, \ldots, e_N \in W$. Then for every $x \in \mathcal{H}$, one has*

$$x = \sum_{j=1}^{N} \langle e_j, x \rangle e_j + y$$

for a unique element $y \in W^\perp$, and

$$\|x\|^2 = \sum_{j=1}^{N} |\langle e_j, x \rangle|^2 + \|y\|^2.$$

PROOF. Defining $y := x - \sum_{j=1}^{N} \langle e_j, x \rangle e_j$, we compute

$$\langle e_j, y \rangle = \langle e_j, x \rangle - \langle e_j, x \rangle = 0$$

for each $j = 1, \ldots, N$, thus $y \in W^\perp$. The second relation then follows from Lemma 3.5. $\square$

THEOREM 3.7 (Cauchy-Schwarz inequality). *In an inner product space $\mathcal{H}$, the relation*

$$|\langle x, y \rangle| \leqslant \|x\| \cdot \|y\|$$

holds for all $x, y \in \mathcal{H}$.

PROOF. Assume $x, y \in \mathcal{H}$ are both nonzero, since the relation is otherwise trivial. Setting $W \subset \mathcal{H}$ to be the 1-dimensional subspace spanned by $e_1 := x/\|x\|$, Theorem 3.6 implies

$$y = \left\langle \frac{x}{\|x\|}, y \right\rangle \frac{x}{\|x\|} + y'$$

for some $y' \in W^\perp$, and thus

$$\|y\|^2 = \left| \left\langle \frac{x}{\|x\|}, y \right\rangle \right|^2 + \|y'\|^2 \geqslant \left| \left\langle \frac{x}{\|x\|}, y \right\rangle \right|^2 = \frac{|\langle x, y \rangle|^2}{\|x\|^2}.$$

$\square$

COROLLARY 3.8. *On any inner product space $\mathcal{H}$, the inner product $\langle \ , \ \rangle : \mathcal{H} \times \mathcal{H} \to \mathbb{K}$ is a continuous function.*

PROOF. Since all topological spaces in this discussion can be regarded as metric spaces, it suffices to prove *sequential* continuity. Suppose $x_n, y_n \in \mathcal{H}$ are sequences with $x_n \to x$ and $y_n \to y$: then

$$\begin{aligned} |\langle x, y \rangle - \langle x_n, y_n \rangle| &= |\langle x - x_n, y \rangle + \langle x_n, y - y_n \rangle| \leqslant |\langle x - x_n, y \rangle| + |\langle x_n, y - y_n \rangle| \\ &\leqslant \|x - x_n\| \cdot \|y\| + \|x_n\| \cdot \|y - y_n\| \to 0 \end{aligned}$$

as $n \to \infty$. $\square$

COROLLARY 3.9. *For any $x \in \mathcal{H}$ in an inner product space $\mathcal{H}$, the formula $\Lambda_x(y) := \langle x, y \rangle$ defines a bounded linear functional $\Lambda_x : \mathcal{H} \to \mathbb{K}$ with $\|\Lambda_x\| = \|x\|$.*

PROOF. The inequality $\|\Lambda_x\| \leqslant \|x\|$ follows directly from Cauchy-Schwarz. To see that it is actually an equality, observe that

$$\frac{|\Lambda_x(x)|}{\|x\|} = \frac{|\langle x, x \rangle|}{\|x\|} = \frac{\|x\|^2}{\|x\|} = \|x\|.$$

$\square$

COROLLARY 3.10. *For any inner product space $\mathcal{H}$ and linear subspace $W \subset \mathcal{H}$, $W^\perp \subset \mathcal{H}$ is a closed subspace.*

PROOF. Suppose $x_n \in W^\perp$ is a sequence with $x_n \to x \in \mathcal{H}$. Then $\langle x_n, w \rangle = 0$ for all $w \in W$, and by the continuity of the inner product, it follows that $\langle x, w \rangle = 0$ as well. $\square$

EXERCISE 3.11. Show that in any inner product space $\mathcal{H}$ with linear subspace $W \subset \mathcal{H}$, W and its closure $\overline{W}$ have the same orthogonal complement.

3.2. Orthogonal complements. You are surely familiar with the fact that every *finite-dimensional* inner product space can be decomposed into the direct sum of any subspace with the orthogonal complement of that subspace.

In infinite dimensions, the following observation should make you a little worried. It is possible for a proper subspace $W \subset \mathcal{H}$ to be dense; for example, we will see later that all L^2-functions on $\mathbb{R}$ can be approximated arbitrarily well (in the L^2-norm) by *smooth* L^2-functions, which therefore form a dense subspace of $L^2(\mathbb{R})$. But if $W \subset \mathcal{H}$ is dense, then according to Exercise 3.11,

$$W^\perp = \overline{W}^\perp = \mathcal{H}^\perp = \{0\}.$$

Here is a direct argument showing the same: for any $x \in W^\perp$, density implies that there exists a sequence $x_n \in W$ with $x_n \to x$, so that $0 = \langle x_n, x \rangle \to \langle x, x \rangle = \|x\|^2$, and thus $x = 0$. (Here we've made implicit use of the Cauchy-Schwarz inequality, namely via the fact that $\langle\ ,\ \rangle : \mathcal{H} \times \mathcal{H} \to \mathbb{K}$ is a continuous function.) In other words, the orthogonal complement of a dense subspace $W \subset \mathcal{H}$ is the trivial subspace, and $\mathcal{H}$ is therefore definitely not $W \oplus W^\perp$.

Here is the good news: the problem with the example of a dense subspace $W \subset \mathcal{H}$ is that W is not a closed subset. If we restrict our attention to *closed* subspaces, then everything is fine:

THEOREM 3.12. *For any closed subspace $W \subset \mathcal{H}$ in a Hilbert space $\mathcal{H}$,*

$$\mathcal{H} = W \oplus W^\perp,$$

i.e. every $x \in \mathcal{H}$ can be written as $x = v + w$ for uniquely-determined elements $v \in W$ and $w \in W^\perp$.

Our proof of this theorem depends on a lemma that I think you will have little trouble believing, but it's actually rather nontrivial, and its proof will form one portion of a larger discussion in the next subsection. It depends in an essential way on the completeness condition, and we shall take it for now as a black box:

LEMMA 3.13 (see Theorem 3.23). *Assume $\mathcal{H}$ is a Hilbert space, $K \subset X$ is a closed convex subset, and $x \in \mathcal{H} \backslash K$. Then K contains a unique point that is closest to x.*

PROOF OF THEOREM 3.12. The uniqueness of the decomposition $x = v + w$ with $v \in W$ and $w \in W^\perp$ is immediate from the nondegeneracy of the inner product: if it were not unique, then two distinct decompositions $x = v + w = v' + w'$ would give rise to a nontrivial vector $v - v' = w' - w \in W \cap W^\perp$, which is impossible since every nonzero $y \in W$ satisfies $\langle y, y \rangle > 0$.

For existence, observe that there is nothing to prove if $x \in W$, so assume $x \in \mathcal{H} \backslash W$. The subspace $W \subset \mathcal{H}$ is a convex set that is closed by assumption. Lemma 3.13 thus implies the existence of an element $v \in W$ that is nearest to x, and we claim that $w := x - v$ then lies in $W^\perp$. Indeed, for any $h \in W$, the fact that $\|x - v\|^2 = \langle x - v, x - v \rangle$ minimizes the distance from x to W implies

$$\begin{aligned} 0 &= \frac{d}{dt} \|x - (v + th)\|^2 \bigg|_{t=0} = \frac{d}{dt} \langle x - (v + th), x - (v + th) \rangle \bigg|_{t=0} = \frac{d}{dt} \langle w - th, w - th \rangle \bigg|_{t=0} \\ &= \frac{d}{dt} \left(\|w\|^2 - 2t \operatorname{Re}\langle w, h \rangle + t^2 \|h\|^2 \right) \bigg|_{t=0} = -2 \operatorname{Re}\langle w, h \rangle, \end{aligned}$$

where the symbol "Re" is redundant in the case $\mathbb{K} = \mathbb{R}$, and the result is then simply $\langle w, h\rangle = 0$. In the complex case, we can plug in $ih \in W$ instead of h, so that the same computation also gives

$$0 = -2\,\mathrm{Re}\langle w, ih\rangle = -2\,\mathrm{Re}\,(i\langle w, h\rangle) = 2\,\mathrm{Im}\langle w, h\rangle,$$

and the conclusion is again $\langle w, h\rangle = 0$ for all $h \in W$, as claimed. □

The following corollary is frequently useful in applications:

COROLLARY 3.14. *In a Hilbert space $\mathcal{H}$, a subspace $W \subset \mathcal{H}$ is dense if and only if $W^\perp = \{0\}$.*

PROOF. In one direction, the statement follows from Exercise 3.11 since density means $\overline{W} = \mathcal{H}$, and $\mathcal{H}^\perp = \{0\}$. Conversely, if $W^\perp = \{0\}$, then $\overline{W} \subset \mathcal{H}$ is a closed subspace with trivial orthogonal complement. Applying Theorem 3.12, we conclude

$$\mathcal{H} = \overline{W} \oplus \overline{W}^\perp = \overline{W},$$

meaning that W is dense. □

EXERCISE 3.15. For $\mathcal{H}$ a Hilbert space and $W \subset \mathcal{H}$ a linear subspace with closure denoted by $\overline{W}$, prove $(W^\perp)^\perp = \overline{W}$. Does this remain true in general if $\mathcal{H}$ is assumed to be an inner product space but not complete?

3.3. Uniform convexity. We now introduce a concept that lies in the background of several important results about Banach and Hilbert spaces, including the theorem above on orthogonal complements.

Recall that a subset K in a vector space X is called **convex** (*konvex*) if K contains the line segment joining any two of its points (see Figure 1), i.e.

$$x, y \in K \quad \Rightarrow \quad tx + (1-t)y \in K \text{ for every } t \in [0, 1].$$

Similarly, a function $f : K \to \mathbb{R}$ on a convex set $K \subset X$ is called **convex** if for every pair of points in its domain, the values of f along the line segment between those points are bounded by the corresponding "convex combinations" of its values at the end points (Figure 2); concretely,

$$\forall x, y \in K \text{ and } t \in [0, 1], \qquad f(tx + (1-t)y) \leqslant tf(x) + (1-t)f(y). \tag{3.1}$$

It is straightforward to show that if f is convex, then $f^{-1}((-\infty, a))$ and $f^{-1}((-\infty, a])$ are convex subsets for every $a \in \mathbb{R}$. We say additionally that f is **strictly convex** (*strikt konvex*) if the inequality in (3.1) is strict for all $t \in (0, 1)$ whenever $x \neq y$.

EXAMPLE 3.16. By a standard exercise in first-year analysis, if $\mathcal{U} \subset \mathbb{R}^n$ is an open convex set, then a C^2-function $f : \mathcal{U} \to \mathbb{R}$ is convex (or strictly convex) if and only if its Hessian at every point is positive semidefinite (or positive definite, respectively).

EXAMPLE 3.17. For any normed vector space $(X, \|\cdot\|)$ and $x_0 \in X$, the triangle inequality implies that the function $X \to [0, \infty) : x \mapsto \|x - x_0\|$ is convex. As a consequence, (closed or open) balls about points in a normed vector space are always convex sets. This remains true if the norm is replaced by a seminorm, and is the reason why a topological vector space with topology generated by a family of seminorms is called a *locally convex* space.

Strict and *uniform* convexity are geometric properties of normed vector spaces that strengthen the observation in Example 3.17 about balls $B \subset X$ being convex—the idea is to require that the segment joining any two points in the ball stays in the *interior* of the ball. This is a nontrivial condition on the "shape" of the unit ball as determined by the norm, and it is not satisfied by every norm (see Exercise 3.21 below). In the following, we denote the closed unit ball and unit sphere in a normed vector space $(X, \|\cdot\|)$ by

$$\bar{B} := \{x \in X \mid \|x\| \leqslant 1\}, \quad \text{and} \quad \partial\bar{B} := \{x \in X \mid \|x\| = 1\}$$

respectively, and denote the distance between two subsets $\mathcal{U}, \mathcal{V} \subset X$ by

$$\mathrm{dist}(\mathcal{U}, \mathcal{V}) := \inf \{ \|x - y\| \mid x \in \mathcal{U},\ x \in \mathcal{V} \} .$$

Definition 3.18. A normed vector space $(X, \|\cdot\|)$ is called **strictly convex** (*strikt konvex*) if

$$x, y \in \bar{B} \text{ with } x \neq y \quad \Rightarrow \quad tx + (1-t)y \in \bar{B} \backslash \partial \bar{B} \quad \forall t \in (0,1).$$

The middle picture in Figure 1 gives an example of something one might imagine the unit ball looking like in a normed vector space that is not strictly convex. The next definition amounts to a quantitative version of strict convexity, in which the distance of the midpoint between x and y to the boundary cannot become arbitrarily small unless x and y are close.

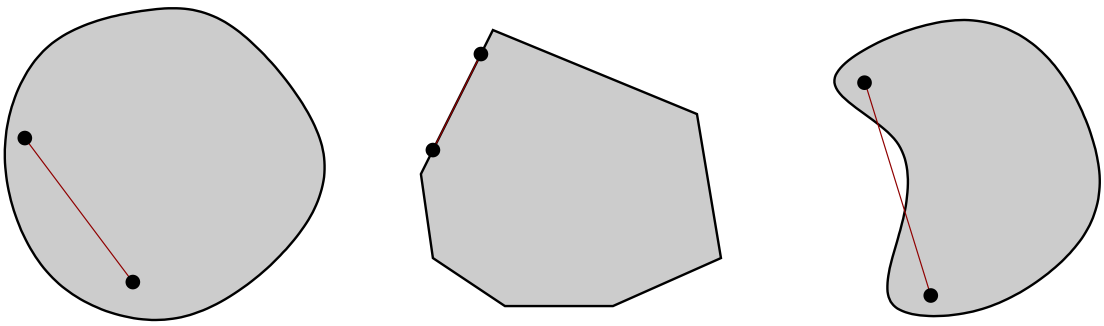

Figure 1. The two sets on the left are convex, while the set on the right is not. The set in the middle is convex but not *strictly convex*, i.e. it contains a segment connecting boundary points that does not stay in the interior. In particular, if this set occurs as the closed unit ball in some normed vector space, it implies that that space is not strictly (and therefore not uniformly) convex.

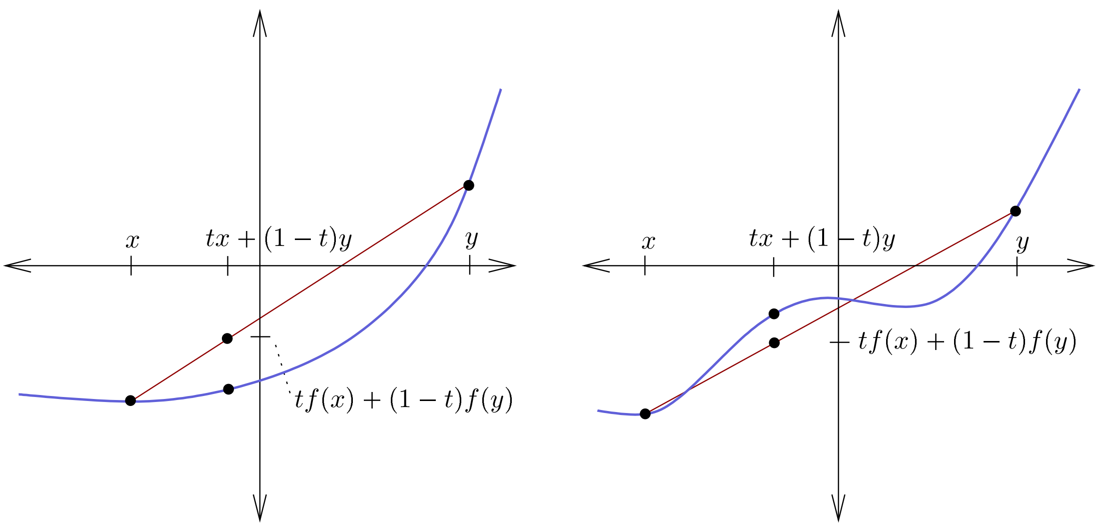

Figure 2. The function $f : \mathbb{R} \to \mathbb{R}$ on the left is convex, and the function on the right is not.

Definition 3.19. A normed vector space $(X, \|\cdot\|)$ is called **uniformly convex** (*gleichmäßig konvex*) if for every $\epsilon > 0$, there exists $\delta > 0$ such that

$$x, y \in \bar{B} \text{ with } \|x - y\| \geqslant \epsilon \qquad \Rightarrow \qquad \operatorname{dist}\left(\frac{x+y}{2}, \partial\bar{B}\right) \geqslant \delta.$$

Observe that every uniformly convex space is clearly also strictly convex.

Remark 3.20. The definition of uniform convexity appears in many references with a weaker condition on x and y, namely that they lie in $\partial\bar{B}$ instead of $\bar{B}$. The resulting notion is equivalent to our definition; for a proof of this, see [Tes, Lemma 5.20]. This detail will not concern us since, for all uniformly convex spaces that we actually encounter, the apparently stronger condition is not any more difficult to prove than the weaker one. On the other hand, our main application of uniform convexity, Theorem 3.23 below, only uses the weaker condition.

Exercise 3.21. For vectors $x = (x_1, \ldots, x_n)$ in $\mathbb{R}^n$, consider the norms

$$|x|_p := \left(\sum_{j=1}^n x_j^p\right)^{1/p} \quad \text{for } 1 \leqslant p < \infty, \qquad |x|_\infty := \max\{|x_1|, \ldots, |x_n|\}.$$

(a) Show (by drawing pictures of the unit ball) that $(\mathbb{R}^n, |\cdot|_1)$ and $(\mathbb{R}^n, |\cdot|_\infty)$ are not strictly convex.

(b) Show that the spaces of real-valued functions of class L^1 or L^∞ on any measure space are not strictly convex. (Note that this implies part (a) if you take the measure space to be $\{1, \ldots, n\}$ with the counting measure.)

We will see in §5.3 that all L^p-spaces for $1 < p < \infty$ are uniformly convex; this of course includes the examples $(\mathbb{R}^n, |\cdot|_p)$ defined in Exercise 3.21. Notice that uniform convexity is not a property of the *equivalence class* of a norm but rather of the norm itself—indeed, all norms on $\mathbb{R}^n$ are equivalent, but some are uniformly convex and some are not.

Proposition 3.22. *Every inner product space $(X, \langle\cdot,\cdot\rangle)$ is uniformly convex.*

Proof. Denoting the norm by $\|\cdot\| := \sqrt{\langle\cdot,\cdot\rangle}$, a straightforward computation yields the **parallelogram identity**,

$$\|v + w\|^2 + \|v - w\|^2 = 2\|v\|^2 + 2\|w\|^2 \qquad \forall v, w \in X, \tag{3.2}$$

which for $\|v\| \leqslant 1$ and $\|w\| \leqslant 1$ implies the relation

$$\frac{1}{4}\|v - w\|^2 \leqslant 1 - \left\|\frac{v+w}{2}\right\|^2.$$

This gives a concrete upper bound on $\|v - w\|$ in terms of the distance from $\frac{v+w}{2}$ to the boundary of the unit ball. □

The following theorem on uniformly convex Banach spaces is a useful source of existence results, and plays a key role in characterizing the duals of Hilbert spaces and L^p-spaces. The special case where X is a Hilbert space appeared already (without proof) as Lemma 3.13, and was used in the proof of the theorem about orthogonal decompositions.

Theorem 3.23. *Assume $(X, \|\cdot\|)$ is a uniformly convex Banach space, $K \subset X$ is a closed convex subset and $x \in X \setminus K$. Then the function $K \to (0, \infty) : k \mapsto \|k - x\|$ attains a unique global minimum.*

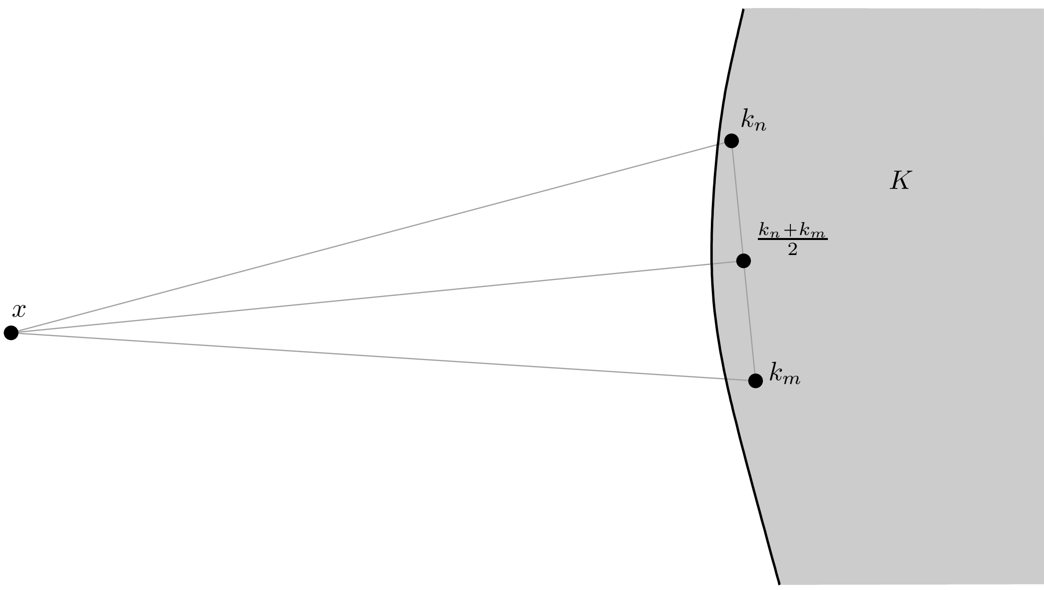

Figure 3. The geometric setup behind the proof of Theorem 3.23.

Note that if $\dim X < \infty$, then Theorem 3.23 follows easily from the fact that since closed and bounded subsets of X are compact, $K \to (0,\infty) : k \mapsto \|k - x\|$ is a proper function: one only has to take a sequence $k_n \in K$ with $\|k_n - x\| \to \inf\{\|k - x\| \mid k \in K\}$ and use compactness to extract a convergent subsequence, whose limit is the desired minimum. This argument completely falls apart if $\dim X = \infty$, because closed bounded subsets are no longer compact. One must instead appeal to the completeness of X, using the idea represented in Figure 3: suppose $k_n, k_m \in K$ both have distances to x that are close to the infimum. After rescaling the whole picture, we can assume without loss of generality that $k_n - x$ and $k_m - x$ are both in the unit ball, in which case so is the midpoint $\frac{(k_n-x)+(k_m-x)}{2} = \frac{k_n+k_m}{2} - x$, where $\frac{k_n+k_m}{2}$ also lies in K since K is convex. By assumption, $\left\|\frac{k_n+k_m}{2} - x\right\|$ cannot be that much smaller than $\|k_n - x\|$ and $\|k_m - x\|$, since both of the latter were already close to the infimum, hence $\frac{k_n+k_m}{2} - x$ cannot be too far away from the boundary of the unit ball. But in that case, uniform convexity implies that $k_n - x$ and $k_m - x$ must be close, or equivalently, k_n and k_m must be close. We will use a version of this argument in the following to show that k_n is a Cauchy sequence, and thus converges to an element that attains the minimum.

Proof of Theorem 3.23. Let $I := \inf\left\{\|k - x\| \;\middle|\; k \in K\right\}$, choose a sequence $k_n \in K$ with $I_n := \|k_n - x\| \to I$, and let

$$z_n := \frac{k_n - x}{I_n} = \frac{k_n - x}{I} + \frac{I - I_n}{I} z_n,$$

which defines a sequence in the unit sphere of X. If $\epsilon > 0$ is given, we can choose $N \in \mathbb{N}$ such that $I_n < I + \epsilon$ for all $n \geqslant N$. For any $m, n \geqslant N$, the fact that K is convex implies $\frac{k_m+k_n}{2} \in K$, thus it satisfies

$$\left\|\frac{(k_m - x) + (k_n - x)}{2}\right\| = \left\|\frac{k_m + k_n}{2} - x\right\| \geqslant I,$$

which implies

$$\left\| \frac{z_m + z_n}{2} \right\| = \left\| \frac{1}{I} \frac{(k_m - x) + (k_n - x)}{2} + \frac{I - I_m}{2I} z_m + \frac{I - I_n}{2I} z_n \right\| \geqslant 1 - \frac{\epsilon}{I}.$$

Since the latter can be made arbitrarily close to 1 by choosing $\epsilon > 0$ small, uniform convexity now implies that $\|z_m - z_n\|$ can be assumed arbitrarily small for N large, so z_n is a Cauchy sequence and therefore converges to some $z_\infty \in X$. It follows that k_n converges to $k_\infty := x + I z_\infty$, and since K is a closed set, $k_\infty \in K$. Clearly $\|x - k_\infty\| = I$.

The uniqueness of the minimum follows almost immediately since, if $k_0, k_1 \in K$ are two minimums, then the argument above shows that $k_0, k_1, k_0, k_1, \ldots$ is a Cauchy sequence, implying $k_0 = k_1$. $\square$

Exercise 3.24. Let $\mathbb{R}^\infty$ denote the vector space of infinite tuples $x = (x_1, x_2, \ldots)$ of real numbers such that at most finitely many of the coordinates $x_n \in \mathbb{R}$ are nonzero. This becomes an inner product space if we define on $\mathbb{R}^\infty$ the obvious generalization of the Euclidean inner product,

$$\langle x, y \rangle := \sum_{n=1}^{\infty} x_n y_n \in \mathbb{R}, \tag{3.3}$$

where the sum always converges since only finitely many of its terms are nonzero. Define a subspace $W \subset \mathbb{R}^\infty$ by $W := \left\{ x \in \mathbb{R}^\infty \;\middle|\; \sum_{n=1}^{\infty} \frac{x_n}{n} = 0 \right\}$.

(a) Prove that $W \subset \mathbb{R}^\infty$ is a closed subspace of codimension 1.

(b) Prove that the orthogonal complement $W^\perp = \left\{ x \in \mathbb{R}^\infty \;\middle|\; \langle x, v \rangle = 0 \text{ for all } v \in W \right\}$ is the trivial subspace of $\mathbb{R}^\infty$.

(c) In §3.2 we proved that for any closed subspace W in a Hilbert space $\mathcal{H}$, $\mathcal{H} = W \oplus W^\perp$. Where does the proof of this theorem go wrong if you try to carry it out with the Hilbert space $\mathcal{H}$ replaced by the *incomplete* inner product space $\mathbb{R}^\infty$?

Hint: $\mathbb{R}^\infty$ is a dense subspace of the Hilbert space ℓ^2 consisting of tuples $x = (x_1, x_2, \ldots)$ that are allowed to have infinitely many nonzero coordinates but must also satisfy $\sum_{n=1}^\infty x_n^2 < \infty$. Equivalently, ℓ^2 is $L^2(\mathbb{N}, \nu)$, the space of square-integrable functions $\mathbb{N} \to \mathbb{R} : n \mapsto x_n$ with the counting measure ν. Notice that $W = \mathbb{R}^\infty \cap z^\perp$ for an element $z \in \ell^2 \backslash \mathbb{R}^\infty$.

3.4. The Riesz representation theorem. We now give an explicit characterization of the dual space of any Hilbert space.

Theorem 3.25 (Riesz representation theorem for Hilbert spaces). *For any Hilbert space $\mathcal{H}$, the map*

$$\mathcal{H} \to \mathcal{H}^* : x \mapsto \Lambda_x, \qquad \Lambda_x(y) := \langle x, y \rangle$$

is a real-linear isomorphism, and it is also an ***isometry****, meaning*

$$\|\Lambda_x\| = \|x\| \qquad \textit{for all } x \in \mathcal{H}.$$

Remark 3.26. The term "real-linear" appears in this statement because if $\mathbb{K} = \mathbb{C}$, the isomorphism $\mathcal{H} \to \mathcal{H}^*$ is not complex linear, but instead complex *antilinear*.

Proof of Theorem 3.25. Corollary 3.9 implies that the map $x \mapsto \Lambda_x$ is a well-defined isometry, and therefore injective. Our task is now to show that it is surjective, i.e. that for any bounded linear functional $\Lambda : \mathcal{H} \to \mathbb{K}$, there exists some $x \in \mathcal{H}$ satisfying $\Lambda(y) = \langle x, y \rangle$ for all $y \in \mathcal{H}$. This is clear if $\Lambda = 0$, so assume $\Lambda \neq 0$. By Exercise 3.27 below, the kernel $K := \ker \Lambda \subset \mathcal{H}$ is then a closed subspace of codimension 1, so by Theorem 3.12, $\mathcal{H} = K \oplus K^\perp$, where $K^\perp \subset \mathcal{H}$ is a 1-dimensional subspace. By choosing any $y \neq 0 \in K^\perp$ and rescaling it by a suitable scalar, we can then find an element $x \in K^\perp$ such that $\Lambda(x) = \|x\|^2 = \Lambda_x(x)$. Since Λ and Λ_x both vanish on K, and $\mathcal{H}$ is spanned by K and x, it follows that $\Lambda = \Lambda_x$. □

EXERCISE 3.27. Assume X is a vector space over $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ and $W \subset X$ is a subspace. One says that W has **codimension** k if the quotient vector space X/W has dimension k. (Note that k may be finite even if X and W are both infinite dimensional.)

(a) Show that the following conditions are equivalent:
 (i) $\operatorname{codim} W = 1$;
 (ii) There exists a vector $w \in X \backslash W$ such that every $x \in X$ can be written as $x = v + \lambda w$ for unique elements $v \in W$ and $\lambda \in \mathbb{K}$;
 (iii) $W = \ker \Lambda$ for some nontrivial linear map $\Lambda : X \to \mathbb{K}$.[4]

(b) Show that if $W = \ker \Lambda = \ker \Lambda'$ for two linear functionals $\Lambda, \Lambda' : X \to \mathbb{K}$, then $\Lambda' = c\Lambda$ for some nonzero scalar $c \in \mathbb{K}$.

(c) Assuming X is a normed vector space and $W = \ker \Lambda$ for a nontrivial linear functional $\Lambda : X \to \mathbb{K}$, show that the following conditions are equivalent:
 (a) $W \subset X$ is closed;
 (b) $W \subset X$ is not dense;
 (c) $\Lambda : X \to \mathbb{K}$ is continuous.

 Hint: The closure of any subspace is also a subspace. If $\Lambda : X \to \mathbb{K}$ is not bounded, there exists a bounded sequence $x_n = v_n + \lambda_n w \in X$ with $v_n \in W$, $w \in X \backslash W$ and $\lambda_n \in \mathbb{K}$ such that $|\Lambda(x_n)| \to \infty$. What can you say about $\frac{v_n}{\lambda_n}$?

COROLLARY 3.28. *For two Hilbert spaces X, Y, every bounded linear operator $A : X \to Y$ has a unique* **adjoint** *(Adjungierte), meaning a bounded linear operator $A^* : Y \to X$ satisfying the relation*

$$\langle y, Ax \rangle = \langle A^* y, x \rangle \qquad \text{for all} \qquad x \in X,\ y \in Y.$$

PROOF. We recall that the transpose $A^\top : Y^* \to X^*$ of A is defined by $(A^\top \lambda)(x) := \lambda(Ax)$. By the Riesz representation theorem, $\lambda \in Y^*$ can always be written as Λ_y for some $y \in Y$, thus

$$\langle y, Ax \rangle = \Lambda_y(Ax) = (A^\top \Lambda_y)(x)$$

will be equal to $\langle A^* y, x \rangle$ for all $x \in X$ if and only if

$$\Lambda_{A^* y} = A^\top \Lambda_y.$$

This formula uniquely determines A^*: writing

$$X \xrightarrow{\Phi_X} X^*, \qquad Y \xrightarrow{\Phi_Y} Y^*$$

for the real-linear isomorphisms given by the Riesz representation theorem, we can define

$$A^* := \Phi_X^{-1} A^\top \Phi_Y.$$

Note that in the case $\mathbb{K} = \mathbb{C}$, the presence of *two* complex-antilinear maps in this composition makes A^* complex linear. Moreover, it is bounded, with

$$\|A^*\| \leqslant \|\Phi_X^{-1}\| \cdot \|A^\top\| \cdot \|\Phi_Y\| = \|A^\top\| = \|A\|.$$

□

DEFINITION 3.29. A bounded linear operator $A : \mathcal{H} \to \mathcal{H}$ on a Hilbert space $\mathcal{H}$ is called **self-adjoint** (*selbstadjungiert*) if $A^* = A$.

[4] Linear maps $X \to \mathbb{K}$ are also called *linear functionals*, and subspaces $W \subset X$ of codimension 1 are also called *hyperplanes*.

EXERCISE 3.30. Assume X and Y are inner product spaces, and $A : X \to Y$ and $A^* : Y \to X$ are linear maps satisfying the adjoint relation

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

Denote the images of these operators by $\operatorname{im} A \subset Y$ and $\operatorname{im} A^* \subset X$.

(a) Prove: $\ker A^* = (\operatorname{im} A)^\perp$ and $\ker A = (\operatorname{im} A^*)^\perp$.
(b) Assume Y is complete, $A : X \to Y$ is continuous and its image is closed. Show that for a given $y \in Y$, the equation $Ax = y$ has solutions $x \in X$ if and only if $\langle y, z \rangle = 0$ for all $z \in \ker A^*$.

EXERCISE 3.31. For an inner product space $\mathcal{H}$ and subspace $W \subset \mathcal{H}$ such that $\mathcal{H} = W \oplus W^\perp$, the **orthogonal projection** (*orthogonale Projektion*) to W is the unique linear map $P : \mathcal{H} \to \mathcal{H}$ such that $P|_W$ is the identity map on W and $\ker P = W^\perp$. Prove:

(a) P is bounded and self-adjoint, and satisfies $P^2 = P$.
(b) The orthogonal projection to $W^\perp$ is given by $\mathbb{1} - P : \mathcal{H} \to \mathcal{H}$.
(c) If $\mathcal{H}$ is complete and $\Pi : \mathcal{H} \to \mathcal{H}$ is a self-adjoint bounded linear operator with $\Pi^2 = \Pi$, then $\operatorname{im} \Pi \subset \mathcal{H}$ is closed and Π is the orthogonal projection onto $\operatorname{im} \Pi$.
Hint: The image of an orthogonal projection is the kernel of another one.

EXERCISE 3.32. Assume $\mathcal{H}$ is a Hilbert space.

(a) Show that the formula $\langle \Lambda_x, \Lambda_y \rangle := \langle y, x \rangle$ defines an inner product on $\mathcal{H}^*$ such that the operator norm $\| \cdot \|$ satisfies $\|\Lambda\|^2 = \langle \Lambda, \Lambda \rangle$ for all $\Lambda \in \mathcal{H}^*$, thus making $\mathcal{H}^*$ into a Hilbert space over $\mathbb{K}$.
(b) Prove that every Hilbert space is reflexive.

3.5. Orthonormal bases and separability. Let us define the **span** of a nonempty subset $S \subset \mathcal{H}$ in an inner product space to be the *closure* of the set of all finite linear combinations of elements of S:

$$\operatorname{Span}(S) := \overline{\left\{ \sum_{j=1}^{N} \lambda_j e_j \;\middle|\; N \in \mathbb{N},\ e_j \in S,\ \lambda_j \in \mathbb{K} \text{ for } j = 1, \dots, N \right\}}.$$

An **orthonormal basis** (*Orthonormalbasis*) of $\mathcal{H}$ is an orthonormal set such that $\operatorname{Span}(S) = \mathcal{H}$. Note that when $\dim \mathcal{H} = \infty$, orthonormal bases are *not* the same thing as Hamel bases: the use of closures in the definition of $\operatorname{Span}(S)$ means that arbitrary elements of $\mathcal{H}$ do not need to be expressible as *finite* linear combinations of basis elements. They only need to be *approximated* by such finite linear combinations. Theorem 3.34 below will give us a useful way to understand in practice what this means; we will see that it is easier to understand if the basis is countable, which turns out to be true for most of the important Hilbert spaces arising in applications.

THEOREM 3.33. *An orthonormal set in a Hilbert space is an orthonormal basis if and only if it is maximal.*

PROOF. If $S \subset \mathcal{H}$ is an orthonormal set with $\operatorname{Span}(S) \neq \mathcal{H}$, then since $\operatorname{Span}(S)$ is closed by definition, it is not dense in $\mathcal{H}$, so that by Corollary 3.14, it has a nontrivial orthogonal complement. That complement contains a unit vector v that can be added to S, implying that S is not a maximal orthonormal set.

Conversely, $S \subset \mathcal{H}$ being an orthonormal basis means $\operatorname{Span}(S) = \mathcal{H}$, thus $\operatorname{Span}(S)^\perp = \{0\}$, implying that there is no nontrivial element $x \in \mathcal{H}$ orthogonal to every $e \in S$, hence S is maximal. $\square$

THEOREM 3.34. *Suppose $\{e_\alpha \in \mathcal{H}\}_{\alpha\in I}$ is an orthonormal basis of a Hilbert space $\mathcal{H}$ with $\dim \mathcal{H} = \infty$. Then for every $x \in \mathcal{H}$, the set*

$$I_x := \{\alpha \in I \mid \langle e_\alpha, x\rangle \neq 0\}$$

is at most countable, and x can be written as an absolutely convergent series

$$x = \sum_{\alpha\in I_x} \langle e_\alpha, x\rangle e_\alpha,$$

where the partial sums are defined via any choice of ordering for the countable set I_x.

PROOF. By assumption, any $x \in \mathcal{H}$ can be written as $x = \lim_{n\to\infty} x_n$ where each $x_n \in \mathcal{H}$ is a finite linear combination of elements from the basis. It follows that the set I_{x_n} is finite for each n. If $\alpha \in I_x$, then

$$\langle e_\alpha, x_n\rangle \to \langle e_\alpha, x\rangle \neq 0$$

implies that α also belongs to I_{x_n} for all n sufficiently large, thus

$$I_x \subset \bigcup_{n=1}^{\infty} I_{x_n},$$

proving that I_x is at most countable. If I_x is finite, then the result now follows from Theorem 3.6. If not, choose an ordering and write $I_x = \{\alpha_1, \alpha_2, \alpha_3, \ldots\}$, and rewrite $e_j := e_{\alpha_j}$ for $j \in \mathbb{N}$. We claim that the series $\sum_{j=1}^{\infty} \langle e_j, x\rangle e_j$ converges absolutely to x. Denote its partial sums by

$$x_N := \sum_{j=1}^{N} \langle e_j, x\rangle e_j, \qquad N \in \mathbb{N}.$$

By Theorem 3.6, for each N, $x - x_N$ is orthogonal to the span of $\{e_1, \ldots, e_N\}$, and

$$\|x\|^2 = \sum_{j=1}^{N} |\langle e_j, x\rangle|^2 + \|x - x_N\|^2,$$

implying

$$\sum_{j=1}^{\infty} \|\langle e_j, x\rangle e_j\|^2 = \sum_{j=1}^{\infty} |\langle e_j, x\rangle|^2 < \infty,$$

which establishes absolute convergence, and since $\mathcal{H}$ is a Banach spce, it follows that the series also converges to some element $x' \in \mathcal{H}$. One can also see this directly as follows: for $N \geqslant M$, we have

$$\|x_N - x_M\|^2 = \left\| \sum_{j=M+1}^{N} \langle e_j, x\rangle e_j \right\|^2 = \sum_{j=M+1}^{N} |\langle e_j, x\rangle|^2,$$

which can be made arbitrarily small by taking both M and N large, implying that x_N is a Cauchy sequence, and thus converges. To see that its limit $x' = \lim_{N\to\infty} x_N$ is actually x, we observe that for all $M \leqslant N$,

$$\langle e_M, x_N\rangle = \langle e_M, x\rangle$$

which implies by taking $N \to \infty$ that $\langle e_M, x'\rangle = \langle e_M, x\rangle$ for all $M \in \mathbb{N}$. Moreover, x is by definition orthogonal to every e_α for $\alpha \notin I_x$, and so is x', since this is also true for each of the partial sums x_N, due to the orthogonality of the basis elements. It follows that $x - x'$ is orthogonal to *all* of the basis elements, and therefore also to their span, which is $\mathcal{H}$, implying $x - x' = 0$. □

It will now be helpful to recall a standard notion from the theory of metric spaces.

DEFINITION 3.35. A metric space is **separable** (*separabel*) if it contains a countable dense subset.

THEOREM 3.36. *Every Hilbert space $\mathcal{H}$ admits an orthonormal basis, and the basis is at most countable if and only if $\mathcal{H}$ is separable.*

PROOF. If $\mathcal{H}$ is separable, then there exists a dense sequence $\{x_1, x_2, x_3, \dots\} \subset \mathcal{H}$, some subsequence of which is linearly independent but has the same span. Applying the Gram-Schmidt orthogonalization procedure then produces a countable orthonormal set that spans $\mathcal{H}$. We leave it as an exercise to show conversely that if $\mathcal{H}$ admits a countable orthonormal basis, then it also admits a countable dense subset. (Hint: Use rational coefficients.)

If $\mathcal{H}$ is not separable, then there is no obvious way to apply the Gram-Schmidt algorithm as described above, but one can still make the set of all orthonormal sets into a partially ordered set and apply Zorn's lemma as in the proof of Theorem 2.40. □

REMARK 3.37. For those who like to keep track of such things, the proof above that *all* Hilbert spaces admit orthonormal bases relies on the axiom of choice, but the proof in the separable case does not. The latter encompasses almost all of the examples that we will actually take interest in.

EXERCISE 3.38. Let ν denote the counting measure on a set I, i.e. every subset $E \subset I$ is ν-measurable and $\nu(E) \in \mathbb{N} \cap \{0, \infty\}$ is the number of points in E. It follows that every function $f : I \to \mathbb{C}$ is ν-measurable, and by a straightforward exercise in measure theory, a ν-integrable function can be nonzero on at most countably many points $\alpha_1, \alpha_2, \alpha_3, \dots \in I$, so that its integral is given by an absolutely convergent series

$$\int_I f\, d\nu = \sum_{\alpha \in I} f(\alpha) := \sum_{n=1}^{\infty} f(\alpha_n) \in \mathbb{C}.$$

All summations appearing in the following should be understood in this sense. The complex Hilbert space $L^2(I, \nu)$ now consists of all functions $f : I \to \mathbb{C}$ that are nonzero on at most countably many points and satisfy $\|f\|_{L^2}^2 = \sum_{\alpha \in I} |f(\alpha)|^2 < \infty$, with the inner product of two functions in this space given by

$$\langle f, g \rangle_{L^2} = \sum_{\alpha \in I} \overline{f(\alpha)} g(\alpha) \in \mathbb{C}.$$

(a) Show that if the set I is finite or countably infinite, then $L^2(I, \nu)$ is separable.
Hint: Show that every $f \in L^2(I, \nu)$ can be approximated arbitrarily well by functions that have real and imaginary parts in $\mathbb{Q}$ at all points and are nonzero on at most finitely many.
(b) Show that if I is uncountable, then $L^2(I, \nu)$ is not separable.
(c) If $\mathcal{H}$ is a complex[5] Hilbert space with orthonormal basis $\{e_\alpha\}_{\alpha \in I}$, show that the map

$$\mathcal{H} \to L^2(I, \nu) : x \mapsto f_x \quad \text{where} \quad f_x(\alpha) := \langle e_\alpha, x \rangle$$

is a unitary isomorphism of Hilbert spaces, i.e. it is an isomorphism and satisfies $\langle f_x, f_y \rangle_{L^2} = \langle x, y \rangle$ for all $x, y \in \mathcal{H}$. Conclude that both this map and its inverse are continuous, and that $\mathcal{H}$ is separable if and only if I is not uncountable.

Comment: Almost all infinite-dimensional Hilbert spaces that one encounters in applications (e.g. $L^2(\mathbb{R})$ or $L^2([0, 1])$ and the related Sobolev spaces that we will study later) turn out to be separable. Thus all of them are unitarily isomorphic to $\ell^2 := L^2(\mathbb{N}, \nu)$.

[5]The analogous statement for a real Hilbert space is obtained by taking functions in $L^2(I, \nu)$ to be real valued and omitting complex conjugation from all formulas.

EXERCISE 3.39. For $\mathcal{H}$ a Hilbert space containing an infinite orthonormal set $e_1, e_2, e_3, \ldots \in \mathcal{H}$, prove that the bounded sequence $\{e_n\}_{n=1}^{\infty}$ has no convergent subsequence. In particular, the closed unit ball in $\mathcal{H}$ is not compact.
Comment: A topological space X is called "locally compact" if for every point $x \in X$, every neighborhood of x contains a compact neighborhood of x, e.g. in a Hilbert space, such a neighborhood could be a sufficiently small closed ball about x. Local compactness in a Hilbert space is in fact equivalent to the condition that the closed unit ball is compact, so this problem in combination with a standard result from first-year analysis proves that a Hilbert space is locally compact if and only if it is finite dimensional. We will later prove that the same is true in Banach spaces; in fact, it is true in arbitrary Hausdorff topological vector spaces. If you're curious to see a proof of the latter statement, see
https://terrytao.wordpress.com/2011/05/24/locally-compact-topological-vector-spaces/

Part 2: Real analysis and L^p-spaces

The second portion of this course deals less with general theorems about Banach or Hilbert spaces than with specific examples of them, whose properties need to be understood *well*, before the full power of the theory can be put into practice. The most important objects in this story are the L^p-spaces, but we will also touch upon Fourier series and transforms, Sobolev spaces, and the theory of "generalized" functions, also known as *distributions*. The main prerequisite is some basic knowledge of measure theory, and it is common in some countries for this material to appear as part of a graduate-level course called "real analysis," which begins with the general theory of the Lebesgue integral and then develops it further.

4. Preliminaries

This section is an interlude before we get to the important part: its purpose is to clarify some notation and basic notions that should already be familiar from the measure-theoretic portion of your introductory analysis courses. In particular, in case you have only seen Lebesgue integration for *real-valued* functions before, I want to be extra sure that you understand the (easy) extension of this theory to complex- or vector-valued functions.

4.1. Integrals of vector-valued functions. We continue with the convention established in §1.5: $\mathbb{K}$ is a field that is assumed to be either $\mathbb{R}$ or $\mathbb{C}$, and we consider functions with values in a fixed finite-dimensional inner product space $(V, \langle\ ,\ \rangle)$ over $\mathbb{K}$, with norm denoted by

$$|\cdot| := \sqrt{\langle\cdot,\cdot\rangle}.$$

The discussion of Fourier analysis starting in §10 will require choosing $\mathbb{K} = \mathbb{C}$, but in most other places, the differences between the real and complex cases will be negligible, e.g. we will sometimes need to use the relation

$$\langle v + w, v + w\rangle = |v|^2 + 2\,\mathrm{Re}\langle v, w\rangle + |w|^2,$$

which is true in both cases, the difference being only that in the real case, the symbol "Re" is redundant. We recall the convention that a complex inner product is antilinear in its first argument and linear in its second:

$$\langle iv, w\rangle = -i\langle v, w\rangle, \qquad \langle v, iw\rangle = i\langle v, w\rangle.$$

We will sometimes make use of the fact that a complex vector space is also a real vector space (of twice the dimension).

CONVENTION 4.1. By the standard definition, a **measure space** (*Messraum*) consists of three pieces of data $(X, \mathcal{A}, \mu)$: a set X, a σ-algebra $\mathcal{A} \subset 2^X$ and a measure $\mu : \mathcal{A} \to [0, \infty]$. Since we will almost never have occasion to talk about the σ-algebra itself, we shall typically omit it from the notation and simply call (X, μ) a measure space, referring when necessary to the elements of $\mathcal{A}$ as the **measurable** (or **μ-measurable**) **sets** (*messbare Mengen*).

Given a measure space (X, μ), a function $f : X \to V$ is considered **measurable** (*messbar*) if the preimage of every open subset of V is μ-measurable in X. It is easy to show that if we choose any real basis $e_1, \ldots, e_n$ of V and write $f = \sum_{j=1}^n f_j e_j$ for functions $f_j : X \to \mathbb{R}$, then f is measurable if and only if all of the f_j are measurable. Similarly, if f is measurable then $|f| : X \to [0, \infty)$ is also measurable, and in this case the component functions f_j are μ-integrable if and only if $\int_X |f| \, d\mu < \infty$. One can then define the vector-valued integral

$$\int_X f \, d\mu = \sum_{j=1}^n \left(\int_X f_j \, d\mu \right) e_j \in V. \tag{4.1}$$

We will sometimes also write

$$\int_X f(x) \, d\mu(x) := \int_X f \, d\mu$$

when we want to specify the name of the variable $x \in X$.

Exercise 4.2. Show that for μ-integrable functions $f : X \to V$, the integral $\int_X f \, d\mu \in V$ defined above is independent of the choice of real basis $e_1, \ldots, e_n \in V$.

Exercise 4.3. Show that for every μ-integrable function $f : X \to V$, $\left|\int_X f \, d\mu\right| \leqslant \int_X |f| \, d\mu$.

The simplest example beyond $V = \mathbb{R}$ is $V = \mathbb{C}$ with the standard inner product $\langle v, w \rangle := \bar{v} w$. Here we can take $e_1 := 1$ and $e_2 := i$ as a real basis of $\mathbb{C}$, so $f : X \to \mathbb{C}$ is measurable/integrable if and only if its real and imaginary parts are both measurable/integrable, and (4.1) becomes

$$\int_X f \, d\mu = \int_X (\operatorname{Re} f) \, d\mu + i \int_X (\operatorname{Im} f) \, d\mu \in \mathbb{C}.$$

Remark 4.4. The assumption $\dim V < \infty$ is inessential for much of what follows, though obviously the definition of $\int_X f \, d\mu \in V$ requires some modification if V has no finite basis. A definition (using approximation by step functions) for the case where V is an arbitrary Banach space may be found in [**Lan93**]. Since many details become more complicated in this more general setting, we will stick to the case $\dim V < \infty$ but give occasional remarks on what needs to be modified in order to lift this assumption.

4.2. Differentiation under the integral sign. The following standard consequence of the dominated convergence theorem will be an essential tool to have at our disposal.

Theorem 4.5. *Suppose (Y, ν) is a measure space, M is a metric space, and $\varphi : M \times Y \to V$ is a function with the following properties:*

(1) For every $x \in M$, the function $\varphi(x, \cdot) : Y \to V$ is measurable and satisfies $|\varphi(x, \cdot)| \leqslant \psi$ for some fixed ν-integrable function $\psi : Y \to [0, \infty]$ independent of x;
(2) For every $y \in Y$, the function $\varphi(\cdot, y) : M \to V$ is continuous.

Then the function $F : M \to V$ given by

$$F(x) := \int_Y \varphi(x, \cdot) \, d\nu$$

is continuous. If additionally M is an open subset of $\mathbb{R}^m$ with coordinates $x = (x_1, \ldots, x_m)$ and the partial derivatives $\frac{\partial \varphi}{\partial x_j} : M \times Y \to V$ exist for every $j = 1, \ldots, m$ and also satisfy the two conditions above, then F is continuously differentiable and satisfies

$$\partial_j F(x) = \int_Y \frac{\partial \varphi}{\partial x_j}(x, \cdot) \, d\nu$$

for every $x \in M$ and $j = 1, \ldots, m$.

PROOF. To prove $F : M \to V$ is continuous at a point $x \in M$, consider a sequence $x_n \in M$ with $x_n \to x$. Since $\varphi(\cdot, y) : M \to V$ is continuous for every $y \in Y$, the sequence of functions $\varphi(x_n, \cdot) : Y \to \mathbb{R}$ converges pointwise to $\varphi(x, \cdot) : Y \to \mathbb{R}$, and by assumption it also satisfies

$$|\varphi(x_n, \cdot)| \leqslant \psi \quad \text{for all } n$$

for a fixed ν-integrable function $\psi : Y \to [0, \infty]$. The dominated convergence theorem thus implies $F(x_n) \to F(x)$.

Now suppose additionally that $M = \mathcal{U} \subset \mathbb{R}^m$ is open and $\frac{\partial \varphi}{\partial x_j}(x, y)$ exists for all $(x, y) \in \mathcal{U} \times Y$ and defines a function that is (for each fixed $y \in Y$) continuous with respect to $x \in \mathcal{U}$ and (for each fixed $x \in \mathcal{U}$) measurable with respect to $y \in Y$, additionally satisfying the bound $\left|\frac{\partial \varphi}{\partial x_j}(x, \cdot)\right| \leqslant \psi$ for all $x \in \mathcal{U}$. Let $e_1, \ldots, e_m$ denote the standard basis of $\mathbb{R}^m$. The partial derivative $\frac{\partial \varphi}{\partial x_j}(x, y)$ is then the limit as $h \to 0$ of the difference quotients

$$D_j^h \varphi(x, y) := \frac{\varphi(x + he_j, y) - \varphi(x, y)}{h} \in V,$$

where for each $x \in \mathcal{U}$, the function $D_j^h \varphi(x, \cdot) : Y \to V$ is defined for all $h \in \mathbb{R} \backslash \{0\}$ sufficiently close to 0. For any sequence $h_n \in \mathbb{R} \backslash \{0\}$ with $h_n \to 0$, we therefore have

$$D_j^{h_n} \varphi(x, \cdot) \to \frac{\partial \varphi}{\partial x_j}(x, \cdot) \quad \text{pointwise on } Y. \tag{4.2}$$

For every $y \in Y$ and $h \in \mathbb{R}$ sufficiently close to 0, the fact that $\varphi(\cdot, y)$ is continuously differentiable with respect to x_j allows us to derive a formula for $D_j^h \varphi(x, y)$ using the fundamental theorem of calculus: we have

$$\varphi(x + he_j, y) = \varphi(x, y) + \int_0^1 \frac{d}{dt} \varphi(x + the_j, y)\, dt = \varphi(x, y) + h \int_0^1 \frac{\partial \varphi}{\partial x_j}(x + the_j, y)\, dt,$$

and thus

$$D_j^h \varphi(x, y) = \int_0^1 \frac{\partial \varphi}{\partial x_j}(x + the_j, y)\, dt, \tag{4.3}$$

giving rise to the bound

$$\left|D_j^h \varphi(x, y)\right| \leqslant \int_0^1 \psi(y)\, dt = \psi(y).$$

Since ψ is integrable, one can again apply the dominated convergence theorem and obtain a convergence result for the corresponding difference quotients of F: for any sequence $h_n \in \mathbb{R} \backslash \{0\}$ with $h_n \to 0$, we have

$$D_j^{h_n} F(x) := \frac{F(x + h_n e_j) - F(x)}{h_n} = \int_Y D_j^{h_n} \varphi(x, \cdot)\, d\nu \to \int_Y \frac{\partial \varphi}{\partial x_j}(x, \cdot)\, d\nu$$

Since the sequence h_n was arbitrary, this proves

$$\frac{\partial F}{\partial x_j}(x) = \lim_{h \to 0} D_j^h F(x) = \int_Y \frac{\partial \varphi}{\partial x_j}(x, \cdot)\, d\nu,$$

and the continuity of $\frac{\partial F}{\partial x_j}$ now follows from the same argument as the continuity of F. $\square$

REMARK 4.6. The hypotheses of Theorem 4.5 can be weakened (at the cost of more cumbersome notation) in various ways that are sometimes useful. Most importantly, since the continuity and differentiability of F are purely local conditions, the bounds $|\varphi(x, \cdot)| \leqslant \psi$ and $\left|\frac{\partial \varphi}{\partial x_j}(x, \cdot)\right| \leqslant \psi$ do not really need to hold with a single function ψ for every $x \in M$; it suffices if every $x_0 \in M$ has a neighborhood $\mathcal{U} \subset M$ and an associated integrable function $\psi_{x_0} : Y \to [0, \infty]$ that bounds these functions for all $x \in \mathcal{U}$. One can also insert the words "almost everywhere" in various places among the hypotheses, so that certain steps in the proof make sense only after deleting sets of measure zero from Y, which is harmless. For more elaborate versions of the statement, see e.g. [**AE01**, Theorems 3.17 and 3.18]) or [**Wen19**].

4.3. Some standard function spaces. Let's review some essential facts that need to be understood about spaces of L^p-functions or C^m-functions on domains in Euclidean space.

Assume (X, μ) is an arbitrary measure space, and $(V, \langle\ ,\ \rangle)$ is again a finite-dimensional inner product space over $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ with norm $|\cdot| := \sqrt{\langle \cdot, \cdot \rangle}$. The L^p-norm of a measurable function $f : X \to V$ is defined for each $p \in [1, \infty)$ by

$$\|f\|_{L^p} := \|f\|_{L^p(X)} := \left(\int_X |f(x)|^p \, d\mu(x) \right)^{1/p} \in [0, \infty],$$

and for the case $p = \infty$,

$$\|f\|_{L^\infty} := \|f\|_{L^\infty(X)} := \operatorname{ess\,sup}_{x \in X} |f(x)| := \inf \left\{ c \geqslant 0 \;\middle|\; |f| \leqslant c \text{ almost everywhere} \right\} \in [0, \infty].$$

We assume the reader is familiar with the standard Minkowski and Hölder inequalities, and the fact that the space $L^p(X, \mu)$ of equivalence classes of measurable functions (defined almost everywhere) with finite L^p-norms is a Banach space. We will typically abbreviate

$$L^p(X) := L^p(X, \mu)$$

when the measure is clear from context. Here is a precise statement of the completeness theorem:

Theorem 4.7 (see e.g. [**Sal16**, §4.2]). *For $1 \leqslant p \leqslant \infty$, every L^p-Cauchy sequence $f_n \in L^p(X)$ is L^p-convergent and also has a pointwise almost everywhere convergent subsequence. In the case $p = \infty$, the original sequence also converges pointwise almost everywhere.* □

The usual Hölder inequality for real-valued functions combines with the Cauchy-Schwarz inequality on $(V, \langle\ ,\ \rangle)$ and Exercise 4.3 to give the relation

$$\left| \int_X \langle f(x), g(x) \rangle \, d\mu(x) \right| \leqslant \int_X |\langle f(x), g(x) \rangle| \, d\mu(x) \leqslant \|f\|_{L^p} \cdot \|g\|_{L^q}$$

for $f \in L^p(X)$ and $g \in L^q(X)$ with $\frac{1}{p} + \frac{1}{q} = 1$.

Exercise 4.8. In case you have only seen $L^p(X)$ defined for real-valued functions before, convince yourself that the proof of Theorem 4.7 still goes through when the functions in $L^p(X)$ take values in an arbitrary (real or complex) finite-dimensional vector space.

Exercise 4.9. Show that for every measurable function $f : X \to V$,

$$\|f\|_{L^\infty} \leqslant \liminf_{p \to \infty} \|f\|_{L^p},$$

and if additionally either $\mu(X) < \infty$ or $f \in L^r(X)$ for some $r \in [1, \infty)$, then $\|f\|_{L^\infty} = \lim_{p \to \infty} \|f\|_{L^p}$. *Hint: For the case with $f \in L^r(X)$ for some $r < \infty$, show that $\|f\|_{L^p} \leqslant \|f\|_{L^r}^{r/p} \cdot \|f\|_{L^\infty}^{1-r/p}$ holds for every $p > r$. (Note that this is not a version of Hölder's inequality—it is easier.) Use this to bound* $\limsup_{p \to \infty} \|f\|_{L^p}$.

When X is an open subset of Euclidean space

$$X := \Omega \subset \mathbb{R}^n \qquad \text{with} \qquad \mu := m \text{ (Lebesgue measure)},$$

it is often useful to consider functions that need not be in $L^p(\Omega)$ but restrict to L^p-functions on all compact subsets. Since compact subsets of $\mathbb{R}^n$ are bounded and therefore have finite measure, this includes for instance the nontrivial constant functions, which are not in $L^p(\Omega)$ unless $m(\Omega) < \infty$. We define the vector space

$$L^p_{\mathrm{loc}}(\Omega) := \left\{ f : \Omega \to V \;\middle|\; f|_K \in L^p(K) \text{ for all } K \subset \Omega \text{ compact} \right\} / \sim,$$

where as usual the equivalence relation $f \sim g$ means $f = g$ almost everywhere on Ω. The functions in $L^p_{\mathrm{loc}}(\Omega)$ are said to be **locally of class** L^p, and in the case $p = 1$, a function $f \in L^1_{\mathrm{loc}}(\Omega)$ is called **locally integrable** on Ω. The space $L^p_{\mathrm{loc}}(\Omega)$ is strictly larger than $L^p(\Omega)$, and it is not a Banach space since there is no single norm to determine whether or not a given function is of class L^p_{loc}. It does however have a natural topology as a locally convex space, defined via the family of seminorms

$$\|f\|_{L^p(K)} = \left\| f|_K \right\|_{L^p}, \tag{4.4}$$

where K ranges over the set of all compact subsets $K \subset \Omega$. Note that these are seminorms rather than norms, because a function $f \in L^p_{\mathrm{loc}}(\Omega)$ may be nontrivial but satisfy $\|f\|_{L^p(K)} = 0$ because it vanishes almost everywhere on K. Convergence of a sequence $f_j \to f$ in $L^p_{\mathrm{loc}}(\Omega)$ means that $\|f - f_j\|_{L^p(K)} \to 0$ is satisfied for all of these seminorms, which is equivalent to saying that the restrictions of f_j to every compact subset $K \subset \Omega$ are convergent in $L^p(K)$ to $f|_K$.

It is possible to derive the topology of $L^p_{\mathrm{loc}}(\Omega)$ from a countable subfamily of the seminorms in (4.4). Indeed, Ω can always be covered by a nested sequence

$$\Omega_1 \subset \bar{\Omega}_1 \subset \Omega_2 \subset \bar{\Omega}_2 \subset \ldots \subset \bigcup_{m \in \mathbb{N}} \Omega_m = \Omega$$

of open subsets $\Omega_m \subset \Omega$ with compact closures $K_m := \bar{\Omega}_m$, so that any compact subset $K \subset \Omega$ is contained in Ω_m for $m \in \mathbb{N}$ sufficiently large. For a concrete construction of Ω_m, one can for instance define $\Omega_m := \left\{ x \in \Omega \;\middle|\; |x| < m \text{ and } \operatorname{dist}(x, \mathbb{R}^n \backslash \Omega) > 1/m \right\}$, where for two subsets $A, B \subset \mathbb{R}^n$, we denote

$$\operatorname{dist}(A, B) := \inf \left\{ |x - y| \;\middle|\; x \in A,\; y \in B \right\}.$$

A sequence $f_j \in L^p_{\mathrm{loc}}(\Omega)$ is then L^p_{loc}-convergent if and only if it converges in each of the seminorms $\|\cdot\|_{L^p(K_m)}$ for $m \in \mathbb{N}$, and similarly, every open subset of $L^p_{\mathrm{loc}}(\Omega)$ is a union of sets of the form $\left\{ f \in L^p_{\mathrm{loc}}(\Omega) \;\middle|\; \|f - f_0\|_{L^p(K_m)} < \epsilon \right\}$ for $f_0 \in L^p_{\mathrm{loc}}(\Omega)$, $m \in \mathbb{N}$ and $\epsilon > 0$. It follows via Theorem 2.26 that $L^p_{\mathrm{loc}}(\Omega)$ is metrizable, with open subsets defined via the metric

$$d(f, g) := \sum_{m=1}^{\infty} \frac{1}{2^m} \frac{\|f - g\|_{L^p(K_m)}}{1 + \|f - g\|_{L^p(K_m)}}.$$

In fact, $L^p_{\mathrm{loc}}(\Omega)$ is a Fréchet space: completeness follows from the completeness of the Banach space $L^p(K_m)$ for every m, as a sequence $f_j \in L^p_{\mathrm{loc}}(\Omega)$ is Cauchy if and only if $f_j|_{K_m}$ is Cauchy in $L^p(K_m)$ for every m.

Continuing under the assumption that $\Omega \subset \mathbb{R}^n$ is an open subset, we shall continue to denote

$$C^m(\Omega) := \{ f : \Omega \to V \; m \text{ times continuously differentiable} \}$$

for integers $m \geqslant 0$. This is not a Banach space, but it can be made into one by imposing an extra boundedness condition, i.e. by defining the C^m-norm

$$\|f\|_{C^m} := \|f\|_{C^m(\Omega)} := \sum_{|\alpha| \leqslant m} \sup_{x \in \Omega} |\partial^\alpha f(x)|,$$

and then defining the Banach space

$$C^m_b(\Omega) := \left\{ f \in C^m(\Omega) \;\middle|\; \|f\|_{C^m} < \infty \right\}.$$

Convergence of a sequence f_j in the C^m-norm means uniform convergence of f_j and all its derivatives up to order m. By standard results of first-year analysis, $C^m_b(\Omega)$ with this norm is a Banach space for every integer $m \geqslant 0$. A useful subspace of $C^m_b(\Omega)$ can be defined by[6]

$$C^m(\bar{\Omega}) := \left\{ f \in C^m_b(\Omega) \;\middle|\; \partial^\alpha f \text{ is uniformly continuous for all multi-indices } \alpha \text{ with } |\alpha| \leqslant m \right\}.$$

The following exercise explains the motivation for this notation.

EXERCISE 4.10. Let $\bar{\Omega} \subset \mathbb{R}^n$ denote the closure of the open subset $\Omega \subset \mathbb{R}^n$.

(a) Show that if $f : \Omega \to \mathbb{R}$ is uniformly continuous, then it admits a (necessarily unique) continuous extension over $\bar{\Omega}$. (Note that the converse is also true if Ω is bounded, since continuous functions on compact sets are always uniformly continuous.)
(b) Show that $C^m(\bar{\Omega})$ is a closed subspace of $C^m_b(\Omega)$, hence it is a Banach space with the C^m-norm.

In particular, $C^m(\bar{\Omega})$ can be characterized as the space of C^m-functions on Ω whose derivatives up to order m all admit bounded continuous extensions to $\bar{\Omega}$. (The word "bounded" is redundant here if Ω itself is bounded, since $\bar{\Omega}$ is then compact.)[7]

For smooth (i.e. infinitely differentiable) functions, we define

$$C^\infty(\Omega) := \bigcap_{m \geqslant 0} C^m(\Omega), \qquad C^\infty_b(\Omega) := \bigcap_{m \geqslant 0} C^m_b(\Omega),$$

and endow the latter with the locally convex topology defined via the entire sequence of norms $\|\cdot\|_{C^m}$ for $m \geqslant 0$, hence a sequence $f_j \in C^\infty_b(\Omega)$ is C^∞-convergent if and only if its derivatives of all orders are uniformly convergent. One could similarly define $C^\infty(\bar{\Omega})$, but this turns out to be the same space as $C^\infty_b(\Omega)$ since the boundedness of the derivatives of order $m+1$ implies uniform continuity for derivatives of order m. Since the family of C^m-norms for $m \geqslant 0$ is countable, one can define a metric on $C^\infty_b(\Omega)$ in the same manner that we did so for $L^p_{\mathrm{loc}}(\Omega)$, and the completeness of $C^m_b(\Omega)$ for each $m \geqslant 0$ implies that $C^\infty_b(\Omega)$ is a Fréchet space.

The C^m-topologies also have local variants, which are defined on $C^m(\Omega)$ without requiring any boundedness condition: we say that a sequence $f_j \in C^m(\Omega)$ is **C^m_{loc}-convergent** to $f \in C^m(\Omega)$ if

$$\|f - f_j\|_{C^m(K)} := \sum_{|\alpha| \leqslant m} \max_{x \in K} |\partial^\alpha f(x) - \partial^\alpha f_j(x)| \to 0$$

for every compact subset $K \subset \Omega$. As with $L^p_{\mathrm{loc}}(\Omega)$, one can use an exhaustion of Ω by a nested sequence of open subsets with compact closure to characterize this notion of convergence via a countable family of seminorms, making $C^m(\Omega)$ into a Fréchet space with the C^m_{loc}-topology. There is similarly a C^∞_{loc}-topology on $C^\infty(\Omega)$, in which sequences converge if and only if their derivatives of *all* orders converge on compact subsets, and this endows $C^\infty(\Omega)$ with a natural Fréchet space structure. Note that for each $m \in \mathbb{N} \cup \{0, \infty\}$, C^m_{loc}-convergence is a much weaker notion than C^m-convergence, i.e. many sequences converge in C^m_{loc} but not in C^m, and the behavior of a C^m_{loc}-convergent sequence "near infinity" can be arbitrarily wild.

The **support** (*Träger*) $\mathrm{supp}(f) \subset \Omega$ of a function $f : \Omega \to V$ is the closure of the set $\{x \in \Omega \mid f(x) \neq 0\}$. We will denote

$$C^m_0(\Omega) := \left\{ f \in C^m(\Omega) \;\middle|\; \mathrm{supp}(\Omega) \subset \Omega \text{ is compact} \right\}.$$

[6]There is potential ambiguity in the notation when $\Omega = \mathbb{R}^n$ since $\mathbb{R}^n$ is its own closure, but $C^m(\bar{\mathbb{R}}^n)$ is nonetheless a smaller space than $C^m(\mathbb{R}^n)$.

[7]If $\bar{\Omega}$ is compact and has a sufficiently "nice" boundary, meaning for instance that the boundary is a C^m-smooth submanifold of $\mathbb{R}^n$, then one can show with somewhat more effort that $C^m(\bar{\Omega})$ is the space of C^m-functions on Ω that admit extensions of class C^m over some open neighborhood of $\bar{\Omega}$; for details, see [**AF03**, §5.19–§5.21].

This is a subspace of both of the Banach spaces $C_b^m(\Omega)$ and $C^m(\bar{\Omega})$, though not a closed subspace in either case, as a sequence of functions with growing compact supports can easily be C^m-convergent to one whose support is not compact.

5. Duality in L^p-spaces

5.1. The pairing of L^p and L^q. For this section, assume (X, μ) is an arbitrary measure space, and $(V, \langle\ ,\ \rangle)$ is again a finite-dimensional inner product space over $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ with norm $|\cdot| := \sqrt{\langle\cdot,\cdot\rangle}$. Our aim is to prove a characterization of the space $(L^p(X))^*$ of bounded linear functionals $L^p(X) \to \mathbb{K}$ which, like Theorem 3.25, is also sometimes called the *Riesz representation theorem*. To prepare the statement, notice that whenever $1 \leqslant p, q \leqslant \infty$ with $\frac{1}{p} + \frac{1}{q} = 1$, Hölder's inequality gives rise to a real-linear map

$$L^q(X) \to (L^p(X))^* : g \mapsto \Lambda_g := \int_X \langle g, \cdot\rangle \, d\mu \tag{5.1}$$

satisfying $\|\Lambda_g\|_{(L^p)^*} \leqslant \|g\|_{L^q}$, where $\|\cdot\|_{(L^p)^*}$ denotes the operator norm on bounded linear operators $L^p(X) \to \mathbb{K}$.

LEMMA 5.1. *Assume $1 \leqslant p, q \leqslant \infty$, $\frac{1}{p} + \frac{1}{q} = 1$ and, additionally, either $p < \infty$ or X is σ-finite.*[8] *Then for every $f \in L^p(X)$,*

$$\sup_{g \in L^q(X)\setminus\{0\}} \frac{\left|\int_X \langle g, f\rangle \, d\mu\right|}{\|g\|_{L^q}} = \|f\|_{L^p},$$

and the ratio on the left hand side attains its maximum in the case $p < \infty$.

PROOF. Hölder's inequality implies that the ratio in question can never be greater than $\|f\|_{L^p}$. There is nothing to prove if $f = 0$, so assume $f \in L^p(X)$ is nontrivial. If $p < \infty$, we define $g : X \to V$ by $g := |f|^{p-2} f$ at points where $f \neq 0$ and $g := 0$ otherwise. Then g satisfies $|g| = 1$ almost everywhere if $p = 1$, and in the other cases, $|g|^q = |f|^{q(p-1)} = |f|^p$, thus $g \in L^q(X)$ and

$$\int_X \langle g, f\rangle \, d\mu = \int_X |f|^p \, d\mu = \|f\|_{L^p}^p = (\|f\|_{L^p}^p)^{\frac{p-1}{p}} \cdot \|f\|_{L^p} = (\|g\|_{L^q}^q)^{1-\frac{1}{p}} \cdot \|f\|_{L^p} = \|g\|_{L^q} \cdot \|f\|_{L^p},$$

so this choice of $g \in L^q(X)$ maximizes the ratio in question.

In the case $p = \infty$ and $q = 1$, we argue by contradiction and suppose that $\|f\|_{L^\infty}$ is strictly greater than the supremum of $\left|\int_X \langle g, f\rangle \, d\mu\right| / \|g\|_{L^1}$ over all $g \in L^1(X)\setminus\{0\}$. Then there exists a constant c strictly greater than this supremum such that the set $A' := \{x \in X \mid |f(x)| \geqslant c\}$ has positive measure. Assuming X is σ-finite, there also exists a subset $A \subset A'$ with $0 < \mu(A) < \infty$, and the function g defined as $f/|f|$ on A and 0 everywhere else is then in $L^1(X)$, with $\|g\|_{L^1} = \mu(A)$. Since $|f| \geqslant c > \left|\int_X \langle g, f\rangle \, d\mu\right| / \|g\|_{L^1}$ on A, we now find the contradiction,

$$\left|\int_X \langle g, f\rangle \, d\mu\right| = \int_A |f| \, d\mu \geqslant \mu(A) \cdot c = \|g\|_{L^1} \cdot c > \left|\int_X \langle g, f\rangle \, d\mu\right|.$$

□

COROLLARY 5.2. *For every $1 \leqslant p, q \leqslant \infty$ with $\frac{1}{p} + \frac{1}{q} = 1$, if either $p > 1$ or X is σ-finite, then the bounded real-linear map (5.1) is injective and satisfies $\|\Lambda_g\|_{(L^p)^*} = \|g\|_{L^q}$ for all $g \in L^q(X)$.* □

EXERCISE 5.3. Show that for any $f \in L^\infty(X)$ satisfying $|f| < \|f\|_{L^\infty}$ almost everywhere, the inequality $\left|\int_X \langle g, f\rangle \, d\mu\right| \leqslant \|g\|_{L^1} \cdot \|f\|_{L^\infty}$ is strict for every $g \in L^1(X)\setminus\{0\}$.

---

[8] Certain measure-theoretic pathologies can arise in the case $p = \infty$ that are excluded if we assume X is σ-finite. This is not the most general assumption possible, but it suffices for all applications we will want to consider. For more general versions of the results in this section involving duality between $L^1(X)$ and $L^\infty(X)$, see [Sal16, §4.5].

Here is the hard part:

THEOREM 5.4 (Riesz representation theorem for L^p). *The map* (5.1) *is bijective for all* $1 < p, q < \infty$ *with* $\frac{1}{p} + \frac{1}{q} = 1$*, and also for* $p = 1$ *and* $q = \infty$ *if* X *is* σ*-finite.*

REMARK 5.5. In the case $\mathbb{K} = \mathbb{C}$, the map $L^q(X) \to (L^p(X))^*$ in (5.1) is complex antilinear and thus is not, strictly speaking, an isomorphism of complex Banach spaces. However, one can also define a space $(L^p(X))'$ consisting of all bounded complex-antilinear functionals $\Lambda : L^p(X) \to \mathbb{C}$ and consider a complex-linear map defined by

$$L^q(X) \to (L^p(X))' : g \mapsto \Lambda'_g := \int_X \langle \cdot, g \rangle \, d\mu. \tag{5.2}$$

It is an easy exercise to check that this map is bijective whenever (5.1) is, so under the same hypotheses as Theorem 5.4, it is a complex Banach space isomorphism.

The proof of Theorem 5.4 given below follows the same strategy as our proof of the corresponding statement about Hilbert spaces in Theorem 3.25. The crucial idea in the latter was that given a nontrivial dual vector $\Lambda \in \mathcal{H}^*$ for a Hilbert space $\mathcal{H}$, the right place to search for elements x with $\Lambda = \langle x, \cdot \rangle$ is in the *orthogonal complement* of the closed hyperplane $\ker \Lambda \subset \mathcal{H}$. While the notion of orthogonality does not make sense in $L^p(X)$ for $p \neq 2$, Hölder's inequality furnishes us with a reasonable substitute in the form of the natural pairing of L^p with L^q for $\frac{1}{p} + \frac{1}{q} = 1$; informally, we can thus regard the orthogonal complement of a subspace in $L^p(X)$ as a subspace of $L^q(X)$. With this notion in mind, the main task is then to prove, as we did for Hilbert spaces in §3.2, that a proper closed subspace $K \subset L^p(X)$ always has a *nontrivial* orthogonal complement. Our proof of this in the Hilbert space setting required two fundamental ingredients:

(1) The uniform convexity of every Hilbert space $\mathcal{H}$;
(2) The differentiability of the function $t \mapsto \|x + tv\|^2$ for any $x, v \in \mathcal{H}$.

Both were easy to prove using the characterization of the Hilbert space norm via an inner product, but since the latter is not available in $L^p(X)$ for $p \neq 2$, we will have to work a bit harder.

Recall that every Banach space $(E, \|\cdot\|)$ has a canonical continuous inclusion into the dual of its dual space, defined by

$$\Phi : E \to E^{**}, \qquad \Phi(v)\Lambda := \Lambda(v) \text{ for } v \in E,\ \Lambda \in E^*.$$

The injectivity of this map for general Banach spaces is not so obvious, though for $E = L^p(X)$ with $p < \infty$, it is an easy consequence of the following corollary of Lemma 5.1. Outside of these special cases, it follows immediately from the Hahn-Banach theorem, which we will cover in Part 3 of this course; its standard proof uses the axiom of choice.

LEMMA 5.6. *For every normed vector space* $(E, \|\cdot\|)$ *and every* $x \in E$*, there exists a dual vector* $\Lambda \in E^*$ *with* $\|\Lambda\| = 1$ *and* $\Lambda(x) = \|x\|$*.*

PROOF FOR $E = L^p(X)$ WITH $p < \infty$. Given $f \in L^p(X)$, choose $\Lambda := \Lambda_g \in (L^p(X))^*$ for $g \in L^q(X)$ as in Lemma 5.1, then normalize g. □

COROLLARY 5.7. *For every Banach space* $(E, \|\cdot\|)$*, the canonical map* $\Phi : E \to E^{**}$ *is an injective isometry, i.e. it satisfies* $\|\Phi(x)\| = \|x\|$ *for every* $x \in E$*.* □

Recall that $(E, \|\cdot\|)$ is called **reflexive** if the inclusion $\Phi : E \hookrightarrow E^{**}$ is also surjective. For $E = L^p(X)$ with $1 < p < \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$, Theorem 5.4 identifies E^* with $L^q(X)$ and then identifies E^{**} in turn with $L^p(X)$, so that under these identifications, $\Phi : E \to E^{**}$ becomes a map $L^p(X) \to L^p(X)$ uniquely determined by[9]

$$\int_X \langle \Phi(f), g\rangle \, d\mu = \int_X \langle f, g\rangle \, d\mu \quad \text{for all} \quad g \in L^q(X).$$

This implies that $\int_X \langle \Phi(f) - f, g\rangle \, d\mu$ vanishes for all $g \in L^q(X)$, proving that the function $\Phi(f) - f \in L^p(X)$ is identified with the trivial element of $(L^q(X))^*$, which makes $\Phi : L^p(X) \to L^p(X)$ the identity map.

COROLLARY 5.8. *For $1 < p < \infty$, $L^p(X)$ is reflexive.* □

REMARK 5.9. Reflexivity is in fact a general property of uniformly convex Banach spaces, by the *Milman-Pettis theorem*; see e.g. [**RS80**, Problem V.15].

Theorem 5.4 is false for $p = \infty$ and $q = 1$; the dual of $L^\infty(X)$ is generally a larger space than can be described via such a pairing. One can see this by comparing Lemma 5.6 with Exercise 5.3: there exist nontrivial functions $f \in L^\infty(X)$ for which an element $\Lambda \in (L^\infty(X))^*$ with $\|\Lambda\|_{(L^\infty)^*} = 1$ satisfying $|\Lambda(f)| = \|f\|_{L^\infty}$ must exist, but the strictness of the inequality in Exercise 5.3 implies that Λ cannot be represented by any function in $L^1(X)$.[10] For more counterexamples, see also [**Rud87**, Chapter 6, Exercise 13] or [**Sal16**, Example 4.36]. It follows that $L^1(X)$ is not reflexive, and by the next exercise, neither is $L^\infty(X)$.

EXERCISE 5.10. For a Banach space E, let $\Phi_E : E \hookrightarrow E^{**}$ and $\Phi_{E^*} : E^* \hookrightarrow E^{***}$ denote the canonical inclusions, and denote by $\Phi_E^\top : E^{***} \to E^*$ the transpose of Φ_E.

(a) Show that $\Phi_E^\top \circ \Phi_{E^*}$ is the identity map on E^*.
(b) Show that the image of Φ_E is always a closed subspace of E^{**}.
(c) Deduce that E^* is reflexive if and only if E is reflexive.
Hint: Another easy consequence of the Hahn-Banach theorem is that if $A : X \to Y$ is a bounded linear operator between Banach spaces such that $\operatorname{im} A \subset Y$ is closed and $A^* : Y^* \to X^*$ is injective, then A is surjective.

5.2. Differentiability of the norm. Let us examine whether the function $\|f + tg\|_{L^p}^p$ can be differentiated with respect to $t \in \mathbb{R}$ for $f, g \in L^p(X)$. Assume in the following

$$1 < p < \infty.$$

For $v, w \in V$ and $t \in \mathbb{R}$ with $v + tw \neq 0$, the differentiability of the function $x \mapsto x^{p/2}$ for $x \neq 0$ implies

$$\begin{aligned} \frac{d}{dt}|v + tw|^p &= \frac{d}{dt}\langle v + tw, v + tw\rangle^{p/2} \\ &= \frac{p}{2}\langle v + tw, v + tw\rangle^{\frac{p}{2}-1} \cdot \frac{d}{dt}\left(|v|^2 + 2t \operatorname{Re}\langle v, w\rangle + t^2|w|^2\right) \\ &= p|v + tw|^{p-2}\left(\operatorname{Re}\langle v, w\rangle + t|w|^2\right) = p|v + tw|^{p-2} \cdot \operatorname{Re}\langle v + tw, w\rangle. \end{aligned} \tag{5.3}$$

[9] One needs to be a bit careful with this argumentation in the case $\mathbb{K} = \mathbb{C}$, because the bijection $E^* \cong L^q(X)$ is then complex antilinear rather than linear, so substituting $L^q(X)$ for E^* identifies E^{**} with the space $(L^q(X))'$ of bounded complex-*antilinear* maps $L^q(X) \to \mathbb{C}$ instead of the actual dual space of $L^q(X)$. As mentioned in Remark 5.5, however, the Riesz representation identifies the latter complex-linearly with $L^p(X)$.

[10] Quoting Lemma 5.6 for $L^\infty(X)$ means we are relying on the Hahn-Banach theorem, which is inherently non-constructive, i.e. it guarantees the existence of an element in $(L^\infty(X))^* \backslash L^1(X)$ as an artefact of the axioms of set theory, but gives no hint how one could ever write one down. In fact, *all* proofs that $(L^\infty(X))^* \backslash L^1(X) \neq \varnothing$ are non-constructive in this sense. Readers who wish to explore this particular set-theoretic rabbit hole may consult [**Sch99**, Chapter 14]; see also https://mathoverflow.net/questions/5351/whats-an-example-of-a-space-that-needs-the-hahn-banach-theorem.

Notice that by the Cauchy-Schwarz inequality on $(V, \langle\ ,\ \rangle)$, the right hand side of this expression satisfies

$$\left| p|v + tw|^{p-2} \cdot \mathrm{Re}\langle v + tw, w\rangle \right| \leqslant p|v + tw|^{p-1} \cdot |w|,$$

whenever $v + tw \neq 0$. Since $p > 1$, one can therefore sensibly define the right hand side of (5.3) to be 0 when $v + tw = 0$, and the relation remains correct, since in this case

$$\frac{d}{dt}|v + tw|^p = \frac{d}{ds}|(v + tw) + sw|^p \bigg|_{s=0} = \frac{d}{ds}|s|^p|w|^p \bigg|_{s=0} = |w|^p \lim_{s \to 0} \frac{|s|^p}{s} = 0.$$

With this understood, for any given $f, g \in L^p(X)$, differentiation under the integral sign now suggests the formula

$$\begin{aligned} \frac{d}{dt}\|f + tg\|_{L^p}^p &= \frac{d}{dt}\int_X |f(x) + tg(x)|^p \, d\mu(x) = \int_X \frac{\partial}{\partial t}|f(x) + tg(x)|^p \, d\mu(x) \\ &= \int_X p|f(x) + tg(x)|^{p-2} \cdot \mathrm{Re}\langle f(x) + tg(x), g(x)\rangle \, d\mu(x), \end{aligned} \tag{5.4}$$

where the same application of the Cauchy-Schwarz inequality interprets the integrand on the right as 0 whenever $f(x) + tg(x) = 0$. Let us use Theorem 4.5 to justify this formula at $t = 0$. Set $(Y, \nu) := (X, \mu)$ and $M := (-1, 1) \subset \mathbb{R}$ and define $\varphi : (-1, 1) \times X \to V$ by $\varphi(t, x) := |f(x) + tg(x)|^p$, so $\frac{\partial \varphi}{\partial t}(t, x)$ is given by the integrand on the right hand side of (5.4). Both φ and $\frac{\partial \varphi}{\partial t}$ are then continuous functions of $t \in (-1, 1)$ for every fixed $x \in X$. For every fixed $t \in (-1, 1)$, they also satisfy

$$|\varphi(t, x)| \leqslant (|f(x)| + |g(x)|)^p \tag{5.5}$$

and

$$\left|\frac{\partial \varphi}{\partial t}(t, x)\right| \leqslant p\,(|f(x)| + |g(x)|)^{p-1} \cdot |g(x)|. \tag{5.6}$$

By Minkowski's inequality,

$$\int_X (|f(x)| + |g(x)|)^p \, d\mu(x) = \big\| |f| + |g| \big\|_{L^p}^p \leqslant (\|f\|_{L^p} + \|g\|_{L^p})^p < \infty,$$

thus the right hand side of (5.5) defines a μ-integrable function on X. It follows in turn that the function $(|f| + |g|)^{p-1}$ is of class $L^{p/(p-1)}$ on X, and since $\frac{p-1}{p} + \frac{1}{p} = 1$, Hölder's inequality implies that the right hand side of (5.6) is also μ-integrable. The hypotheses of Theorem 4.5 are thus satisfied, and we conclude:

LEMMA 5.11. *For any $f, g \in L^p(X)$ with $1 < p < \infty$, the function $\mathbb{R} \to [0, \infty) : t \mapsto \|f + tg\|_{L^p}^p$ is differentiable and satisfies*

$$\frac{d}{dt}\|f + tg\|_{L^p}^p \bigg|_{t=0} = p \int_X |f|^{p-2} \cdot \mathrm{Re}\langle f, g\rangle \, d\mu.$$

□

5.3. Uniform convexity of L^p. In order to prove that $L^p(X)$ is uniformly convex for $1 < p < \infty$, we begin with the observation that the function

$$V \to \mathbb{R} : v \mapsto |v|^p$$

is strictly convex for all $p \in (1, \infty)$. One can show this by computing that its Hessian is positive definite everywhere outside of the origin; at the origin it may fail to have second derivatives, but it is then easy enough to check the convexity condition along segments connecting 0 to any other point. It follows that the function

$$\psi : V \times V \to \mathbb{R} : (v, w) \mapsto \frac{|v|^p + |w|^p}{2} - \left| \frac{v+w}{2} \right|^p \tag{5.7}$$

is nonnegative everywhere, and strictly positive whenever $v \neq w$. For any constant $\epsilon > 0$, its restriction to the compact subset

$$K := \left\{ (v, w) \in V \times V \;\middle|\; |v - w|^p \geqslant \epsilon \text{ and } |v|^p + |w|^p \leqslant 1 \right\}$$

therefore satisfies $\psi|_K \geqslant \delta$ for some constant $\delta > 0$.[11] Now if $v, w \in V$ are any elements with $v \neq w$, set $\tau := (|v|^p + |w|^p)^{1/p} > 0$, $v' := v/\tau$ and $w' := w/\tau$, so $|v'|^p + |w'|^p = 1$, and the condition $|v' - w'|^p \geqslant \epsilon$ is equivalent to $|v - w|^p \geqslant \epsilon\tau^p$. Under this condition, $\psi(v', w') \geqslant \delta$ becomes $\psi(v, w) \geqslant \delta\tau^p$, which proves:

LEMMA 5.12. *Given any $p \in (1, \infty)$ and $\epsilon > 0$, there exists $\delta > 0$ such that the function ψ in (5.7) satisfies*

$$|v - w|^p \geqslant \epsilon \left(|v|^p + |w|^p \right) \quad \Rightarrow \quad \psi(v, w) \geqslant \delta \left(|v|^p + |w|^p \right) \qquad \forall v, w \in V.$$

□

EXERCISE 5.13. Extract from Lemma 5.12 a new proof that $(V, \langle\ ,\ \rangle)$ is uniformly convex.

The uniform convexity of $L^p(X)$ is an easy application of the following estimate.

THEOREM 5.14. *Given any $p \in (1, \infty)$ and $\epsilon > 0$, there exists $\delta > 0$ such that for all $f, g \in L^p(X)$,*

$$\frac{\|f\|_{L^p}^p + \|g\|_{L^p}^p}{2} - \left\| \frac{f+g}{2} \right\|_{L^p}^p \geqslant \delta \left[\|f - g\|_{L^p}^p - \epsilon \left(\|f\|_{L^p}^p + \|g\|_{L^p}^p \right) \right].$$

PROOF. Given $f, g \in L^p(X)$ and $\epsilon > 0$, decompose X into the subsets

$$A := \left\{ x \in X \;\middle|\; |f(x) - g(x)|^p \geqslant \epsilon \left(|f(x)|^p + |g(x)|^p \right) \right\}, \qquad A^c = X \backslash A.$$

For $x \in A$, we have $\psi(f(x), g(x)) \geqslant \delta_0 \left(|f(x)|^p + |g(x)|^p \right)$ for some constant $\delta_0 > 0$ provided by Lemma 5.12. Now using the fact that $|f - g|^p < \epsilon \left(|f|^p + |g|^p \right)$ on A^c, while $\psi(f, g) \geqslant 0$ and $\left| \frac{f-g}{2} \right|^p = \left| \frac{f+(-g)}{2} \right|^p \leqslant \frac{|f|^p + |g|^p}{2}$ hold everywhere, we estimate

$$\begin{aligned}
\frac{\|f\|_{L^p}^p + \|g\|_{L^p}^p}{2} - \left\| \frac{f+g}{2} \right\|_{L^p}^p &\geqslant \int_A \psi(f, g)\, d\mu \geqslant \delta_0 \int_A \left(|f|^p + |g|^p \right) d\mu \geqslant \frac{\delta_0}{2^{p-1}} \int_A |f - g|^p\, d\mu \\
&= \frac{\delta_0}{2^{p-1}} \left(\|f - g\|_{L^p}^p - \int_{A^c} |f - g|^p\, d\mu \right) \\
&\geqslant \frac{\delta_0}{2^{p-1}} \left(\|f - g\|_{L^p}^p - \epsilon \int_{A^c} \left(|f|^p + |g|^p \right) d\mu \right) \\
&\geqslant \frac{\delta_0}{2^{p-1}} \left(\|f - g\|_{L^p}^p - \epsilon \left(\|f\|_{L^p}^p + \|g\|_{L^p}^p \right) \right).
\end{aligned}$$

[11]Recall from Remark 4.4 that we are assuming $\dim V < \infty$, and we are using that assumption here in order to say that K is compact. However, if V is an infinite-dimensional Hilbert space, then one can fix an orthonormal basis, single out two basis vectors $e_1, e_2 \in V$ and then argue as follows: if $(v_n, w_n) \in K$ is a sequence such that $\psi(v_n, w_n) \to 0$, then by choosing suitable new orthonormal bases for each n, we can transform each (v_n, w_n) by isometries of $(V, \langle\ ,\ \rangle)$ (which leave both K and ψ invariant) so that without loss of generality, each v_n and w_n lies in the span of e_1 and e_2. It follows now that the sequence (v_n, w_n) lives in a compact subset of V, so a subsequence converges to some $(v, w) \in K$ with $\psi(v, w) = 0$, which cannot exist. The estimate $\psi|_K \geqslant \delta > 0$ therefore also holds in this case.

Set $\delta := \delta_0/2^{p-1}$. □

COROLLARY 5.15. *For* $1 < p < \infty$, $L^p(X)$ *is uniformly convex.*

PROOF. Assuming f and g in the inequality of Theorem 5.14 lie in the unit ball in $L^p(X)$, the left hand side is bounded above by $1 - \|(f+g)/2\|_{L^p}^p$, and the right hand side bounded below by $\delta\left(\|f-g\|_{L^p}^p - 2\epsilon\right)$. It follows that for any $\epsilon > 0$, there exists $\delta > 0$ such that all f, g in the unit ball satisfy

$$1 - \left\|\frac{f+g}{2}\right\|_{L^p}^p \geqslant \delta\left(\|f-g\|_{L^p}^p - 2\epsilon\right).$$

To show uniform convexity, assume $\epsilon' > 0$ is given: we then need to find $\delta' > 0$ such that bounding f and g away from each other by ϵ' forces $(f+g)/2$ away from the boundary of the ball by distance δ'. Set $\epsilon := (\epsilon')^p/4$, so that $\|f-g\|_{L^p} \geqslant \epsilon'$ implies $\|f-g\|_{L^p}^p - 2\epsilon \geqslant 2\epsilon$, and then choose $\delta > 0$ sufficiently small so that the inequality is satisfied. It follows that if $\|f-g\|_{L^p} \geqslant \epsilon'$, then $1 - \|(f+g)/2\|_{L^p}^p \geqslant 2\delta\epsilon = \delta(\epsilon')^p/2$, so setting $\delta' := \delta(\epsilon')^p/2$ does the trick. □

REMARK 5.16. The notion of uniform convexity and Corollary 5.15 are originally due to Clarkson [**Cla36**], and the literature contains many other proofs based on more powerful inequalities than in Theorem 5.14; see for instance [**LL01**, §2.5], which uses *Hanner's inequality.* Our proof has been adapted from [**Shi18**].

5.4. Proof of the representation theorem. As in the Hilbert space case, the idea for finding a function $g \in L^q(X)$ to represent any given $\Lambda \in (L^p(X))^*$ is to look for nontrivial functions whose pairing with $L^p(X)$ annihilates $\ker\Lambda$. We do this by finding the closest point in $\ker\Lambda$ to some $h \in L^p(X)\setminus\ker\Lambda$.

PROOF OF THEOREM 5.4 FOR $1 < p < \infty$. Assume $p, q \in (1, \infty)$ with $\frac{1}{p} + \frac{1}{q} = 1$. Given $\Lambda \in (L^p(X))^*$, we need to find $g \in L^q(X)$ such that $\int_X \langle g, f\rangle\, d\mu = \Lambda(f)$ for all $f \in L^p(X)$. Assume $\Lambda \neq 0$ since the problem is otherwise trivial, let $K := \ker\Lambda \subset L^p(X)$ and choose $h \in L^p(X)\setminus K$; after multiplication by a scalar, we may assume $\Lambda(h) = 1$. Then K is a closed convex subset, and since $L^p(X)$ is uniformly convex, Theorem 3.23 provides an element $k_0 \in K$ minimizing the distance to h. For any $k \in K$, Lemma 5.11 then gives

$$0 = \left.\frac{d}{dt}\|h - (k_0 - tk)\|_{L^p}^p\right|_{t=0} = p\int_X |h - k_0|^{p-2} \cdot \operatorname{Re}\langle h - k_0, k\rangle\, d\mu,$$

where the integral on the right hand side is well defined due to Hölder's inequality. The symbol "Re" in this formula is redundant in the case $\mathbb{K} = \mathbb{R}$, while if $\mathbb{K} = \mathbb{C}$, replacing $k \in K$ with $ik \in K$ in this relation shows that the same thing holds with the imaginary part instead of the real part, implying that the function $\tilde{g} := |h - k_0|^{p-2}(h - k_0)$ satisfies

$$\int_X \langle \tilde{g}, k\rangle\, d\mu = 0 \quad \text{for all} \quad k \in K.$$

Observe that since $h - k_0 \in L^p(X)$ and $|\tilde{g}| \leqslant |h - k_0|^{p-1}$, $\tilde{g} \in L^q(X)$. Now let $g := c\tilde{g} \in L^q(X)$ for a constant $c > 0$ to be determined momentarily. The relation above implies $\int_X \langle g, f\rangle\, d\mu = \Lambda(f)$ holds for all $f \in K$, and moreover,

$$\int_X \langle g, h - k_0\rangle\, d\mu = c\int_X |h - k_0|^{p-2}\langle h - k_0, h - k_0\rangle\, d\mu = c\|h - k_0\|_{L^p}^p > 0,$$

so the latter matches $\Lambda(h - k_0) = \Lambda(h) = 1$ if we set $c := 1/\|h - k_0\|_{L^p}^p$. Clearly $h - k_0 \notin K$, so $L^p(X)$ is spanned by K and $h - k_0$, thus we have proved that $\int_X \langle g, f\rangle\, d\mu = \Lambda(f)$ holds for all $f \in L^p(X)$. □

The case $p = 1$ is easily derived from the case $p > 1$ if X has finite measure, and we will then use σ-finiteness to extend to the case $\mu(X) = \infty$. We will need to know that L^1-functions can be approximated by L^p-functions for $p > 1$.

LEMMA 5.17. *For every* $p \in (1, \infty]$, $L^p(X) \cap L^1(X)$ *is dense in* $L^1(X)$.

PROOF. Given $f \in L^1(X)$ and $n \in \mathbb{N}$, denote

$$A_n := \{x \in X \mid |f(x)| \leqslant n\}$$

and define $f_n : X \to V$ as the product of f with the characteristic function of A_n. Since $f \in L^1(X)$ and $|f| > 1$ on $X\backslash A_1$, we have $\mu(X\backslash A_1) \leqslant \int_{X\backslash A_1} |f|\, d\mu \leqslant \int_X |f|\, d\mu < \infty$, i.e. $X\backslash A_1$ has finite measure. Clearly $|f_n| \leqslant n$ everywhere for each $n \in \mathbb{N}$, and since $|f|^p \leqslant |f|$ on A_1,

$$\|f_n\|_{L^p}^p = \int_{X\backslash A_1} |f_n|^p\, d\mu + \int_{A_1} |f|^p\, d\mu \leqslant n^p \mu(X\backslash A_1) + \int_{A_1} |f|\, d\mu \leqslant n^p \mu(X\backslash A_1) + \|f\|_{L^1} < \infty,$$

so $f_n \in L^p(X)$ for all p. The complements of the sets A_n form a nested sequence with empty intersection

$$X\backslash A_1 \supset X\backslash A_2 \supset X\backslash A_3 \supset \ldots \supset \bigcap_{n\in\mathbb{N}} (X\backslash A_n) = \varnothing,$$

and they all have finite measure since $\mu(X\backslash A_1) < \infty$, thus

$$\|f - f_n\|_{L^1} = \int_{X\backslash A_n} |f|\, d\mu \to 0 \quad \text{as} \quad n \to \infty,$$

proving $f_n \to f$ in $L^1(X)$. □

PROOF OF THEOREM 5.4 FOR $p = 1$ AND $\mu(X) < \infty$. The advantage of having finite measure is that for every $p' > p \geqslant 1$, $L^{p'}(X)$ is contained in $L^p(X)$, and the inclusion $L^{p'}(X) \hookrightarrow L^p(X)$ is a continuous linear map. This follows from Hölder's inequality, which for $r \geqslant p$ with $\frac{1}{p'} + \frac{1}{r} = \frac{1}{p}$ gives

$$\|f\|_{L^p} \leqslant \|1\|_{L^r} \cdot \|f\|_{L^{p'}} = \mu(X)^{1/r} \cdot \|f\|_{L^{p'}}.$$

Now if $\Lambda \in (L^1(X))^*$, then for $f \in L^p(X)$ with $1 < p < \infty$,

$$|\Lambda(f)| \leqslant \|\Lambda\|_{(L^1)^*} \cdot \|f\|_{L^1} \leqslant \mu(X)^{1/q} \cdot \|\Lambda\|_{(L^1)^*} \cdot \|f\|_{L^p}, \tag{5.8}$$

where $q \in (1, \infty)$ is determined by $\frac{1}{p} + \frac{1}{q} = 1$. This means Λ also belongs to $(L^p(X))^*$, so by the $p > 1$ case of Theorem 5.4, there exists a function $g_p \in L^q(X)$ such that $\Lambda(f) = \int_X \langle g_p, f\rangle\, d\mu$ for all $f \in L^p(X)$. Notice that if $p < p' < \infty$, then $g_{p'} \in L^{q'}(X)$ with $\frac{1}{p'} + \frac{1}{q'} = 1$, where $q' < q$, thus $L^{p'}(X) \subset L^p(X)$ and $L^q(X) \subset L^{q'}(X)$. It follows that g_p is also in $L^{q'}(X)$ and satisfies

$$\int_X \langle g_p - g_{p'}, f\rangle\, d\mu = \Lambda(f) - \Lambda(f) = 0 \quad \text{for all} \quad f \in L^{p'}(X),$$

$g_p - g_{p'} \in L^{q'}(X)$ defines the trivial element of $(L^{p'}(X))^*$, implying $g_p - g_{p'} = 0$ almost everywhere. For this reason we will now drop p from the notation and write g_p for every $p \in (1, \infty)$ as a single function g, which belongs to $L^q(X)$ for every $q \in (1, \infty)$. By (5.8) and Corollary 5.2, it satisfies

$$\|g\|_{L^q} = \|\Lambda\|_{(L^p)^*} \leqslant \mu(X)^{1/q} \cdot \|\Lambda\|_{(L^1)^*} \quad \text{for every} \quad q \in (1, \infty).$$

We claim that this implies $g \in L^\infty(X)$ with $\|g\|_{L^\infty} \leqslant \|\Lambda\|_{(L^1)^*}$. Indeed, for each $c > 0$, let $A_c := \{x \in X \mid |g(x)| \geqslant c\}$; then fixing $p, q \in (1, \infty)$ with $\frac{1}{p} + \frac{1}{q} = 1$, we have

$$c\mu(A_c)^{1/q} \leqslant \|g\|_{L^q} \leqslant \mu(X)^{1/q} \cdot \|\Lambda\|_{(L^1)^*}.$$

Taking the limit $q \to \infty$ then yields $c \leqslant \|\Lambda\|_{(L^1)^*}$ unless $\mu(A_c) = 0$, thus proving the claim.

We have now found a function $g \in L^\infty(X)$ such that $\Lambda(f) = \int_X \langle g, f\rangle \, d\mu$ holds for all $f \in L^p(X)$ with $1 < p < \infty$. For an arbitrary $f \in L^1(X)$, Lemma 5.17 then provides a sequence $f_n \in L^p(X)$ with $f_n \to f_1$ in L^1, and Hölder's inequality implies

$$\left| \int_X \langle g, f\rangle \, d\mu - \int_X \langle g, f_n\rangle \, d\mu \right| \leqslant \int_X |\langle g, f - f_n\rangle| \, d\mu \leqslant \|g\|_{L^\infty} \cdot \|f - f_n\|_{L^1} \to 0,$$

thus

$$\Lambda(f) = \lim_{n\to\infty} \Lambda(f_n) = \lim_{n\to\infty} \int_X \langle g, f_n\rangle \, d\mu = \int_X \langle g, f\rangle \, d\mu.$$

$\square$

Proof of Theorem 5.4 for $p = 1$ and $\mu(X) = \infty$. We assume X is σ-finite, so $X = \bigcup_{n\in\mathbb{N}} X_n$ for subsets $X_n \subset X$ with $\mu(X_n) < \infty$, and without loss of generality

$$X_1 \subset X_2 \subset X_3 \subset \dots .$$

Any $\Lambda \in (L^1(X))^*$ gives rise to functionals $\Lambda_n \in (L^1(X_n))^*$ for every $n \in \mathbb{N}$, defined by

$$\Lambda_n(f) := \Lambda(f_n), \qquad \text{where} \qquad f_n := \begin{cases} f & \text{on } X_n, \\ 0 & \text{on } X \backslash X_n, \end{cases}$$

and they satisfy

$$\|\Lambda_n\|_{(L^1)^*} = \sup_{f\in L^1(X_n)\backslash\{0\}} \frac{|\Lambda(f_n)|}{\|f_n\|_{L^1}} \leqslant \sup_{f\in L^1(X)\backslash\{0\}} \frac{|\Lambda(f)|}{\|f\|_{L^1}} = \|\Lambda\|_{(L^1)^*}.$$

Applying the theorem for the case of finite measure, we obtain functions $g_n \in L^\infty(X_n)$ such that $\Lambda(f) = \int_{X_n} \langle g_n, f\rangle \, d\mu$ for every $f \in L^1(X)$ that vanishes outside of X_n, with norms satisfying $\|g_n\|_{L^\infty} \leqslant \|\Lambda\|_{(L^1)^*}$ for all n. Notice that for $n > m \geqslant 1$ and a function $f \in L^1(X)$ that vanishes outside of X_m, f also vanishes outside of X_n and thus satisfies

$$\int_{X_m} \langle g_m, f\rangle \, d\mu = \Lambda(f) = \int_{X_n} \langle g_n, f\rangle \, d\mu = \int_{X_m} \langle g_n, f\rangle \, d\mu,$$

implying $\int_{X_m} \langle g_m - g_n, f\rangle \, d\mu = 0$ for all $f \in L^1(X_m)$. It follows that $g_m - g_n|_{X_m} \in L^\infty(X_m)$ defines the trivial element of $(L^1(X_m))^*$ and therefore vanishes almost everywhere. This shows that each g_n can in fact be regarded as the restriction to X_n of a single function $g : X \to V$, and since $\|g_n\|_{L^\infty} \leqslant \|\Lambda\|_{(L^1)^*}$ for every n, the set on which $|g| > \|\Lambda\|_{(L^1)^*}$ is the union of countably many sets of measure zero, implying $g \in L^\infty(X)$ with $\|g\|_{L^\infty} \leqslant \|\Lambda\|_{(L^1)^*}$.

We claim finally that $\Lambda(f) = \int_X \langle g, f\rangle \, d\mu$ holds for every $f \in L^1(X)$. To see this, for each $n \in \mathbb{N}$ define $h_n \in L^1(X)$ as the product of f with the characteristic function of X_n, so $\|f - h_n\|_{L^1} = \int_{X\backslash X_n} |f| \, d\mu \to 0$ as $n \to \infty$. Using the continuity of Λ and Hölder's inequality, we now conclude

$$\Lambda(f) = \lim_{n\to\infty} \Lambda(h_n) = \lim_{n\to\infty} \int_X \langle g, h_n\rangle \, d\mu = \int_X \langle g, f\rangle \, d\mu.$$

$\square$

6. Separability of L^p

Recall that a topological space is called **separable** if it contains a countable dense subset. The simplest examples that come to mind are finite-dimensional vector spaces, e.g. $\mathbb{Q}^n$ is a countable dense subset of $\mathbb{R}^n$. In this section, we would like to prove that $L^p(X)$ is also separable when $1 \leqslant p < \infty$. This requires some measure-theoretic assumptions on X, so in order to avoid over-complicating the problem, we shall restrict ourselves to the case where X is a subset Ω of $\mathbb{R}^n$. (See [Sal16, §4.3] for a treatment of more general situations.)

THEOREM 6.1. *For any $p \in [1, \infty)$ and any Lebesgue-measurable set $\Omega \subset \mathbb{R}^n$ endowed with the Lebesgue measure m, the space $L^p(\Omega)$ is separable.*

We have an immediate corollary involving Hilbert spaces, thanks to Theorem 3.36:

COROLLARY 6.2. *For any Lebesgue-measurable set $\Omega \subset \mathbb{R}^n$, the Hilbert space $L^2(\Omega)$ admits a countable orthonormal basis.* □

We shall prove Theorem 6.1 by constructing an explicit countable set of functions $Q(\mathbb{R}^n) \subset L^p(\mathbb{R}^n)$ that is dense in $L^p(\mathbb{R}^n)$. Given any $f \in L^p(\Omega)$ for $\Omega \subset \mathbb{R}^n$, one can then extend f to a function $\hat{f} \in L^p(\mathbb{R}^n)$ that vanishes outside of Ω, find a sequence $\hat{f}_k \in Q(\mathbb{R}^n)$ converging to $\hat{f}$ in L^p, and observe that the restrictions $f_k := \hat{f}_k|_\Omega$ therefore converge in L^p to f, proving that the countable set $Q(\Omega) := \left\{ f|_\Omega \;\middle|\; f \in Q(\mathbb{R}^n) \right\}$ is dense in $L^p(\Omega)$.

The set $Q(\mathbb{R}^n) \subset L^p(\mathbb{R}^n)$ is easy to describe. In the following, we denote the characteristic function of a subset $A \subset \mathbb{R}^n$ by

$$\chi_A : \mathbb{R}^n \to \mathbb{R}, \qquad \chi_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{otherwise.} \end{cases}$$

Let us first fix a countable dense subset

$$V_0 \subset V$$

in the vector space V where our functions take their values; this is clearly possible since $\dim V < \infty$. (If we were allowing V to be an infinite-dimensional Banach space, then we would now add the assumption that V is separable.) We refer to a set $Q \subset \mathbb{R}^n$ as a **dyadic cube** if Q is of the form

$$Q = \left[\frac{m_1}{2^N}, \frac{m_1 + 1}{2^N} \right] \times \ldots \times \left[\frac{m_n}{2^N}, \frac{m_n + 1}{2^N} \right] \subset \mathbb{R}^n$$

for some $m_1, \ldots, m_n, N \in \mathbb{Z}$ with $N \geqslant 0$. Observe that the set of all dyadic cubes is countable, and so therefore is the set of characteristic functions $\chi_Q : \mathbb{R}^n \to \mathbb{R}$ of dyadic cubes. It follows that for every $k \in \mathbb{N}$, the set of k-tuples of dyadic cubes is countable, and thus so is the set of all finite tuples of dyadic cubes. Finally, for each individual tuple $(Q_1, \ldots, Q_k)$ of dyadic cubes, there is a countable set of functions $f : \mathbb{R}^n \to V$ of the form

$$f = \sum_{j=1}^{k} \chi_{Q_j} v_j, \qquad v_1, \ldots, v_k \in V_0.$$

We define $Q(\mathbb{R}^n)$ to be the set of all functions of this type, i.e. all finite linear combinations (with coefficients in the countable set V_0) of characteristic functions of dyadic cubes. All of these functions are bounded and have compact support, so they belong to $L^p(\mathbb{R}^n)$ for every $p \in [1, \infty]$. Our goal is to prove:

PROPOSITION 6.3. *For every $p \in [1, \infty)$, the countable set $Q(\mathbb{R}^n)$ is dense in $L^p(\mathbb{R}^n)$.*

We will use the following fundamental fact from the theory of Lebesgue integration. Recall that a function is called **simple** (or sometimes a **step function**) if it takes only finitely many values. A simple function on a measure space (X, μ) is measurable if and only if it is a finite linear combination of characteristic functions of measurable sets, and it is then integrable if and only if all of those sets have finite measure, which is equivalent to saying that the function's support has finite measure. The integrable simple functions form a linear subspace of $L^p(X)$ for every $p \in [1, \infty]$, and we shall denote it by

$$S(X) \subset L^p(X).$$

LEMMA 6.4. *For every measure space (X, μ) and $1 \leqslant p < \infty$, $S(X)$ is dense in $L^p(X)$.*

PROOF. Depending on your definition of integration, the $p = 1$ case may be understood as either a theorem or a tautology; e.g. [Lan93] defines $L^1(X)$ to be a quotient (modulo equality almost everywhere) of the L^1-closure of $S(X)$. Let us take the more common definition as in [Sal16], where $\int_X f\,d\mu \in [0,\infty]$ for a measurable function $f : X \to [0,\infty]$ is the supremum of $\int_X s\,d\mu$ for all measurable simple functions with $0 \leqslant s \leqslant f$, and for $f : X \to \mathbb{R}$, $\int_X f\,d\mu := \int_X f^+\,d\mu - \int_X f^-\,d\mu$ with $f^\pm : X \to [0,\infty)$ such that $f = f^+ - f^-$ and $|f| = f^+ + f^-$. Then given $f : X \to \mathbb{R}$ of class L^p, there exist increasing sequences of measurable simple functions $0 \leqslant f_1^\pm \leqslant f_2^\pm \leqslant \ldots \leqslant f^\pm$ such that $f_n^\pm \to f^\pm$ pointwise as $n \to \infty$. Since $(x+y)^p \geqslant x^p + y^p$ for all $x, y \geqslant 0$ and $p \geqslant 1$,[12]

$$\int_X |f^+|^p\,d\mu + \int_X |f^-|^p\,d\mu = \int_X (|f^+|^p + |f^-|^p)\,d\mu \leqslant \int_X |f^+ + f^-|^p\,d\mu = \int_X |f|^p\,d\mu < \infty,$$

so $\int_X |f_n^\pm|^p\,d\mu \leqslant \int_X |f|^p\,d\mu < \infty$. This implies that every $|f_n^\pm|^p$ (and therefore also every $f_n^\pm$) is a finite linear combination of characteristic functions of sets with finite measure, so $f_n^\pm \in S(X)$, and thus $f_n := f_n^+ - f_n^- \in S(X)$. Now $|f - f_n|^p \to 0$ pointwise, and using the convexity of the function $x \mapsto x^p$,

$$|f - f_n|^p = |(f^+ - f_n^+) - (f^- - f_n^-)|^p \leqslant 2^{p-1}|f^+ - f_n^+|^p + 2^{p-1}|f^- - f_n^-|^p \leqslant 2^p\left(|f^+|^p + |f^-|^p\right),$$

where the function on the right hand side is integrable, so the dominated convergence theorem implies $\int_X |f - f_n|^p\,d\mu \to 0$. The result for real-valued functions now easily extends to functions valued in the finite-dimensional vector space V by choosing a real basis as in §4.1. $\square$

EXERCISE 6.5. Show that $S(X)$ is dense in $L^\infty(X)$ if and only if $\mu(X) < \infty$.

With Lemma 6.4 in hand, our goal is now to show that $Q(\mathbb{R}^n)$ is dense in $S(\mathbb{R}^n)$.

LEMMA 6.6. *Every open subset $A \subset \mathbb{R}^n$ is a union of a sequence of dyadic cubes $Q_1, Q_2, Q_3, \ldots$ whose interiors are all pairwise disjoint.*

PROOF. Let $\mathcal{O}$ denote the set of all dyadic cubes that are contained in A. Since dyadic cubes can be arbitrarily small, $A = \bigcup_{Q \in \mathcal{O}} Q$, and the set $\mathcal{O}$ is countable since there are only countably many dyadic cubes in total. Write $\mathcal{O} = \{\hat{Q}_1, \hat{Q}_2, \ldots\}$; this is not the desired sequence since it contains pairs $\hat{Q}_j, \hat{Q}_k$ whose interiors intersect, but observe that for any such pair, the part of $\hat{Q}_j$ disjoint from $\hat{Q}_k$ can be covered by finitely many smaller dyadic cubes whose interiors are disjoint from each other and from $\hat{Q}_k$. We can therefore construct a new sequence $Q_1, Q_2, \ldots$ by setting $Q_1 := \hat{Q}_1$ and then replacing each $\hat{Q}_k$ for $k \geqslant 2$ with a finite collection of dyadic cubes with interiors that are disjoint from each other and from $\bigcup_{j=1}^{k-1} \hat{Q}_j$. $\square$

LEMMA 6.7. *For every open subset $A \subset \mathbb{R}^n$ with $m(A) < \infty$ and every $v \in V$, $\epsilon > 0$ and $p \in [1,\infty)$, $Q(\mathbb{R}^n)$ contains a function f with $\|\chi_A v - f\|_{L^p} < \epsilon$.*

PROOF. Pick $v_0 \in V_0$ with $|v - v_0|^p < \epsilon^p / m(A)$ and let $Q_1, Q_2, Q_3, \ldots$ denote the sequence of dyadic cubes provided by Lemma 6.6 to cover A. Since $\sum_{k=1}^\infty m(Q_k) = m(A) < \infty$, we have

[12] This inequality is an easy algebraic exercise when $p \in \mathbb{N}$, but when p is not an integer, one can argue as follows. Assume $y > 0$ since otherwise the result is obvious. Dividing by y^p, it is then equivalent to prove $(1+x)^p \geqslant 1 + x^p$ for all $x \geqslant 0$ and $p \geqslant 1$. Differentiating with respect to x, it is easy to show that $(1+x)^p - 1 - x^p$ is an increasing function on $\{x \geqslant 0\}$ if $p \geqslant 1$, and since it vanishes at $x = 0$, it is therefore nonnegatve.

$\lim_{k\to\infty}\sum_{j=k}^{\infty} m(Q_j) = 0$, so the functions $f_k := \left(\sum_{j=1}^{k} \chi_{Q_j}\right) v_0$ satisfy

$$\begin{aligned}\|\chi_A v - f_k\|_{L^p}^p &= \sum_{j=1}^{k} |v - v_0|^p m(Q_j) + \sum_{j=k+1}^{\infty} |v|^p m(Q_j) \\ &\leqslant |v - v_0|^p \cdot m(A) + |v|^p \sum_{j=k+1}^{\infty} m(Q_j) \to |v - v_0|^p \cdot m(A) \quad \text{as} \quad k \to \infty,\end{aligned}$$

thus $\|\chi_A v - f_k\|_{L^p} < \epsilon$ for k sufficiently large. ☐

We next appeal to the fact that the Lebesgue measure m is **outer regular**, meaning that for every Lebesgue-measurable set $A \subset \mathbb{R}^n$,

$$m(A) = \inf\left\{ m(A') \;\middle|\; A \subset A' \subset \mathbb{R}^n,\ A' \text{ open}\right\}.$$

It follows that whenever $m(A) < \infty$, there exists a nested sequence of open sets $A_1 \supset A_2 \supset A_3 \supset \ldots \supset A' := \bigcap_{k\in\mathbb{N}} A_k \supset A$ such that $m(A) = m(A')$. The set A' is not generally open, but it is a Borel set, a so-called G_δ. In this situation, $|\chi_A - \chi_{A_n}|^p$ converges almost everywhere to 0, and since it is clearly also bounded by a fixed integrable function for every n, the dominated convergence theorem implies $\chi_{A_n} \to \chi_A$ in L^p. Since Lemma 6.7 provides arbitrarily good approximations $f_n \in Q(\mathbb{R}^n)$ for each $\chi_{A_n} v \in L^p(\mathbb{R}^n)$, we've proved:

LEMMA 6.8. *Lemma 6.7 remains true with A replaced by an arbitrary Lebesgue-measurable set in $\mathbb{R}^n$ with finite measure.* ☐

PROOF OF PROPOSITION 6.3 (AND THUS THEOREM 6.1). By Lemma 6.4, it suffices to prove that $Q(\mathbb{R}^n)$ is dense in $S(\mathbb{R}^n)$ in the L^p-norm. Elements of $S(\mathbb{R}^n)$ are of the form $\sum_{j=1}^{k} \chi_{A_j} v_j$, where each $A_j \subset \mathbb{R}^n$ is Lebesgue measurable with finite measure and $v_j \in V$. By Lemma 6.8, each $\chi_{A_j} v_j$ can be approximated arbitrarily well in the L^p-norm by functions in $Q(\mathbb{R}^n)$, so we are done. ☐

EXERCISE 6.9. The following shows that for almost any interesting measure space (X, μ), $L^\infty(X)$ is not separable:

(a) Show that if E is a Banach space containing an uncountable discrete subset, then E is not separable.

(b) Suppose (X, μ) is a measure space containing infinitely many disjoint subsets with positive measure. Show that $L^\infty(X)$ contains an uncountable subset $S \subset L^\infty(X)$, consisting of functions that take only the values 0 and 1, such that $\|f - g\|_{L^\infty} = 1$ for any two distinct $f, g \in S$.
Hint: If you've forgotten or never seen the proof via Cantor's diagonal argument that $\mathbb{R}$ is uncountable, looking it up may help.

EXERCISE 6.10. Here is another nonseparable Banach space that sometimes arises naturally. Assume $\mathcal{H}$ is an infinite-dimensional separable Hilbert space, and let $\mathscr{L}(\mathcal{H})$ denote the Banach space of bounded linear operators $\mathcal{H} \to \mathcal{H}$. Use an orthonormal basis of $\mathcal{H}$ to find a continuous embedding of $L^\infty(X)$ into $\mathscr{L}(\mathcal{H})$ for a suitable measure space X, and deduce from this that $\mathscr{L}(\mathcal{H})$ cannot be separable.

7. Weak convergence

In finite dimensions, a sequence $x_k \in \mathbb{K}^n$ converges to $x_\infty \in \mathbb{K}^n$ if and only if the n sequences formed by the coordinates of these vectors all converge to the corresponding coordinates of x_∞.

Writing $e_1, \dots, e_n \in \mathbb{K}^n$ for the standard orthonormal basis, the latter condition can be expressed equivalently as

$$\lim_{k\to\infty} \langle e_j, x_k\rangle = \langle e_j, x_\infty\rangle \quad \text{for all} \quad j = 1, \dots, n.$$

There is an obvious way to generalize this condition for a sequence x_k in an infinite-dimensional Hilbert space $\mathcal{H}$, though the resulting notion of convergence turns out to depend on a choice of orthonormal basis (see Exercise 7.2 below). A stronger condition that is clearly independent of any choice of basis is

$$\lim_{k\to\infty} \langle v, x_k\rangle = \langle v, x_\infty\rangle \quad \text{for all} \quad v \in \mathcal{H}.$$

In light of the Riesz representation theorem, this can be expressed equivalently as:

$$\lim_{k\to\infty} \Lambda(x_k) = \Lambda(x_\infty) \quad \text{for all} \quad \Lambda \in \mathcal{H}^*.$$

In this form, the condition also makes sense in arbitrary normed vector spaces, leading to the following important definition.

Definition 7.1. A sequence x_n in a normed vector space E is said to **converge weakly** to $x \in E$, written

$$x_n \rightharpoonup x,$$

if $\Lambda(x_n) \to \Lambda(x)$ for all $\Lambda \in E^*$.

With this definition in mind, the usual notion of convergence in a normed vector space (written "$x_n \to x$") is sometimes also called **strong convergence**. If $\dim E < \infty$, then it is easy to check that there is no difference between weak and strong convergence. In infinite-dimensional spaces, strong convergence clearly implies weak convergence due to the continuity of the functionals $\Lambda \in E^*$, but the following exercise shows that the converse is false.

Exercise 7.2. Suppose $\mathcal{H}$ is a Hilbert space containing an infinite orthonormal set $\{e_n \in \mathcal{H}\}_{n=1}^\infty$. Prove:

(a) The sequence e_n converges weakly to 0 but has no strongly convergent subsequence.
(b) For any bounded sequence $\lambda_n \in \mathbb{K}$, the sequence $x_n := \lambda_n e_n \in \mathcal{H}$ converges weakly to 0.
(c) For any unbounded sequence $\lambda_n \in \mathbb{K}$, $x_n := \lambda_n e_n \in \mathcal{H}$ satisfies $\lim_{n\to\infty}\langle e_j, x_n\rangle = 0$ for every $j \in \mathbb{N}$, but is nonetheless not weakly convergent.
Hint: Given a subsequence λ_{n_j} with $|\lambda_{n_j}| \geqslant j$ for $j = 1, 2, 3, \dots$, find a convergent series of the form $v := \sum_{j=1}^\infty a_j e_{n_j} \in \mathcal{H}$ for suitable scalars $a_j \in \mathbb{K}$ such that $\langle v, x_{n_j}\rangle \nrightarrow 0$ as $j \to \infty$.

Whenever we discuss a notion of convergence, there should be a topology in the background. Every normed vector space E comes with a natural topology, usually called the **norm topology** (sometimes also the **strong topology**), for which a set is open if and only if it is a union of open balls. The **weak topology** on E is generally different: it is the locally convex topology defined via the uncountably infinite family of seminorms

$$\{\|\cdot\|_\Lambda : E \to [0,\infty)\}_{\Lambda\in E^*}, \quad \text{where} \quad \|x\|_\Lambda := |\Lambda(x)|.$$

Notice that these are not norms since $\Lambda(x) = 0$ does not imply $x = 0$, but they are seminorms due to the linearity of Λ. The weak topology on E is thus the topology generated by all subsets of the form $\{x \in E \mid |\Lambda(x) - \Lambda(x_0)| < \epsilon\}$ for $x_0 \in E$, $\epsilon > 0$ and $\Lambda \in E^*$, and a sequence $x_n \in E$ converges to $x \in E$ in the weak topology if and only if it converges in all the seminorms, which means precisely that $\Lambda(x_n) \to \Lambda(x)$ for all $\Lambda \in E^*$, i.e. $x_n \rightharpoonup x$. A subset $\mathcal{U} \subset E$ that belongs to the weak topology is sometimes called **weakly open**. We will see below (see Remark 7.5) that all weakly open sets are also open in the usual sense, but the converse is generally false.

Remark 7.3. On a locally convex space E with topology generated by a family of seminorms $\{\|\cdot\|_\alpha : E \to [0,\infty)\}_{\alpha\in I}$, it is conventional to require that no nonzero $x \in E$ can satisfy $\|x\|_\alpha = 0$ for every $\alpha \in I$. This guarantees that a convergent sequence in E can only have one limit, and is equivalent to the condition that the topology defined by the seminorms on E is Hausdorff. The weak topology does satisfy this condition on every normed vector space, but this fact is not always obvious: it depends on the knowledge that for every nonzero $x \in E$ there exists a dual vector $\Lambda \in E^*$ with $\Lambda(x) \neq 0$. In all of the explicit examples that we deal with, it will be clear that this is true, e.g. Lemma 5.1 guarantees it for the L^p-spaces. For arbitrary Banach spaces, it follows from the Hahn-Banach theorem (see Lemma 5.6).

Exercise 7.4. This exercise gives an alternative characterization of the weak topology on a normed vector space E as the smallest topology for which the map $E \to \mathbb{K}$ defined by every dual vector $\Lambda \in E^*$ is continuous. In other words, the weak topology contains exactly the sets that must be considered open in order for these maps to be called continuous, but no more.

(a) Show that for every $\Lambda \in E^*$, the map $\Lambda : E \to \mathbb{K}$ is continuous in the weak topology.

Continuity of the maps $\Lambda : E \to \mathbb{K}$ means that for every $\Lambda \in E^*$ and every open set $\mathcal{U} \subset \mathbb{K}$, the set $\Lambda^{-1}(\mathcal{U}) \subset E$ needs to be open. Let $\mathcal{T}$ denote the smallest topology on E that contains all sets of this form, which means that a set is in $\mathcal{T}$ if and only if it is a union of finite intersections of sets of the form $\Lambda^{-1}(\mathcal{U})$ for arbitrary dual vectors $\Lambda \in E^*$ and open sets $\mathcal{U} \subset \mathbb{K}$. Part (a) shows that the weak topology contains $\mathcal{T}$. We now aim to show that these two topologies are the same.

(b) Show that for every $y \in E$, the translation map $\tau_y : E \to E : x \mapsto x + y$ is continuous with respect to the topology $\mathcal{T}$.
(c) Show that for every $\Lambda \in E^*$, $x_0 \in E$ and $\epsilon > 0$, the set $\{x \in E \mid |\Lambda(x - x_0)| < \epsilon\}$ is in $\mathcal{T}$, and conclude that $\mathcal{T}$ contains the weak topology.

Remark 7.5. Since every bounded linear functional $\Lambda : E \to \mathbb{K}$ is continuous in the norm topology on E, it follows from Exercise 7.4 that the norm topology contains the weak topology, i.e. every weakly open set is also open with respect to the norm. In general, however, the norm topology is strictly larger, e.g. if E is an infinite-dimensional Hilbert space, then Exercise 7.2 exhibits a sequence $x_n \in E$ that converges to 0 in the weak topology but not in the norm topology—the reason being that the norm topology has too many open neighborhoods of 0 for x_n to lie in all of them for n large. Relatedly, the fact that $\|x_n\| = 1$ for all n but $x_n \rightharpoonup 0$ in that exercise demonstrates that the norm $\|\cdot\| : E \to [0,\infty)$ is not a continuous function in the weak topology, though it is of course continuous in the norm topology.

Exercise 7.6. In the setting of Exercise 7.2, show that every neighborhood of $0 \in \mathcal{H}$ in the weak topology contains infinitely many of the vectors $x_n := \sqrt{n}e_n$ for $n \in \mathbb{N}$. In particular, the closure of the set $\{e_1, \sqrt{2}e_2, \sqrt{3}e_3, \dots\} \subset \mathcal{H}$ contains 0.
Remark: In a topological space, a set is closed if and only if its complement is open, and the closure of a set is by definition the intersection of all closed sets containing that set. Exercise 7.2 shows that the sequence $\sqrt{n}e_n$ has no subsequence weakly convergent to 0, so the present exercise demonstrates that the notion of the "closure" of a discrete set in the weak topology does not match your intuition from the theory of metric spaces—this shows in fact that the weak topology on $\mathcal{H}$ is not metrizable.

Combining Definition 7.1 with the Riesz representation theorem leads naturally to the following notion:

Definition 7.7. For a measure space (X,μ) and $1 \leqslant p < \infty$ such that either X is σ-finite or $p > 1$, we say that a sequence $f_n \in L^p(X)$ is **weakly L^p-convergent** to a function $f \in L^p(X)$ and write $f_n \xrightarrow{L^p} f$ if for every $g \in L^q(X)$ with $\frac{1}{p} + \frac{1}{q} = 1$,

$$\int_X \langle g, f_n \rangle \, d\mu \to \int_X \langle g, f \rangle \, d\mu.$$

For $p = \infty$, the notion of convergence in Definition 7.7 still makes sense but cannot be called "weak convergence" since the dual space of $L^\infty(X)$ is generally larger than $L^1(X)$. But we could instead view $L^\infty(X)$ as the dual space of $L^1(X)$ and fit this notion into the following context. For a normed vector space E, there is a natural topology on its dual space E^* that is generally even weaker[13] than the weak topology. The **weak* topology** on E^* is, namely, the locally convex topology defined via the family of seminorms

$$\{\|\cdot\|_x : E^* \to [0, \infty)\}_{x \in E}, \qquad \text{where} \qquad \|\Lambda\|_x := |\Lambda(x)|.$$

In light of the natural inclusion $E \hookrightarrow E^{**}$, this family of seminorms is a subset of the family that defines the weak topology, though the two families are exactly the same whenever E is a reflexive Banach space, that is:

PROPOSITION 7.8. *If E is a reflexive Banach space, then the weak and weak* topologies on E^* are identical.* □

We observe that since the space $L^p(X)$ for $1 < p < \infty$ is reflexive and can be identified with the dual space of $L^q(X)$ for $\frac{1}{p} + \frac{1}{q} = 1$, $L^p(X)$ has a natural weak* topology which is the same as its weak topology. On the other hand, the analogue of Definition 7.7 for $p = \infty$ describes convergence in the weak* topology on $L^\infty(X)$, which is strictly weaker than the weak topology, due to the fact that the dual of $L^\infty(X)$ is strictly larger than $L^1(X)$.

In analogy with Exercise 7.4, one can show that the weak* topology is the smallest topology such that for every $x \in E$, the function $E^* \to \mathbb{K} : \Lambda \mapsto \Lambda(x)$ is continuous. A sequence $\Lambda_n \in E^*$ is **weak* convergent** if and only if for every $x \in E$, $\Lambda_n(x) \to \Lambda(x)$, i.e. the functionals $\Lambda_n : E \to \mathbb{K}$ converge *pointwise* to $\Lambda : E \to \mathbb{K}$. Notice that for every nonzero $\Lambda \in E^*$, there necessarily exists a vector $x \in E$ for which $\|\Lambda\|_x \neq 0$, thus limits of weak* convergent sequences are unique and the weak* topology is Hausdorff (cf. Remark 7.3). This provides an easy proof (without requiring the Hahn-Banach theorem) that the weak topology on E^* is also Hausdorff, since every weak* open subset of E^* is also weakly open; or in terms of convergence, every weakly convergent sequence also converges in the weak* topology.

REMARK 7.9. The definitions above do not require E to be complete, but there is a subtlety to be aware of when considering normed vector spaces that are not Banach spaces. If E is a Banach space and $F \subset E$ is a proper dense subspace, then $F^* = E^*$ since every bounded linear functional on F extends uniquely to one on E. The norms on F^* and E^* are also the same, so as Banach spaces they are identical, but their weak* topologies may nonetheless be different. In practice, we will only consider examples in which E is complete, in which case the reader may feel free to ignore this remark.

The next result demonstrates that the weak* topology is often, indeed, *much* weaker than the norm topology on E^*. Having fewer open sets means that sequences can more easily converge, so they are more likely to have convergent subsequences.

[13]When comparing two topologies $\mathcal{T}_1$ and $\mathcal{T}_2$ on the same set, one says that $\mathcal{T}_1$ is **weaker** than $\mathcal{T}_2$ if $\mathcal{T}_1 \subset \mathcal{T}_2$. In this context, "weaker" is a synonym for "smaller," and the word **coarser** is also sometimes used with the same meaning, while in the other direction, one says that $\mathcal{T}_2$ is **stronger** / **finer** / **larger** than $\mathcal{T}_1$. Weakening a topology makes it easier for sequences to converge, i.e. every $\mathcal{T}_2$-convergent sequence is also $\mathcal{T}_1$-convergent, but there may also be $\mathcal{T}_1$-convergent sequences that are not $\mathcal{T}_2$-convergent. Similarly, weakening the topology makes it easier for maps from other spaces into X to be continuous, but harder for functions defined on X to be continuous.

THEOREM 7.10 (Banach-Alaoglu theorem, separable case). *Assume E is a separable normed vector space. Then every bounded sequence in E^* has a weak* convergent subsequence.*

PROOF. Fix a sequence $\Lambda_n \in E^*$ satisfying $\|\Lambda_n\| \leqslant C$ for some constant $C > 0$.

Claim 1: If $F \subset E$ is a countable subset, then after replacing Λ_n with a suitable subsequence, we can assume $\Lambda_n(x)$ converges for every $x \in F$. We prove this via the Cantor diagonal argument. Let $F = \{x_1, x_2, x_3, \dots\}$, and observe that for each $k, n \in \mathbb{N}$, $|\Lambda_n(x_k)| \leqslant C\|x_k\|$, thus for every fixed $k \in \mathbb{N}$ the sequence $\{\Lambda_n(x_k)\}_{n=1}^{\infty}$ is bounded in $\mathbb{K}$. Let $\Lambda_n^{(1)}$ denote a subsequence of Λ_n such that the sequence $\Lambda_n^{(1)}(x_1)$ converges in $\mathbb{K}$. Then choose $\Lambda_n^{(2)}$ to be a subsequence of $\Lambda_n^{(1)}$ such that the sequence $\Lambda_n^{(2)}(x_2)$ also converges in $\mathbb{K}$. Continuing in this way, we obtain a sequence of sequences such that the diagonal subsequence $\Lambda_n^{(n)}$ has the desired property.

Claim 2: If $F \subset E$ is a dense subset such that $\Lambda_n(x)$ converges for every $x \in F$, then $\Lambda_n(x)$ also converges for every $x \in E$. Indeed, for any given $x \in E$, one can choose $x' \in F$ arbitrarily close to x and then estimate

$$\begin{aligned} |\Lambda_m(x) - \Lambda_n(x)| &\leqslant |\Lambda_m(x) - \Lambda_m(x')| + |\Lambda_m(x') - \Lambda_n(x')| + |\Lambda_n(x') - \Lambda_n(x)| \\ &\leqslant 2C\|x - x'\| + |\Lambda_m(x') - \Lambda_n(x')|. \end{aligned}$$

Since $\Lambda_n(x')$ is a Cauchy sequence in $\mathbb{K}$, this shows that $\Lambda_n(x)$ is also a Cauchy sequence.

Finally, since E is separable, we are free to assume the two subsets denoted by $F \subset E$ in claims 1 and 2 are the same set, so both claims together allow us to replace Λ_n with a subsequence such that $\Lambda_n(x)$ is convergent for every $x \in E$. Define $\Lambda : E \to \mathbb{K}$ by

$$\Lambda(x) := \lim_{n \to \infty} \Lambda_n(x).$$

It is easy to check that Λ is linear and satisfies $|\Lambda(x)| \leqslant C\|x\|$, thus $\Lambda \in E^*$ and Λ_n is weak* convergent to Λ. □

Since $L^p(\Omega)$ is separable and reflexive for $1 < p < \infty$ and $\Omega \subset \mathbb{R}^n$, this implies:

COROLLARY 7.11. *Assume $\Omega \subset \mathbb{R}^n$ is a Lebesgue-measurable subset and $1 < p < \infty$. Then every L^p-bounded sequence $f_k \in L^p(\Omega)$ has a weakly L^p-convergent subsequence.* □

EXERCISE 7.12. Find a sequence $f_n \in L^p(\mathbb{R})$ for $1 < p < \infty$ that converges weakly to 0 but satisfies $\|f_n\|_{L^p} = 1$ for all n, and deduce that f_n has no L^p-convergent subsequence.

REMARK 7.13. $L^\infty(\Omega)$ is also the dual space of a separable Banach space, namely $L^1(\Omega)$, so Theorem 7.10 implies that L^∞-bounded sequences have weak* convergent subsequences. This case was not included in Corollary 7.11 since the weak and weak* topologies on $L^\infty(\Omega)$ are not the same.

EXAMPLE 7.14. There are two troubles with the case $p = 1$ in Corollary 7.11, one more serious than the other. The less serious problem is that $L^1(\Omega)$ is not the dual space of $L^\infty(\Omega)$, though since it is *contained* in the dual of $L^\infty(\Omega)$, one could still deduce from Theorem 7.10 a result about weakly L^1-convergent subsequences if $L^\infty(\Omega)$ were separable. The lack of separability is the more serious problem, and the following example shows that it cannot be overcome. For $n \in \mathbb{N}$, define $f_n \in L^1(\mathbb{R})$ to be the characteristic function of the interval $[n-1, n]$, so clearly $\|f_n\|_{L^1} = 1$ for every n. Consider an arbitrary subsequence f_{n_k} for some $1 \leqslant n_1 < n_2 < n_3 < \dots$, and define a function $g \in L^\infty(\mathbb{R})$ such that $g = (-1)^k$ on $[n_k - 1, n_k]$ for each $k \in \mathbb{N}$ and $g = 0$ everywhere else. Then the sequence $\int_{-\infty}^{\infty} g(x) f_{n_k}(x)\, dx = (-1)^k$ does not converge, thus f_{n_k} cannot be weakly convergent. The problem here is in essence that $L^\infty(\mathbb{R})$ is just too large a space, and as a consequence, weak L^1-convergence is harder to achieve than in the case $p > 1$.

The Banach-Alaoglu theorem implies that even though the unit sphere in $L^p(X)$ for $1 < p < \infty$ is not compact, it is *weakly* compact: arbitrary sequences with unit norm need not have accumulation points with respect to the L^p-norm, but in the weak topology they do. You may be wondering *which* points can arise as accumulation points in this scenario, e.g. must they also lie in the unit sphere? Let us show that they are at least bounded:

PROPOSITION 7.15. *In any normed vector space $(E, \|\cdot\|)$, if $x_n \rightharpoonup x$, then $\|x\| \leqslant \liminf_{n\to\infty} \|x_n\|$.*

PROOF. Using Lemma 5.6, choose $\Lambda \in E^*$ with $\|\Lambda\| = 1$ and $\Lambda(x) = \|x\|$. Then since $\Lambda(x_n) \to \Lambda(x)$ and $|\Lambda(x_n)| \leqslant \|x_n\|$,

$$\|x\| = \Lambda(x) = \liminf_{n\to\infty} \Lambda(x_n) \leqslant \liminf_{n\to\infty} \|x_n\|.$$

□

Recall from §5.3 that $L^p(X)$ is uniformly convex for $1 < p < \infty$. The next result therefore gives a useful criterion for strong L^p-convergence in terms of weak convergence. But in light of the Banach-Alaoglu theorem, it also says something that your geometric intuition may find shocking: the unit sphere is not closed in the weak topology. In particular, every sequence in the unit sphere that fails to have a strongly convergent subsequence has one that converges weakly to something in the *interior* of the unit sphere. It is known in fact that for any infinite-dimensional normed vector space E, the weak* closure of the unit sphere in E^* is the entire closed unit ball (cf. [**BS18**, Corollary 3.28]).

THEOREM 7.16. *If $(E, \|\cdot\|)$ is a uniformly convex Banach space and $x_n \in E$ is a sequence with $x_n \rightharpoonup x$ and $\|x_n\| \to \|x\|$, then $x_n \to x$.*

PROOF. We can assume $x \neq 0$ since the statement is otherwise trivial. Since the norms converge, we can also replace x_n and x with $x_n/\|x_n\|$ and $x/\|x\|$ respectively in order to assume $\|x_n\| = \|x\| = 1$ for all n without loss of generality. The weak convergence $x_n \rightharpoonup x$ implies $x_n + x \rightharpoonup 2x$, so combining Proposition 7.15 with the triangle inequality now gives

$$2 = \|2x\| \leqslant \liminf_{n\to\infty} \|x_n + x\| \leqslant \limsup_{n\to\infty} \|x_n + x\| \leqslant \limsup_{n\to\infty} (\|x_n\| + \|x\|) = 2,$$

and hence $\|x_n + x\| \to 2$, or equivalently, $\left\|\frac{x_n+x}{2}\right\| \to 1$. The conclusion $\|x_n - x\| \to 0$ then follows from uniform convexity. □

Just out of interest, let us state the more general version of the Banach-Alaoglu theorem, which does not require E to be separable. Its meaning is a bit harder to interpret, since the weak* topology is not generally first countable, so compactness need not imply sequential compactness.[14] We will neither prove nor make use of this version of the theorem, but proofs may be found e.g. in [**RS80**, §IV.5] or [**BS18**, §3.2]; it is a consequence of *Tychonoff's theorem* on the compactness of arbitrary products of compact topological spaces, which is equivalent to the axiom of choice (see [**Wen23**, §6]).

THEOREM 7.17 (Banach-Alaoglu theorem, general case). *For any normed vector space E, the closed unit ball in E^* is compact in the weak* topology.* □

[14] A topological space X is called **first countable** if for every $x \in X$, there is a countable sequence $\mathcal{U}_n \subset X$ of neighborhoods of x such that every neighborhood of x contains $\mathcal{U}_n$ for some $n \in \mathbb{N}$. First countability is a sufficient condition for the compactness of a subset to imply that all of its sequences have convergent subsequences (see e.g. [**Wen23**, §5]). It is easy to show that all metrizable topologies have this property, but scenarios like that of Exercise 7.6 reveal that the weak and weak* topologies generally do not.

8. Mollification

For this section, we consider functions defined on Lebesgue-measurable sets $\Omega \subset \mathbb{R}^n$ and define all integrals with respect to the Lebesgue measure m. For a Lebesgue-integrable function $f : \Omega \to V$, we write the integral as

$$\int_\Omega f\, dm :=: \int_\Omega f(x)\, dx :=: \int_\Omega f(x_1, \dots, x_n)\, dx_1 \dots dx_n.$$

We saw in §6 that the space $Q(\mathbb{R}^n)$ of functions that take constant values on finitely many dyadic cubes is dense in $L^p(\mathbb{R}^n)$ for every $p \in [1, \infty)$. It is not hard to convince oneself that every function in $Q(\mathbb{R}^n)$ can in turn be approximated arbitrarily well in the L^p-norm (again for $p < \infty$) by a compactly supported *continuous* function, thus proving that the space of continuous functions with compact support is dense in $L^p(\Omega)$.[15] We would now like to prove something more ambitious, and far more useful in applications:

THEOREM 8.1. *For every $p \in [1, \infty)$, $C^\infty(\mathbb{R}^n) \cap L^p(\mathbb{R}^n)$ is a dense subspace of $L^p(\mathbb{R}^n)$.*

Two important generalizations of Theorem 8.1 follow almost immediately. First: one can replace $\mathbb{R}^n$ by an arbitrary open subset $\Omega \subset \mathbb{R}^n$ and show that $C^\infty(\Omega) \cap L^p(\Omega)$ is dense in $L^p(\Omega)$. For the proof, one extends any given function $f \in L^p(\Omega)$ to $\tilde{f} \in L^p(\mathbb{R}^n)$ via

$$\tilde{f} := \begin{cases} f & \text{on } \Omega, \\ 0 & \text{on } \mathbb{R}^n \backslash \Omega, \end{cases}$$

and then approximates f with $f_\epsilon|_\Omega$ for smooth functions $f_\epsilon \in C^\infty(\mathbb{R}^n) \cap L^p(\mathbb{R}^n)$ that approximate $\tilde{f}$ in $L^p(\mathbb{R}^n)$. Further: $C^\infty(\Omega) \cap L^p(\Omega)$ in this statement can be replaced with

$$C_0^\infty(\Omega) := \left\{ f \in C^\infty(\Omega) \;\middle|\; f \text{ has compact support in } \Omega \right\}.$$

To see this, one first chooses for any given $f \in L^p(\Omega)$ and $\epsilon > 0$ an approximation $f_\epsilon \in C^\infty(\Omega) \cap L^p(\Omega)$ with $\|f - f_\epsilon\|_{L^p} < \frac{\epsilon}{2}$, and then replaces f_ϵ with βf_ϵ for a smooth compactly supported function $\beta : \Omega \to [0, 1]$ that satisfies $\beta|_{\mathcal{U}} \equiv 1$ for a sufficiently large open subset $\mathcal{U} \subset \Omega$. Taking a sequence of such cutoff functions β_N and subsets $\mathcal{U}_N$ such that $\bigcup_{N \in \mathbb{N}} \mathcal{U}_N = \Omega$, one can arrange that

$$\|f_\epsilon - \beta_N f_\epsilon\|_{L^p} < \frac{\epsilon}{2} \quad \text{and therefore} \quad \|f - \beta_N f_\epsilon\|_{L^p} < \epsilon$$

for $N \gg 0$ sufficiently large. For more details on this generalization, see e.g. [**LL01**, §2.19]; we summarize the result as follows:

COROLLARY 8.2. *For every $p \in [1, \infty)$ and every open subset $\Omega \subset \mathbb{R}^n$, $C_0^\infty(\Omega)$ is dense in $L^p(\Omega)$.* □

EXERCISE 8.3. Show that the space of bounded continuous functions is not dense in $L^\infty(\mathbb{R})$.

We prove Theorem 8.1 in the next several subsections using the *convolution*, a construction that is worth getting to know well, as it has a multitude of applications beyond this one theorem.

8.1. Continuity under translation. For $v \in \mathbb{R}^n$ and a function $f : \mathbb{R}^n \to V$, the **translation operator** τ_v produces a new function $\tau_v f : \mathbb{R}^n \to V$ defined by

$$(\tau_v f)(x) := f(x + v).$$

Clearly τ_v defines a bounded and norm-preserving linear map $L^p(\mathbb{R}^n) \to L^p(\mathbb{R}^n)$. Continuity of f is equivalent to the condition that for every convergent sequence $v_k \to v_\infty$ in $\mathbb{R}^n$, the functions $\tau_{v_k} f$ converge pointwise to $\tau_{v_\infty} f$. This is not true in general for functions $f \in L^p(\mathbb{R}^n)$ since they are not

[15] For a discussion of the density of $C^0(X)$ in $L^p(X)$ on more general measure spaces X, see [**Sal16**, §4.3].

generally continuous, but it will be useful to know that it becomes true if pointwise convergence is replaced by L^p-convergence:

THEOREM 8.4. *If* $1 \leqslant p < \infty$ *and* $f \in L^p(\mathbb{R}^n)$*, then the map* $\mathbb{R}^n \to L^p(\mathbb{R}^n) : v \mapsto \tau_v f$ *is continuous.*

PROOF. Since every τ_v defines a bounded linear operator $L^p(\mathbb{R}^n) \to L^p(\mathbb{R}^n)$ with $\|\tau_v\|_{\mathscr{L}(L^p)} = 1$, we have $\|\tau_{w+v} f - \tau_w f\|_{L^p} = \|\tau_w(\tau_v f - f)\|_{L^p} \leqslant \|\tau_v f - f\|_{L^p}$. It will thus suffice to prove that $\|\tau_v f - f\|_{L^p} \to 0$ as $v \to 0$ for all f belonging to some dense subset of $L^p(\mathbb{R}^n)$. Let $\widehat{Q}(\mathbb{R}^n)$ denote the space of all finite linear combinations $\sum_j \chi_{Q_j} f_j : \mathbb{R}^n \to V$, where $f_j \in V$ and each $Q_j \subset \mathbb{R}^n$ is a cube, i.e. any set of the form $[a_1, a_1 + d] \times \ldots \times [a_n, a_n + d]$ for $(a_1, \ldots, a_n) \in \mathbb{R}^n$ and $d > 0$. Then $\widehat{Q}(\mathbb{R}^n)$ contains the set $Q(\mathbb{R}^n)$ spanned by characteristic functions of dyadic cubes, and having proved in Proposition 6.3 that the latter is dense in $L^p(\mathbb{R}^n)$, it follows that $\widehat{Q}(\mathbb{R}^n)$ is also dense. For an individual cube $Q = [a_1 + d, \ldots, a_n + d]$, we have

$$\|\tau_v \chi_Q - \chi_Q\|_{L^p}^p = \int_{\mathbb{R}^n} |\tau_v \chi_Q - \chi_Q|^p \, dm = m\big((v + Q)\backslash Q\big) + m\big(Q\backslash(v + Q)\big) \to 0 \quad \text{as} \quad v \to 0,$$

thus for any $f = \sum_j \chi_{Q_j} f_j \in \widehat{Q}(\mathbb{R}^n)$, Minkowski's inequality gives

$$\|\tau_v f - f\|_{L^p} \leqslant \sum_j \big\|\tau_v \chi_{Q_j} - \chi_{Q_j}\big\|_{L^p} \cdot |f_j| \to 0 \quad \text{as} \quad v \to 0.$$

□

8.2. Convolution and regularity. The **convolution** of two scalar-valued functions $f, g : \mathbb{R}^n \to \mathbb{K}$ is a scalar-valued function $f * g$ defined by

$$(f * g)(x) := \int_{\mathbb{R}^n} f(x - y) g(y) \, dy. \tag{8.1}$$

More generally, one can also allow one of f or g to take values in the vector space V, so that $f * g$ also takes values in V; we will generally assume this in the following without further commentary. The domain of $f * g$ is the set of all points $x \in \mathbb{R}^n$ for which the integrand on the right hand side of (8.1) is a Lebesgue-integrable function of y. It may happen that $(f * g)(x)$ is defined for some but not all $x \in \mathbb{R}^n$. In practice, we will only consider situations in which $(f * g)(x)$ is defined for *almost every* x; the function $f * g$ is then defined almost everywhere on $\mathbb{R}^n$. Since $f * g$ is defined via an integral, it does not change if either f or g is changed on a set of measure zero; it can therefore make sense to speak of the convolution $f * g$ of two elements $f \in L^p(\mathbb{R}^n)$ and $g \in L^q(\mathbb{R}^n)$, and in such discussions we will typically not distinguish between actual functions and equivalence classes of functions defined almost everywhere.

REMARK 8.5. In many situations, it can also make sense to define $f * g$ on a suitable subset of $\mathbb{R}^n$ for two functions f and g that are not defined everywhere on $\mathbb{R}^n$. One case that often arises is when f is defined on some open subset $\Omega \subset \mathbb{R}^n$ and g is defined on $\mathbb{R}^n$ but has compact support in the r-ball $B_r \subset \mathbb{R}^n$ about the origin for some small $r > 0$. If x belongs to the set

$$\Omega_r := \left\{ x \in \Omega \;\middle|\; \operatorname{dist}(x, \mathbb{R}^n \backslash \Omega) \geqslant r \right\},$$

then either $x - y \in \Omega$ or $g(y) = 0$ holds for every $y \in \mathbb{R}^n$, thus one can make sense of the right hand side of (8.1) by interpreting the integrand to be 0 whenever $g(y) = 0$. The convolution $f * g$ is thus defined on all points of Ω_r for which this integrand (suitably interpreted) is integrable.

EXERCISE 8.6. Use a change of variables to prove $f * g = g * f$.

An important property of the convolution is that $f * g$ is in general at least as "nice" as the nicest function among f and g.[16] In particular, if either f or g is of class C^1, then Exercise 8.6 allows us to relabel the functions so that f is in C^1 without loss of generality, and we can then try to prove the formula

$$\partial_k(f * g)(x) = \frac{\partial}{\partial x_k}\int_{\mathbb{R}^n} f(x-y)g(y)\,dy = \int_{\mathbb{R}^n} \partial_k f(x-y)g(y)\,dy = (\partial_k f * g)(x).$$

This will be valid whenever f and g satisfy suitable conditions to apply Theorem 4.5 and justify differentiating under the integral sign. In practice, it is often easy to verify these conditions, and importantly, they do not require g to be differentiable, nor even continuous. For example:

THEOREM 8.7. *For any $f \in C_0^\infty(\mathbb{R}^n)$ and $g \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$, the function $f * g$ is smooth on $\mathbb{R}^n$, and for every multi-index α,*

$$\partial^\alpha(f * g) = (\partial^\alpha f) * g.$$

PROOF. By assumption f is smooth and vanishes outside of a compact subset $K \subset \mathbb{R}^n$, which implies that f is bounded. For every $x \in \mathbb{R}^n$, the integrand $y \mapsto f(x-y)g(y)$ can then only be nontrivial on the compact subset $K_x := \{x - k \in \mathbb{R}^n \mid k \in K\}$, and g is integrable on this domain, implying that the whole integrand is integrable on $\mathbb{R}^n$ and $(f * g)(x)$ is therefore defined for every $x \in \mathbb{R}^n$.

The function $x \mapsto (f * g)(x)$ is now defined as a parameter-dependent integral, where in the integrand only $f(x-y)$ depends on the parameter x. The result thus follows from Theorem 4.5, since:

- The integrand is Lebesgue integrable for every $x \in \mathbb{R}^n$;
- The integrability is also "locally uniform" in the sense that to every $x_0 \in \mathbb{R}^n$, one can associate a neighborhood $\mathcal{U} \subset \mathbb{R}^n$ of x_0 and an integrable function that bounds the integrand from above for every $x \in \mathcal{U}$.
- The function $x \mapsto f(x-y)g(y)$ is smooth for every $y \in \mathbb{R}^n$ and has partial derivative with respect to x_j given by $x \mapsto \partial_j f(x-y)g(y)$, which is again a continuous function of x.

Theorem 4.5 now implies $\partial_j(f * g) = (\partial_j f) * g$, and the generalization to arbitrary multi-indices follows by induction. ☐

8.3. Young's inequality. The following result is an elegant application of Fubini's theorem and Hölder's inequality.[17]

THEOREM 8.8. *For arbitrary functions $f \in L^1(\mathbb{R}^n)$ and $g \in L^p(\mathbb{R}^n)$ with $1 \leqslant p \leqslant \infty$, $f * g$ is defined almost everywhere on $\mathbb{R}^n$, belongs to $L^p(\mathbb{R}^n)$ and satisfies*

$$\|f * g\|_{L^p} \leqslant \|f\|_{L^1} \cdot \|g\|_{L^p}.$$

PROOF. The case $p = \infty$ is an easy exercise, so consider the case $1 \leqslant p < \infty$. Let $q \in [1, \infty]$ with $\frac{1}{p} + \frac{1}{q} = 1$; then

$$|f(x-y)g(y)| = |f(x-y)|^{1/p}|g(y)| \cdot |f(x-y)|^{1/q},$$

[16] The technical term for this notion of "niceness" is *regularity*, e.g. proving regularity of a function typically means proving that it is differentiable or smooth etc.

[17] For various more general forms of Young's inequality, see [**Sal16**, Theorem 7.33] or [**LL01**, §4.2].

and Hölder's inequality implies for every $x \in \mathbb{R}^n$,

$$\begin{aligned}\varphi(x) &:= \int_{\mathbb{R}^n} |f(x-y)g(y)|\, dy \\ &\leqslant \left(\int_{\mathbb{R}^n} |f(x-y)| \cdot |g(y)|^p \, dy\right)^{1/p} \cdot \left(\int_{\mathbb{R}^n} |f(x-y)|\, dy\right)^{1/q} \\ &\leqslant \|f\|_{L^1}^{1/q} \left(\int_{\mathbb{R}^n} |f(x-y)| \cdot |g(y)|^p \, dy\right)^{1/p}.\end{aligned}$$

Now apply Fubini's theorem for nonnegative measurable functions to

$$\mathbb{R}^n \times \mathbb{R}^n \to [0, \infty] : (x, y) \mapsto |f(x-y)| \cdot |g(y)|^p;$$

it follows that φ^p is a measurable function and

$$\begin{aligned}\|\varphi\|_{L^p}^p &= \int_{\mathbb{R}^n} [\varphi(x)]^p \, dx \leqslant \int_{\mathbb{R}^n} \|f\|_{L^1}^{p/q} \left(\int_{\mathbb{R}^n} |f(x-y)| \cdot |g(y)|^p \, dy\right) dx \\ &= \|f\|_{L^1}^{p/q} \int_{\mathbb{R}^n \times \mathbb{R}^n} |f(x-y)| \cdot |g(y)|^p \, dx\, dy \\ &= \|f\|_{L^1}^{p/q} \int_{\mathbb{R}^n} |g(y)|^p \left(\int_{\mathbb{R}^n} |f(x-y)|\, dx\right) dy = \|f\|_{L^1}^{p/q+1} \cdot \|g\|_{L^p}^p \\ &= \|f\|_{L^1}^p \cdot \|g\|_{L^p}^p < \infty.\end{aligned} \tag{8.2}$$

The function φ^p must therefore satisfy $\varphi^p < \infty$ almost everywhere, implying that $\varphi < \infty$ also holds almost everywhere, from which it follows that the convolution $f * g$ is defined almost everywhere.

As a further application of Fubini's theorem, one can show that $f * g$ is also a measurable function; in fact, the convolution of two Lebesgue-measurable functions is always Borel measurable. We'll skip the proof of this, though see [Sal16, Theorem 7.32(iii)]. Since $|f * g| \leqslant \varphi$, the estimate $\|f * g\|_{L^p} \leqslant \|f\|_{L^1} \cdot \|g\|_{L^p}$ now follows. □

Exercise 8.9. Prove as a corollary of Theorem 8.8 that the convolution defines a continuous bilinear operator

$$L^1(\mathbb{R}^n) \times L^p(\mathbb{R}^n) \to L^p(\mathbb{R}^n) : (f, g) \mapsto f * g.$$

8.4. Approximate identities. We can now prove Theorem 8.1, and in the process explain a useful general trick called **mollification**, by which non-smooth functions can be approximated by smooth ones. One of the motivating ideas in the background is that of the "Dirac δ-function,", a fictional function $\delta : \mathbb{R}^n \to \mathbb{R}$ that one imagines being defined by $\delta(x) = 0$ for $x \neq 0$ and $\delta(0) = \infty$ so that

$$\int_{\mathbb{R}^n} \varphi(x)\delta(x)\, dx = \varphi(0)$$

for all φ in some reasonable class of functions on $\mathbb{R}^n$. While δ cannot be defined as an actual function, it can easily be *approximated* by smooth functions—such an approximation is sometimes called a **mollifier**.

Definition 8.10. An **approximate identity** on $\mathbb{R}^n$ is a sequence of smooth functions $\rho_j : \mathbb{R}^n \to [0, \infty)$ such that for every smooth compactly supported function φ on $\mathbb{R}^n$,

$$\int_{\mathbb{R}^n} \varphi(x)\rho_j(x)\, dx \to \varphi(0) \quad \text{as} \quad j \to \infty.$$

The functions ρ_j in Definition 8.10 are not required to have compact support, and it will be important when we prove the Fourier inversion formula in §11.5 to be able to choose specific examples that are not compactly supported but have other nice properties. For applications involving the convolution, however, it is useful to impose the following stricter condition.

DEFINITION 8.11. A sequence of functions ρ_j on $\mathbb{R}^n$ will be said to have **shrinking support** if for every $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that the support of ρ_j is contained in the ϵ-ball about $0 \in \mathbb{R}^n$ for every $j \geqslant N$.

LEMMA 8.12. *A sequence of smooth functions $\rho_j : \mathbb{R}^n \to [0, \infty)$ with shrinking support is an approximate identity if and only if $\int_{\mathbb{R}^n} \rho_j \, dm \to 1$ as $j \to \infty$, and in this case, the condition in Definition 8.10 is also satisfied for all (not necessarily smooth or compactly supported) measurable functions φ on $\mathbb{R}^n$ that are continuous at the origin.*

PROOF. Assume $\operatorname{supp}(\rho_j)$ is contained in the ball $B_{r_j} \subset \mathbb{R}^n$ of radius $r_j > 0$ for some sequence $r_j \to 0$. If ρ_j is an approximate identity, then we can choose $N \in \mathbb{N}$ and a smooth compactly supported function $\varphi : \mathbb{R}^n \to [0, 1]$ that equals 1 on B_{r_j} for all $j \geqslant N$, and write

$$\int_{\mathbb{R}^n} \rho_j \, dm = \int_{B_{r_j}} \rho_j \, dm = \int_{B_{r_j}} \varphi \rho_j \, dm = \int_{\mathbb{R}^n} \varphi \rho_j \, dm \longrightarrow \varphi(0) = 1 \quad \text{as} \quad j \to \infty.$$

Conversely, if $\int_{\mathbb{R}^n} \rho_j \, dm \to 1$, then for any function φ on $\mathbb{R}^n$ that is continuous at 0,

$$\begin{aligned} \left| \varphi(0) - \int_{\mathbb{R}^n} \varphi \, \rho_j \, dm \right| &= \left| \varphi(0) \left(1 - \int_{\mathbb{R}^n} \rho_j \, dm \right) + \int_{\mathbb{R}^n} [\varphi(0) - \varphi(x)] \, \rho_j(x) \, dx \right| \\ &\leqslant |\varphi(0)| \cdot \left| 1 - \int_{\mathbb{R}^n} \rho_j \, dm \right| + \sup_{x \in B_{r_j}} |\varphi(0) - \varphi(x)| \int_{\mathbb{R}^n} \rho_j \, dm \to 0. \end{aligned}$$

$\square$

EXAMPLE 8.13. Choose a smooth function $\rho : \mathbb{R}^n \to [0, \infty)$ with compact support in the unit ball B_1 such that $\int_{\mathbb{R}^n} \rho \, dm = 1$. For $j \in \mathbb{N}$, the functions $\rho_j : \mathbb{R}^n \to [0, \infty)$ defined by $\rho_j(x) := j^n \rho(jx)$ then satisfy $\int_{\mathbb{R}^n} \rho_j \, dm = 1$ and have compact support in $B_{1/j}$ for all j, so this sequence forms an approximate identity with shrinking support.

THEOREM 8.14. *Fix an approximate identity ρ_j with shrinking support, and given $f \in L^p(\mathbb{R}^n)$ with $1 \leqslant p < \infty$, let $f_j := \rho_j * f = f * \rho_j$ for $j \in \mathbb{N}$, that is,*

$$f_j(x) := \int_{\mathbb{R}^n} f(x - y) \rho_j(y) \, dy. \tag{8.3}$$

Then:

(1) *f_j is a smooth function on $\mathbb{R}^n$ for every $j \in \mathbb{N}$.*
(2) *$\|f_j\|_{L^p} \leqslant C \|f\|_{L^p}$ for every $j \in \mathbb{N}$ and a constant $C > 0$, which may be assumed arbitrarily close to 1 for sufficiently large j.*
(3) *f_j converges in $L^p(\mathbb{R}^n)$ to f as $j \to \infty$.*

REMARK 8.15. The formula (8.3) can be interpreted as defining $f_j(x)$ to be a *weighted average* of the values of f in a neighborhood of x, where the size of the neighborhood becomes arbitrarily small as j becomes large. The latter follows from the assumption that ρ_j has shrinking support.

REMARK 8.16. The motivation for the term "approximate identity" is that if the δ-function existed as an actual function, it would satisfy $\delta * f = f * \delta = f$ for all reasonable functions f, making it an identity element in the algebra defined via the convolution product. We will see in §13 that this notion can be made rigorous by interpreting δ as a so-called *generalized function*, or *distribution*.

Proof of Theorem 8.14. The first two statements in the theorem follow from Theorems 8.7 and 8.8 since, by Lemma 8.12, $\|\rho_j\|_{L^1} = \int_{\mathbb{R}^n} \rho_j \, dm \to 1$. Let us write

$$\operatorname{supp}(\rho_j) \subset B_{r_j} \qquad \text{and} \qquad \left| \int_{\mathbb{R}^n} \rho_j \, dm - 1 \right| < \epsilon_j$$

for a pair of sequences $r_j, \epsilon_j > 0$ that converge to zero. For the third statement in the theorem, we first give a proof under the additional assumption that f is almost everywhere bounded and has compact support, i.e. assume there exists a constant $R > 0$ such that

$$\|f\|_{L^\infty} \leqslant R \quad \text{and} \quad f|_{\mathbb{R}^n \setminus B_R} \equiv 0. \tag{8.4}$$

Since $\|\rho_j\|_{L^1}$ is bounded, Young's inequality (Theorem 8.8) now implies that f_j satisfies a uniform L^∞-bound for all j, and since $\operatorname{supp}(\rho_j) \subset B_{r_j}$ with $r_j \to 0$, we can also assume for large j that f_j has compact support in B_{R+1}. It follows that f and f_j are in $L^1(\mathbb{R}^n)$, and we claim: $f_j \to f$ in $L^1(\mathbb{R}^n)$. To prove this, we use (8.3) and estimate

$$\begin{aligned} |f_j(x) - f(x)| &= \left| \int_{\mathbb{R}^n} [f(x-y) - f(x)] \, \rho_j(y) \, dy + f(x) \left(\int_{\mathbb{R}^n} \rho_j \, dm - 1 \right) \right| \\ &\leqslant \int_{\mathbb{R}^n} |f(x-y) - f(x)| \, \rho_j(y) \, dy + \epsilon_j |f(x)|, \end{aligned} \tag{8.5}$$

so by Fubini's theorem,

$$\begin{aligned} \|f_j - f\|_{L^1} &\leqslant \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^n} |f(x-y) - f(x)| \cdot \rho_j(y) \, dy \right) dx + \epsilon_j \|f\|_{L^1} \\ &= \int_{\mathbb{R}^n} \rho_j(y) \left(\int_{\mathbb{R}^n} |f(x-y) - f(x)| \, dx \right) dy + \epsilon_j \|f\|_{L^1} \\ &= \int_{B_{r_j}} \rho_j(y) \|\tau_{-y} f - f\|_{L^1} \, dy + \epsilon_j \|f\|_{L^1} \leqslant \sup_{y \in B_{r_j}} \|\tau_{-y} f - f\|_{L^1} \cdot \|\rho_j\|_{L^1} + \epsilon_j \|f\|_{L^1}. \end{aligned}$$

This goes to 0 as $j \to \infty$ since $\epsilon_j, r_j \to 0$ and (by Theorem 8.4), $y \mapsto \tau_y f$ is a continuous map $\mathbb{R}^n \to L^1(\mathbb{R}^n)$.

Having established $f_j \to f$ in L^1, we also know that f_j has a subsequence for which $|f_j - f|^p$ converges pointwise almost everywhere to 0, and $|f_j - f|^p$ is also uniformly bounded by a constant multiple of the characteristic function of B_{R+1}, which is integrable. The dominated convergence theorem then implies

$$\|f_j - f\|_{L^p}^p = \int_{\mathbb{R}^n} |f_j - f|^p \, dm \to \int_{\mathbb{R}^n} 0 \, dm = 0.$$

This conclusion applies at first to a subsequence, but if f_j were not convergent to f in $L^p(\mathbb{R}^n)$, then we could now find a subsequence that stays a positive distance away from f in the L^p-norm, and the L^1-convergence would then give a contradiction via the argument above, thus we have actually proved the convergence $f_j \xrightarrow{L^p} f$.

Without the additional conditions (8.4), one can instead argue as follows: for a given function $f \in L^p(\mathbb{R}^n)$ and a constant $R > 0$, define

$$f^R(x) := \begin{cases} f(x) & \text{if } x \in B_R \text{ and } |f(x)| \leqslant R, \\ 0 & \text{otherwise.} \end{cases}$$

It is not hard to show that $\|f - f^R\|_{L^p}$ can be made arbitrarily small by choosing $R > 0$ sufficiently large. Then f^R satisfies the conditions (8.4) and can therefore be approximated arbitrarily well in the L^p-norm by $f_j^R := \rho_j * f^R$. By Young's inequality,

$$\|f_j - f_j^R\|_{L^p} = \|\rho_j * (f - f^R)\|_{L^p} \leqslant \|\rho_j\|_{L^1} \cdot \|f - f^R\|_{L^p}$$

can then also be made arbitrarily small, thus $\|f - f_j\|_{L^p}$ becomes arbitrarily small for j sufficiently large. $\square$

While we are on this subject, we can prove a similar result on approximation of C^m-functions that will be useful when we talk about distributions in §13. The statement requires a slight expansion of the notion of C^m_{loc}-convergence defined in §4.3. Observe that if

$$\Omega_1 \subset \Omega_2 \subset \ldots \subset \bigcup_{j \in \mathbb{N}} \Omega_j = \Omega \subset \mathbb{R}^n$$

is a nested sequence of open subsets in $\mathbb{R}^n$, then every compact set $K \subset \Omega$ belongs to Ω_j for $j \in \mathbb{N}$ sufficiently large. A sequence of C^m-functions $f_j : \Omega_j \to V$ is said to be **convergent in** $C^m_{\text{loc}}(\Omega)$ to a function $f : \Omega \to V$ if for every compact subset $K \subset \Omega$ and $N \in \mathbb{N}$ such that $K \subset \Omega_N$, the sequence of functions $f_N, f_{N+1}, f_{N+2}, \ldots$ restricted to K is C^m-convergent to $f|_K$. The only difference between this and the definition in §4.3 is that the limit function f may be defined on a strictly larger domain than any function in the sequence.

THEOREM 8.17. *Suppose $\Omega \subset \mathbb{R}^n$ is an open subset, $f \in C^m(\Omega)$ for some integer $m \geqslant 0$, and $\rho_j : \mathbb{R}^n \to [0, \infty)$ for $j \in \mathbb{N}$ is an approximate identity with shrinking support. Then there exists a nested sequence of open subsets $\Omega_1 \subset \Omega_2 \subset \ldots \subset \bigcup_{j\in\mathbb{N}} \Omega_j = \Omega$ such that for each $j \in \mathbb{N}$, $f_j := \rho_j * f$ is defined (in the sense of Remark 8.5) and smooth on Ω_j, and the sequence f_j converges to f in $C^m_{\text{loc}}(\Omega)$.*

PROOF. Assume $\text{supp}(\rho_j) \subset B_{r_j}$ with $r_j \to 0$, and define

$$\Omega_j := \left\{ x \in \Omega \;\middle|\; \text{dist}(x, \mathbb{R}^n \backslash \Omega) > 2r_j \right\}.$$

Then $f_j(x) = \int_{\mathbb{R}^n} \rho_j(x - y) f(y)\, dy$ can be defined for all $x \in \Omega_j$ since $y \in \Omega$ whenever $x - y \in \text{supp}(\rho_j)$. Smoothness follows by differentiating under the integral sign as in Theorem 8.7 to prove $\partial^\alpha f_j(x) = (\partial^\alpha \rho_j * f)(x)$ for all multi-indices α and $x \in \Omega_j$; here Theorem 4.5 is applicable because ρ_j is bounded and f is integrable on the region $B_{r_j}(x)$ where $\rho_j(x - \cdot)$ can be nonzero. To prove $f_j \to f$ in C^m_{loc}, suppose $K \subset \Omega$ is compact, and pick $N \in \mathbb{N}$ large enough so that $K \subset \Omega_N$ and the slightly larger compact set

$$K' := \left\{ x \in \mathbb{R}^n \;\middle|\; \text{dist}(x, K) \leqslant r_N \right\}$$

is also contained in Ω. Then for $x \in K$ and $j \geqslant N$, (8.5) gives

$$|f_j(x) - f(x)| \leqslant \sup_{y \in B_{r_j}} |f(x - y) - f(x)| \cdot \|\rho_j\|_{L^1} + \epsilon_j \|f\|_{C^0(K)}.$$

Since x and $x - y$ in this expression both belong to K' and f is uniformly continuous on K', this implies uniform convergence $f_j \to f$ on K. To prove the same for derivatives up to order m, we observe that for any multi-index α with $|\alpha| \leqslant m$, $x \in K'$ and j sufficiently large,

$$\partial^\alpha f(x) = \frac{\partial^{|\alpha|}}{\partial x^\alpha} \int_{B_{r_j}} f(x - y) \rho_j(y)\, dy = \int_{B_{r_j}} \partial^\alpha f(x - y) \rho_j(y)\, dy = (\partial^\alpha f * \rho_j)(x),$$

where Theorem 4.5 justifies differentiation under the integral sign since $\partial^\alpha f$ is well defined and bounded on $B_{r_j}(x)$ while ρ_j is integrable. The same argument that was used for f_j then implies uniform convergence $\partial^\alpha f_j \to \partial^\alpha f$ on K. $\square$

9. Absolute continuity

9.1. The fundamental theorem of calculus. Let us consider the following question.

QUESTION 9.1. *What is the largest class of functions f on a compact interval $[a, b] \subset \mathbb{R}$ such that the formula $f(x) = f(a) + \int_a^x f'(t)\,dt$ holds?*

Here we regard $\int_a^b f(t)\,dt$ as alternative notation for the Lebesgue integral $\int_{[a,b]} f\,dm$ if $a \leqslant b$, or $-\int_{[b,a]} f\,dm$ if $b \leqslant a$. The formula is easy to prove under the assumption that f is continuously differentiable, but we already know it is valid somewhat more generally than this, e.g. it clearly also holds if f is continuous and only piecewise C^1, and it is not hard to think up examples in which f is non-differentiable on a countably infinite subset but the formula still holds. In order for the right hand side to make sense at all, f only needs to be differentiable almost everywhere on $[a, b]$, and its (almost everywhere well-defined) derivative needs to be in $L^1([a, b])$. Is that enough? No:

EXAMPLE 9.2. The **Cantor function** is a continuous, surjective and monotone increasing function $f : [0, 1] \to [0, 1]$ whose derivative is well defined and vanishes on a subset of full measure, namely the complement of the Cantor ternary set $C \subset [0, 1]$. In particular, f is defined to be constant on each of the intervals that are removed in order to define C:

$$f|_{(1/3,2/3)} := \frac{1}{2},$$
$$f|_{(1/9,2/9)} := \frac{1}{4}, \qquad f|_{(7/9,8/9)} := \frac{3}{4},$$
$$f|_{(1/27,2/27)} := \frac{1}{8}, \quad f|_{(7/27,8/27)} := \frac{3}{8}, \quad f|_{(19/27,20/27)} := \frac{5}{8}, \quad f|_{(25/27,26/27)} := \frac{7}{8}$$

and so forth (see Figure 4). The easiest way to define f at all other points is as the uniform limit of a sequence of piecewise affine, continuous, increasing and surjective functions $f_n : [0, 1] \to [0, 1]$. Such a sequence is uniquely determined by the following conditions (Figure 5):

- $f_0(x) := x$;
- For each $n \in \mathbb{N}$, f_n takes the same constant values as f on each of the 2^{n-1} intervals of length $1/3^n$ that are removed in the definition of C, and has constant slope on all other subintervals of $[0, 1]$.

It is easy to check from this definition that $|f_n - f_{n-1}| \leqslant c/2^n$ for some constant $c > 0$ and all $n \in \mathbb{N}$, thus the sequence f_n is uniformly Cauchy and therefore converges to a continuous function f, which is automatically monotone and surjective.[18]

Since the Cantor function has values on the entire interval $[0, 1]$ in spite of its derivative vanishing almost everywhere, it clearly lacks whatever property is needed for the fundamental theorem of calculus to hold. Let us reformulate the question slightly: suppose $f \in L^1([a, b])$, and consider the function F defined on $[a, b]$ by

$$F(x) := \int_a^x f(t)\,dt.$$

[18] A more precise formula for f can be deduced from the fact that it is continuous and constant on a sequence of intervals whose union is dense. It is easiest to express in terms of base-3 and base-2 expansions: since all points $x \in C$ have unique base-3 expansions $0.a_1a_2a_3\ldots$ with $a_n \in \{0, 2\}$ for all $n = 1, 2, 3, \ldots$, one can write $f(x) \in [0, 1]$ so that its base-2 expansion is $0.b_1b_2b_3\ldots$ with $b_n := a_n/2$ for all n. In other words, $f\left(\sum_{n=1}^\infty \frac{2a_n}{3^n}\right) = \sum_{n=1}^\infty \frac{a_n}{2^n}$, assuming $a_n \in \{0, 1\}$ for all $n \in \mathbb{N}$.

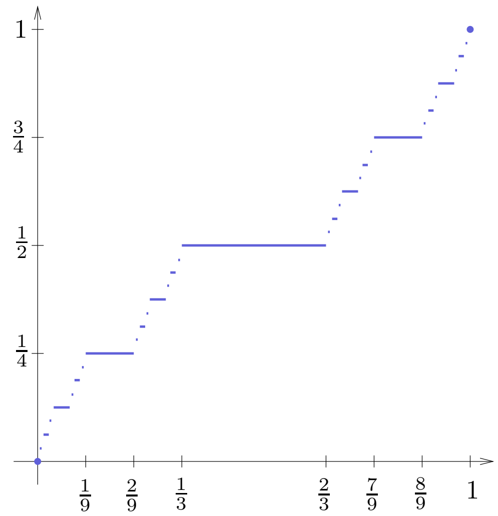

FIGURE 4. An imperfect picture of the Cantor function. Despite the appearance of jump discontinuities in the approximate graph drawn here, it is continuous.

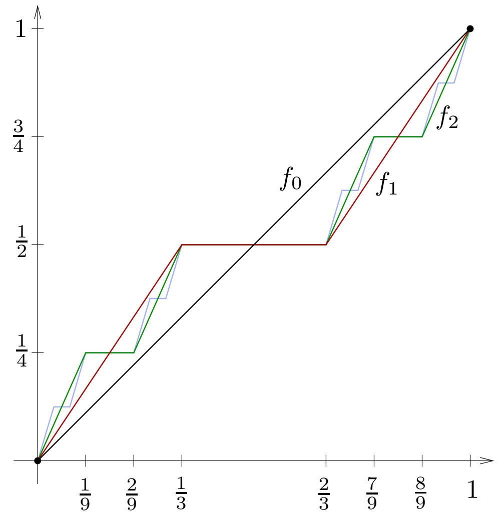

FIGURE 5. A sequence of piecewise affine functions converging uniformly to the Cantor function.

One of the main results of this section (Corollary 9.12 below) will show that F must be differentiable almost everywhere and its derivative is f. The Cantor function also has the first property, but since it is evidently not the integral of its derivative, we deduce that the Cantor function cannot be written as an integral of *any* Lebesgue-integrable function on $[0, 1]$. So, how do we tell the difference, i.e. what properties does the function F have that the Cantor function does not? Both are continuous, but it turns out that F satisfies a stronger condition than continuity.

LEMMA 9.3. *For any measure space (X, μ) and any $f \in L^1(X)$, given $\epsilon > 0$, there exists $\delta > 0$ such that for all measurable subsets $A \subset X$,*

$$\mu(A) < \delta \quad \Rightarrow \quad \int_A |f| \, d\mu < \epsilon.$$

PROOF. If the result is not true, then there exists a number $\epsilon > 0$ and a sequence of measurable sets $A_n \subset X$ such that

$$\mu(A_n) < \frac{1}{2^n} \quad \text{but} \quad \int_{A_n} |f| \, d\mu \geqslant \epsilon.$$

Define $B_n := \bigcup_{k=n}^{\infty} A_k$, so we have

$$B_1 \supset B_2 \supset B_3 \supset \ldots B := \bigcap_{n \in \mathbb{N}} B_n,$$

and $\mu(B_n) \leqslant \sum_{k=n}^{\infty} \mu(A_k) < \sum_{k=n}^{\infty} \frac{1}{2^k} = \frac{1}{2^{n-1}}$, thus $\mu(B) = \lim_{n\to\infty} \mu(B_n) = 0$. This implies $\lim_{n\to\infty} \int_{B_n} |f| \, d\mu = \int_B |f| \, d\mu = 0$, which is a contradiction since $B_n \supset A_n$ for every n and thus $\int_{B_n} |f| \, d\mu \geqslant \int_{A_n} |f| \, d\mu \geqslant \epsilon > 0$. □

Returning to the function $F(x) = \int_a^x f(t) \, dt$ with $f \in L^1([a, b])$, consider the consequences of Lemma 9.3 for subsets $A \subset [a, b]$ defined as finite unions of intervals $A = \bigcup_{j=1}^N [a_j, b_j]$ with $a \leqslant a_1 \leqslant b_1 \leqslant \ldots \leqslant a_N \leqslant b_N \leqslant b$. The lemma provides for every $\epsilon > 0$ a $\delta > 0$ such that whenever $m(A) = \sum_{j=1}^N (b_j - a_j) < \delta$, it follows that

$$\sum_{j=1}^N |F(b_j) - F(a_j)| = \sum_{j=1}^N \left| \int_{a_j}^{b_j} f(t) \, dt \right| \leqslant \sum_{j=1}^N \int_{[a_j, b_j]} |f| \, dm = \int_A |f| \, dm < \epsilon.$$

In other words, F satisfies the following condition:

DEFINITION 9.4. A function F on an interval $I \subset \mathbb{R}$ is called **absolutely continuous** if for every $\epsilon > 0$ there exists $\delta > 0$ such that for all finite sequences $a_1 \leqslant b_1 \leqslant \ldots \leqslant a_N \leqslant b_N$ of points in I,

$$\sum_{j=1}^N (b_j - a_j) < \delta \quad \Rightarrow \quad \sum_{j=1}^N |F(b_j) - F(a_j)| < \epsilon.$$

This definition would be the same as uniform continuity for functions on $I \subset \mathbb{R}$ if one only allowed $N = 1$, but the extension to all finite unions of intervals makes it a strictly stronger condition than uniform continuity. The Cantor function, for example, is uniformly continuous (as are all continuous functions on compact intervals), but the next exercise shows that it is not absolutely continuous:

EXERCISE 9.5. Show that if $F : [a, b] \to \mathbb{R}$ is absolutely continuous, then it maps every set of measure zero in $[a, b]$ to a set of measure zero in $\mathbb{R}$.

EXERCISE 9.6. Show that every Lipschitz continuous function on a compact interval $[a, b]$ is also absolutely continuous.

Here is the answer to Question 9.1:

THEOREM 9.7 (Fundamental theorem of calculus for the Lebesgue integral). *For a nontrivial compact interval $[a, b] \subset \mathbb{R}$ and a function f on $[a, b]$, the following conditions are equivalent:*

(1) f is absolutely continuous;
(2) f is differentiable almost everywhere, its derivative f' is in $L^1([a, b])$, and $f(x) = f(a) + \int_a^x f'(t) \, dt$ for all $x \in [a, b]$.

We have already proved the easy direction of this theorem, as a consequence of Lemma 9.3. We will show in Corollary 9.12 that for any $g \in L^1([a,b])$, the absolutely continuous function given by $F(x) = c + \int_a^x g(t)\,dt$ for a constant $c = F(a)$ is almost everywhere differentiable and its derivative is g. This statement is a consequence of the *Lebesgue differentiation theorem*, introduced in the next subsection. What then remains to be proved is that *every* absolutely continuous function f on $[a,b]$ can be written in the form $f(x) = f(a) + \int_a^x g(t)\,dt$ for some $g \in L^1([a,b])$. We will prove this in §9.5 as a consequence of a simple version of the *Radon-Nikodým theorem*, proved in §9.4.

Combining Exercise 9.6 with Theorem 9.7 produces a slightly surprising consequence:

COROLLARY 9.8. *Every Lipschitz continuous function on a compact interval $[a,b] \subset \mathbb{R}$ with $b > a$ is differentiable almost everywhere.* □

Corollary 9.8 also holds for functions on open domains in $\mathbb{R}^n$, and is known in that level of generality as *Rademacher's theorem*. For a concise proof built on top of the one-dimensional case, see [**Hei05**].

9.2. The Lebesgue differentiation theorem. Here is another natural question, which we will need to answer before we learn how to differentiate integrals of L^1-functions.

QUESTION 9.9. *For locally integrable functions f on $\mathbb{R}^n$, what relation is there between $f(x)$ and the "average" value of f on arbitrarily small balls about x?*

Let us denote by

$$B_r(x) \subset \mathbb{R}^n$$

the open ball of radius $r > 0$ about a point $x \in \mathbb{R}^n$.

DEFINITION 9.10. For a function $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$, $x \in \mathbb{R}^n$ is called a **Lebesgue point** of f if the average value of $|f - f(x)|$ on $B_r(x)$ converges to zero as $r \to 0$, i.e.

$$\lim_{r\to 0^+} \frac{1}{m(B_r(x))} \int_{B_r(x)} |f(y) - f(x)|\, dy = 0.$$

Whenever x is a Lebesgue point of f, one has

$$\lim_{r\to 0^+} \frac{1}{m(B_r(x))} \int_{B_r(x)} f(y)\, dy = f(x) \tag{9.1}$$

since

$$\left| \frac{1}{m(B_r(x))} \int_{B_r(x)} f\, dm - f(x) \right| = \left| \frac{1}{m(B_r(x))} \int_{B_r(x)} (f - f(x))\, dm \right|$$
$$\leqslant \frac{1}{m(B_r(x))} \int_{B_r(x)} |f - f(x)|\, dm.$$

Clearly x is a Lebesgue point whenever f is continuous at x, but Lebesgue-integrable functions can easily be discontinuous everywhere. Moreover, changing f on a set of measure zero changes the right hand side of (9.1) at some points but not the left hand side, so the most one could hope for in general is for (9.1) to be true for almost every x. That turns out to be true, and thus gives the best possible answer to Question 9.9:

THEOREM 9.11 (Lebesgue differentiation theorem). *For any $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$, almost every point of $\mathbb{R}^n$ is a Lebesgue point of f.*

To see why this is called a *differentiation* theorem, consider the case $n = 1$. If $f \in L^1([a,b])$, extend f to a function in $L^1(\mathbb{R})$ that vanishes outside $[a,b]$, and consider the function $F(x) := \int_a^x f(t)\,dt$. If $x \in (a,b)$ is a Lebesgue point of f, then for all $h > 0$ sufficiently small, we have

$$\begin{aligned}\left|\frac{F(x+h)-F(x)}{h} - f(x)\right| &= \left|\frac{1}{h}\int_x^{x+h} f(t)\,dt - f(x)\right| \leqslant \frac{1}{h}\int_x^{x+h} |f(t)-f(x)|\,dt \\ &\leqslant 2\frac{1}{m((x-h,x+h))}\int_{x-h}^{x+h} |f(t)-f(x)|\,dt,\end{aligned}$$

and the latter becomes arbitrarily small when $h > 0$ is small. A similar statement is proved in the same manner for $h < 0$ and shows that at every Lebesgue point x, $F'(x) = f(x)$.

COROLLARY 9.12. *For every $f \in L^1([a,b])$, the function $F(x) := \int_a^x f(t)\,dt$ is differentiable almost everywhere on (a,b) and satisfies $F' = f$.* □

The proof of Theorem 9.11 requires a result called the *Hardy-Littlewood maximal inequality*, which we will discuss in the next subsection. In order to see what is needed, let us set up the general framework of the proof first.

We begin with two easy observations:

(1) If f is a continuous function on $\mathbb{R}^n$, then every point in $\mathbb{R}^n$ is a Lebesgue point.
(2) If almost every point is a Lebesgue point for all $f \in L^1(\mathbb{R}^n)$, then the same holds for all $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$.

The second statement follows from the purely local nature of the Lebesgue point condition, i.e. it depends on f only in arbitrarily small neighborhoods of x. Then if we cut off the values of $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ outside the ball $B_k(0) \subset \mathbb{R}^n$ to produce a function in $L^1(\mathbb{R}^n)$ whose set of non-Lebesgue points in $B_k(0)$ we can prove has measure zero, it follows that the set of non-Lebesgue points of f will be the union of these sets for all $k \in \mathbb{N}$, and thus also has measure zero.

With this understood, let us associate to any $f \in L^1(\mathbb{R}^n)$ and $r \geqslant 0$ the functions $f^r : \mathbb{R}^n \to [0,\infty]$ defined by

$$f^r(x) := \frac{1}{m(B_r(x))}\int_{B_r(x)} |f - f(x)|\,dm \quad \text{for } r > 0, \qquad f^0(x) := \limsup_{r\to 0} f^r(x).$$

The goal is to prove that $f^0 = 0$ almost everywhere. For each $N \in \mathbb{N}$, let

$$A_N := \left\{x \in \mathbb{R}^n \;\middle|\; f^0(x) > 1/N\right\}.$$

We will deduce the desired result from:

LEMMA 9.13. *For every $N \in \mathbb{N}$, A_N is contained in a Lebesgue-measurable set of measure less than $\frac{1}{N}$.*

Indeed, if this lemma holds, then since $A_1 \subset A_2 \subset A_3 \subset \ldots$, it follows that every A_N is a set of measure zero. Their union therefore also has measure zero, and that is precisely the set on which $f^0 > 0$.

In order to estimate the measure of A_N, we appeal to the density of continuous functions in $L^1(\mathbb{R}^n)$ and choose a sequence $f_1, f_2, f_3, \ldots$ of continuous functions on $\mathbb{R}^n$ such that $f_k \to f$ in L^1.

We can then pick k large and use f_k to estimate $f^r(x)$ for $r > 0$ small:

$$\begin{aligned} f^r(x) &= \frac{1}{m(B_r(x))} \int_{B_r(x)} |f - f(x)|\, dm \\ &\leqslant \frac{1}{m(B_r(x))} \int_{B_r(x)} \left(|f - f_k| + |f_k - f_k(x)| + |f_k(x) - f(x)|\right) dm \\ &= \frac{1}{m(B_r(x))} \int_{B_r(x)} |f - f_k|\, dm + f_k^r(x) + |f_k(x) - f(x)|. \end{aligned} \tag{9.2}$$

We cannot assume $f_k \to f$ uniformly, so in this last expression, the third term might not become arbitrarily small for all x as $k \to 0$, but it is easy to show that it does so outside of a set of small measure. Indeed, we can associate to any given measurable function g on a measure space (X, μ) the sets $A_t := \{x \in X \mid |g(x)| > t\}$ for $t > 0$, and then estimate $\|g\|_{L^1} \geqslant \int_{A_t} |g|\, d\mu \geqslant \mu(A_t)t$. The result is known as **Chebyshev's inequality**:

$$\mu\big(\{x \in X \mid |g(x)| > t\} \big) \leqslant \frac{\|g\|_{L^1}}{t}. \tag{9.3}$$

Applying this to $f - f_k \in L^1(\mathbb{R}^n)$, we can arrange by choosing k sufficiently large to make $|f_k(x) - f(x)|$ arbitrarily small for all x outside of a set that has arbitrarily small measure. Having chosen k in this way, the second term in the last line of (9.2) also becomes arbitrarily small as $r \to 0$. However, estimating the first term requires some non-obvious input: we would like to claim that since $|f - f_k|$ has a small L^1-norm, its average value over $B_r(x)$ also satisfies some small bound as $r \to 0$. If we were first fixing $r > 0$ and then letting $k \to \infty$, it would be obvious that this term vanishes in the limit, but unfortunately, the order of quantifiers is the other way around: we have already fixed k and need to estimate the average for all small $r > 0$ in terms of $\|f - f_k\|_{L^1}$. This is what the Hardy-Littlewood maximal inequality is for.

9.3. Maximal functions and weak L^1. We now introduce a missing ingredient in the proof of Lemma 9.13.

DEFINITION 9.14. For $f \in L^1_{\text{loc}}(\mathbb{R}^n)$, the **maximal function** $Mf : \mathbb{R}^n \to [0, \infty]$ is defined by

$$Mf(x) := \sup_{r>0} \frac{1}{m(B_r(x))} \int_{B_r(x)} |f|\, dm.$$

LEMMA 9.15. *For every $f \in L^1_{\text{loc}}(\mathbb{R}^n)$, the maximal function $Mf : \mathbb{R}^n \to [0, \infty]$ is Borel measurable.*

PROOF. It suffices to show that $(Mf)^{-1}((t, \infty))$ is a Borel set for every $t \in \mathbb{R}$; we will show in fact that it is open, i.e. Mf is lower semicontinuous. The condition $Mf(x) > t$ implies that for some $r > 0$, $\frac{1}{m(B_r(x))} \int_{B_r(x)} |f|\, dm > t$. This remains true after replacing $m(B_r(x))$ with $m(B_{r'}(x))$ for some slightly larger $r' > r$, and since $B_r(x) \subset B_{r'}(x')$ for all $x' \in \mathbb{R}^n$ sufficiently close to x, we then have

$$t < \frac{1}{m(B_{r'}(x))} \int_{B_r(x)} |f|\, dm \leqslant \frac{1}{m(B_{r'}(x))} \int_{B_{r'}(x)} |f|\, dm \leqslant Mf(x').$$

□

It is clearly not possible to achieve a general pointwise bound on Mf for $f \in L^1(\mathbb{R}^n)$, e.g. if f is defined as $1/\sqrt{|x|}$ on $[-1, 1]$ and 0 on the rest of $\mathbb{R}$, then $f \in L^1(\mathbb{R})$ but its average values on $[-r, r]$ diverge to ∞ as $r \to 0$, giving $Mf(0) = \infty$. A realistic hope, however, would be to bound the measure of sets on which Mf exceeds any given value. If $Mf \in L^1(\mathbb{R}^n)$, then such a bound follows from the Chebyshev inequality (9.3). In general, it would be too much to hope for Mf to be globally integrable, but there also exist functions that are not in $L^1(\mathbb{R}^n)$ and nonetheless satisfy a bound of the form (9.3), with the L^1-norm replaced by some other constant. A simple example is $f(x) := 1/x$, which is not in $L^1(\mathbb{R})$ but satisfies $m\big(\{x \in X \mid |f(x)| > t\}\big) = 2/t$, thus it belongs to the following class of functions.

DEFINITION 9.16. A measurable function f on (X, μ) is called **weakly integrable** if there exists a constant $C > 0$ such that

$$\mu\big(\{x \in X \mid |f(x)| > t\}\big) \leqslant \frac{C}{t} \quad \text{for all} \quad t > 0.$$

We will denote the space of such functions by $L^1_{\text{weak}}(X)$.

One can define a "norm" on $L^1_{\text{weak}}(X)$ by

$$\|f\|_{L^1_{\text{weak}}} := \sup_{t>0} t\mu\big(\{x \in X \mid |f(x)| > t\}\big).$$

Just one caveat: $\|\cdot\|_{L^1_{\text{weak}}}$ satisfies $\|cf\|_{L^1_{\text{weak}}} = |c| \cdot \|f\|_{L^1_{\text{weak}}}$ for $c \in \mathbb{K}$ and $\|f\|_{L^1_{\text{weak}}} = 0$ if and only if f vanishes almost everywhere, but it does not satisfy the triangle inequality. Instead it satisfies (see [Sal16, Lemma 6.2])

$$\|f + g\|_{L^1_{\text{weak}}} \leqslant \frac{\|f\|_{L^1_{\text{weak}}}}{\lambda} + \frac{\|g\|_{L^1_{\text{weak}}}}{1 - \lambda} \quad \text{for} \quad 0 < \lambda < 1,$$

and

$$\sqrt{\|f + g\|_{L^1_{\text{weak}}}} \leqslant \sqrt{\|f\|_{L^1_{\text{weak}}}} + \sqrt{\|g\|_{L^1_{\text{weak}}}}.$$

As a consequence, $L^1_{\text{weak}}(X)$ is not a normed vector space, but one can regard it as a topological vector space with respect to the metric $\operatorname{dist}(f, g) := \sqrt{\|f - g\|_{L^1_{\text{weak}}}}$. The inequality (9.3) can now be interpreted as saying that there is a natural continuous inclusion $L^1(X) \hookrightarrow L^1_{\text{weak}}(X)$.

THEOREM 9.17 (Hardy-Littlewood). *There exists a constant $C > 0$ depending only on the dimension n such that the estimate $\|Mf\|_{L^1_{\text{weak}}} \leqslant C\|f\|_{L^1}$ holds for all $f \in L^1(\mathbb{R}^n)$.*

The proof requires a simple version of a result known as the *Vitali covering lemma*.

LEMMA 9.18 (Vitali). *For every finite collection of open balls $B_{r_1}(x_1), \dots, B_{r_N}(x_N) \subset \mathbb{R}^n$, there exists a subset $I \subset \{1, \dots, N\}$ such that $B_{r_i}(x_i) \cap B_{r_j}(x_j) = \varnothing$ for every $i, j \in I$ with $i \neq j$, and*

$$\bigcup_{i=1}^N B_{r_i}(x_i) \subset \bigcup_{j \in I} B_{3r_j}(x_j).$$

In particular, $m\left(\bigcup_{i=1}^N B_{r_i}(x_i)\right) \leqslant 3^n \sum_{j \in I} m(B_{r_j}(x_j))$.

PROOF. Abbreviate $B_i := B_{r_i}(x_i)$, and reorder the balls so that, without loss of generality, $r_1 \geqslant \ldots \geqslant r_N$. Define $I = \{i_1, \dots, i_\ell\} \subset \{1, \dots, N\}$ such that $i_1 := 1$ and, for each $j \geqslant 1$, i_{j+1} is the smallest number greater than i_j such that $B_{i_{j+1}}$ is disjoint from $B_{i_1} \cup \ldots \cup B_{i_j}$. This process terminates after finitely many steps, and if $k \notin I$, it means that B_k intersects B_i for some $i \in I$ with $i < k$. Since $r_k \leqslant r_i$, it follows that $B_k \subset B_{3r_i}(x_i)$. □

PROOF OF THEOREM 9.17. We shall prove that the stated inequality holds with $C = 3^n$: in other words, for every $f \in L^1(\mathbb{R}^n)$ and $t > 0$,

$$m(A_t) \leqslant \frac{3^n \cdot \|f\|_{L^1}}{t}, \quad \text{where} \quad A_t := \{x \in \mathbb{R}^n \mid Mf(x) > t\}. \tag{9.4}$$

By the inner regularity of the Lebesgue measure (see [Sal16, Theorem 2.13]), it will suffice to prove that every compact set $K \subset \mathbb{R}^n$ with $Mf|_K > t$ satisfies $m(K) \leqslant \frac{3^n \cdot \|f\|_{L^1}}{t}$. For each $x \in K$, the condition $Mf(x) > t$ means there exists a ball $B(x) \subset \mathbb{R}^n$ about x such that

$$\frac{1}{m(B(x))} \int_{B(x)} |f| \, dm > t. \tag{9.5}$$

Using the compactness of K, choose a finite subcollection $B_1, \ldots, B_N$ of such balls so that $K \subset \bigcup_{j=1}^N B_j$. After reordering the balls, we can then apply Lemma 9.18 to assume that $B_1, \ldots, B_\ell$ (for some $\ell \leqslant N$) are all disjoint and, using (9.5),

$$m(K) \leqslant m\left(\bigcup_{j=1}^N B_j\right) \leqslant 3^n \sum_{j=1}^{\ell} m(B_j) < 3^n \sum_{j=1}^{\ell} \frac{1}{t} \int_{B_j} |f| \, dm \leqslant \frac{3^n}{t} \|f\|_{L^1}.$$

□

We now have enough tools to complete the proof of the Lebesgue differentiation theorem.

Proof of Lemma 9.13 (and thus Theorem 9.11). The estimate (9.2) implies

$$f^r(x) \leqslant M(f - f_k)(x) + f_k^r(x) + |f_k(x) - f(x)|. \tag{9.6}$$

Given $N \in \mathbb{N}$, choose k large enough so that

$$\|f - f_k\|_{L^1} < \frac{1}{3^n \cdot 4N^2}.$$

Then Chebyshev's inequality (9.3) implies

$$m\left(\left\{x \in \mathbb{R}^n \;\middle|\; |f_k(x) - f(x)| > 1/2N\right\}\right) \leqslant 2N\|f - f_k\|_{L^1} < \frac{1}{3^n \cdot 2N} < \frac{1}{2N},$$

and by Theorem 9.17 (in particular (9.4)),

$$m\left(\left\{x \in \mathbb{R}^n \;\middle|\; M(f - f_k)(x) > 1/2N\right\}\right) \leqslant 2N \cdot 3^n \|f - f_k\|_{L^1} < \frac{1}{2N},$$

thus for all $x \in \mathbb{R}^n$ outside a set of measure at most $1/N$, (9.6) becomes

$$f^r(x) \leqslant \frac{1}{2N} + f_k^r(x) + \frac{1}{2N} = f_k^r(x) + \frac{1}{N}.$$

Letting r go to 0, we conclude $f^0(x) \leqslant \frac{1}{N}$ since $f_k^0(x) = 0$ by the continuity of f_k. □

9.4. The Radon-Nikodým theorem. Recall that if (X, μ) is a measure space and $f : X \to [0, \infty]$ a measurable function, then one can define another measure λ on the same σ-algebra by

$$\lambda(A) := \int_A f \, d\mu. \tag{9.7}$$

In order to show that absolutely continuous functions can always be written as integrals, we will first answer the following question, which turns out to be easier:

Question 9.19. *Given two measures μ and λ defined on the same measurable space X, does there exist a measurable function $f : X \to [0, \infty]$ such that (9.7) holds for all measurable sets A?*

The function $f : X \to [0, \infty]$ in this relation, if it exists, is sometimes called the **Radon-Nikodým derivative** of λ with respect to μ, and written as

$$\frac{d\lambda}{d\mu} := f.$$

It is not hard to think up necessary conditions for the existence of such a function. For example, there clearly is no such function if μ is the counting measure on $\mathbb{R}$ and λ is the Lebesgue measure, as $\int_A f\,d\mu = \sum_{x\in A} f(x)$ then can only be finite when $A \subset \mathbb{R}$ is a countable set, whose Lebesgue measure is therefore zero. It is also impossible if one takes the Lebesgue measure on $\mathbb{R}^n$ as μ and the **Dirac measure**

$$\delta(A) := \begin{cases} 1 & \text{if } 0 \in A, \\ 0 & \text{otherwise} \end{cases}$$

as λ; this is just the statement that the "Dirac δ-function" so popular among physicists does not actually exist. One thing both of these counterexamples have in common is that one can find measurable sets $A \subset X$ for which $\mu(A) = 0$ but $\lambda(A) \neq 0$. This possibility clearly needs to be excluded since $\int_A f\,d\mu = 0$ for every function when $\mu(A) = 0$.

Definition 9.20. Given a measure space (X, μ), a measure λ defined on the same σ-algebra is called **absolutely continuous** with respect to μ (written "$\lambda \ll \mu$") if the implication

$$\mu(A) = 0 \quad \Rightarrow \quad \lambda(A) = 0$$

holds for all measurable sets $A \subset X$.

The following exercise is not logically necessary for our exposition, but it demonstrates that there are nontrivial connections between Definition 9.20 and the notion of absolutely continuous functions.

Exercise 9.21 (cf. Lemma 9.3). Show that if $\lambda \ll \mu$ and $\lambda(X) < \infty$, then for every $\epsilon > 0$ there exists a $\delta > 0$ such that for measurable sets $A \subset X$,

$$\mu(A) < \delta \quad \Rightarrow \quad \lambda(A) < \epsilon.$$

Theorem 9.22 (Radon-Nikodým). *If μ and λ are two σ-finite measures on the same measurable space X, then the following conditions are equivalent:*

(1) $\lambda \ll \mu$;
(2) There exists a measurable function $f : X \to [0, \infty]$ satisfying $\lambda(A) = \int_A f\,d\mu$ for all measurable sets $A \subset X$.

The implication (2) $\Rightarrow$ (1) is immediate. We will prove (1) $\Rightarrow$ (2) using the natural isomorphism between L^∞ and the dual space of L^1. To see how this arises, note that if λ is given by (9.7), then for every real-valued λ-integrable function $g \in L^1(X, \lambda)$,

$$\int_X g\,d\lambda = \int_X gf\,d\mu. \tag{9.8}$$

The non-obvious trick is to view $\Lambda(g) := \int_X g\,d\lambda \in \mathbb{R}$ as defining a bounded linear functional $\Lambda : L^1(X, \lambda + \mu) \to \mathbb{R}$, which makes sense since $(\lambda + \mu)(A) := \lambda(A) + \mu(A)$ defines yet another measure on the same σ-algebra as λ and μ, and we have

$$|\Lambda(g)| = \left|\int_X g\,d\lambda\right| \leqslant \int_X |g|\,d\lambda \leqslant \int_X |g|\,d(\lambda + \mu) = \|g\|_{L^1(\lambda+\mu)}. \tag{9.9}$$

Since $\lambda + \mu$ is σ-finite, it follows then from the Riesz representation theorem that there exists a real-valued function $h \in L^\infty(X, \lambda + \mu)$ with $\|h\|_{L^\infty(\lambda+\mu)} \leqslant 1$ such that

$$\int_X g\,d\lambda = \int_X hg\,d(\lambda + \mu) \quad \text{for all} \quad g \in L^1(X, \lambda + \mu). \tag{9.10}$$

This is enough information to derive a formula for f in terms of h: indeed, combining (9.8) and (9.10) gives

$$\int_X gf\,d\mu = \int_X hg\,d(\lambda+\mu) = \int_X hg\,d\lambda + \int_X hg\,d\mu = \int_X (hgf + hg)\,d\mu = \int_X gh(f+1)\,d\mu$$

for all $g \in L^1(X, \lambda+\mu) \subset L^1(X, \lambda)$. This suggests the relation $f = h(f+1)$, or equivalently

$$f = \frac{h}{1-h}.$$

There are a few subtle issues to check before we can call this a proof—you may notice for instance that we have not yet used the condition $\lambda \ll \mu$.

Proof of Theorem 9.22. Following the trick described above, we note that since λ and μ are both σ-finite, $\lambda+\mu$ is also a σ-finite measure, so that the Riesz representation theorem (Theorem 5.4) gives a natural isomorphism between $L^\infty(X, \lambda+\mu)$ and the dual space of $L^1(X, \lambda+\mu)$. The bounded linear functional $\Lambda : L^1(X, \lambda+\mu) \to \mathbb{R}$ defined by $\Lambda(g) := \int_X f\,d\lambda$ therefore gives rise to a unique (up to equality almost everywhere) real-valued function $h \in L^\infty(X, \lambda+\mu)$ satisfying (9.10).

We claim that h satisfies $0 \leqslant h < 1$ almost everywhere with respect to the measure μ. Indeed, for $n \in \mathbb{N}$, let $A_n := \{x \in X \mid h \leqslant -1/n\}$, and suppose $A'_n \subset A_n$ is any subset for which $\lambda(A'_n) + \mu(A'_n) < \infty$. The function $g := \chi_{A'_n}$ is then in $L^1(X, \lambda+\mu)$, so plugging it into (9.10) gives

$$0 \leqslant \lambda(A'_n) = \int_X g\,d\lambda = \int_X hg\,d(\lambda+\mu) = \int_{A'_n} h\,d(\lambda+\mu) \leqslant -\frac{1}{n}\left[\lambda(A'_n) + \mu(A'_n)\right] \leqslant 0,$$

implying $\mu(A'_n) = \lambda(A'_n) = 0$. Since λ and μ are both σ-finite, A_n is a union of countably many subsets on which λ and μ are both finite, so having shown that μ and λ vanish on all of these subsets, it follows that $\mu(A_n) = \lambda(A_n) = 0$. The set on which $h < 0$ is now the countable union of the sets A_n for $n \in \mathbb{N}$, and therefore also has measure zero with respect to both μ and λ. The other bound follows similarly by setting $A := \{x \in X \mid h(x) \geqslant 1\}$: for any subset $A' \subset A$ with $\lambda(A')$ and $\mu(A')$ both finite, we can plug $g := \chi_{A'}$ into (9.10) and find

$$\lambda(A') = \int_X g\,d\lambda = \int_X hg\,d(\lambda+\mu) = \int_{A'} h\,d\lambda + \int_{A'} h\,d\mu \geqslant \lambda(A') + \mu(A'),$$

implying $\mu(A') = 0$. (Notice that this time, we do not immediately also obtain $\lambda(A') = 0$; the latter follows since $\lambda \ll \mu$, but it need not be true without the absolute continuity assumption.) Appealing once more to the σ-finiteness of λ and μ, this implies $\mu(A) = 0$.

The function $f := \frac{h}{1-h}$ therefore satisfies $0 \leqslant f < \infty$ almost everywhere with respect to μ, so we can define a measure μ_f by

$$\mu_f(A) := \int_A f\,d\mu.$$

We claim $\mu_f = \lambda$. To see this, let us rewrite the relation (9.10) in the form

$$\int_X (1-h)g\,d\lambda = \int_X hg\,d\mu \quad \text{for all} \quad g \in L^1(X, \lambda+\mu).$$

Now if $A \subset X$ is any measurable subset for which the function $g := \frac{1}{1-h}\chi_A$ is in $L^1(X, \lambda+\mu)$, we obtain

$$\lambda(A) = \int_X \chi_A\,d\lambda = \int_X (1-h)g\,d\lambda = \int_X hg\,d\mu = \int_X \frac{h}{1-h}\chi_A\,d\mu = \int_A f\,d\mu = \mu_f(A).$$

To extend this to an arbitrary measurable subset $A \subset X$, we can again appeal to σ-finiteness and write

$$X = \bigcup_{n \in \mathbb{N}} X_n = \bigcup_{n \in \mathbb{N}} Y_n$$

for two sequences of subsets $X_1 \subset X_2 \subset \ldots \subset X$ and $Y_1 \subset Y_2 \subset \ldots \subset X$ with $\mu(X_n) < \infty$ and $\lambda(Y_n) < \infty$. For $n \in \mathbb{N}$, let

$$A_n := \{x \in A \mid 1 - h(x) \geqslant 1/n\} \cap X_n \cap Y_n \subset X.$$

Then A_n has finite $(\mu + \lambda)$-measure and $\frac{1}{1-h} \leqslant n$ on A_n, thus $\frac{1}{1-h}\chi_{A_n} \in L^1(X, \lambda + \mu)$, so the calculation above proves $\lambda(A_n) = \mu_f(A_n)$. To finish, observe that since $h < 1$ almost everywhere with respect to μ, absolute continuity $\lambda \ll \mu$ implies that this is also true with respect to λ, and thus

$$\lambda\Big(A \backslash \bigcup_{n \in \mathbb{N}} A_n\Big) = \mu_f\Big(A \backslash \bigcup_{n \in \mathbb{N}} A_n\Big) = 0.$$

This justifies the following limit computation:

$$\lambda(A) = \lim_{n \to \infty} \lambda(A_n) = \lim_{n \to \infty} \mu_f(A_n) = \mu_f(A).$$

□

REMARK 9.23. Without the condition $\lambda \ll \mu$, the function h constructed in the proof above may satisfy $h = 1$ on a set with positive λ-measure, in which case $\lim_{n \to \infty} \lambda(A_n) \leqslant \lambda(A)$ in the last step, producing an inequality

$$\int_A f \, d\mu \leqslant \lambda(A)$$

which may in general be strict. The argument still proves that equality holds for every measurable set $A \subset X$ such that the function $\frac{1}{1-h}\chi_A \in L^1(X, \lambda + \mu)$, but in pathological examples, there may be no interesting sets with this property.

EXERCISE 9.24. Find the function $f : \mathbb{R}^n \to [0, \infty]$ that is constructed in the proof of Theorem 9.22 for the case where $\mu := m$ is the Lebesgue measure and $\lambda := \delta$ the Dirac measure. On which sets $A \subset \mathbb{R}^n$ is equality achieved in $\int_A f \, dm \leqslant \delta(A)$?

REMARK 9.25. There are more general versions of the Radon-Nikdým theorem for so-called *signed measures* and *complex measures*, in which f in the formula $\lambda(A) = \int_A f \, d\mu$ may be a real or complex-valued μ-integrable function. See for example [**Rud87**, Chapter 6] or [**Sal16**, §5.4].

9.5. Absolutely continuous functions are integrals. In light of Corollary 9.12, the hard direction of the fundamental theorem of calculus for the Lebesgue integral now follows from:

LEMMA 9.26. *Every absolutely continuous function F on $[a, b] \subset \mathbb{R}$ is given by*

$$F(x) = F(a) + \int_a^x f(t) \, dt \tag{9.11}$$

for some $f \in L^1([a, b])$.

The lemma is valid for functions $F : [a, b] \to V$ with values in an arbitrary finite-dimensional vector space, but we will focus on the case $V = \mathbb{R}$, which immediately implies the general case after choosing a real basis of V. For real-valued functions, we will deduce it from the Radon-Nikodým theorem. Some alternative approaches and interesting related facts are outlined in §9.5.3.

9.5.1. *The case of strictly increasing functions.* Let us assume $F : [a, b] \to \mathbb{R}$ is absolutely continuous and strictly increasing, i.e. it satisfies

$$F(y) > F(x) \quad \text{whenever} \quad y > x.$$

The following lemma produces another connection between the notions of absolute continuity for functions and measures.

LEMMA 9.27. *F maps every set of measure zero in $[a, b]$ to a set of measure zero in $\mathbb{R}$.*

PROOF. If $A \subset [a, b]$ has measure zero, then for any given $\delta > 0$, A is contained in the union of a sequence of disjoint intervals (a_i, b_i) such that $\sum_{i=1}^{\infty}(b_i - a_i) < \delta$. Absolute continuity guarantees that if $\epsilon > 0$ is given, $\delta > 0$ in the previous sentence can be chosen so that for every $k \in \mathbb{N}$, $\sum_{i=1}^{k} |f(b_i) - f(a_i)| < \epsilon$, and consequently,

$$\sum_{i=1}^{\infty} |f(b_i) - f(a_i)| \leqslant \epsilon.$$

The image $F(A)$ is therefore contained in a countable union of open intervals $(F(a_i), F(b_i))$ whose total measure is at most ϵ. ☐

EXERCISE 9.28. Find a set of measure zero whose image under the Cantor function of Example 9.2 has positive measure. (Note that in Lemma 9.27, we did not actually need to assume that F is *strictly* increasing.)

PROOF OF LEMMA 9.26 FOR F STRICTLY INCREASING. We claim that the formula

$$\lambda(A) := m(F(A))$$

defines a measure on $[a, b]$ with $\lambda \ll m$.

We need to check first that the image under F of every Lebesgue measurable set $A \subset [a, b]$ is Lebesgue measurable. By the inner regularity of the Lebesgue measure (see [Sal16, Theorem 2.13]), we can choose a sequence of compact subsets $K_1 \subset K_2 \subset K_3 \subset \ldots \subset A$ such that A is the disjoint union of a set A_0 of measure zero with $K_\infty := \bigcup_{n \in \mathbb{N}} K_n$. Since F is continuous, $F(K_n) \subset \mathbb{R}$ is also compact for each n, which makes $F(K_\infty)$ a countable union of compact sets and thus a Borel set. By Lemma 9.27, $F(A_0)$ is another set of measure zero, thus $F(A) = F(A_0) \cup F(K_\infty)$ is Lebesgue measurable.

Since $F : [a, b] \to \mathbb{R}$ is strictly increasing, it is also injective, so disjoint measurable subsets $A_1, A_2, A_3, \ldots \subset [a, b]$ have disjoint images, and it follows that λ as defined above is σ-additive. It is clearly also finite since F has bounded image, so this proves that λ is a measure, and Lemma 9.27 implies that it is absolutely continuous with respect to the Lebesgue measure.

The Radon-Nikodým theorem now provides a measurable function $f : [a, b] \to [0, \infty]$ such that for all measurable subsets $A \subset [a, b]$,

$$m(F(A)) = \int_A f \, dm.$$

In particular for $A := [a, x]$, this gives $F(x) - F(a) = \int_a^x f(t) \, dt$, and the global integrability of f follows from this by setting $x = b$. ☐

9.5.2. *The general case.* If $F : [a, b] \to \mathbb{R}$ is increasing but not strictly, then there is a cheap trick to reduce Lemma 9.26 to the strictly increasing case: we consider the function

$$G(x) := x + F(x),$$

which *is* strictly increasing (and also absolutely continuous), even if F is constant on some subinterval. The strictly increasing case therefore provides a function $g \in L^1([a,b])$ such that $G(x) = G(a) + \int_a^x g(t)\,dt$, and it follows that

$$F(x) = F(a) + \int_a^x [g(t) - 1]\,dt.$$

The proof for increasing functions is thus complete.

The conclusion of the proof now follows from an important general observation about all functions of *bounded variation*, which includes the absolutely continuous functions. Given any function f on $[a,b]$, we define the **total variation** of f by

$$TV(f) := \sup\left\{ \sum_{i=1}^N |f(x_i) - f(x_{i-1})| \;\middle|\; N \geqslant 1 \text{ and } a = x_0 < x_1 < \ldots < x_N = b \right\} \in [0,\infty],$$

and say that f is of **bounded variation** if $TV(f) < \infty$. Notice that if f is of bounded variation, then its restriction to every compact subinterval of $[a,b]$ is also of bounded variation.

LEMMA 9.29. *If f is absolutely continuous on $[a,b]$, then it is of bounded variation, and the function $Vf : [a,b] \to [0,\infty)$ defined by $Vf(x) := TV(f|_{[a,x]})$ is also absolutely continuous.*

PROOF. We note first that if $a \leqslant x < c < y \leqslant b$, then $|f(y)-f(x)| \leqslant |f(y)-f(c)|+|f(c)-f(x)|$, thus if $a \leqslant x_0 < x_1 < \ldots < x_N = b$ is any partition of the interval $[a,b]$ that does not include $c \in (a,b)$, adding c to the partition can only increase the value of the sum in the definition of $TV(f)$. It follows that we lose no generality if we modify the definition of $TV(f)$ so that the supremum ranges only over partitions that include c, which gives rise to the relation

$$TV(f) = TV(f|_{[a,c]}) + TV(f|_{[c,b]}).$$

In particular, this implies

$$Vf(y) = Vf(x) + TV(f|_{[x,y]}) \quad \text{for every } y > x \text{ in } [a,b], \tag{9.12}$$

so the function $Vf : [a,b] \to [0,\infty]$ is increasing. Now choose $\epsilon > 0$ and $\delta > 0$ as in the definition of absolute continuity, and choose a partition $a = t_0 < t_1 < \ldots < t_N = b$ such that $t_i - t_{i-1} < \delta$ for every i. Any partition of $[t_{i-1}, t_i]$ is then a finite collection of closures of disjoint open intervals with total length less than δ, implying $TV(f|_{[t_{i-1},t_i]}) < \epsilon$ and thus

$$TV(f) = \sum_{i=1}^N TV(f|_{[t_{i-1},t_i]}) < N\epsilon < \infty.$$

Keeping the same ϵ and δ, if $a \leqslant a_1 < b_1 < \ldots < a_n < b_n \leqslant b$ satisfy $\sum_{i=1}^n (b_i - a_i) < \delta$, then (9.12) implies

$$\sum_{i=1}^n |Vf(b_i) - Vf(a_i)| = \sum_{i=1}^n TV(f|_{[a_i,b_i]}).$$

The latter is the supremum of sums $\sum_j |f(t_j) - f(t_{j-1})|$ over finite collections of intervals $[t_{j-1}, t_j]$ with disjoint interiors whose lengths add up to $\sum_{i=1}^n (b_i - a_i) < \delta$, hence the sum is less than ϵ. ☐

LEMMA 9.30. *For any function $f : [a,b] \to \mathbb{R}$ of bounded variation, the functions Vf, $Vf + f$ and $Vf - f$ on $[a,b]$ are increasing.*

PROOF. That Vf is increasing follows already from (9.12), and in fact for $y > x$,

$$Vf(y) - Vf(x) = TV(f|_{[x,y]}) \geqslant |f(y) - f(x)|$$

since $x < y$ is a particular example of a partition of $[x, y]$. It follows that $Vf(y) - Vf(x) \geqslant f(y) - f(x)$ and $Vf(y) - Vf(x) \geqslant f(x) - f(y)$, thus

$$Vf(y) - f(y) \geqslant Vf(x) - f(x) \quad \text{and} \quad Vf(y) + f(y) \geqslant Vf(x) + f(x).$$

$\square$

Conclusion of the proof of Lemma 9.26. An arbitrary absolutely continuous function $F : [a, b] \to \mathbb{R}$ can be decomposed as

$$F = \frac{1}{2}(VF + F) - \frac{1}{2}(VF - F),$$

where by Lemmas 9.29 and 9.30, $VF + F$ and $VF - F$ are each absolutely continuous and increasing. We have already proved therefore that both can be represented as integrals of L^1-functions on $[a, b]$, and the same thus follows for F. $\square$

9.5.3. *Alternative approaches.* There are other ways of proving Lemma 9.26 without using the Radon-Nikodým theorem. Since we expect an absolutely continuous function F to be the integral of its derivative, one method for finding the function f in (9.11) is to prove directly that F is differentiable almost everywhere. This can be deduced from the following somewhat surprising classical result of Lebesgue:

Theorem 9.31 (Lebesgue). *Every monotone function $f : [a, b] \to \mathbb{R}$ is differentiable almost everywhere.*

A proof of this theorem using a more elaborate version of the Vitali covering lemma (cf. Lemma 9.18) may be found in [**Roy88**, Chapter 5]; see also [**RSN90**] for a slightly different exposition. It takes only slightly more effort to see that for a monotone function $f : [a, b] \to \mathbb{R}$, f' belongs to $L^1([a, b])$. The argument goes as follows: consider the case where $f : [a, b] \to \mathbb{R}$ is increasing, and for convenience, extend f over $\mathbb{R}$ with constant values $f(x) = f(a)$ for $x \leqslant a$ and $f(x) = f(b)$ for $x \geqslant b$. The **difference quotients**

$$D_h f(x) := \frac{f(x + h) - f(x)}{h}$$

are then well-defined functions $D_h f : [a, b] \to \mathbb{R}$ for every $h \in \mathbb{R}\backslash\{0\}$, and they are clearly measurable functions since f is measurable. By Theorem 9.31, there exists a function $f' : [a, b] \to \mathbb{R}$ which can be defined as $f'(x) = \lim_{h \to 0} D_h f(x)$ wherever this limit exists and zero everywhere else; this means $D_h f \to f'$ almost everywhere as $h \to 0$, thus f' is measurable. The difference quotients are easily seen to satisfy a "discrete" variant of the fundamental theorem of calculus: instead of $\int_a^b f'(t)\, dt = f(b) - f(b)$, one has

$$\begin{aligned}\int_{[a,b]} D_h f\, dm &= \frac{1}{h}\int_a^b [f(x + h) - f(x)]\, dx = \frac{1}{h}\int_{a+h}^{b+h} f(x)\, dx - \frac{1}{h}\int_a^b f(x)\, dx \\ &= \frac{1}{h}\int_{[b,b+h]} f\, dm - \frac{1}{h}\int_{[a,a+h]} f\, dm.\end{aligned} \tag{9.13}$$

Taking $h > 0$, in the situation at hand we have defined f to be constant on $[b, b + h]$, and $-f$ is bounded above by $-f(a)$ on $[a, a + h]$, so this computation implies

$$\int_{[a,b]} D_h f\, dm \leqslant f(b) - f(a) \quad \text{for } h > 0.$$

Now if we consider the sequence of functions $D_{1/n}f$ for $n \in \mathbb{N}$, they have nonnegative values since f is increasing, and they converge almost everywhere to f', thus Fatou's lemma (see [**Sal16**, Theorem 1.41]) gives

$$\int_{[a,b]} f'\,dm = \int_{[a,b]} \liminf_{n\to\infty} D_{1/n}h\,dm \leqslant \liminf_{n\to\infty} \int_{[a,b]} D_{1/n}h\,dm \leqslant f(b) - f(a).$$

COROLLARY 9.32. *For every monotone function $f : [a,b] \to \mathbb{R}$, f' is measurable and satisfies*

$$\|f'\|_{L^1} \leqslant |f(b) - f(a)|.$$

□

REMARK 9.33. Corollary 9.32 is analogous to the inequality in Remark 9.23 for a measure λ on $[a,b]$ that need not satisfy $\lambda \ll m$. It comes with the caveat that without an assumption of absolute continuity, the inequality may fail to carry any interesting information, e.g. the Cantor function (Example 9.2) shows that the left hand side can simply vanish, even when the function f is far from being constant.

Theorem 9.31 and Corollary 9.32 have an immediate consequence for the class of functions $f : [a,b] \to \mathbb{R}$ that can be written as the difference $f_+ - f_-$ between two increasing functions $f_\pm : [a,b] \to \mathbb{R}$. This is precisely the class of functions with bounded variation that we saw in §9.5.2, which includes all absolutely continuous functions, thus:

COROLLARY 9.34. *Every absolutely continuous function $f : [a,b] \to \mathbb{R}$ is differentiable almost everywhere and its derivative belongs to $L^1([a,b])$.* □

With this result in place, we can compare any given absolutely continuous function F on $[a,b]$ with the function $F_1(x) := \int_a^x F'(t)\,dt$. The latter is also absolutely continuous and differentiable almost everywhere, with $F_1' = F'$ by Corollary 9.12, so looking at $F - F_1$ reduces the problem to proving:

LEMMA 9.35. *Every absolutely continuous function on $[a,b]$ whose derivative vanishes almost everywhere is constant.*

For a fairly short proof of this, based again on the Vitali covering lemma, see [**Roy88**, §5.4, Lemma 13].

From a functional-analytic perspective, there is a more interesting alternative argument to be found in the most recent edition [**RF10**] of Royden's classic textbook, based originally on the article [**FH15**]. It makes use of the following notion, which should seem natural in light of Lemma 9.3:

DEFINITION 9.36. A collection $\mathcal{F}$ of integrable functions on a measure space (X,μ) is called **uniformly integrable** (or **equi-integrable**) if for every $\epsilon > 0$, there exists $\delta > 0$ such that for all $f \in \mathcal{F}$ and measurable subsets $A \subset X$,

$$\mu(A) < \delta \quad \Rightarrow \quad \int_A |f|\,d\mu < \epsilon.$$

You should think of this definition as a close analogue of *equicontinuity*: the point is that the correspondence between ϵ and δ is not allowed to depend on the choice of the function $f \in \mathcal{F}$.

EXAMPLE 9.37. If $\mathcal{F}$ is any collection of measurable functions f satisfying $|f| \leqslant g$ for some fixed $g \in L^1(X)$, then Lemma 9.3 implies that $\mathcal{F}$ is uniformly integrable.

The relevance of uniform integrability to this discussion arises from the following observation:

LEMMA 9.38 ([**FH15**] or [**RF10**, §6.4]). *A continuous function f on $[a,b]$ is absolutely continuous if and only if its family of difference quotients $\{D_h f\}_{0<h\leqslant 1}$ is uniformly integrable on $[a,b]$.*

With this understood, one can now appeal to a useful generalization of the dominated convergence theorem:

THEOREM 9.39 (Vitali's convergence theorem; see [RF10, §4.6]). *For a Lebesgue-measurable subset $X \subset \mathbb{R}$ with finite measure, if $\{f_n\}_{n\in\mathbb{N}}$ is a uniformly integrable collection of functions on X such that $f_n \to f$ pointwise almost everywhere, then $f \in L^1(X)$ and $\int_X f_n\,dm \to \int_X f\,dm$.*

REMARK 9.40. The convergence theorem is also true on general finite measure spaces if one adds the condition $f \in L^1(X)$ to the hypotheses; in our situation, that version would also suffice in light of Corollary 9.34. There is also a version for spaces with infinite measure; see [RF10, §18.3].

Knowing that an absolutely continuous function f is differentiable almost everywhere, one can now deduce $\int_{[a,b]} f'\,dt = f(b) - f(a)$ as follows: setting $h = 1/n$ in (9.13) gives a sequence of relations whose left hand sides converge as $n \to \infty$ to $\int_{[a,b]} f'\,dt$ due to Lemma 9.38 and Vitali's convergence theorem. At the same time, the right hand sides converge to $f(b) - f(a)$ since f is continuous, so we are done.

10. Fourier series

10.1. Fully periodic functions. In this section we consider functions f on $\mathbb{R}^n$ that are 1-periodic in every variable, meaning that the relation

$$f(x_1, \dots, x_{j-1}, x_j + 1, x_{j+1}, \dots, x_n) = f(x_1, \dots, x_n)$$

holds for every $j = 1, \dots, n$. We shall refer to functions with this property as **fully periodic** functions. Some obvious examples include the trigonometric functions

$$\begin{aligned} &\sin(2\pi k x_j), \quad \text{for} \quad k = 1, 2, 3, \dots \text{ and } j = 1, \dots, n, \\ &\cos(2\pi k x_j), \quad \text{for} \quad k = 0, 1, 2, \dots \text{ and } j = 1, \dots, n, \end{aligned} \tag{10.1}$$

plus all products and linear combinations of these functions. The idea of a Fourier series is to express arbitrary fully periodic functions as (possibly infinite) sums of products of precisely these functions. Algebraically, it is much easier to work with complex exponentials than trigonometric functions, thus we shall allow all our fully periodic functions to take values in a *complex* vector space and, instead of writing them in terms of the functions in (10.1), try to express them as linear combinations of products of the functions

$$e^{2\pi i k x_j}, \quad \text{for} \quad k \in \mathbb{Z} \text{ and } j = 1, \dots, n. \tag{10.2}$$

Notice that an arbitrary finite product of such functions takes the form

$$\varphi_k(x) := e^{2\pi i k \cdot x}, \quad \text{for} \quad k \in \mathbb{Z}^n, \tag{10.3}$$

where $k \cdot x$ denotes the standard Euclidean inner product of two vectors $k, x \in \mathbb{R}^n$. We use this notation to distinguish the inner product on $\mathbb{R}^n$ from the *complex* inner product $\langle\ ,\ \rangle$ on the finite-dimensional vector space V in which our functions will take their values. For this discussion, we explicitly set

$$\mathbb{K} := \mathbb{C},$$

and since it will often be relevant, we remind the reader that the standing convention for the complex inner product on V is

$$\langle iv, w \rangle = -i\langle v, w \rangle, \qquad \langle v, iw \rangle = i\langle v, w \rangle.$$

The main theorem on Fourier series states that every function in a sufficiently reasonable class of fully periodic functions $f : \mathbb{R}^n \to V$ can be expressed as a convergent sum of the form

$$f(x) = \sum_{k\in\mathbb{Z}^n} e^{2\pi i k \cdot x} \hat{f}_k \tag{10.4}$$

for a unique set of coefficients $\widehat{f}_k \in V$, called the **Fourier coefficients** of f. The right hand side of (10.4) is called the **Fourier series** of f. Since complex exponentials are linear combinations of trigonometric functions, it is always possible (though often tiresome) to rewrite a Fourier series as a sum of products of the trigonometric functions appearing in (10.1); in particular, the Fourier series of a *real*-valued fully periodic function $f : \mathbb{R}^n \to \mathbb{R}$ can always be re-expressed as a *real*-linear combination of products of real-valued trigonometric functions, so that complex numbers need not be mentioned. In applications, the complex numbers typically have no intrinsic meaning but make calculations much easier.

10.2. Function spaces on the torus and the lattice. A fully periodic function on $\mathbb{R}^n$ can equivalently be regarded as a funtion on an n-dimensional torus $\mathbb{T}^n$, which is by definition an n-fold Cartesian product of circles. The most convenient definition of $\mathbb{T}^n$ for our purposes is as follows. The **lattice** $\mathbb{Z}^n \subset \mathbb{R}^n$ is a subgroup of $\mathbb{R}^n$ with respect to the operation of vector addition, and since $\mathbb{R}^n$ is an abelian group, the subgroup is normal. We define $\mathbb{T}^n$ to be the quotient group

$$\mathbb{T}^n := \mathbb{R}^n / \mathbb{Z}^n,$$

so in other words, elements of $\mathbb{T}^n$ are equivalence classes of vectors in $\mathbb{R}^n$, such that two vectors $x, y \in \mathbb{R}^n$ are in the same equivalence class if and only if $x - y \in \mathbb{Z}^n$. In the case $n = 1$, the map

$$\mathbb{R}/\mathbb{Z} \ni [t] \mapsto (\cos(2\pi t), \sin(2\pi t)) \in \mathbb{R}^2$$

gives a natural bijection between $\mathbb{T}^1$ and the unit circle in $\mathbb{R}^2$, which is also often denoted by $S^1 \subset \mathbb{R}^2$ since it is a "1-dimensional sphere". Through this bijection, one can identify $\mathbb{T}^n$ with the n-fold product of copies of S^1.

We can make $\mathbb{T}^n$ into a metric space by defining

$$d([x], [y]) := \inf_{(x,y) \in [x] \times [y]} |x - y|,$$

where $|\cdot|$ denotes the standard Euclidean norm on $\mathbb{R}^n$. You should take a moment to convince yourself that this expression really defines a metric on $\mathbb{T}^n$. Moreover, the natural projection map

$$\pi : \mathbb{R}^n \to \mathbb{T}^n : x \mapsto [x]$$

is continuous with respect to this metric, and since $\mathbb{T}^n$ is the image under π of the compact subset $[0, 1]^n \subset \mathbb{R}^n$, it follows that $\mathbb{T}^n$ is compact.

The Lebesgue measure m on $\mathbb{R}^n$ also determines a natural measure on $\mathbb{T}^n$. Let $\mathcal{L}(\mathbb{R}^n)$ denote the σ-algebra of Lebesgue-measurable subsets of $\mathbb{R}^n$, and define $\mathcal{L}(\mathbb{T}^n) \subset 2^{\mathbb{T}^n}$ to consist of all sets $A \subset \mathbb{T}^n$ with the property that $\pi^{-1}(A) \in \mathcal{L}(\mathbb{R}^n)$. In other words, $\mathcal{L}(\mathbb{T}^n)$ is the largest σ-algebra on $\mathbb{T}^n$ for which the projection map $\pi : \mathbb{R}^n \to \mathbb{T}^n$ is measurable. For $A \in \mathcal{L}(\mathbb{T}^n)$, we then define

$$m(A) := m(\pi^{-1}(A) \cap [0, 1)^n) \geqslant 0.$$

It is straightforward to check that $(\mathbb{T}^n, \mathcal{L}(\mathbb{T}^n), m)$ by this definition is a measure space, and moreover, since $[0, 1)^n \subset \mathbb{R}^n$ has finite Lebesgue measure, $m(\mathbb{T}^n)$ is finite; indeed, $m(\mathbb{T}^n) = 1$.

EXERCISE 10.1. Show that every fully periodic function $f : \mathbb{R}^n \to V$ corresponds to a unique function $F : \mathbb{T}^n \to V$ such that

$$f(x) = F([x]) \quad \text{for all} \quad x \in \mathbb{R}^n,$$

and conversely, every function $F : \mathbb{T}^n \to V$ determines a fully periodic function $f : \mathbb{R}^n \to V$ via this same relation. Show moreover that f is continuous/measurable if and only if F is continuous/measurable, respectively, and for an integrable function $F : \mathbb{T}^n \to V$,

$$\int_{\mathbb{T}^n} F(x)\, dx := \int_{\mathbb{T}^n} F\, dm = \int_{[0,1)^n} F \circ \pi\, dm.$$

Since $\mathbb{T}^n$ is now both a compact metric space and a finite measure space, Exercise 10.1 has the following useful consequences. First, every continuous fully periodic function is equivalent to a continuous function on a compact metric space, and is therefore *bounded.* Second, a function $f : \mathbb{T}^n \to V$ can be integrable even if $f \circ \pi : \mathbb{R}^n \to V$ is not, as it is only the integral of $|f \circ \pi|$ over the cube $[0,1)^n$ that needs to be finite. In fact, periodicity guarantees that fully periodic functions $f : \mathbb{R}^n \to V$ can *never* be Lebesgue integrable on $\mathbb{R}^n$ unless they vanish almost everywhere, but this only happens because the function $f : \mathbb{R}^n \to V$ contains too much redundant information. Integrating f instead over the finite measure space $\mathbb{T}^n$ circumvents this problem.

In the following, we will keep Exercise 10.1 in mind and typically blur the distinction between arbitrary functions $\mathbb{T}^n \to V$ and fully periodic functions $\mathbb{R}^n \to V$. We will also drop the equivalence classes from the notation and denote elements of $\mathbb{T}^n$ simply as vectors $x \in \mathbb{R}^n$ when there is no ambiguity; when this notational shortcut is used, it means that *any* representative $x \in \mathbb{R}^n$ of the given element in $\mathbb{T}^n$ may be chosen, and no important results will depend on this choice.

Notice that if a fully periodic function is differentiable, then its partial derivatives are also periodic functions, thus we can sensibly speak of differentiable functions on $\mathbb{T}^n$ and define the hierarchy of function spaces

$$C^0(\mathbb{T}^n) \supset C^1(\mathbb{T}^n) \supset C^2(\mathbb{T}^n) \supset \ldots \supset C^\infty(\mathbb{T}^n),$$

where for each $k = 0, 1, 2, \ldots, \infty$ we define $C^k(\mathbb{T}^n)$ to be the vector space of fully periodic functions $\mathbb{R}^n \to V$ that are k-times continuously differentiable. For $k < \infty$, these spaces all admit natural Banach space structures, of which only the case $k = 0$ will be especially important for our purposes: the norm on $C^0(\mathbb{T}^n)$ is defined by

$$\|f\|_{C^0} := \max_{x \in \mathbb{R}^n} |f(x)|,$$

where the existence of the maximum is guaranteed by the fact that $\mathbb{T}^n$ is compact. Similarly, for each $p \in [1, \infty]$ the measure on $\mathbb{T}^n$ gives rise to a Banach space of V-valued functions (defined almost everywhere),

$$L^p(\mathbb{T}^n) := L^p(\mathbb{T}^n, m), \qquad \|f\|_{L^p} := \begin{cases} \left(\int_{\mathbb{T}^n} |f|^p \, dm\right)^{1/p} & \text{for } p < \infty, \\ \operatorname{ess\,sup} |f| & \text{for } p = \infty. \end{cases}$$

Note that since $\mathbb{T}^n$ is the image of the compact and finite-measure subset $[0,1]^n$ under the projection $\pi : \mathbb{R}^n \to \mathbb{T}^n$, a function $f : \mathbb{T}^n \to V$ will belong to $L^p(\mathbb{T}^n)$ if and only if $f \circ \pi : \mathbb{R}^n \to V$ is *locally* of class L^p on $\mathbb{R}^n$, i.e. its restriction to every compact subset must be of class L^p, but $f \circ \pi$ itself will not usually belong to $L^p(\mathbb{R}^n)$. The space $L^2(\mathbb{T}^n)$ has a natural complex inner product defined by

$$\langle f, g \rangle_{L^2} := \int_{\mathbb{T}^n} \langle f(x), g(x) \rangle \, dx,$$

which makes $L^2(\mathbb{T}^n)$ into a Hilbert space.

Since the continuous functions on $\mathbb{T}^n$ are bounded and $\mathbb{T}^n$ has finite measure, $C^0(\mathbb{T}^n)$ is a subspace of $L^p(\mathbb{T}^n)$ for every $p \in [1, \infty]$; so, therefore, is $C^\infty(\mathbb{T}^n)$. In fact:

Proposition 10.2. *For every $p \in [1, \infty)$, $C^\infty(\mathbb{T}^n)$ is a dense linear subspace of $L^p(\mathbb{T}^n)$.*

Proof. We shall deduce this from the result in §8 that $C^\infty(\mathbb{R}^n)$ is dense in $L^p(\mathbb{R}^n)$. Given $f \in L^p(\mathbb{T}^n)$, define $F : \mathbb{R}^n \to V$ by

$$F = f \circ \pi \text{ on } [0,1)^n, \qquad F = 0 \text{ elsewhere},$$

where $\pi : \mathbb{R}^n \to \mathbb{T}^n$ is the quotient projection. Then $F \in L^p(\mathbb{R}^n)$, so for every $\epsilon > 0$, there exists a smooth function $F_\epsilon \in C^\infty(\mathbb{R}^n)$ with

$$\|F - F_\epsilon\|_{L^p} < \epsilon.$$

Given any $\delta > 0$, we can also choose a smooth function $\beta_\delta : \mathbb{R}^n \to [0,1]$ that has compact support in $(0,1)^n$ and satisfies

$$\beta_\delta \equiv 1 \quad \text{on} \quad [\delta, 1-\delta]^n.$$

The function $\beta_\delta F_\epsilon : \mathbb{R}^n \to V$ is then smooth and has compact support in $(0,1)^n$, so it gives rise to a uniquely determined fully periodic smooth function $G_\epsilon^\delta : \mathbb{R}^n \to V$ such that $G_\epsilon^\delta = \beta_\delta F_\epsilon$ on $[0,1)^n$. Let $g_\epsilon^\delta : \mathbb{T}^n \to V$ denote the corresponding function on the n-torus such that $G_\epsilon^\delta = g_\epsilon^\delta \circ \pi$.

We claim that $\|f - g_\epsilon^\delta\|_{L^p}$ can be made arbitrarily small if ϵ and δ are each chosen sufficiently small. Indeed, abbreviate $Q := [0,1)^n$ and $Q_\delta := [\delta, 1-\delta]^n$. Then

$$\begin{aligned}\|f - g_\epsilon^\delta\|_{L^p}^p &= \int_{\mathbb{T}^n} |f - g_\epsilon^\delta|^p \, dm = \int_Q |F - G_\epsilon^\delta|^p \, dm \\ &= \int_{Q_\delta} |F - G_\epsilon^\delta|^p \, dm + \int_{Q \backslash Q_\delta} |F - G_\epsilon^\delta|^p \, dm.\end{aligned}$$

Since $G_\epsilon^\delta = \beta_\delta F_\epsilon = F_\epsilon$ on Q_δ, the first term in the second line is

$$\int_{Q_\delta} |F - G_\epsilon^\delta|^p \, dm = \int_{Q_\delta} |F - F_\epsilon|^p \, dm \leqslant \int_{\mathbb{R}^n} |F - F_\epsilon|^p \, dm = \|F - F_\epsilon\|_{L^p}^p < \epsilon^p,$$

which is made arbitrarily small by choosing $\epsilon > 0$ small. To estimate the other term in the second line, we can use the fact that β_δ takes values in $[0,1]$ and write

$$|F - G_\epsilon^\delta| = |F - \beta_\delta F_\epsilon| = |F - F_\epsilon + F_\epsilon(1 - \beta_\delta)| \leqslant |F - F_\epsilon| + |F_\epsilon|,$$

hence by Minkowski's inequality,

$$\begin{aligned}\left(\int_{Q\backslash Q_\delta} |F - G_\epsilon^\delta|^p \, dm\right)^{1/p} &\leqslant \left(\int_{Q\backslash Q_\delta} (|F - F_\epsilon| + |F_\epsilon|)^p \, dm\right)^{1/p} \\ &\leqslant \left(\int_{Q\backslash Q_\delta} |F - F_\epsilon|^p \, dm\right)^{1/p} + \left(\int_{Q\backslash Q_\delta} |F_\epsilon|^p \, dm\right)^{1/p}.\end{aligned}$$

Here, the first term in the second line is bounded above by $\|F - F_\epsilon\|_{L^p} < \epsilon$, while if $\epsilon > 0$ is fixed, the second term can be made arbitrarily small for sufficiently small $\delta > 0$ since $|F_\epsilon|^p$ is Lebesgue integrable and $\bigcap_{\delta>0}(Q\backslash Q_\delta)$ is a set of measure zero, so that $\lim_{\delta\to 0}\int_{Q\backslash Q_\delta} |F_\epsilon|^p \, dm = 0$. This proves the claim. □

Since the Fourier coefficients of a function $f : \mathbb{T}^n \to V$ are meant to be a collection of vectors $\hat{f}_k$ associated to elements $k \in \mathbb{Z}^n$, it will be useful to regard the collection of all these coefficients as a function

$$\hat{f} : \mathbb{Z}^n \to V.$$

There is no meaningful notion of continuity or differentiability for such functions, but we *can* speak of L^p-spaces on $\mathbb{Z}^n$ with respect to the **counting measure**, i.e. let $\nu : 2^{\mathbb{Z}^n} \to [0, \infty]$ denote the measure such that $\nu(A)$ for each $A \subset \mathbb{Z}^n$ is the number of elements in A. The L^p-spaces with respect to this measure are conventionally denoted by

$$\ell^p(\mathbb{Z}^n) := L^p(\mathbb{Z}^n, \nu), \qquad 1 \leqslant p \leqslant \infty,$$

and since nonempty subsets in $\mathbb{Z}^n$ always have positive measure, the elements in these spaces are actual functions, not just equivalence classes of functions. The counting measure identifies integrals with infinite series and integrability with absolute summability, so for each $p \in [1, \infty)$, the ℓ^p-norm of a function $f : \mathbb{Z}^n \to V$ is

$$\|f\|_{\ell^p} := \left(\sum_{k \in \mathbb{Z}^n} |f(k)|^p \right)^{1/p},$$

while

$$\|f\|_{\ell^\infty} := \sup_{k \in \mathbb{Z}^n} |f(k)|.$$

There is one more space of functions $f : \mathbb{Z}^n \to V$ that we will need to consider, called the space of **rapidly decreasing** coefficients and denoted by $\mathscr{S}(\mathbb{Z}^n)$. A function $f : \mathbb{Z}^n \to V$ is defined to be in $\mathscr{S}(\mathbb{Z}^n)$ if and only if for every n-variable polynomial function $P : \mathbb{R}^n \to \mathbb{R}$, the function

$$\mathbb{Z}^n \to V : k \mapsto P(k) f(k)$$

is bounded. Equivalently, this means that for every $m \in \mathbb{N}$, the function $k \mapsto |k|^m f(k)$ is bounded on $\mathbb{Z}^n$. Since m in this expression can be chosen arbitrarily large, it is clear that functions $f(k)$ in $\mathscr{S}(\mathbb{Z}^n)$ always decay to 0 as $|k| \to \infty$. In fact:

EXERCISE 10.3. Show that $\mathscr{S}(\mathbb{Z}^n)$ is a dense linear subspace of $\ell^p(\mathbb{Z}^n)$ for every $p \in [1, \infty)$. *Hint: All functions $\mathbb{Z}^n \to V$ with bounded support are in $\mathscr{S}(\mathbb{Z}^n)$.*

10.3. The transformations $\mathscr{F}$ and $\mathscr{F}^*$. Suppose for the moment that $(V, \langle\ ,\ \rangle)$ is $\mathbb{C}$ with its standard Hermitian inner product. The functions $\varphi_k(x) := e^{2\pi i k \cdot x}$ defined in (10.3) for $k \in \mathbb{Z}^n$ can then be regarded as elements of $C^\infty(\mathbb{T}^n)$ since they are fully periodic and smooth. Since they are bounded and $\mathbb{T}^n$ has finite measure, they can also be regarded as elements of $L^2(\mathbb{T}^n)$, and as it turns out, they form an orthonormal set:

EXERCISE 10.4. Show that $\|\varphi_k\|_{L^2} = 1$ for every $k \in \mathbb{Z}^n$ and $\langle \varphi_k, \varphi_{k'} \rangle_{L^2} = 0$ for every $k \neq k' \in \mathbb{Z}^n$.

If we now assume $\hat{f}_k \in \mathbb{C}$ are coefficients such that the sum $\sum_{k \in \mathbb{Z}^n} \hat{f}_k \varphi_k$ converges in the L^2-norm to $f \in L^2(\mathbb{T}^n)$, then Exercise 10.4 makes it easy to compute the Fourier coefficients in terms of f: we have

$$\hat{f}_k = \sum_{p \in \mathbb{Z}^n} \langle \varphi_k, \hat{f}_p \varphi_p \rangle_{L^2} = \langle \varphi_k, f \rangle_{L^2} = \int_{\mathbb{T}^n} e^{-2\pi i k \cdot x} f(x)\, dx.$$

This computation generalizes in a straightforward way to functions valued in a general finite-dimensional complex inner product space $(V, \langle\ ,\ \rangle)$ if we engage in a slight abuse of notation: let us define

$$\langle \varphi, \psi v \rangle_{L^2} := \langle \varphi, \psi \rangle_{L^2} v \in V \qquad \text{for } \varphi, \psi : \mathbb{T}^n \to \mathbb{C},$$

which by linearity gives rise to a natural pairing $\langle \varphi, f \rangle \in V$ for any pair of L^2-functions $\varphi : \mathbb{T}^n \to \mathbb{C}$ and $f : \mathbb{T}^n \to V$. The computation of $\hat{f}_k$ above then becomes valid for vector-valued functions. We shall take this formula as a definition of a transformation $\mathscr{F}$, which sends functions $f : \mathbb{T}^n \to V$ to functions

$$\mathscr{F} f := \hat{f} : \mathbb{Z}^n \to V : k \mapsto \hat{f}_k$$

whenever the integral on the right hand side of the following expression is well defined for all k:

$$(\mathscr{F} f)_k := \hat{f}_k := \int_{\mathbb{T}^n} e^{-2\pi i k \cdot x} f(x)\, dx. \tag{10.5}$$

It is clear, for instance, that if $f \in L^1(\mathbb{T}^n)$, then all of the coefficients $\hat{f}_k$ are well defined and they satisfy a uniform bound

$$|\hat{f}_k| \leqslant \|f\|_{L^1},$$

hence $\mathscr{F}$ defines a bounded linear operator

$$\mathscr{F} : L^1(\mathbb{T}^n) \to \ell^\infty(\mathbb{Z}^n). \tag{10.6}$$

There is a similar transformation $\mathscr{F}^*$ that associates to a function $g : \mathbb{Z}^n \to V : k \mapsto g_k$ a function

$$\mathscr{F}^* g := \check{g} : \mathbb{T}^n \to V$$

defined by

$$(\mathscr{F}^* g)(x) := \check{g}(x) := \sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} g_k. \tag{10.7}$$

As with the definition of $\mathscr{F}$ in (10.5), this definition comes with the caveat that at first glance, it only makes sense if the sum converges for every x. So for instance, it makes sense whenever $g \in \ell^1(\mathbb{Z}^n)$, as the sum then converges absolutely and uniformly; since the partial sums are all finite sums of continuous functions, it follows in this case that $\check{g} : \mathbb{T}^n \to V$ is a continuous function and satisfies $|\check{g}(x)| \leqslant \sum_{k \in \mathbb{Z}^n} |g_k| = \|g\|_{\ell^1}$ for all $x \in \mathbb{T}^n$, thus $\mathscr{F}^*$ defines a bounded linear operator

$$\mathscr{F}^* : \ell^1(\mathbb{Z}^n) \to C^0(\mathbb{T}^n). \tag{10.8}$$

We have already seen that under certain circumstances, the operators $\mathscr{F}$ and $\mathscr{F}^*$ are inverse to each other, e.g. the computation following Exercise 10.4 above shows that if $g : \mathbb{Z}^n \to \mathbb{C}$ is a function such that $\mathcal{F}^* g \in L^2(\mathbb{T}^n)$ and the series in the definition of $\mathcal{F}^* g$ converges in the L^2-norm, then $\mathcal{F}\mathcal{F}^* g = g$. The next two theorems are the main results we need to prove about Fourier series.

THEOREM 10.5. *The transformations $\mathscr{F}$ and $\mathscr{F}^*$ defined in* (10.5) *and* (10.7) *respectively have the following properties:*

1. *$\mathscr{F}$ maps $C^\infty(\mathbb{T}^n)$ bijectively onto $\mathscr{S}(\mathbb{Z}^n)$.*
2. *$\mathscr{F}^*$ maps $\mathscr{S}(\mathbb{Z}^n)$ bijectively onto $C^\infty(\mathbb{T}^n)$.*
3. *The bijections $\mathscr{F} : C^\infty(\mathbb{T}^n) \to \mathscr{S}(\mathbb{Z}^n)$ and $\mathscr{F}^* : \mathscr{S}(\mathbb{Z}^n) \to C^\infty(\mathbb{T}^n)$ are inverse to each other.*
4. *For every $f \in C^\infty(\mathbb{T}^n)$, the series $\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} \hat{f}_k$ converges absolutely and uniformly with all derivatives to f.*

THEOREM 10.6 (Parseval's identity). *For every $f, g \in C^\infty(\mathbb{T}^n)$,*

$$\langle \hat{f}, \hat{g} \rangle_{\ell^2} = \langle f, g \rangle_{L^2}.$$

Since $C^\infty(\mathbb{T}^n)$ is dense in $L^2(\mathbb{T}^n)$, Parseval's identity gives rise to a unique bounded linear extension of the operator $\mathscr{F} : C^\infty(\mathbb{T}^n) \to \mathscr{S}(\mathbb{Z}^n)$ to an operator

$$\mathscr{F} : L^2(\mathbb{T}^n) \to \ell^2(\mathbb{Z}^n). \tag{10.9}$$

In other words, for each $f \in L^2(\mathbb{T}^n)$, we can choose an approximating sequence $f_j \in C^\infty(\mathbb{T}^n)$ with $f_j \xrightarrow{L^2} f$ as $j \to \infty$, and define $\hat{f} = \mathscr{F} f \in \ell^2(\mathbb{Z}^n)$ as the ℓ^2-limit of the ℓ^2-Cauchy sequence $\hat{f}_j \in \mathscr{S}(\mathbb{Z}^n)$. This description makes $\mathscr{F} : L^2(\mathbb{T}^n) \to \ell^2(\mathbb{Z}^n)$ sound more abstract than it really is: in fact, since $\mathbb{T}^n$ has finite measure, $L^2(\mathbb{T}^n)$ is a subspace of $L^1(\mathbb{T}^n)$, so the operator in (10.9) is just the restriction of $\mathscr{F} : L^1(\mathbb{T}^n) \to \ell^\infty(\mathbb{Z}^n)$ to this subspace. In the other direction, the density of $\mathscr{S}(\mathbb{Z}^n)$ in $\ell^2(\mathbb{Z}^n)$ implies that $\mathscr{F}^* : \mathscr{S}(\mathbb{Z}^n) \to C^\infty(\mathbb{T}^n)$ extends uniquely to an operator

$$\mathscr{F}^* : \ell^2(\mathbb{Z}^n) \to L^2(\mathbb{T}^n),$$

defined similarly by choosing for any $g \in \ell^2(\mathbb{Z}^n)$ an approximating sequence $g_j \in \mathscr{S}(\mathbb{Z}^n)$ with $g_j \xrightarrow{\ell^2} g$ and writing $\mathscr{F}^* g = \check{g} \in L^2(\mathbb{T}^n)$ for the L^2-limit of the L^2-Cauchy sequence $\check{g}_j \in C^\infty(\mathbb{T}^n)$. Here there is an obvious choice of approximating sequence g_j available, namely

$$g_j(k) := \begin{cases} g(k) & \text{if } |k| \leqslant j, \\ 0 & \text{otherwise.} \end{cases}$$

This makes $\check{g}_j \in C^\infty(\mathbb{T}^n)$ a sequence of partial sums for the Fourier series $\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} g(k)$, so the conclusion is that this series converges to $\check{g}$ in the L^2-norm. It clearly cannot be expected to converge uniformly since $\check{g} \in L^2(\mathbb{T}^n)$ is not generally continuous, and there is also no guarantee of pointwise convergence, not even almost everywhere. The compositions $\mathscr{F}^*\mathscr{F} : C^\infty(\mathbb{T}^n) \to C^\infty(\mathbb{T}^n)$ and $\mathscr{F}\mathscr{F}^* : \mathscr{S}(\mathbb{Z}^n) \to \mathscr{S}(\mathbb{Z}^n)$, of course, each extend to $L^2(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$ and $\ell^2(\mathbb{Z}^n) \to \ell^2(\mathbb{Z}^n)$ respectively as the identity map. We summarize this discussion as follows.

Corollary 10.7. *The transformations $\mathscr{F}$ and $\mathscr{F}^*$ defined in (10.5) and (10.7) give well-defined unitary maps*[19] *$L^2(\mathbb{T}^n) \to \ell^2(\mathbb{Z}^n)$ and $\ell^2(\mathbb{Z}^n) \to L^2(\mathbb{T}^n)$ respectively, where in the latter case, the series should be interpreted as an L^2-convergent (but not necessarily pointwise convergent) series of functions in $L^2(\mathbb{T}^n)$. Moreover, these two transformations are inverse to each other.* $\square$

In light of Exercise 10.4, another way to say this is as follows:

Corollary 10.8. *For any orthonormal basis $v_1, \dots, v_m$ of $(V, \langle\ ,\ \rangle)$, the functions*

$$\left\{ e^{2\pi i k \cdot x} v_j \right\}_{k \in \mathbb{Z}^n,\ j = 1, \dots, m}$$

form an orthonormal basis of the Hilbert space $L^2(\mathbb{T}^n)$. $\square$

The remainder of §10 is concerned with the proofs of Theorems 10.5 and 10.6, and along the way, we will prove some relations between the Fourier operations and differentiation which are frequently useful in applications.

10.4. Fourier series and derivatives. If one ignores the words "bijectively" and "onto," then the first statement in Theorem 10.5 becomes an easy consequence of the following exercise.

Exercise 10.9. Use integration by parts to show that for every $f \in C^1(\mathbb{T}^n)$, $k = (k_1, \dots, k_n) \in \mathbb{Z}^n$ and $j = 1, \dots, n$,

$$\widehat{\partial_j f}_k = 2\pi i k_j \hat{f}_k.$$

Recall from §4.3 that a **multi-index** in n variables is an n-tuple $\alpha = (\alpha_1, \dots, \alpha_n)$ of nonnegative integers, and we denote its **order** by $|\alpha| := \alpha_1 + \ldots + \alpha_n$. This gives rise to the differential operator

$$\partial^\alpha := \partial_1^{\alpha_1} \ldots \partial_n^{\alpha_n}$$

of order $|\alpha|$ for functions on $\mathbb{R}^n$, as well as a complex-valued polynomial function of $z = (z_1, \dots, z_n) \in \mathbb{C}^n$ defined by

$$z^\alpha := z_1^{\alpha_1} \ldots z_n^{\alpha_n}.$$

For $f \in C^\infty(\mathbb{T}^n)$, repeating the formula in Exercise 10.9 finitely many times now yields

$$\widehat{\partial^\alpha f}_k = (2\pi i k)^\alpha \hat{f}_k \tag{10.10}$$

for any multi-index α.

[19]A linear map $T : \mathcal{H} \to \mathcal{H}'$ between two Hilbert spaces is called **unitary** if it is an isometry, i.e. $\langle Tx, Ty \rangle = \langle x, y \rangle$ for all $x, y \in \mathcal{H}$. Such maps also satisfy $\|Tx\| = \|x\|$ for all $x \in \mathcal{H}$, hence they are continuous.

PROOF OF THEOREM 10.5, PART 1. Assume $f \in C^\infty(\mathbb{T}^n)$, and choose any multi-index α. Since $\partial^\alpha f$ is bounded and $\mathbb{T}^n$ has finite measure, $\partial^\alpha f$ also belongs to $L^1(\mathbb{T}^n)$, implying in light of (10.6) that $\widehat{\partial^\alpha f} \in \ell^\infty(\mathbb{Z}^n)$. The relation (10.10) then implies that

$$k^\alpha \hat{f}_k = \frac{k^\alpha}{(2\pi i)^{|\alpha|} k^\alpha} \widehat{\partial^\alpha f}_k = \frac{1}{(2\pi i)^{|\alpha|}} \widehat{\partial^\alpha f}_k$$

is bounded independently of $k \in \mathbb{Z}^n$. Since this is true for every multi-index α, it follows that $k \mapsto P(k)\hat{f}_k$ is a bounded function $\mathbb{Z}^n \to V$ for every polynomial P, hence $\hat{f} \in \mathscr{S}(\mathbb{Z}^n)$.

We've proved:

$$\mathscr{F}(C^\infty(\mathbb{T}^n)) \subset \mathscr{S}(\mathbb{Z}^n).$$

$\square$

The next exercise is an easy application of the standard theorem on term-by-term differentiation of infinite series—the point is that whenever $g \in \ell^1(\mathbb{Z}^n)$, the partial sums of the series $\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} g(k)$ converge uniformly with respect to $x \in \mathbb{T}^n$.

EXERCISE 10.10. Given a function $g : \mathbb{Z}^n \to V$ and $j \in \{1, \dots, n\}$, let $g_j : \mathbb{Z}^n \to V$ denote the function defined by $g_j(k) := k_j g(k)$ for $k = (k_1, \dots, k_n) \in \mathbb{Z}^n$. Show that if g and g_j both belong to $\ell^1(\mathbb{Z}^n)$, then $\check{g} : \mathbb{T}^n \to V$ is continuous and has a continuous partial derivative $\partial_j \check{g} : \mathbb{T}^n \to V$ given by

$$\partial_j \check{g}(x) = \widecheck{2\pi i g_j}(x).$$

PROOF OF THEOREM 10.5, PART 2. We consider the second statement in the theorem: suppose $g \in \mathscr{S}(\mathbb{Z}^n)$. Then the function $k \mapsto k^\alpha g(k)$ also belongs to $\mathscr{S}(\mathbb{Z}^n)$ for every multi-index α, and is therefore in $\ell^1(\mathbb{Z}^n)$. Iterating the result of Exercise 10.10 finitely many times then proves that for every multi-index α, $\partial^\alpha \check{g}$ exists and is continuous and is given by

$$\partial^\alpha \check{g} = (2\pi i)^{|\alpha|} \check{g}_\alpha, \tag{10.11}$$

where $g_\alpha : \mathbb{Z}^n \to V$ is given by $g_\alpha(k) := k^\alpha g(k)$. In particular, $\check{g} : \mathbb{T}^n \to V$ is smooth.

We've proved:

$$\mathscr{F}^*(\mathscr{S}(\mathbb{Z}^n)) \subset C^\infty(\mathbb{T}^n).$$

$\square$

The main remaining step in the proof of Theorem 10.5 is to show that

$$\mathscr{F}\mathscr{F}^*|_{\mathscr{S}(\mathbb{Z}^n)} = \mathrm{Id}_{\mathscr{S}(\mathbb{Z}^n)} \qquad \text{and} \qquad \mathscr{F}^*\mathscr{F}|_{C^\infty(\mathbb{T}^n)} = \mathrm{Id}_{C^\infty(\mathbb{T}^n)}.$$

We have already proved the first relation, as a consequence of the L^2-orthonormality of the functions φ_k. We shall prove in §10.6 that the relation $\mathscr{F}^*\hat{f} = f$ holds for $f \in C^\infty(\mathbb{T}^n)$. As preparation for the latter, we first need a quick digression on the topic of approximate identities.

10.5. Approximate identities. In §8.4, we considered sequences of smooth functions $\rho_j : \mathbb{R}^n \to [0, \infty)$ that approximate the so-called "Dirac δ-function". In the context of fully periodic functions, the analogous object to $\delta : \mathbb{R}^n \to [0, \infty)$ would be a nonnegative function δ on $\mathbb{T}^n$ that satisfies

$$\int_{\mathbb{T}^n} \varphi(x)\delta(x)\, dx = \varphi(0) \quad \text{for all} \quad \varphi \in C^\infty(\mathbb{T}^n).$$

If such a function existed, it would need to be identically zero on $\mathbb{T}^n \backslash \{0\}$ and have an infinite value at 0, so δ cannot be understood as a function in the classical sense, though one can make sense of it as either a measure or a distribution (i.e. a "generalized function", see §13). What is perhaps more important in many applications is that one can *approximate* it with actual smooth functions.

Definition 10.11. An **approximate identity** on $\mathbb{T}^n$ is a sequence $\rho_j : \mathbb{T}^n \to [0, \infty)$ of nonnegative smooth functions such that for every $\varphi \in C^\infty(\mathbb{T}^n)$,

$$\lim_{j\to\infty} \int_{\mathbb{T}^n} \rho_j(x)\varphi(x)\, dx = \varphi(0).$$

Remark 10.12. The convolution of two functions on $\mathbb{T}^n$ is defined analogously to functions on $\mathbb{R}^n$, by

$$(f * g)(x) := \int_{\mathbb{T}^n} f(x-y)g(y)\, dy,$$

where $x - y \in \mathbb{T}^n = \mathbb{R}^n/\mathbb{Z}^n$ makes sense for $x, y \in \mathbb{T}^n$ since the lattice $\mathbb{Z}^n$ is a subgroup of $\mathbb{R}^n$ with respect to vector addition. One can again use a change of variables to show $f * g = g * f$ (cf. Exercise 8.6). The defining property of an approximate identity thus implies that for any $f \in C^\infty(\mathbb{T}^n)$, $(\rho_j * f)(x) = (f * \rho_j)(x) = \int_{\mathbb{T}^n} f(x-y)\rho_j(y)\, dy \to f(x)$, so

$$\rho_j * f \to f \text{ pointwise for } f \in C^\infty(\mathbb{T}^n). \tag{10.12}$$

The term "approximate identity" refers to the ring structure on $L^1(\mathbb{T}^n)$ defined via the convolution operator. If a δ-function "$\delta := \lim_{j\to\infty} \rho_j$" existed, then it would satisfy $\delta * f = f * \delta = f$ for every smooth function f, thus it would define an identity element in the convolution ring.

The next result describes one of several simple tricks for finding examples of approximate identities.

Proposition 10.13. *Suppose $\rho : \mathbb{T}^n \to [0, \infty)$ is a smooth function satisfying $\rho(0) = 1$ and $\rho(x) < 1$ for all $x \ne 0 \in \mathbb{T}^n$, and for each $j \in \mathbb{N}$, let $c_j := \int_{\mathbb{T}^n} [\rho(x)]^j\, dx > 0$. Then the sequence $\rho_j : \mathbb{T}^n \to [0, \infty)$ defined by*

$$\rho_j(x) := \frac{1}{c_j}[\rho(x)]^j$$

is an approximate identity.

Proof. Let $B_\delta(0) \subset \mathbb{T}^n$ denote the open ball of radius $\delta > 0$ about $0 \in \mathbb{T}^n$. We claim that for every $\delta > 0$,

$$\int_{\mathbb{T}^n \backslash B_\delta(0)} \rho_j(x)\, dx \to 0$$

as $j \to \infty$. Indeed, $\rho < 1$ on the compact set $\mathbb{T}^n \backslash B_\delta(0)$, thus $\rho \leqslant b$ on this set for some constant $b \in (0, 1)$. Choose $a \in (b, 1)$: then since $\rho(0) = 1$, there also exists a $\delta' \in (0, \delta)$ such that $\rho \geqslant a$ on $B_{\delta'}(0)$. This implies

$$c_j = \int_{\mathbb{T}^n} [\rho(x)]^j\, dx \geqslant \int_{B_{\delta'}(0)} [\rho(x)]^j\, dx \geqslant a^j m(B_{\delta'}(0)),$$

thus

$$\begin{aligned}\int_{\mathbb{T}^n \backslash B_\delta(0)} \rho_j(x)\, dx &= \frac{1}{c_j} \int_{\mathbb{T}^n \backslash B_\delta(0)} [\rho(x)]^j\, dx \leqslant \frac{b^j m(\mathbb{T}^n \backslash B_\delta(0))}{a^j m(B_{\delta'}(0))} \\ &= \left(\frac{b}{a}\right)^j \frac{m(\mathbb{T}^n \backslash B_\delta(0))}{m(B_{\delta'}(0))} \to 0.\end{aligned}$$

Now given $f \in C^\infty(\mathbb{T}^n)$ and any $\epsilon > 0$, choose $\delta > 0$ such that $|f(x) - f(0)| < \epsilon$ for all $x \in B_\delta(0)$. Since $\int_{\mathbb{T}^n} \rho_j(x)\,dx = 1$ for all j by construction, we then have

$$\begin{aligned}
\left|\int_{\mathbb{T}^n} \rho_j(x) f(x)\,dx - f(0)\right| &= \left|\int_{\mathbb{T}^n} \rho_j(x)\,[f(x) - f(0)]\,dx\right| \leqslant \int_{\mathbb{T}^n} \rho_j(x)\,|f(x) - f(0)|\,dx \\
&= \int_{B_\delta(0)} \rho_j(x)\,|f(x) - f(0)|\,dx + \int_{\mathbb{T}^n\setminus B_\delta(0)} \rho_j(x)\,|f(x) - f(0)|\,dx \\
&\leqslant \epsilon \int_{B_\delta(0)} \rho_j(x)\,dx + 2\max_{x\in\mathbb{T}^n}|f(x)| \int_{\mathbb{T}^n\setminus B_\delta(0)} \rho_j(x)\,dx \\
&\leqslant \epsilon + 2\max_{x\in\mathbb{T}^n}|f(x)| \int_{\mathbb{T}^n\setminus B_\delta(0)} \rho_j(x)\,dx \to \epsilon \quad \text{as } j \to \infty.
\end{aligned}$$

Since $\epsilon > 0$ can be chosen arbitrarily small, this proves $\int_{\mathbb{T}^n} \rho_j(x) f(x)\,dx \to f(0)$. □

If the δ-function existed, its Fourier coefficients would have to be $\delta_k = \int_{\mathbb{T}^n} e^{-2\pi i k\cdot x}\delta(x)\,dx = 1$ for all $k \in \mathbb{Z}^n$, giving rise to the formal expression

$$\delta(x) = \sum_{k\in\mathbb{Z}^n} e^{2\pi i k\cdot x}. \tag{10.13}$$

Both sides of this formula are nonsense mathematically, but it is worth remembering anyway, as it encapsulates two rigorously provable statements about Fourier series of approximate identities:

LEMMA 10.14. *For any approximate identity $\rho_j : \mathbb{T}^n \to [0,\infty)$, the Fourier coefficients $(\hat\rho_j)_k \in \mathbb{C}$ satisfy a uniform bound $|(\hat\rho_j)_k| \leqslant C$ for some constant $C > 0$ independent of $j \in \mathbb{N}$ and $k \in \mathbb{Z}^n$, and $\lim_{j\to\infty}(\hat\rho_j)_k = 1$ for all k.*

PROOF. The convergence $(\hat\rho_j)_k \to 1$ as $j \to \infty$ follows immediately from the formula $(\hat\rho_j)_k = \int_{\mathbb{T}^n} e^{-2\pi i k\cdot x}\rho_j(x)\,dx$ and the defining property of an approximate identity. In particular for $k = 0 \in \mathbb{Z}^n$, we have $\lim_{j\to\infty}(\hat\rho_j)_0 = 1$, so there exists a bound $(\hat\rho_j)_0 = \int_{\mathbb{T}^n} \rho_j(x)\,dx \leqslant C$ independent of j. Then

$$|(\hat\rho_j)_k| \leqslant \int_{\mathbb{T}^n} \left|e^{-2\pi i k\cdot x}\rho_j(x)\right|\,dx \leqslant \int_{\mathbb{T}^n} \rho_j(x)\,dx \leqslant C$$

holds for every $j \in \mathbb{N}$ and $k \in \mathbb{Z}^n$. □

LEMMA 10.15. *There exists an approximate identity $\rho_j : \mathbb{T}^n \to [0,\infty)$ that is equal to its own Fourier series for every j, i.e. it satisfies $\mathscr{F}^*\mathscr{F}\rho_j = \rho_j$.*

PROOF. Define $\beta : \mathbb{T}^1 \to [0,\infty)$ by $\beta(t) := \frac{\cos(2\pi t)+1}{2}$ and $\rho : \mathbb{T}^n \to [0,\infty)$ by

$$\rho(x_1,\dots,x_n) := \beta(x_1)\dots\beta(x_n),$$

and let ρ_j denote the approximate identity described by Proposition 10.13 in terms of this particular choice of ρ. Since β is a complex linear combination of $e^{2\pi i t}$ and $e^{-2\pi i t}$, ρ is a finite linear combination of functions from the orthonormal set $\{\varphi_k\}_{k\in\mathbb{Z}^n}$, and the same is therefore true of all its powers $[\rho(x)]^j$ for $j \in \mathbb{N}$. This proves that ρ_j is equal to a finite Fourier series for every j. □

10.6. Completeness. We are now ready to prove that $\mathscr{F}^*\hat f = f$ for every $f \in C^\infty(\mathbb{T}^n)$. Let us first describe an informal "physicist's version" of the proof, in which we refuse to worry about annoying analytical issues like integrability, convergence, and whether the δ-function actually exists. The main tool needed for this is Fubini's theorem, which we apply for functions on $\mathbb{T}^n \times \mathbb{Z}^n$

with the product of our Lebesgue-type measure m on $\mathbb{T}^n$ with the counting measure ν on $\mathbb{Z}^n$. For $f \in C^\infty(\mathbb{T}^n)$, we compute:

$$\begin{aligned}
(\mathscr{F}^* \widehat{f})(x) &= \sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} \widehat{f}_k \, dx = \int_{\mathbb{Z}^n} e^{2\pi i k \cdot x} \left(\int_{\mathbb{T}^n} e^{-2\pi i k \cdot y} f(y) \, dy \right) d\nu(k) \\
&= \int_{\mathbb{T}^n \times \mathbb{Z}^n} e^{2\pi i k \cdot (x-y)} f(y) \, d(m(y) \otimes \nu(k)) = \int_{\mathbb{T}^n} \left(\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot (x-y)} f(y) \right) dy \\
&= \int_{\mathbb{T}^n} \left(\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot (x-y)} \right) f(y) \, dy = \int_{\mathbb{T}^n} \delta(x-y) f(y) \, dy = \int_{\mathbb{T}^n} f(x-y) \delta(y) \, dy \\
&= f(x).
\end{aligned}$$

Several steps in this derivation are formal manipulations that cannot be taken literally. The interchange of the integral and the summation is meant to be a result of applying Fubini's theorem to the function $(y,k) \mapsto e^{2\pi i k \cdot (x-y)} f(y)$ on $\mathbb{T}^n \times \mathbb{Z}^n$, though unfortunately, the latter is not $(m \otimes \nu)$-integrable. The δ-function then appears due to (10.13), and from there we apply a straightforward change of variables followed by the defining property of the δ-function.

The way to make all this mathematically precise is by introducing the Fourier coefficients of an approximate identity ρ_j in the second line. This will make the function on $\mathbb{T}^n \times \mathbb{Z}^n$ integrable and thus produce a mathematically correct formula, which converges to the desired formula as $j \to \infty$.

Proof of Theorem 10.5, part 3. Assume $f \in C^\infty(\mathbb{T}^n)$ is given. By Lemma 10.15, we can choose an approximate identity $\rho_j : \mathbb{T}^n \to [0, \infty)$ that equals its own Fourier series for every j, and by Lemma 10.14, its Fourier coefficients are uniformly bounded and converge to 1. Since ρ_j is smooth, the function $k \mapsto (\widehat{\rho}_j)_k$ on $\mathbb{Z}^n$ belongs to $\mathscr{S}(\mathbb{Z}^n) \subset \ell^1(\mathbb{Z}^n)$, so that the function $F : \mathbb{T}^n \times \mathbb{Z}^n \to V$ given by

$$F(y,k) := e^{2\pi i k \cdot x} e^{-2\pi i k \cdot y} (\widehat{\rho}_j)_k f(y)$$

satisfies

$$|F(y,k)| \leqslant |(\widehat{\rho}_j)_k| \cdot |f(y)|,$$

and is therefore $(m \otimes \nu)$-integrable as a consequence of Fubini's theorem for nonnegative measurable functions. We can then apply Fubini's theorem for integrable functions, giving

$$\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} (\widehat{\rho}_j)_k \left(\int_{\mathbb{T}^n} e^{-2\pi i k \cdot y} f(y) \, dy \right) = \int_{\mathbb{T}^n} \left(\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot (x-y)} (\widehat{\rho}_j)_k \right) f(y) \, dy.$$

The left hand side of this expression is $\sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} (\widehat{\rho}_j)_k \widehat{f}_k$, and since $\widehat{f} \in \mathscr{S}(\mathbb{Z}^n) \subset \ell^1(\mathbb{Z}^n)$, Lemma 10.14 implies via the dominated convergence theorem (applied on $\mathbb{Z}^n$ with the counting measure) that this converges to $(\mathscr{F}^* \widehat{f})(x)$ as $j \to \infty$. Since each ρ_j is equal to its Fourier series, the right hand side is

$$\int_{\mathbb{T}^n} \rho_j(x-y) f(y) \, dy = (\rho_j * f)(x),$$

which converges in turn to $f(x)$ by (10.12).

We've proved:

$$\mathscr{F}^* \mathscr{F} f = f \quad \text{for all} \quad f \in C^\infty(\mathbb{T}^n).$$

□

Proof of Theorem 10.5, conclusion. The first three statements in the theorem have already been established, so it only remains to verify that for $g \in \mathscr{S}(\mathbb{Z}^n)$, the Fourier series

$\sum_{k\in\mathbb{Z}^n} e^{2\pi i k\cdot x} g(k)$ converges uniformly and so do its derivatives of all orders. The uniform convergence is clear since $\mathscr{S}(\mathbb{Z}^n) \subset \ell^1(\mathbb{Z}^n)$. Applying an arbitrary differential operator ∂^α to the terms of the series changes the coefficients to $(-2\pi i k)^\alpha g(k)$, and this function of k is still in $\mathscr{S}(\mathbb{Z}^n)$, so the resulting series also converges uniformly. □

10.7. Parseval's identity. The proof of Theorem 10.6 is based mainly on the observation that $\mathscr{F}$ and $\mathscr{F}^*$ are adjoint operations.

LEMMA 10.16. *For every $f \in C^\infty(\mathbb{T}^n)$ and $g \in \mathscr{S}(\mathbb{Z}^n)$,*

$$\langle g, \mathscr{F} f\rangle_{\ell^2} = \langle \mathscr{F}^* g, f\rangle_{L^2}.$$

PROOF. We again use Fubini's theorem for a function on $\mathbb{T}^n \times \mathbb{Z}^n$ with the product measure $m \otimes \nu$:

$$\begin{aligned}
\langle g, \mathscr{F} f\rangle_{\ell^2} &= \sum_{k\in\mathbb{Z}^n} \left\langle g(k), \int_{\mathbb{T}^n} e^{-2\pi i k\cdot x} f(x)\, dx \right\rangle = \sum_{k\in\mathbb{Z}^n} \left(\int_{\mathbb{T}^n} \langle g(k), e^{-2\pi i k\cdot x} f(x)\rangle\, dx \right) \\
&= \int_{\mathbb{T}^n\times\mathbb{Z}^n} e^{-2\pi i k\cdot x} \langle g(k), f(x)\rangle\, d(m(x)\otimes\nu(k)) \\
&= \int_{\mathbb{T}^n} \left(\sum_{k\in\mathbb{Z}^n} \langle e^{2\pi i k\cdot x} g(k), f(x)\rangle \right) dx = \int_{\mathbb{T}^n} \left\langle \sum_{k\in\mathbb{Z}^n} e^{2\pi i k\cdot x} g(k), f(x) \right\rangle dx \\
&= \int_{\mathbb{T}^n} \langle \check{g}(x), f(x)\rangle\, dx = \langle \mathscr{F}^* g, f\rangle_{L^2}.
\end{aligned}$$

Here the use of Fubini's theorem is justified since f is smooth and g is rapidly decreasing, so $(x,k) \mapsto |e^{-2\pi i k\cdot x}\langle g(k), f(x)\rangle| \leqslant |g(k)| \cdot |f(x)|$ defines an integrable function on $\mathbb{T}^n \times \mathbb{Z}^n$. □

PROOF OF THEOREM 10.6. For $f, g \in C^\infty(\mathbb{T}^n)$, we have $\hat{f}, \hat{g} \in \mathscr{S}(\mathbb{Z}^n)$, so Lemma 10.16 and the fact that $\mathscr{F}$ and $\mathscr{F}^*$ are inverses gives

$$\langle \hat{f}, \hat{g}\rangle_{\ell^2} = \langle \mathscr{F}^* f, \mathscr{F}^* g\rangle_{\ell^2} = \langle f, \mathscr{F}\mathscr{F}^* g\rangle_{L^2} = \langle f, g\rangle_{L^2}.$$

□

11. The Fourier transform

11.1. The Fourier transform on the Schwartz space. In this section we again assume $(V, \langle\ ,\ \rangle)$ is a finite-dimensional complex inner product space, but we now consider functions $f : \mathbb{R}^n \to V$ that are not periodic. One cannot expect these to be expressible in terms of the fully periodic functions $\varphi_k(x) := e^{2\pi i k\cdot x}$ for $k \in \mathbb{Z}^n$. On the other hand, if the periodicity condition is dropped, then the oscillatory function φ_k is well defined on $\mathbb{R}^n$ for *every* $k \in \mathbb{R}^n$, and it is natural to wonder whether arbitrary functions on $\mathbb{R}^n$ can be regarded in some sense as linear combinations of oscillatory functions of this type. Since k can now take uncountably many distinct values, our notion of a "linear combination" needs to be expanded for this discussion: instead of trying to write $f(x)$ as a series, we would now like to write it as an integral

$$f(x) = \int_{\mathbb{R}^n} e^{2\pi i p\cdot x} \hat{f}(p)\, dp \tag{11.1}$$

for some function $\hat{f} : \mathbb{R}^n \to V$, called the **Fourier transform** of f. Our discussion of the Fourier transform in this section will closely parallel that of the Fourier series, but it is in some respects more elegant, as the theory of the Fourier transform exhibits a certain symmetry that is lacking in the periodic case. This is evident when one sees the formulas for the transformations $\mathscr{F}$ and $\mathscr{F}^*$, each of which converts a function $\mathbb{R}^n \to V$ into another function $\mathbb{R}^n \to V$: for any class of functions

$f, g : \mathbb{R}^n \to V$ such that the following integrals converge, we define the **Fourier transform** of f and **Fourier inverse transform** of g respectively by[20]

$$(\mathscr{F} f)(p) := \hat{f}(p) := \int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} f(x)\, dx, \tag{11.2}$$

and

$$(\mathscr{F}^* g)(x) := \check{g}(x) := \int_{\mathbb{R}^n} e^{2\pi i p \cdot x} g(p)\, dp. \tag{11.3}$$

Both are clearly well defined if $f, g \in L^1(\mathbb{R}^n)$, in which case $\hat{f}$ and $\check{g}$ are both bounded functions; in fact, one can easily show via Theorem 4.5 that in this case they are continuous, so that $\mathscr{F}$ and $\mathscr{F}^*$ each define bounded linear operators

$$\mathscr{F}, \mathscr{F}^* : L^1(\mathbb{R}^n) \to C^0_b(\mathbb{R}^n).$$

Recall from §4.3 that $C^0_b(\mathbb{R}^n)$ is the Banach space of *bounded* continuous functions on $\mathbb{R}^n$, with the usual sup-norm[21]

$$\|f\|_{C^0} := \sup_{x \in \mathbb{R}^n} |f(x)|.$$

Before we can discuss in what sense these two operators are inverse to each other, we must introduce suitable function spaces on which they will both be bijective. In the setting of Fourier series, this role was played by the spaces $C^\infty(\mathbb{T}^n)$ and $\mathscr{S}(\mathbb{Z}^n)$. In the present setting, we need a single space of functions on $\mathbb{R}^n$ that combines features of both of these.

DEFINITION 11.1. The **Schwartz space** $\mathscr{S}(\mathbb{R}^n)$, also known as the space of **smooth and rapidly decreasing** functions, consists of all smooth functions $f : \mathbb{R}^n \to V$ with the property that for every pair of multi-indices α and β, the function $\mathbb{R}^n \to V$ given by $x^\alpha \partial^\beta f(x)$ is bounded.

EXERCISE 11.2. Show that a smooth function f on $\mathbb{R}^n$ belongs to $\mathscr{S}(\mathbb{R}^n)$ if and only if for every multi-index α and every $k \in \mathbb{N}$, there exists a constant $C > 0$ dependent on α and k such that

$$|\partial^\alpha f(x)| \leqslant \frac{C}{1 + |x|^k} \quad \text{for all} \quad x \in \mathbb{R}^n.$$

EXERCISE 11.3. Show that $\mathscr{S}(\mathbb{R}^n) \subset L^p(\mathbb{R}^n)$ for every $p \in [1, \infty]$, and for every $f \in \mathscr{S}(\mathbb{R}^n)$ and every multi-index α, the functions $\partial^\alpha f$ and $x \mapsto x^\alpha f(x)$ also belong to $\mathscr{S}(\mathbb{R}^n)$.

The next two theorems are the main results we aim to prove in this section about the Fourier transform.

THEOREM 11.4. *The transformations $\mathscr{F}$ and $\mathscr{F}^*$ each map $\mathscr{S}(\mathbb{R}^n)$ bijectively to itself, and they are inverse to each other.*

THEOREM 11.5 (Plancherel's theorem). *For every $f, g \in \mathscr{S}(\mathbb{R}^n)$, $\langle f, g \rangle_{L^2} = \langle \hat{f}, \hat{g} \rangle_{L^2}$.*

[20] The literature contains several differing opinions on where the factor of 2π should appear in (11.2) and (11.3). Our convention is the same as in [**LL01**, **DM72**], but many books omit it from the exponent, at the cost of having to insert some power $1/2\pi$ (depending on the dimension) in front of one or both integrals. A professor of mine once told of a lecture on Fourier analysis in which the speaker had solved this problem right at the beginning by saying, "Let $2\pi = 1$."

[21] Unlike the norm for continuous functions on the compact space $\mathbb{T}^n$, the supremum in the definition of $\|f\|_{C^0}$ need not be achieved for continuous functions on $\mathbb{R}^n$, and $C^0_b(\mathbb{R}^n)$ does not contain *all* continuous functions on $\mathbb{R}^n$, since not all of them are bounded.

In particular, the linear operators $\mathscr{F}, \mathscr{F}^* : \mathscr{S}(\mathbb{R}^n) \to \mathscr{S}(\mathbb{R}^n)$ are isometries (and are therefore continuous) with respect to the L^2-norm. The Schwartz space contains the space of smooth compactly supported functions, which is dense in $L^2(\mathbb{R}^n)$, thus $\mathscr{S}(\mathbb{R}^n)$ is also dense in $L^2(\mathbb{R}^n)$, so this result implies:

COROLLARY 11.6. *The operators $\mathscr{F}, \mathscr{F}^* : \mathscr{S}(\mathbb{R}^n) \to \mathscr{S}(\mathbb{R}^n)$ admit unique extensions to unitary isomorphisms $L^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ such that $\mathscr{F}^* = \mathscr{F}^{-1}$.* □

PROPOSITION 11.7. *For $f \in L^1(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$, the definitions of $\mathscr{F}f$ and $\mathscr{F}^*f$ in (11.2) and (11.3) respectively agree (up to equality almost everywhere) with their definitions as described in Corollary 11.6 via Plancherel's theorem and the density of $\mathscr{S}(\mathbb{R}^n) \subset L^2(\mathbb{R}^n)$.*

PROOF. To avoid confusion, let us denote by $\hat{f}_{L^1}$ and $\hat{f}_{L^2}$ the two possible definitions of $\hat{f}$ as defined via (11.2) or via the density of $\mathscr{S}(\mathbb{R}^n) \subset L^2(\mathbb{R}^n)$. We claim first that there exists a sequence of smooth compactly supported functions $f_j \in C_0^\infty(\mathbb{R}^n)$ that converge to f in both the L^1- and L^2-norms. Indeed, choosing an approximate identity $\rho_j : \mathbb{R}^n \to [0, \infty)$ with shrinking support as in §8, the smooth functions $\rho_j * f$ converge to f in both L^1 and L^2 according to Theorem 8.14, and one can then define f_j by multiplying these by suitable compactly supported cutoff functions as in the discussion preceding Corollary 8.2. With this sequence chosen, the functions $f_j \in C_0^\infty(\mathbb{R}^n)$ also belong to $\mathscr{S}(\mathbb{R}^n)$, so the L^2-convergence $f_j \to f$ implies that $\hat{f}_j$ converges in L^2 to $\hat{f}_{L^2}$, and it follows that $\hat{f}_j$ also has a subsequence converging pointwise almost everywhere to $\hat{f}_{L^2}$. But since $\mathscr{F} : L^1(\mathbb{R}^n) \to C_b^0(\mathbb{R}^n)$ as defined by (11.2) is a bounded linear map, the L^1-convergence $f_j \to f$ implies additionally that $\hat{f}_j$ converges uniformly to the continuous function $\hat{f}_{L^1}$. This can only be true if $\hat{f}_{L^1} = \hat{f}_{L^2}$ almost everywhere. A completely analogous argument works for $\mathscr{F}^*$. □

For a function $f \in L^2(\mathbb{R}^n)$ that is not in $L^1(\mathbb{R}^n)$, the formula for $\hat{f}$ in (11.2) does not strictly make sense, because the integral does not converge, but the continuity of $\mathscr{F} : L^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ means that one can define $\hat{f}$ as the L^2-limit of the L^2-Cauchy sequence $\hat{f}_j \in C_b^0(\mathbb{R}^n)$ for any sequence $f_j \in L^1(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$ converging in L^2 to f. Exercise 11.9 below describes a reasonable trick for carrying this out in practice.

REMARK 11.8. If $f \in L^2(\mathbb{R}^n)$ but $f \notin L^1(\mathbb{R}^n)$, then $\hat{f}$ and $\check{f}$ are not functions, strictly speaking, but rather equivalence classes of functions up to equality almost everywhere, so their values $\hat{f}(p)$ and $\check{f}(p)$ at individual points $p \in \mathbb{R}^n$ are not well defined. In contrast, $\hat{f}(p)$ and $\check{f}(p)$ are well defined for every $p \in \mathbb{R}^n$ via the integrals (11.2) or (11.3) if $f \in L^1(\mathbb{R}^n)$.

EXERCISE 11.9. Show that for $f, g \in L^2(\mathbb{R}^n)$, the following conditions are equivalent:

(1) $\hat{f} = g$ almost everywhere;
(2) There exists a sequence $R_j \to \infty$ such that $\lim_{j\to\infty} \int_{B_{R_j}} e^{-2\pi i p\cdot x} f(x)\, dx = g(p)$ for almost every $p \in \mathbb{R}^n$. (Here $B_R \subset \mathbb{R}^n$ denotes the ball of radius R about the origin.)

Hint: Multiply f by characteristic functions to define L^2-close approximations that are also in $L^1(\mathbb{R}^n)$.

11.2. Fourier transforms and derivatives. For $f \in L^1(\mathbb{R}^n)$, the function $e^{-2\pi p\cdot x} f(x)$ is continuous with respect to p and also, as a function of x, bounded for every $p \in \mathbb{R}^n$ by the fixed Lebesgue-integrable function $|f| : \mathbb{R}^n \to [0, \infty)$. Viewing $\hat{f}(p)$ as a parameter-dependent integral and applying Theorem 4.5 (and similarly for $\check{f}(x)$) thus proves:

PROPOSITION 11.10. *If $f \in L^1(\mathbb{R}^n)$, then $\mathscr{F}f$ and $\mathscr{F}^*f$ belong to $C_b^0(\mathbb{R}^n)$, and the resulting maps $\mathscr{F}, \mathscr{F}^* : L^1(\mathbb{R}^n) \to C_b^0(\mathbb{R}^n)$ are bounded linear operators.* □

EXERCISE 11.11. Use integration by parts and/or Theorem 4.5 to establish the following analogues of Exercises 10.9 and 10.10:

(a) Suppose $f \in L^1(\mathbb{R}^n)$, f has a continuous partial derivative $\partial_j f : \mathbb{R}^n \to V$ that also belongs to $L^1(\mathbb{R}^n)$ for some $j = 1, \dots, n$, and f "decays at infinity" in the sense that $\lim_{R\to\infty} \sup_{x\in\mathbb{R}^n\setminus B_R} |f(x)| = 0$, where $B_R \subset \mathbb{R}^n$ denotes the ball of radius R about $0 \in \mathbb{R}^n$. Then
$$\widehat{\partial_j f}(p) = 2\pi i p_j \hat{f}(p) \quad \text{and} \quad \widecheck{\partial_j f}(x) = -2\pi i x_j \check{f}(x)$$
for each $p = (p_1, \dots, p_n) \in \mathbb{R}^n$ and $x = (x_1, \dots, x_n) \in \mathbb{R}^n$.

(b) Given $f : \mathbb{R}^n \to V$ and $j \in \{1, \dots, n\}$, let $f_j : \mathbb{R}^n \to V$ denote the function $f_j(x) := x_j f(x)$ for $x = (x_1, \dots, x_n) \in \mathbb{R}^n$. If f and f_j both belong to $L^1(\mathbb{R}^n)$, then $\check{f}, \hat{f} : \mathbb{R}^n \to V$ are continuous and have continuous partial derivatives $\partial_j \check{f}, \partial_j \hat{f} : \mathbb{R}^n \to V$ given by
$$\partial_j \check{f}(x) = \widecheck{2\pi i f_j}(x) \quad \text{and} \quad \partial_j \hat{f}(p) = \widehat{-2\pi i f_j}(p).$$

If $f \in \mathscr{S}(\mathbb{R}^n)$, then the conditions in both parts of Exercise 11.11 are satisfied and the formulas may be iterated arbitrarily many times, proving that for every multi-index α,

$$\begin{aligned} \widehat{\partial^\alpha f}(p) &= (2\pi i p)^\alpha \hat{f}(p), & \widecheck{\partial^\alpha f}(x) &= (-2\pi i x)^\alpha \check{f}(x), \\ \partial^\alpha \check{f}(x) &= (2\pi i)^{|\alpha|} \widecheck{f_\alpha}(x), & \partial^\alpha \hat{f}(p) &= (-2\pi i)^{|\alpha|} \widehat{f_\alpha}(p), \end{aligned} \tag{11.4}$$

where $f_\alpha(x) := x^\alpha f(x)$. Implicit in the last two formulas is that $\partial^\alpha \check{f}$ and $\partial^\alpha \hat{f}$ exist for every α, i.e. in this case, $\check{f}$ and $\hat{f}$ are also smooth.

PROOF OF THEOREM 11.4, PART 1. For $f \in \mathscr{S}(\mathbb{R}^n)$, we have already shown above that $\hat{f}$ is smooth, and for each pair of multi-indices α, β $\partial^\alpha \hat{f}$ satisfies
$$p^\beta \partial^\alpha \hat{f}(p) = p^\beta (-2\pi i)^{|\alpha|} \widehat{f_\alpha}(p) = \frac{(-2\pi i)^{|\alpha|}}{(2\pi i)^{|\beta|}} (2\pi i p)^\beta \widehat{f_\alpha}(p) = \frac{(-2\pi i)^{|\alpha|}}{(2\pi i)^{|\beta|}} \widehat{\partial^\beta f_\alpha}(p).$$
By the definition of the Schwartz space, f_α and $\partial^\beta f_\alpha$ also belong to $\mathscr{S}(\mathbb{R}^n)$, so in particular, the latter is in $L^1(\mathbb{R}^n)$ and its Fourier transform is therefore bounded. This proves $\hat{f} \in \mathscr{S}(\mathbb{R}^n)$. One shows in the same manner that $\check{f} \in \mathscr{S}(\mathbb{R}^n)$, so we have proved:
$$\mathscr{F}(\mathscr{S}(\mathbb{R}^n)) \subset \mathscr{S}(\mathbb{R}^n) \quad \text{and} \quad \mathscr{F}^*(\mathscr{S}(\mathbb{R}^n)) \subset \mathscr{S}(\mathbb{R}^n).$$
□

11.3. The Gaussian. One class of functions in $\mathscr{S}(\mathbb{R}^n)$ whose Fourier transforms can be computed explicitly are the *Gaussians*, i.e. functions of the form $Ae^{-c|x|^2}$ for constants $A, c > 0$. The computation carried out in this subsection is more than just an amusing exercise: the proof of the inversion formula in §11.5 will require an approximate identity with particular properties, and Gaussians furnish the most convenient construction of such an object.

PROPOSITION 11.12. *For any constant $a > 0$, the function $f(x) := e^{-a^2|x|^2}$ on $\mathbb{R}^n$ satisfies*
$$\hat{f}(x) = \check{f}(x) = \frac{\pi^{n/2}}{a^n} e^{-(\pi/a)^2|x|^2}.$$

PROOF. By Fubini's theorem,
$$\begin{aligned} \hat{f}(p) &= \int_{\mathbb{R}^n} e^{-a^2(x_1^2+\dots+x_n^2)} e^{-2\pi i(p_1x_1+\dots+p_nx_n)}\, dx = \int_{\mathbb{R}^n} e^{-a^2x_1^2} \dots e^{-a^2x_n^2} e^{-2\pi i p_1 x_1} \dots e^{-2\pi i p_n x_n}\, dx \\ &= \left(\int_{-\infty}^{\infty} e^{-a^2x_1^2} e^{-2\pi i p_1 x_1}\, dx_1 \right) \dots \left(\int_{-\infty}^{\infty} e^{-a^2x_n^2} e^{-2\pi i p_n x_n}\, dx_n \right), \end{aligned}$$

thus it will suffice to prove that the stated formula for $\hat{f}$ is correct in the case $n = 1$. Consider $f(x) := e^{-a^2x^2}$ on $\mathbb{R}$. Instead of computing the integral

$$\hat{f}(p) = \int_{-\infty}^{\infty} e^{-a^2x^2} e^{-2\pi ipx}\, dx$$

explicitly for every $p \in \mathbb{R}$, we shall identify the function $\hat{f}$ as the unique solution to a certain initial value problem. For $p = 0$, we have

$$\hat{f}(0) = \int_{-\infty}^{\infty} e^{-a^2x^2}\, dx = \frac{1}{a}\int_{-\infty}^{\infty} e^{-u^2}\, du = \frac{\sqrt{\pi}}{a},$$

which follows via the substitution $u = ax$ and the well-known formula $\int_{-\infty}^{\infty} e^{-u^2}\, du = \sqrt{\pi}$. Applying (11.4) and then integrating by parts, we also have

$$\begin{aligned}\hat{f}'(p) &= -2\pi i\widehat{xf}(p) = -2\pi i \int_{-\infty}^{\infty} xe^{-a^2x^2}e^{-2\pi ipx}\, dx = \frac{i\pi}{a^2}\int_{-\infty}^{\infty} \frac{d}{dx}\left(e^{-a^2x^2}\right)\cdot e^{-2\pi ipx}\, dx \\ &= -\frac{i\pi}{a^2}\int_{-\infty}^{\infty} e^{-a^2x^2}\frac{d}{dx}\left(e^{-2\pi ipx}\right)\, dx = \frac{i\pi}{a^2}\cdot 2\pi ip\int_{-\infty}^{\infty} e^{-a^2x^2}e^{-2\pi ipx}\, dx = -\frac{2\pi^2}{a^2}p\hat{f}(p),\end{aligned}$$

in other words, $\hat{f} : \mathbb{R} \to \mathbb{C}$ satisfies the initial value problem

$$\begin{cases} \dfrac{d\hat{f}}{dp} & = -2(\pi/a)^2 p\hat{f}, \\ \hat{f}(0) & = \sqrt{\pi}/a. \end{cases}$$

The unique solution to this problem is $\hat{f}(p) = \frac{\sqrt{\pi}}{a}e^{-(\pi/a)^2p^2}$.

Since f is a real-valued function, $\check{f}$ is the complex conjugate of $\hat{f}$, which is $\hat{f}$ itself. □

COROLLARY 11.13. *The Gaussian $f(x) = e^{-a^2|x|^2}$ with $a > 0$ satisfies $\mathscr{F}^*\mathscr{F}f = \mathscr{F}\mathscr{F}^*f = f$.* □

11.4. Approximate identities revisited. If the Dirac δ-function $\delta : \mathbb{R}^n \to [0, \infty]$ were an actual function in $\mathscr{S}(\mathbb{R}^n)$, its Fourier transform would be

$$\hat{\delta}(p) = \int_{\mathbb{R}^n} e^{-2\pi ip\cdot x}\delta(x)\, dx = 1,$$

leading to the slightly nonsensical formula

$$\delta(x) = \int_{\mathbb{R}^n} e^{2\pi ip\cdot x}\, dp. \tag{11.5}$$

As with most things involving the δ-function, one can make mathematical sense of this formula in terms of approximate identities, and the proof of the Fourier inversion formula in the next subsection will require the existence of an approximate identity for which the inversion formula is known to hold. For our purposes in this context, "approximate identity" means the following:

DEFINITION 11.14. A **tempered approximate identity** on $\mathbb{R}^n$ is a sequence $\rho_j : \mathbb{R}^n \to [0, \infty)$ of nonnegative functions in $\mathscr{S}(\mathbb{R}^n)$ such that for every $\varphi \in \mathscr{S}(\mathbb{R}^n)$,

$$\lim_{j\to\infty}\int_{\mathbb{R}^n} \rho_j(x)\varphi(x)\, dx = \varphi(0). \tag{11.6}$$

Note that the assumption $\rho_j \in \mathscr{S}(\mathbb{R}^n)$ implies that the convolution $\rho_j * f$ is a well-defined function $\mathbb{R}^n \to V$ for every $f \in \mathscr{S}(\mathbb{R}^n)$, and (11.6) then implies $(\rho_j * f)(x) = (f * \rho_j)(x) = \int_{\mathbb{R}^n} f(x-y)\rho_j(y)\,dy \to f(x)$, hence

$$\rho_j * f \to f \text{ pointwise for } f \in \mathscr{S}(\mathbb{R}^n). \tag{11.7}$$

LEMMA 11.15. *Suppose $\rho : \mathbb{R}^n \to [0,\infty)$ is a smooth function satisfying $\int_{\mathbb{R}^n} \rho(x)\,dx = 1$. Then the sequence $\rho_j : \mathbb{R}^n \to [0,\infty)$ defined by*

$$\rho_j(x) := j^n \rho(jx)$$

satisfies (11.6) for every bounded continuous function $\varphi : \mathbb{R}^n \to V$.

PROOF. We use the change of variables $y := jx$ to write

$$\int_{\mathbb{R}^n} \rho_j(x)\varphi(x)\,dx = \int_{\mathbb{R}^n} \rho(y)\varphi(y/j)\,dy.$$

Since φ is bounded and continuous, the integrands on the right converge pointwise as $j \to \infty$ to $\varphi(0)\rho$ and are uniformly bounded by a constant multiple of the integrable function ρ. The dominated convergence theorem thus implies that the integrals converge to $\int_{\mathbb{R}^n} \varphi(0)\rho\,dm = \varphi(0)$. □

LEMMA 11.16. *There exists a tempered approximate identity $\rho_j \in \mathscr{S}(\mathbb{R}^n)$ with the following properties:*

(1) $\mathscr{F}^*\widehat{\rho}_j = \rho_j$ for every j;
(2) The functions $\widehat{\rho}_j$ satisfy a uniform bound $|\widehat{\rho}_j| \leqslant C$ for all j and converge pointwise to 1.

PROOF. Set $\rho(x) := \frac{1}{\sqrt{\pi}}e^{-|x|^2}$ and use this to define ρ_j as in Lemma 11.15. Then ρ_j is a Gaussian for every j, so both ρ_j and $\widehat{\rho}_j$ are in $\mathscr{S}(\mathbb{R}^n)$, and Corollary 11.13 implies $\mathscr{F}^*\widehat{\rho}_j = \rho_j$. Applying Lemma 11.15 with the bounded continuous function $f(x) = e^{-2\pi i p\cdot x}$ for each $p \in \mathbb{R}^n$, we also have

$$\lim_{j\to\infty} \widehat{\rho}_j(p) = \lim_{j\to\infty} \int_{\mathbb{R}^n} e^{-2\pi i p\cdot x}\rho_j(x)\,dx = 1$$

and

$$|\widehat{\rho}_j(p)| \leqslant \|\rho_j\|_{L^1} = \|\rho\|_{L^1} = 1,$$

where a quick computation via change of variables gives $\int_{\mathbb{R}^n} \rho_j\,dm = \int_{\mathbb{R}^n} \rho\,dm$. Alternatively, these last two statements also follow from the explicit computation of $\widehat{\rho}_j$ in Proposition 11.12. □

11.5. The Fourier inversion formula. We can now prove that the operators $\mathscr{F}$ and $\mathscr{F}^*$ on $\mathscr{S}(\mathbb{R}^n)$ are inverse to each other.

The "physicist's proof" that $\mathscr{F}^*\widehat{f} = f$ works via the following adventurous application of Fubini's theorem:

$$\begin{aligned}(\mathscr{F}^*\widehat{f})(x) &= \int_{\mathbb{R}^n} e^{2\pi i p\cdot x}\widehat{f}(p)\,dp = \int_{\mathbb{R}^n} e^{2\pi i p\cdot x}\left(\int_{\mathbb{R}^n} e^{-2\pi i p\cdot y} f(y)\,dy\right)dp \\ &= \int_{\mathbb{R}^n\times\mathbb{R}^n} e^{2\pi i p\cdot(x-y)} f(y)\,dy\,dp = \int_{\mathbb{R}^n}\left(\int_{\mathbb{R}^n} e^{2\pi i p\cdot(x-y)} f(y)\,dp\right)dy \\ &= \int_{\mathbb{R}^n}\left(\int_{\mathbb{R}^n} e^{2\pi i p\cdot(x-y)}\,dp\right) f(y)\,dy = \int_{\mathbb{R}^n} \delta(x-y)f(y)\,dy = \int_{\mathbb{R}^n} f(x-y)\delta(y)\,dy \\ &= f(x).\end{aligned}$$

Here the δ-function appears due to the formal relation (11.5), and something clearly must be modified to justify the use of Fubini's theorem since $(y,p) \mapsto e^{2\pi i p\cdot(x-y)}f(y)$ is not an integrable function on $\mathbb{R}^n \times \mathbb{R}^n$ for any $x \in \mathbb{R}^n$. In analogy with §10.6, the remedy is to multiply this function by the Fourier transform of a tempered approximate identity $\rho_j \in \mathscr{S}(\mathbb{R}^n)$, and then take the limit of the resulting relation as $j \to \infty$.

Proof of Theorem 11.4, conclusion. Given $f \in \mathscr{S}(\mathbb{R}^n)$, we need to show $\mathscr{F}^* \hat{f} = f$. Choose a tempered approximate identity ρ_j with the properties listed in Lemma 11.16. We then have $\hat{\rho}_j \in L^1(\mathbb{R}^n)$, and for every $x \in \mathbb{R}^n$, the function $F : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{C}$ given by

$$F(y,p) := e^{2\pi i p\cdot x} e^{-2\pi i p\cdot y} \hat{\rho}_j(p) f(y)$$

is therefore integrable. Applying Fubini's theorem to the integral of F over $\mathbb{R}^n \times \mathbb{R}^n$ now gives

$$\int_{\mathbb{R}^n} e^{2\pi i p\cdot x} \hat{\rho}_j(p) \left(\int_{\mathbb{R}^n} e^{-2\pi i p\cdot y} f(y) dy \right) dp = \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^n} e^{2\pi i p\cdot(x-y)} \hat{\rho}_j(p)\, dp \right) f(y)\, dy.$$

The left hand side is $\int_{\mathbb{R}^n} e^{2\pi i p\cdot x} \hat{\rho}_j(p) \hat{f}(p)\, dp$, which converges via the dominated convergence theorem as $j \to \infty$ to $\int_{\mathbb{R}^n} e^{2\pi i p\cdot x} \hat{f}(p)\, dp$ since $\hat{f} \in \mathscr{S}(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$ while (by Lemma 11.16) $\hat{\rho}_j$ is uniformly bounded and converges pointwise to 1. The right hand side is likewise $\int_{\mathbb{R}^n} \rho_j(x-y) f(y)\, dy = (\rho_j * f)(x)$, which converges to $f(x)$ by (11.7). We've proved:

$$\mathscr{F}^* \mathscr{F} f = f \quad \text{for all} \quad f \in \mathscr{S}(\mathbb{R}^n).$$

An almost identical argument proves $\mathscr{F}\mathscr{F}^* f = f$ for all $f \in \mathscr{S}(\mathbb{R}^n)$. □

11.6. Plancherel's theorem. With the Fourier inversion formula in hand, Plancherel's theorem will follow easily from the observation that $\mathscr{F}$ and $\mathscr{F}^*$ are adjoints:

Lemma 11.17. *For every $f, g \in \mathscr{S}(\mathbb{R}^n)$,*

$$\langle g, \mathscr{F} f \rangle_{L^2} = \langle \mathscr{F}^* g, f \rangle_{L^2}.$$

Proof. Since f and g are both in $\mathscr{S}(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$, the function $(x,p) \mapsto |e^{-2\pi i p\cdot x}\langle g(p), f(x)\rangle| \leqslant |g(p)| \cdot |f(x)|$ is integrable on $\mathbb{R}^n \times \mathbb{R}^n$, so Fubini's theorem gives

$$\begin{aligned}
\langle g, \mathscr{F} f \rangle_{L^2} &= \int_{\mathbb{R}^n} \left\langle g(p), \left(\int_{\mathbb{R}^n} e^{-2\pi i p\cdot x} f(x)\, dx \right) \right\rangle dp = \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^n} \langle g(p), e^{-2\pi i p\cdot x} f(x) \rangle\, dx \right) dp \\
&= \int_{\mathbb{R}^n \times \mathbb{R}^n} e^{-2\pi i p\cdot x} \langle g(p), f(x) \rangle\, d(m(x) \otimes m(p)) = \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^n} \langle e^{2\pi i p\cdot x} g(p), f(x) \rangle\, dp \right) dx \\
&= \int_{\mathbb{R}^n} \left\langle \int_{\mathbb{R}^n} e^{2\pi i p\cdot x} g(p)\, dp, f(x) \right\rangle dx = \int_{\mathbb{R}^n} \langle \check{g}(x), f(x) \rangle\, dx = \langle \mathscr{F}^* g, f \rangle_{L^2}.
\end{aligned}$$

□

Proof of Theorem 11.5. For $f, g \in \mathscr{S}(\mathbb{R}^n)$, the Fourier transforms $\hat{f}, \hat{g}$ are also in $\mathscr{S}(\mathbb{R}^n)$, so Lemma 11.17 together with the relation $\mathscr{F}\mathscr{F}^* = \mathrm{Id}$ gives

$$\langle \hat{f}, \hat{g} \rangle_{L^2} = \langle \mathscr{F}^* f, \mathscr{F}^* g \rangle_{L^2} = \langle f, \mathscr{F}\mathscr{F}^* g \rangle_{L^2} = \langle f, g \rangle_{L^2}.$$

□

11.7. Convolutions. If f and g are functions of class L^1 and L^2 respectively on $\mathbb{R}^n$ and at least one of them is assumed to be scalar valued (so that pointwise products $f(x)g(x)$ are well defined), then Young's inequality (Theorem 8.8) implies that the convolution $f * g$ is a well-defined function in $L^2(\mathbb{R}^n)$. Since $\hat{f}, \check{f} \in C_b^0(\mathbb{R}^n)$ and $\hat{g}, \check{g} \in L^2(\mathbb{R}^n)$ in this situation, the pointwise products $\hat{f}\hat{g}$ and $\check{f}\check{g}$ are also well-defined functions in $L^2(\mathbb{R}^n)$, and the formulas in the following result therefore make sense:

THEOREM 11.18. *If $f \in L^1(\mathbb{R}^n)$ and $g \in L^2(\mathbb{R}^n)$, then $\mathscr{F}(f * g) = \hat{f}\hat{g}$ and $\mathscr{F}^*(f * g) = \check{f}\check{g}$ almost everywhere.*

PROOF. We focus on the formula for $\mathscr{F}(f * g)$, as the same argument works for $\mathscr{F}^*(f * g)$. If $f, g \in L^1(\mathbb{R}^n)$, then the formula is a straightforward application of Fubini's theorem, which we leave as a (highly recommended!) exercise. To extend this result to general $g \in L^2(\mathbb{R}^n)$, one can choose a sequence $g_k \in \mathscr{S}(\mathbb{R}^n)$ with $g_k \to g$ in L^2: then $\widehat{f * g_k} = \hat{f}\hat{g}_k$ since $\mathscr{S}(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$, and Young's inequality implies $f * g_k \to f * g$ in L^2, hence by Plancherel's theorem, $\hat{f}\hat{g}_k = \widehat{f * g_k} \xrightarrow{L^2} \widehat{f * g}$. At the same time, $\hat{f} \in C^0_b(\mathbb{R}^n)$ and Plancherel's theorem also implies $\hat{g}_k \to \hat{g}$ in L^2, thus $\hat{f}\hat{g}_k$ also converges in L^2 to $\hat{f}\hat{g}$. □

There are two analogues of Theorem 11.18 for fully periodic functions and Fourier series. We defined in Remark 10.12 the convolution of two periodic functions f and g as an integral over the torus $\mathbb{T}^n$. There is a similar definition for functions on $\mathbb{Z}^n$, with Lebesgue integration replaced by summation (i.e. integration with respect to the counting measure): for two functions f, g on $\mathbb{Z}^n$ such that at least one is scalar valued, we write

$$(f * g)(k) := \sum_{j \in \mathbb{Z}^n} f(k - j)g(j).$$

This is considered well-defined for a given $k \in \mathbb{Z}^n$ if and only if the sum on the right hand side converges absolutely.

EXERCISE 11.19. Adapt the proof of Theorem 8.8 to show that Young's inequality also holds for functions on $\mathbb{T}^n$ and $\mathbb{Z}^n$, that is:

(a) For any $f \in L^1(\mathbb{T}^n)$ and $g \in L^p(\mathbb{T}^n)$ with $1 \leqslant p \leqslant \infty$, $(f * g)(x)$ is defined for almost every $x \in \mathbb{T}^n$ and determines a function $f * g \in L^p(\mathbb{T}^n)$ such that $\|f * g\|_{L^p} \leqslant \|f\|_{L^1} \cdot \|g\|_{L^p}$.
(b) For any $f \in \ell^1(\mathbb{Z}^n)$ and $g \in \ell^p(\mathbb{Z}^n)$ with $1 \leqslant p \leqslant \infty$, $(f * g)(k)$ is defined for all $k \in \mathbb{Z}^n$ and satisfies $\|f * g\|_{\ell^p} \leqslant \|f\|_{\ell^1} \cdot \|g\|_{\ell^p}$.

EXERCISE 11.20. Prove the following analogues of Theorem 11.18 for Fourier series:

(a) For any $f, g \in L^1(\mathbb{T}^n)$, the Fourier coefficients of $f * g \in L^1(\mathbb{T}^n)$ are given by $\widehat{f * g}_k = \hat{f}_k \hat{g}_k$.[22]
(b) If $f \in C^0(\mathbb{T}^n)$ and $g \in L^2(\mathbb{T}^n)$ have Fourier coefficients $\hat{f} \in \ell^1(\mathbb{Z}^n)$ and $\hat{g} \in \ell^2(\mathbb{Z}^n)$, then the Fourier coefficients of fg are given by $\widehat{fg}_k = (\hat{f} * \hat{g})_k$.

The following amusing variation on Theorem 11.18 will be useful in our discussion of nowhere differentiable functions in the next subsection. Suppose f is a fully periodic function on $\mathbb{R}^n$, expressed as a Fourier series $f(x) = \sum_{k \in \mathbb{Z}^n} e^{2\pi i k \cdot x} \hat{f}_k$. This function does not belong to $L^p(\mathbb{R}^n)$ for any $p < \infty$ unless it is zero almost everywhere, thus we cannot define a Fourier transform for f in the usual sense.[23] In the following paragraph, we shall ignore this difficulty as we did in the initial "physicist's proofs" of Theorems 10.5 and 11.4, thus the reader is asked to temporarily suspend all skepticism about issues like convergence, interchange of summation and integration, and the existence of the Dirac δ-function. The logical gaps will be filled in subsequently.

With this understood, let us pretend that $\hat{f}$ is a well-defined function on $\mathbb{R}^n$ given by the usual formula $\hat{f}(p) = \int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} f(x)\, dx$. To write it down more precisely, observe that the inverse

[22] We are not mentioning the case $g \in L^2(\mathbb{T}^n)$ here because it is redundant: since $\mathbb{T}^2$ has finite measure, $L^2(\mathbb{T}^n) \subset L^1(\mathbb{T}^n)$.

[23] A function $f \notin L^2(\mathbb{R}^n)$ may nonetheless have a well-defined Fourier transform that is not a function but a *tempered distribution*; see §13.6. This notion can be used to give rigorous meaning to formulas like $\hat{1} = \delta$, though it is not required for the present discussion.

Fourier transform of the (fictional) Dirac δ-function is given by $\check{\delta}(x) = \int_{\mathbb{R}^n} e^{2\pi i p\cdot x}\delta(p)\,dp = 1$, so applying the Fourier inversion formula gives the formal relation $\widehat{1} = \delta$, or in verbose form (cf. (11.5)),

$$\delta(p) = \int_{\mathbb{R}^n} e^{-2\pi i p\cdot x}\,dx. \tag{11.8}$$

This suggests the formula

$$\hat{f}(p) = \int_{\mathbb{R}^n} e^{-2\pi i p\cdot x}\sum_{k\in\mathbb{Z}^n} e^{2\pi i k\cdot x}\hat{f}_k\,dx = \sum_{k\in\mathbb{Z}^n}\left(\int_{\mathbb{R}^n} e^{-2\pi i(p-k)\cdot x}\,dx\right)\hat{f}_k = \sum_{k\in\mathbb{Z}^n}\delta(p-k)\hat{f}_k. \tag{11.9}$$

The support of this "function" is $\mathbb{Z}^n$ since $\delta(p) = 0$ for all $p \neq 0$. Now for a given $k \in \mathbb{Z}^n$, choose a smooth compactly supported function $\hat{\psi} : \mathbb{R}^n \to [0,1]$ that satisfies $\hat{\psi}(k) = 1$ and has no other points of $\mathbb{Z}^n$ in its support. We have labeled it $\hat{\psi}$ because, as an element of $\mathscr{S}(\mathbb{R}^n)$, $\hat{\psi}$ is the Fourier transform of another function $\psi \in \mathscr{S}(\mathbb{R}^n)$. The product of $\hat{\psi}$ with the right hand side of (11.9) is $\delta(p-k)\hat{f}_k$, which is formally the Fourier transform of $e^{2\pi i k\cdot x}\hat{f}_k$, i.e. a single term in the Fourier series for f. Since products of Fourier transforms are Fourier transforms of convolutions according to Theorem 11.18, we can take this formal discussion as motivation for the formula

$$(\psi * f)(x) = e^{2\pi i k\cdot x}\hat{f}_k. \tag{11.10}$$

In contrast with several other questionable things that have been said in this paragraph, (11.10) does not look at all implausible, e.g. both sides are smooth bounded functions on $\mathbb{R}^n$ (for the left hand side this follows from Theorem 8.7 and Young's inequality since $\psi \in \mathscr{S}(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$ and $f \in L^\infty(\mathbb{R}^n)$). Let us now give a rigorous proof.

LEMMA 11.21. *Suppose f is a continuous fully periodic function on $\mathbb{R}^n$ with absolutely summable Fourier coefficients $\hat{f} \in \ell^1(\mathbb{Z}^n)$, and $\psi : \mathbb{R}^n \to \mathbb{C}$ is the inverse Fourier transform of a function $\hat{\psi} \in \mathscr{S}(\mathbb{R}^n)$ with $\hat{\psi}(k) = 1$ for some $k \in \mathbb{Z}^n$ and $\hat{\psi}(k') = 0$ for all $k' \in \mathbb{Z}^n\backslash\{k\}$. Then (11.10) holds.*

PROOF. The reversal of summation and integration in the following computation is justified by the dominated convergence theorem since $\left|\psi(y)\sum_{j\in\mathbb{Z}^n} e^{2\pi i j\cdot(x-y)}\hat{f}_j\right| \leqslant |\psi(y)|\cdot\|\hat{f}\|_{\ell^1}$ and $\psi \in \mathscr{S}(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$:

$$\begin{aligned}(\psi * f)(x) = (f * \psi)(x) &= \int_{\mathbb{R}^n} f(x-y)\psi(y)\,dy = \int_{\mathbb{R}^n}\psi(y)\Big(\sum_{j\in\mathbb{Z}^n} e^{2\pi i j\cdot(x-y)}\hat{f}_j\Big)\,dy \\ &= \sum_{j\in\mathbb{Z}^n} e^{2\pi i j\cdot x}\left(\int_{\mathbb{R}^n}\psi(y)e^{-2\pi i j\cdot y}\,dy\right)\hat{f}_j = \sum_{j\in\mathbb{Z}^n} e^{2\pi i j\cdot x}\hat{\psi}(j)\hat{f}_j = e^{2\pi i k\cdot x}\hat{f}_k.\end{aligned}$$

□

11.8. Nowhere-differentiable functions. Fix constants $a, b > 1$ and consider the function $f : \mathbb{R} \to \mathbb{C}$ defined by

$$f(x) := \sum_{k=0}^\infty \frac{1}{a^k}e^{2\pi i b^k x}. \tag{11.11}$$

Since $\sum_{k=0}^\infty \frac{1}{a^k} < \infty$, the partial sums of this series converge uniformly to a continuous function. If $b \in \mathbb{N}$, then f is periodic, and (11.11) is an expression of its Fourier series. Differentiating it term by term gives

$$f'(x) = 2\pi i \sum_{k=0}^{\infty} \frac{b^k}{a^k} e^{2\pi i b^k x}, \tag{11.12}$$

a formula that should be taken with a grain of salt until we have investigated whether the right hand side converges. In fact, the series converges absolutely and uniformly if $b < a$, and it follows in this case that f is indeed continuously differentiable. The interesting question is what happens when $b \geqslant a$.

THEOREM 11.22. *If $b \geqslant a > 1$, then the function f in (11.11) is not differentiable at any point.*

Up to unimportant details such as the factor of 2π in the exponent, the real part of f is the function that was introduced by Weierstrass in 1872 as the first published example of a continuous but nowhere differentiable function. It was later [**Ban31**,**Maz31**] shown that, while such functions are typically not so easy to write down, they are not at all unusual, e.g. the subset of $C^0([0,1])$ consisting of nowhere differentiable functions is dense, and even better, it is *comeager*, meaning it is a countable intersection of open and dense subsets.[24] In other words, "almost all" continuous functions are nowhere differentiable in some quantifiable sense.

The version of Theorem 11.22 proved by Weierstrass included the extra conditions that b is an odd integer and $b/a > 1 + \frac{3}{2}\pi$, which are not necessary. In the form stated here, Theorem 11.22 is due to Hardy [**Har16**], and our proof below is adapted from [**Joh10**].

Some initial intuition for Theorem 11.22 comes from (11.12), as we have already learned to expect some correspondence between the differentiability of a function and the rate at which its Fourier coefficients decay. This correspondence typically goes in only one direction, e.g. absolute summability of the series $\sum_k \hat{f}_k$ or $\sum_k |k|\hat{f}_k$ implies continuity of f or f' respectively, but not every continuous function has summable Fourier coefficients. The challenge in Theorem 11.22 is similar, as we need to show that if f is differentiable at some point, then the coefficients on the right hand side of (11.12) must indeed be absolutely summable. The Weierstrass function has a special property that makes proving results like this more feasible: its Fourier series is **lacunary**, meaning that most of its Fourier coefficients are zero, and the gaps between its nonzero terms become wider (at an exponential rate) as the series continues. We will not give a more precise definition of this property here, nor mention it explicitly in the proof below, but you may recognize where it is used implicitly if you pay careful attention. A similar result worth mentioning is that for any function g on $S^1 := \mathbb{T}^1$ with a lacunary Fourier series, g is bounded if *and only if* its Fourier coefficients are absolutely summable; see [**Kat04**, §V.1.4].

PROOF OF THEOREM 11.22. As already mentioned, the absolute summability of $\sum_{k=0}^{\infty} \frac{1}{a^k}$ for $a > 1$ implies that f is continuous and bounded. Let us assume that for some $x_0 \in \mathbb{R}$, the difference quotients

$$F(h) := D_h f(x_0) := \frac{f(x_0 + h) - f(x_0)}{h} \quad \text{for} \quad h \in \mathbb{R}\backslash\{0\}$$

have a well-defined limit $f'(x_0) = \lim_{h \to 0} F(h)$. Since f is bounded, it follows that F extends to a bounded continuous function on $\mathbb{R}$. We will show that this assumption implies $\lim_{k\to\infty}(b/a)^k = 0$, and thus $b < a$.

In order to estimate $(b/a)^k$ for large $k \in \mathbb{N}$, we will use the convolution formula (11.10). Choose a smooth function $\hat{\psi} : \mathbb{R} \to [0,1]$ with $\hat{\psi}(1) = 1$ and compact support in the interval $(1/b, b)$, and

[24] Comeager sets are the complements of *meager* sets, which are countable unions of nowhere dense sets. Since there is no meaningful notion of "Lebesgue measure" on an infinite-dimensional vector space, meager sets often play the role of "sets of measure zero" in the infinite-dimensional context.

for each $k \in \mathbb{Z}$, let

$$\widehat{\psi}_k(p) := \widehat{\psi}(p/b^k),$$

which satisfies

$$\operatorname{supp}(\widehat{\psi}_k) \subset (b^{k-1}, b^{k+1}), \qquad \text{thus} \qquad \widehat{\psi}_k(b^n) = \begin{cases} 1 & \text{if } n = k, \\ 0 & \text{if } n \in \mathbb{Z}\backslash\{k\}. \end{cases}$$

Since $\widehat{\psi}_k \in C_0^\infty(\mathbb{R}) \subset \mathscr{S}(\mathbb{R})$, these functions are Fourier transforms of Schwartz-class functions $\psi_k \in \mathscr{S}(\mathbb{R})$, and an easy change of variables in the Fourier inversion formula gives

$$\psi_k(x) = \int_{-\infty}^{\infty} e^{2\pi i p x} \widehat{\psi}(p/b^k)\, dp = b^k \int_{-\infty}^{\infty} e^{2\pi i b^k p x} \widehat{\psi}(p)\, dp = b^k \psi(b^k x).$$

Notice also that since $0 \in \mathbb{R}$ lies outside the support of $\widehat{\psi}_k$ for each $k \in \mathbb{Z}$, we have

$$0 = \widehat{\psi}_k(0) = \int_{-\infty}^{\infty} \psi_k(x)\, dx \tag{11.13}$$

and

$$0 = \widehat{\psi}_k'(0) = -2\pi i \int_{-\infty}^{\infty} x \psi_k(x)\, dx. \tag{11.14}$$

The first of these two relations implies $\int_{-\infty}^{\infty} f(x_0)\psi_k(x)\, dx = 0$, so we now plug in (11.10) and compute:

$$\begin{aligned} \frac{1}{a^k} e^{2\pi i b^k x_0} &= (f * \psi_k)(x_0) = \int_{-\infty}^{\infty} f(x_0 - x)\psi_k(x)\, dx = \int_{-\infty}^{\infty} [f(x_0 - x) - f(x_0)]\, \psi_k(x)\, dx \\ &= -\int_{-\infty}^{\infty} x F(-x)\psi_k(x)\, dx = -b^k \int_{-\infty}^{\infty} x F(-x)\psi(b^k x)\, dx = -\int_{-\infty}^{\infty} \frac{x}{b^k} F(-x/b^k)\psi(x)\, dx, \end{aligned}$$

implying

$$\left(\frac{b}{a}\right)^k e^{2\pi i b^k x_0} = -\int_{-\infty}^{\infty} F(-x/b^k) x \psi(x)\, dx.$$

Since $\psi \in \mathscr{S}(\mathbb{R})$ and F is bounded, the integrand on the right hand side is bounded for every $k \geqslant 0$ by a constant times $|x|\psi \in L^1(\mathbb{R})$, and it converges pointwise as $k \to \infty$ to $F(0)x\psi(x) = f'(x_0)x\psi(x)$. Applying the dominated convergence theorem and the $k = 0$ case of (11.14), we conclude

$$\lim_{k\to\infty} \left(\frac{b}{a}\right)^k e^{2\pi i b^k x_0} = -f'(x_0) \int_{-\infty}^{\infty} x\psi(x)\, dx = 0,$$

thus $b < a$. □

Exercise 11.23. Show that the Weierstrass function (11.11) with arbitrary constants $a, b > 1$ is of class C^m but has no derivative of order $m+1$ at any point, where $m \geqslant 0$ is the unique integer such that $m < \log_b a \leqslant m + 1$.

12. Sobolev spaces via Fourier analysis

12.1. The general idea of Sobolev spaces. In order to study PDEs via functional-analytic methods, one needs function spaces on which derivatives can be defined as bounded linear operators. For instance, the spaces $C^m(\mathbb{R}^n)$ and $C^m(\mathbb{T}^n)$ of bounded functions on $\mathbb{R}^n$ (or in the latter case fully periodic functions on $\mathbb{R}^n$) that have bounded partial derivatives up to order m is a Banach space with respect to the norm

$$\|f\|_{C^m} := \sum_{0 \leqslant |\alpha| \leqslant m} \sup_x |\partial^\alpha f(x)|, \tag{12.1}$$

where the sum ranges over all multi-indices of order at most m. For each $j = 1, \dots, n$, the operation of taking the partial derivative with respect to coordinate x_j then defines a bounded linear operator

$$\partial_j : C^1(\mathbb{R}^n) \to C^0(\mathbb{R}^n) \quad \text{or} \quad \partial_j : C^1(\mathbb{T}^n) \to C^0(\mathbb{T}^n),$$

and similarly, any multi-index α of order $|\alpha| = m$ defines $\partial^\alpha : C^m(\mathbb{R}^n) \to C^0(\mathbb{R}^n)$ or $\partial^\alpha : C^m(\mathbb{T}^n) \to C^0(\mathbb{T}^n)$. That is all fine, but unfortunately the Banach spaces $C^m(\mathbb{R}^n)$ and $C^m(\mathbb{T}^n)$ do not have enough nice properties to be very useful in technical arguments. They are, for example, not reflexive, and their dual spaces are not easy to describe, e.g. by the *Riesz-Markov theorem* (see [**Sal16**, §3.3]), the dual of the space of continuous functions on a compact domain can be identified with a space of *measures*, which is inconveniently much larger than a space of functions. In this sense, the L^p-spaces are much nicer, but they have the obvious drawback that functions of class L^p are typically not even continuous, much less differentiable, so operators like ∂_j cannot be defined on $L^p(\mathbb{R}^n)$ or $L^p(\mathbb{T}^n)$.

The theory of Sobolev spaces, which is indispensable for the modern theory of PDEs, provides a means of keeping the good properties of the L^p-spaces while also permitting differentiation to be a bounded linear operator. Let us suppose first that we want to be able to handle first-order differential operators for functions on an open domain $\Omega \subset \mathbb{R}^n$. There are a few ways that one can imagine defining a suitable generalization of $L^p(\Omega)$ for this purpose:

Idea 1. Define $X_1(\Omega)$ to be the space of functions $f \in L^p(\Omega)$ that are differentiable almost everywhere and satisfy $\partial_j f \in L^p(\Omega)$ for every $j = 1, \dots, n$. A natural choice of norm on this space is

$$\|f\|_{X_1} := \|f\|_{L^p} + \sum_{j=1}^n \|\partial_j f\|_{L^p}. \tag{12.2}$$

Unfortunately, it will turn out that this space is not complete, i.e. it is a reasonable normed vector space, but not a Banach space.

Idea 2. Since $X_1(\Omega)$ as defined above is not complete, one could define $X_2(\Omega)$ to be the closure of $X_1(\Omega) \subset L^p(\Omega)$ with respect to the X_1-norm. This is a reasonable definition, but not convenient to work with—we would prefer to be able to say precisely what the elements of $X_2(\Omega)$ are, rather than just calling it the closure of a dense subspace whose elements we can explicitly describe.

Idea 3. In the case $n = 1$ with $\Omega = (a, b) \subset \mathbb{R}$, one can consider the space $X_3(\Omega)$ of functions that have absolutely continuous extensions to $[a, b]$ such that their (almost everywhere defined) derivatives are of class L^p on (a, b). This is also a reasonable definition, but it only makes sense for functions of one real variable—on domains in $\mathbb{R}^n$, the notion of absolute continuity can be defined for measures, but not functions. It also doesn't give much of a hint how we should handle higher-order derivatives.

The general solution to these problems will be to generalize the notion of the derivative and thus talk about "weakly differentiable" functions; we will do this in §13 by introducing the theory of distributions. But before that, we observe that in the setting of $\Omega = \mathbb{R}^n$ with $p = 2$, a simpler solution is available using the properties of the Fourier transform.

12.2. The spaces $H^m(\mathbb{R}^n)$ and $H^m(\mathbb{T}^n)$. Let us start by writing down a norm that measures derivatives up to order $m \geqslant 0$ by integrating them instead of taking suprema (as the C^m-norm does). The case $m = 1$ appeared already in (12.2), and it generalizes naturally to

$$\|f\|_{W^{m,p}} := \sum_{|\alpha| \leqslant m} \|\partial^\alpha f\|_{L^p}, \tag{12.3}$$

where the summation ranges over all multi-indices α of order at most m; note that this includes the trivial multi-index with $|\alpha| = 0$, so the L^p-norm of f is one of the terms in the sum. If $p = 2$, we can use Plancherel's theorem and (11.4) to rewrite this norm as

$$\sum_{|\alpha| \leqslant m} \|(2\pi i p)^\alpha \hat{f}\|_{L^2}.$$

Up to equivalence of norms, the factors of $2\pi i$ in this expression clearly make no difference, and every monomial p^α with order $|\alpha| \leqslant m$ satisfies $|p^\alpha| \leqslant c(1 + |p|^2)^{m/2}$ for some constant $c > 0$, thus an equivalent norm is given by the simpler expression

$$\|f\|_{H^m} := \left\|(1 + |p|^2)^{m/2} \hat{f}\right\|_{L^2} = \left(\int_{\mathbb{R}^n} (1 + |p|^2)^m |\hat{f}(p)|^2 \, dp\right)^{1/2} \in [0, \infty]. \tag{12.4}$$

Notice that this formula does not require f to be differentiable, nor even continuous; it is defined for *all* L^2-functions on $\mathbb{R}^n$, though we have no guarantee in general that it will be finite. Finiteness of this norm determines a subspace

$$H^m(\mathbb{R}^n) := \left\{ f \in L^2(\mathbb{R}^n) \;\middle|\; \|f\|_{H^m} < \infty \right\}.$$

We now observe two interesting things about this definition: first, it does not actually mention any derivatives of f, so $\|f\|_{H^m}$ might potentially be finite even if f is not differentiable or continuous. We plan to interpret $H^m(\mathbb{R}^n)$ nonetheless as the space of L^2-functions whose derivatives up to order m are also of class L^2, and this interpretation will turn out to be correct as soon as we enlarge our notion of what the word "derivative" can mean in §13. Second, the stated definition of the H^m-norm does not actually require m to be a nonnegative integer. It makes sense in fact for any $m \in \mathbb{R}$ as long as f belongs to a class of functions whose Fourier transforms can be defined, e.g. one can even allow $m < 0$ and drop the condition $f \in L^2$ by allowing f to be a so-called *tempered distribution* (see §13.6). We will not discuss the case $m < 0$ here, but the case of nonnegative real numbers other than integers gives rise to a notion of *fractional differentiability* that is sometimes useful in applications.

THEOREM 12.1. *For every $m \geqslant 0$, $H^m(\mathbb{R}^n)$ is a Hilbert space with respect to the inner product*

$$\langle f, g\rangle_{H^m} := \int_{\mathbb{R}^n} (1 + |p|^2)^m \langle \hat{f}(p), \hat{g}(p)\rangle \, dp.$$

PROOF. The map $H^m(\mathbb{R}^n) \to L^2(\mathbb{R}^n) : f \mapsto (1 + |p|^2)^{m/2} \hat{f}$ is a bijective isometry, so completeness of $H^m(\mathbb{R}^n)$ follows from completeness of $L^2(\mathbb{R}^n)$.[25] □

For fully periodic functions, there is a natural analogue of the space $H^m(\mathbb{R}^n)$ whose definition uses Fourier series instead of the Fourier transform. We define for each $f \in L^2(\mathbb{T}^n)$ the norm

$$\|f\|_{H^m} := \left\|(1 + |k|^2)^{m/2} \hat{f}\right\|_{\ell^2} = \left(\sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^m |\hat{f}_k|^2\right)^{1/2} \in [0, \infty],$$

and set

$$H^m(\mathbb{T}^n) := \left\{ f \in L^2(\mathbb{T}^n) \;\middle|\; \|f\|_{H^m} < \infty \right\}.$$

[25]Theorems 12.1 and 12.2 are also true and can be proved in the same way for $m < 0$, but we are not stating them for that case because our definition of H^m as a subspace of L^2 is only correct for $m \geqslant 0$. For a more general discussion, see e.g. [Tay96].

The proof of the next statement is an easy adaptation of Theorem 12.1.

THEOREM 12.2. *For every $m \geqslant 0$, $H^m(\mathbb{T}^n)$ is a Hilbert space with respect to the inner product*

$$\langle f, g\rangle_{H^m} := \sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^m \langle \hat{f}_k, \hat{g}_k\rangle.$$

$\square$

EXERCISE 12.3. Show that $\mathscr{S}(\mathbb{R}^n) \subset H^m(\mathbb{R}^n)$ and $C^\infty(\mathbb{T}^n) \subset H^m(\mathbb{T}^n)$ for every $m \geqslant 0$.

By construction, any L^2-function f on $\mathbb{R}^n$ or $\mathbb{T}^n$ that has continuous derivatives up to order $m \geqslant 0$ which are also of class L^2 belongs to $H^s(\mathbb{R}^n)$ or $H^s(\mathbb{T}^n)$ respectively for every $s \leqslant m$. In particular, every smooth function with compact support is of class H^m for every m. If α is a multi-index with $|\alpha| \leqslant m$, then for f of class C_0^∞, (10.10) and (11.4) determine formulas for $\partial^\alpha f$ in terms of the Fourier series or transform of f. These formulas also make sense if f is not smooth but is of class H^m, and in this way one also obtains a bound on $\|\partial^\alpha f\|_{L^2}$ in terms of $\|f\|_{H^m}$, proving:

PROPOSITION 12.4. *For any multi-index α of order $|\alpha| = m \in \mathbb{N}$, the operator ∂^α on smooth functions with compact support has a natural extension to a bounded linear map $\partial^\alpha : H^m(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ or $\partial^\alpha : H^m(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$.* $\square$

EXERCISE 12.5. Extend Proposition 12.4 to define ∂^α as a bounded linear map $H^{s+m}(\mathbb{R}^n) \to H^s(\mathbb{R}^n)$ or $H^{s+m}(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$ for every $s \geqslant 0$ whenever $|\alpha| = m$.

EXERCISE 12.6. Show that for $a, b > 1$ with $b \in \mathbb{N}$, the Weierstrass function $f(x) = \sum_{k=0}^\infty \frac{1}{a^k} e^{2\pi i b^k x}$ belongs to $H^m(S^1) := H^m(\mathbb{T}^1)$ if and only if $m < \log_b a$.

For the Weierstrass functions, comparing Exercises 11.23 and 12.6 reveals a fairly straightforward correspondence: for each integer $m \geqslant 0$, $f \in C^m(S^1) := C^m(\mathbb{T}^1)$ if and only if $f \in H^m(S^1)$, and C^m-functions must also belong to $H^s(S^1)$ for some $s \in (m, m+1)$. In particular, f can be nowhere differentiable but will still belong to $H^s(S^1)$ for $s \in (0, 1)$ sufficiently small. But the simplicity of this correspondence is slightly misleading. Beyond the Weierstrass functions, it cannot be true in general that every function of class H^m for an integer $m \geqslant 0$ is also of class C^m; this is clearly false for $m = 0$, since not all L^2-functions are continuous. The following exercises exhibit some less obvious examples.

EXERCISE 12.7. Two simple examples of discontinuous real-valued periodic functions on $\mathbb{R}$ are the *square* and *sawtooth* waves, defined respectiely as the obvious periodic extensions of

$$f(x) := \begin{cases} 1 & \text{if } 0 \leqslant x < 1/2, \\ -1 & \text{if } 1/2 \leqslant x < 1, \end{cases} \qquad \text{and} \qquad g(x) := x \text{ for } 0 \leqslant x < 1.$$

Show that both belong to $H^m(S^1)$ for all $m < 1/2$ but not for $m \geqslant 1/2$.

EXERCISE 12.8. The goal of this exercise is to show that the improper integral

$$f(x) := \int_2^\infty \frac{e^{2\pi i p x}}{p \ln p}\, dp := \lim_{N \to \infty} \int_2^N \frac{e^{2\pi i p x}}{p \ln p}\, dp, \qquad x \in \mathbb{R}\backslash\{0\} \tag{12.5}$$

defines a discontinuous function in $H^{1/2}(\mathbb{R})$. Note that the integrand $\frac{e^{2\pi i p x}}{p \ln p}$ is not a Lebesgue-integrable function of $p \in \mathbb{R}$, so the limit is necessary in order to define the integral, and its convergence is not obvious.

(a) Show that there exists a function $g \in L^2(\mathbb{R})$ whose Fourier transform is given almost everywhere by

$$\widehat{g}(p) = \begin{cases} \frac{1}{p \ln p} & \text{if } p \geqslant 2, \\ 0 & \text{if } p < 2, \end{cases}$$

and that this function belongs to $H^m(\mathbb{R})$ if and only if $m \leqslant 1/2$.

(b) Show that the function g in part (a) is the L^2-limit of the functions $f_N(x) := \int_2^N \frac{e^{2\pi i p x}}{p \ln p}\, dp$ as $N \to \infty$.

(c) Use integration by parts to prove that for every $M \geqslant 2$ and $x \in \mathbb{R}\backslash\{0\}$, the limit $\int_M^\infty \frac{e^{2\pi i p x}}{p \ln p}\, dp := \lim_{N\to\infty} \int_M^N \frac{e^{2\pi i p x}}{p \ln p}\, dp$ exists, depends continuously on x, and satisfies

$$\left| \int_M^\infty \frac{e^{2\pi i p x}}{p \ln p}\, dp \right| \leqslant \frac{1}{\pi |x| \cdot M \ln M}.$$

Deduce from this that the function g in part (a) matches (almost everywhere) the function f defined in (12.5), which is continuous on $\mathbb{R}\backslash\{0\}$.
Hint: Recall that L^2-convergence implies pointwise almost everywhere convergence of a subsequence.

(d) Prove that $\lim_{x\to 0} |f(x)| = \infty$.
Hint: Break up the integral over the intervals $[2, \epsilon/|x|]$ and $[\epsilon/|x|, \infty)$ for some small $\epsilon > 0$ with $|x| < \epsilon/2$. The estimate in part (c) will bound it on the second interval, while on the first, its absolute value should be larger than some positive multiple of $\int_2^{\epsilon/|x|} \frac{dp}{p \ln p}$ whenever ϵ is sufficiently small. Now let $|x| \to 0$ and use the fact that $\int_2^\infty \frac{dp}{p \ln p} = \infty$.

EXERCISE 12.9. Adapt the argument of Exercise 12.8 to show that the L^2-convergent Fourier series $f(x) := \sum_{k=2}^\infty \frac{e^{2\pi i k x}}{k \ln k}$ defines a discontinuous function in $H^{1/2}(S^1) := H^{1/2}(\mathbb{T}^1)$.
Hint: Proving a bound on $\left|\sum_{k=M}^\infty \frac{e^{2\pi i k x}}{k \ln k}\right|$ for $x \neq 0$ requires an analogue of integration by parts for summations, which is easy to prove if you regard the "derivative" of a sequence a_k as the sequence $a_k' := a_{k+1} - a_k$. If you need more inspiration, see [**Rud76**, pp. 70–71].

12.3. The Sobolev embedding theorem. Exercises 12.7, 12.8 and 12.9 demonstrate that functions of class H^m for $m \leqslant 1/2$ on S^1 or $\mathbb{R}$ need not be continuous, though it seems that discontinuous examples for the case $m = 1/2$ are not so easy to construct. We will now show that it becomes impossible for $m > 1/2$, and in fact, such a threshold also exists for functions on $\mathbb{T}^n$ or $\mathbb{R}^n$ and depends on the dimension n. Recall that $H^m(\mathbb{R}^n)$ and $H^m(\mathbb{T}^n)$ were defined as subspaces of $L^2(\mathbb{R}^n)$ and $L^2(\mathbb{T}^n)$ respectively, so their elements are not actually functions, but rather *equivalence classes* of functions defined almost everywhere. This is different from the Banach spaces $C^m(\mathbb{R}^n)$ and $C^m(\mathbb{T}^n)$, whose elements are actual functions. We will say that there exists a **continuous inclusion**

$$H^s(\mathbb{R}^n) \hookrightarrow C^m(\mathbb{R}^n)$$

whenever the following is true: every $f \in H^s(\mathbb{R}^n)$ is equal almost everywhere to a unique function $\widetilde{f} \in C^m(\mathbb{R}^n)$, and the resulting map $H^s(\mathbb{R}^n) \to C^m(\mathbb{R}^n) : f \mapsto \widetilde{f}$ is a bounded linear operator. The existence of a continuous inclusion thus comes with an estimate of the form

$$\|f\|_{C^m} \leqslant c \|f\|_{H^s} \quad \text{for some constant } c > 0 \text{ independent of } f,$$

where we abuse notation by forgetting the distinction between the C^m-function f and the equivalence class in $H^s(\mathbb{R}^n)$ that it represents. There is an obvious similar definition for the spaces of fully periodic functions $H^s(\mathbb{T}^n)$ and $C^m(\mathbb{T}^n)$.

THEOREM 12.10 (Sobolev embedding theorem, case $p = 2$). *Assume $n \in \mathbb{N}$ and $s > 0$ satisfy $2s > n$. Then there exist continuous inclusions*

$$H^{s+m}(\mathbb{R}^n) \hookrightarrow C^m(\mathbb{R}^n) \quad \text{and} \quad H^{s+m}(\mathbb{T}^n) \hookrightarrow C^m(\mathbb{T}^n)$$

for every integer $m \geqslant 0$.

PROOF. We first consider functions $f \in H^s(\mathbb{R}^n)$ with $2s > n$. The main step is to establish a bound on $\|\widehat{f}\|_{L^1}$, as f is then equal almost everywhere to $\mathscr{F}^*\widehat{f}$, which is continuous since $\mathscr{F}^*$ defines a bounded linear operator $L^1(\mathbb{R}^n) \to C^0(\mathbb{R}^n)$. We use the Cauchy-Schwarz inequality:

$$\begin{aligned} \|\widehat{f}\|_{L^1} &= \int_{\mathbb{R}^n} \frac{1}{(1+|p|^2)^{s/2}} \cdot \left|(1+|p|^2)^{s/2}\widehat{f}\right| dp \leqslant \left\| \frac{1}{(1+|p|^2)^{s/2}} \right\|_{L^2} \cdot \left\|(1+|p|^2)^{s/2}\widehat{f}\right\|_{L^2} \\ &\leqslant \left(\int_{\mathbb{R}^n} \frac{1}{(1+|p|^2)^s}\, dp \right)^{1/2} \cdot \|f\|_{H^s} \end{aligned}$$

Using n-dimensional polar coordinates, we see that the integral in the second line converges if and only if $\int_1^\infty \frac{r^{n-1}}{(1+r^2)^s}\, dr < \infty$. For large $r > 0$, the latter integrand behaves like $r^{n-1}/r^{2s} = r^{n-2s-1}$, so the integral converges if and only if $n - 2s < 0$, which is exactly the condition $2s > n$. This proves the continuous inclusion of $H^s(\mathbb{R}^n)$ into $C^0(\mathbb{R}^n)$.

If $f \in H^{s+m}(\mathbb{R}^n)$ with $m \in \mathbb{N}$, then the same argument bounds the L^1-norm of the function $p \mapsto p^\alpha \widehat{f}(p)$ for each multi-index α with $|\alpha| \leqslant m$ in terms of $\|f\|_{H^{s+m}}$, so the argument of Exercise 11.11 shows that the partial derivatives $\partial^\alpha f$ up to order m exist and are continuous. Moreover, their C^0-norms are bounded in terms of the L^1-norm of $p^\alpha \widehat{f}$, which gives a bound for $\|f\|_{C^m}$ in terms of $\|f\|_{H^{s+m}}$.

The result for fully periodic functions follows by essentially the same argument, except that the version of the Cauchy-Schwarz inequality one needs is $\|fg\|_{\ell^1} \leqslant \|f\|_{\ell^2} \cdot \|g\|_{\ell^2}$ for functions $f, g : \mathbb{Z}^n \to [0, \infty)$. The crucial detail is then the convergence of the series

$$\sum_{k \in \mathbb{Z}^n} \frac{1}{(1+|k|^2)^s} < \infty \quad \text{for } 2s > n,$$

which can be established by comparing it with the integral $\int_{\mathbb{R}^n} \frac{1}{(1+|p|^2)^s}\, dp$. □

COROLLARY 12.11. *Any function belonging to $H^s(\mathbb{R}^n)$ for all $s \geqslant 0$ is (after changing its values on a set of measure zero) smooth, and its derivatives of all orders are bounded. Similarly, $\bigcap_{s \geqslant 0} H^s(\mathbb{T}^n) = C^\infty(\mathbb{T}^n)$.* □

Theorem 12.10 leads to the intuition that functions of class H^s have "$s - \frac{n}{2}$ continuous derivatives," where in general the number $s - n/2$ need not be an integer, but should be assumed positive in order for the statement to carry any meaning. We will make this more precise for the case $0 < s - n/2 < 1$ in §12.8.

12.4. Approximation by smooth functions. The following result says that $H^s(\mathbb{T}^n)$ and $H^s(\mathbb{R}^n)$ could just as well have been defined as the closures of the subspaces $C^\infty(\mathbb{T}^n) \subset L^2(\mathbb{T}^n)$ and $\mathscr{S}(\mathbb{R}^n) \subset L^2(\mathbb{R}^n)$ with respect to the H^s-norm. As a first application, it implies that for each $s \geqslant 0$ and each multi-index α with order $|\alpha| = m$, the extension of the classical differential operator ∂^α to a bounded linear operator $H^{s+m}(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$ or $H^{s+m}(\mathbb{R}^n) \to H^s(\mathbb{R}^n)$ is unique (cf. Proposition 12.4 and Exercise 12.5).

THEOREM 12.12. *The subspaces $C^\infty(\mathbb{T}^n) \subset H^s(\mathbb{T}^n)$ and $\mathscr{S}(\mathbb{R}^n) \subset H^s(\mathbb{R}^n)$ are dense for every $s \geqslant 0$.*

PROOF. We begin with the easiest case: suppose $f \in H^s(\mathbb{T}^n)$, and for $j \in \mathbb{N}$, let

$$f_j(x) := \sum_{|k| \leqslant j} e^{2\pi i k \cdot x} \hat{f}_j.$$

Since f is a finite sum of smooth functions, it is smooth, and we have

$$\|f - f_j\|_{H^s}^2 = \sum_{|k|>j} (1 + |k|^2)^s |\hat{f}_k|^2 \to 0 \quad \text{as} \quad j \to \infty$$

since $\sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^s |\hat{f}_k|^2 = \|f\|_{H^s}^2 < \infty$.

For $f \in H^s(\mathbb{R}^n)$, we have $(1+|p|^2)^{s/2} \hat{f} \in L^2(\mathbb{R}^n)$, and the density of $C_0^\infty(\mathbb{R}^n)$ in $L^2(\mathbb{R}^n)$ implies that there exists a sequence $h_j \in C_0^\infty(\mathbb{R}^n)$ with

$$h_j \xrightarrow{L^2} (1 + |p|^2)^{s/2} \hat{f}.$$

The functions $g_j(p) := \frac{h_j(p)}{(1+|p|^2)^{s/2}}$ are then also in $C_0^\infty(\mathbb{R}^n)$, and they satisfy

$$(1 + |p|^2)^{s/2} g_j \xrightarrow{L^2} (1 + |p|^2)^{s/2} \hat{f}. \tag{12.6}$$

Since $C_0^\infty(\mathbb{R}^n) \subset \mathscr{S}(\mathbb{R}^n)$, each g_j is then the Fourier transform of a unique function $f_j \in \mathscr{S}(\mathbb{R}^n)$, and (12.6) implies $f_j \to f$ in H^s. □

REMARK 12.13. A stronger result is true for functions on $\mathbb{R}^n$: the space of smooth *compactly supported* functions $C_0^\infty(\mathbb{R}^n)$, which is a subspace of $\mathscr{S}(\mathbb{R}^n)$, is also dense in $H^s(\mathbb{R}^n)$. For a proof of this in the more general setting of $W^{m,p}$-spaces (assuming $m \in \mathbb{Z}$), see [**AF03**, Theorem 3.22].

Recall from §8 that the density of smooth functions in L^p is proved by taking convolutions of $f \in L^p(\mathbb{R}^n)$ with an approximate identity ρ_j, a trick often referred to as **mollification**. For most purposes, Theorem 12.12 can also be placed into this context: for instance, the approximating sequence $f_j \to f \in H^s(\mathbb{T}^n)$ in the proof above was constructed by defining its Fourier coefficients to be $\hat{f}_j = \chi_{\bar{B}_j} \hat{f} : \mathbb{Z}^n \to V$, where $\chi_{\bar{B}_j} : \mathbb{Z}^n \to [0, 1]$ denotes the characteristic function of the intersection of $\mathbb{Z}^n$ with the closed ball of radius j in $\mathbb{R}^n$. Clearly $\chi_{\bar{B}_j} \in \mathscr{S}(\mathbb{Z}^n)$, so $\chi_{\bar{B}_j}$ defines the Fourier coefficients of a smooth function, namely

$$\rho_j(x) := \sum_{|k| \leqslant j} e^{2\pi i k \cdot x}.$$

Since this function belongs to $L^1(\mathbb{T}^n)$, Exercise 11.20 implies

$$f_j = \rho_j * f.$$

For $f \in H^s(\mathbb{R}^n)$, if we wanted to approximate f with smooth functions in $H^s(\mathbb{R}^n)$ but did not care whether they are rapidly decreasing, we could use a similar trick:

EXERCISE 12.14. Suppose $\rho \in \mathscr{S}(\mathbb{R}^n)$ satisfies $\int_{\mathbb{R}^n} \rho(x)\, dx = 1$, and define $\rho_j(x) := j^n \rho(jx)$.

(a) Show that for any $s \geqslant 0$ and $f \in H^s(\mathbb{R}^n)$, the sequence $\rho_j * f \in C^\infty(\mathbb{R}^n)$ satisfies

$$\|\rho_j * f\|_{H^s} \leqslant \|f\|_{H^s} \qquad \text{and} \qquad \rho_j * f \xrightarrow{H^s} f \text{ as } j \to \infty.$$

Hint: Compute $\hat{\rho}_j$ in terms of $\hat{\rho}$, then use change of variables and dominated convergence to prove $\|f - \rho_j * f\|_{H^s} \to 0$.

(b) Show that the same result holds if $\rho_j \in \mathscr{S}(\mathbb{R}^n)$ is instead defined as $\check{\psi}_j$ for a sequence of smooth functions $\psi_j : \mathbb{R}^n \to [0, 1]$ with compact support in B_{j+1} and $\psi_j|_{B_j} \equiv 1$.

EXERCISE 12.15. Suppose α is a multi-index of order $|\alpha| = m \in \mathbb{N}$.

(a) Use the definition of $\partial^\alpha : H^m(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ in Proposition 12.4 to prove that for every $\psi \in L^1(\mathbb{R}^n)$ and $f \in H^m(\mathbb{R}^n)$, $\partial^\alpha(\psi * f) = \psi * \partial^\alpha f \in L^2(\mathbb{R}^n)$.

(b) Use the result of part (a) to give an alternative proof that for any $f \in H^m(\mathbb{R}^n)$ with $m \in \mathbb{N}$ and any approximate identity ρ_j as in §8.4, $\rho_j * f \to f$ in H^m.

12.5. The Arzelà-Ascoli theorem. We will see in §12.8 below that the inclusion of $H^{s+m}(\mathbb{T}^n)$ into $C^m(\mathbb{T}^n)$ for $2s > n$ is more than just continuous: it also has the useful property that any *bounded* sequence in $H^{s+m}(\mathbb{T}^n)$ has a subsequence that *converges* in $C^m(\mathbb{T}^n)$. Compactness results of this type are ubiquitous in analysis—the Banach-Alaoglu theorem (Theorem 7.10) is another nice example. They are often applied toward proving existence results, e.g. if one has a sequence of functions that *approximately* satisfy a PDE, with an error term converging to zero, then one can expect the limit of any convergent subsequence to be an exact solution. As a prelude to this discussion, we shall review a more basic phenomenon that also leads to many useful compactness results.

You may have seen the Arzelà-Ascoli theorem before in various contexts: it is typically used for instance in proving the Riemann mapping theorem in complex analysis, as well as the Cauchy-Peano theorem (on existence of solutions to ODEs in the absence of a Lipschitz condition). If you haven't seen it before, you'll see it now, and we will subsequently use it to prove the compactness of various natural inclusions of function spaces.

Definition 12.16. Suppose X is a set, and Y is a metric space. A collection $\mathcal{F}$ of maps $X \to Y$ is called **uniformly bounded** if for every $x \in X$, the set $\{f(x) \mid f \in \mathcal{F}\} \subset Y$ is bounded.

Definition 12.17. Suppose (X, d_X) and (Y, d_Y) are metric spaces. A collection $\mathcal{F}$ of maps $X \to Y$ is called **equicontinuous** if for every $\epsilon > 0$, there exists a $\delta > 0$ such that the implication

$$d_X(x, y) < \delta \qquad \Rightarrow \qquad d_Y(f(x), f(y)) < \epsilon$$

holds for all $f \in \mathcal{F}$.

Remark 12.18. It is important to pay attention to the order of quantifiers in Definition 12.17; if you get it wrong, then the definition seems to say merely that the functions $f \in \mathcal{F}$ are all uniformly continuous, but in fact it says more than that. The point is that for any given $\epsilon > 0$, one can choose $\delta > 0$ *independently of* $f \in \mathcal{F}$; the same δ is required to work for all maps in the collection. We will not need it for our discussion, but one can also formulate weaker variants of the definition in which the maps are not required to be uniformly continuous, but also in that case, it is important to allow $\delta > 0$ to be chosen independently of $f \in \mathcal{F}$.

Theorem 12.19 (Arzelà-Ascoli). *Assume (X, d) is a separable metric space, V is a finite-dimensional inner product space, and $\mathcal{F}$ is a uniformly bounded and equicontinuous collection of functions $X \to V$. Then any sequence $f_n \in \mathcal{F}$ has a subsequence converging uniformly on compact subsets to a continuous function $f : X \to V$.*

Proof. Fix a countable subset $E \subset X$. By the same argument as in Claim 1 in the proof of Theorem 7.10 (i.e. using the Cantor diagonal trick), we can first replace the sequence f_n with a subsequence such that $f_n(x) \in V$ converges for every $x \in E$. This depends only on the fact that E is countable and bounded sequences in V have convergent subsequences (because $\dim V < \infty$).

Since X is separable, we are free to assume additionally that E is dense, and the result then follows from the following claim: For any dense subset $E \subset X$ and any sequence $f_n \in \mathcal{F}$ that converges pointwise on E, the restriction of f_n to any compact subset $K \subset X$ is uniformly Cauchy. Indeed, in light of equicontinuity, we can associate to each $\epsilon > 0$ a $\delta > 0$ such that $d(x, y) < \delta$ implies $|f_n(x) - f_n(y)| < \epsilon/3$ for all n. By the compactness of K, we can then cover K with a finite collection of open balls

$$K \subset B_1 \cup \ldots \cup B_N \subset X$$

of radius $\delta/2$, and since E is dense, we can also pick a point $x_i \in E \cap B_i$ for each $i = 1, \dots, N$. Now for any $x \in K$, pick i such that $x \in B_i$, and observe

$$|f_m(x) - f_n(x)| \leqslant |f_m(x) - f_m(x_i)| + |f_m(x_i) - f_n(x_i)| + |f_n(x_i) - f_n(x)|.$$

Since the x_i are only finitely many points and the sequence f_n converges at these points, we can assume the second term on the right is less than $\epsilon/3$ as soon as m and n are sufficiently large. The first and third terms are also less than $\epsilon/3$ due to equicontinuity, since $d(x, x_i) < \delta$. This proves $|f_m(x) - f_n(x)| < \epsilon$ for all m, n sufficiently large, and the threshold required for this depends on ϵ, but not on x. □

REMARK 12.20. We have stated Theorem 12.19 in the form that will be most useful for our purposes, but one can generalize it in various ways. The crucial property of the finite-dimensional vector space V used in the proof was that bounded sequences in this space have convergent subsequences. One sometimes sees the Arzelà-Ascoli theorem stated as a result about equicontinuous collections of maps $X \to Y$ between two metric spaces, the first being separable. In that setting, uniform boundedness is not generally a strong enough hypothesis—one needs an additional condition to ensure that sequences $f_n(x) \in Y$ have convergent subsequences (see e.g. [**Roy88**]).

The following easy application is the first of several such results to be discussed below.

COROLLARY 12.21. *Any C^1-bounded sequence of functions $f_n : [0,1] \to V$ has a C^0-convergent subsequence, and thus a continuous limit.*

PROOF. A C^1-bound $\|f_n\|_{C^1} \leqslant C$ gives rise to both uniform boundedness $|f_n(t)| \leqslant \|f_n\|_{C^0} \leqslant \|f_n\|_{C^1} \leqslant C$ for all $t \in [0,1]$ and equicontinuity

$$|f_n(s) - f_n(t)| \leqslant \max_x |f_n'(x)| \cdot |s-t| \leqslant \|f_n\|_{C^1} \cdot |s-t| \leqslant C|s-t|.$$

□

REMARK 12.22. It must be understood that since $C^1([0,1])$ is infinite dimensional, bounded sequences in $C^1([0,1])$ cannot be expected to have C^1-*convergent* subsequences, and the limit obtained in Corollary 12.21 will not necessarily be differentiable, though it certainly is continuous.

12.6. Compact operators. One way of interpreting Corollary 12.21 above is that the obvious inclusion $C^1([0,1]) \hookrightarrow C^0([0,1])$ is not only a bounded linear operator but also satisfies the following stronger condition:

DEFINITION 12.23. A bounded linear operator $A : X \to Y$ between Banach spaces is called a **compact operator** if it maps every bounded subset of X to a precompact subset of Y, or equivalently, for every bounded sequence $x_n \in X$, the sequence $Ax_n \in Y$ has a convergent subsequence.

EXERCISE 12.24. Use the Arzelà-Ascoli theorem to show that for any integers $m > k \geqslant 0$, the inclusion $C^m(\mathbb{T}^n) \hookrightarrow C^k(\mathbb{T}^n)$ is a compact operator.
Hint: Compositions of compact operators are also compact, thus it suffices to prove that $C^{k+1}(\mathbb{T}^n) \hookrightarrow C^k(\mathbb{T}^n)$ is compact for every $k \geqslant 0$. Start with $k = 0$, and notice that any C^1-bounded sequence in $C^1(\mathbb{T}^n)$ is equicontinuous.

EXERCISE 12.25. Find a bounded sequence in $C^1_b(\mathbb{R}^n)$ that converges pointwise to 0 but does not have any C^0-convergent subsequence.
Hint: Translations!

REMARK 12.26. Exercise 12.25 demonstrates the importance of the words "uniformly *on compact subsets*" in the statement of the Arzelà-Ascoli theorem. The theorem works well for proving compactness of inclusions of function spaces on compact domains, but more is typically needed if domains are allowed to be noncompact.

For an easy example beyond Exercise 12.25 of an operator that is *not* compact, take the identity map $\mathcal{H} \to \mathcal{H}$ on an infinite-dimensional Hilbert space $\mathcal{H}$. In this setting, any countably-infinite orthonormal set produces a bounded sequence that cannot have a convergent subsequence. On the other hand, this example works specifically because $\dim \mathcal{H} = \infty$; since closed and bounded subsets in finite-dimensional vector spaces are always compact, *all* linear operators in finite dimensions are compact. Here is a useful generalization of that statement to infinite dimensions:

PROPOSITION 12.27. *If X and Y are Banach spaces and $A : X \to Y$ is a bounded linear operator with finite-dimensional image, then A is compact.*

PROOF. The image of a bounded sequence $x_n \in X$ is a bounded sequence $Ax_n \in Y$, but by assumption it also belongs to a finite-dimensional subspace $\operatorname{im} A \subset Y$. The result thus follows from the fact that all bounded sequences in finite-dimensional vector spaces have convergent subsequences. □

It will also be useful to know that within the space $\mathscr{L}(X, Y)$ of all bounded linear operators between two Banach spaces, the set of compact operators forms a *closed* subspace:

PROPOSITION 12.28. *If X and Y are Banach spaces and $A_n : X \to Y$ is a sequence of compact operators that converge in the operator norm to an operator $A : X \to Y$, then A is also compact.*

PROOF. Suppose $x_n \in X$ is a bounded sequence. Since $A_1 : X \to Y$ is compact, x_n has a subsequence $x_n^{(1)}$ such that $A_1 x_n^{(1)}$ converges. We can then use the compactness of A_2 to extract from $x_n^{(1)}$ a further subsequence $x_n^{(2)}$ such that $A_2 x_n^{(2)}$ converges. Continuing in this manner, one obtains a sequence of subsequences $x_n^{(j)}$ such that $A_j x_n^{(j)}$ converges as $n \to \infty$ for every $j \in \mathbb{N}$. The diagonal subsequence

$$x_n^{(\infty)} := x_n^{(n)}$$

then has the property that $A_j x_n^{(\infty)}$ converges as $n \to \infty$ for every j.

We claim now that $Ax_n^{(\infty)}$ also converges, which will imply that $A : X \to Y$ is compact. Since Y is complete, it suffices to show that $Ax_n^{(\infty)}$ is a Cauchy sequence. Given $\epsilon > 0$, choose $M \in \mathbb{N}$ such that

$$\|A - A_M\| < \frac{\epsilon}{3 \sup_{n \in \mathbb{N}} \|x_n\|},$$

and then choose $N \in \mathbb{N}$ such that $\|A_M x_n^{(\infty)} - A_M x_m^{(\infty)}\| < \epsilon/3$ for all $m, n \geqslant N$; the latter is possible since $A_M x_n^{(\infty)}$ is a Cauchy sequence. It follows that for all $m, n \geqslant N$,

$$\|Ax_n^{(\infty)} - Ax_m^{(\infty)}\| \leqslant \|(A - A_M)x_n^{(\infty)}\| + \|A_M(x_n^{(\infty)} - x_m^{(\infty)})\| + \|(A_M - A)x_m^{(\infty)}\| \leqslant \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon,$$

thus proving the claim. □

COROLLARY 12.29. *Any bounded linear operator in the closure (with respect to the operator norm) of the space of finite-rank operators is compact.* □

REMARK 12.30. We will not need this at present, but if $\mathcal{H}$ is a separable Hilbert space, then the converse of Corollary 12.29 is also true for bounded linear operators $\mathcal{H} \to \mathcal{H}$, i.e. they are compact if and only if they can be approximated arbitrarily well in the operator norm by operators with finite rank. The proof is not hard; see [**RS80**, Theorem VI.13].

12.7. Compactness of Sobolev inclusions. A much more obvious fact than the Sobolev embedding theorem is that for every $t > s \geqslant 0$, there are continuous inclusions $H^t(\mathbb{R}^n) \hookrightarrow H^s(\mathbb{R}^n)$ and $H^t(\mathbb{T}^n) \hookrightarrow H^s(\mathbb{T}^n)$. If we think of functions of class H^s as being $(s-n/2)$-times differentiable, then these inclusions are analogous to the obvious continuous inclusions $C^m \hookrightarrow C^k$ for $m > k$. Let us focus for this subsection on fully periodic functions, which can be regarded as functions on the compact metric space $\mathbb{T}^n$. We saw in Exercise 12.24 that the inclusion $C^m(\mathbb{T}^n) \hookrightarrow C^k(\mathbb{T}^n)$ for $m > k$ is a compact operator, due to the Arzelà-Ascoli theorem. The analogue of Exercise 12.24 for Sobolev spaces is known as the *Rellich-Kondrachov compactness theorem*:

THEOREM 12.31 (Rellich-Kondrachov for $p = 2$). *For every $t > s \geqslant 0$, the natural inclusion $H^t(\mathbb{T}^n) \hookrightarrow H^s(\mathbb{T}^n)$ is compact.*

EXERCISE 12.32. Adapt Exercise 12.25 to show that the inclusion $H^s(\mathbb{R}^n) \hookrightarrow L^2(\mathbb{R}^n)$ for $s > 0$ is not compact; in particular, there exists an H^s-bounded sequence that converges pointwise to 0 but stays a fixed positive distance away from 0 in the L^2-norm, thus it has no L^2-convergent subsequence. This illustrates that the compactness of the domain $\mathbb{T}^n$ in Theorem 12.31 is relevant.

We will derive Theorem 12.31 from Corollary 12.29: the task is to show that the inclusion $H^t(\mathbb{T}^n) \hookrightarrow H^s(\mathbb{T}^n)$ is in the closure of the space of finite-rank bounded linear operators $H^t(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$.

PROOF OF THEOREM 12.31. Fix $t > s \geqslant 0$, and consider for each $N \in \mathbb{N}$ the operator

$$A_N : H^t(\mathbb{T}^n) \to H^s(\mathbb{T}^n) : f \mapsto \sum_{|k| \leqslant N} e^{2\pi i k \cdot x} \hat{f}_k.$$

The image of A_N is finite dimensional since there are only finitely many lattice points $k \in \mathbb{Z}^n$ satisfying $|k| \leqslant N$. The goal is now to show that A_N converges in the operator norm as $N \to \infty$ to the inclusion $A : H^t(\mathbb{T}^n) \hookrightarrow H^s(\mathbb{T}^n)$, hence the latter is a limit of finite-rank operators and is therefore compact.

To prove $\|A - A_N\| \to 0$, we observe that for each $f \in H^t(\mathbb{T}^n)$, the functions $(A - A_N)f$ have the same Fourier coefficients as f except that every coefficient for $k \in \mathbb{Z}^n$ with $|k| \leqslant N$ is set to zero, hence

$$\begin{aligned}\|(A - A_N)f\|_{H^s}^2 &= \sum_{|k|>N} (1 + |k|^2)^s |\hat{f}_k|^2 = \sum_{|k|>N} \frac{1}{(1 + |k|^2)^{t-s}} (1 + |k|^2)^t |\hat{f}_k|^2 \\ &\leqslant \frac{1}{(1 + N^2)^{t-s}} \sum_{k \in \mathbb{Z}^n} (1 + |k|^2)^t |\hat{f}_k|^2 = \frac{1}{(1 + N^2)^{t-s}} \cdot \|f\|_{H^t}^2.\end{aligned}$$

This proves $\|A - A_N\|^2 \leqslant \frac{1}{(1+N^2)^{t-s}}$, and the latter converges to 0 as $N \to \infty$ since $t > s$. ☐

The compactness of the inclusions $H^t(\mathbb{T}^n) \hookrightarrow H^s(\mathbb{T}^n)$ has an interesting consequence related to the Sobolev embedding theorem: if $2s > n$, then there also exists some $t < s$ such that $2t > n$, and the continuous inclusion $H^{s+m}(\mathbb{T}^n) \hookrightarrow C^m(\mathbb{T}^n)$ thus factors into a composition of two inclusions,

$$H^{s+m}(\mathbb{T}^n) \hookrightarrow H^{t+m}(\mathbb{T}^n) \hookrightarrow C^m(\mathbb{T}^n).$$

The first of these is compact, and therefore so is the composition:[26]

COROLLARY 12.33. *For $2s > n$, the continuous inclusions $H^{s+m}(\mathbb{T}^n) \hookrightarrow C^m(\mathbb{T}^n)$ in Theorem 12.10 are also compact.* ☐

[26]Lemma: Any composition of a compact operator with a bounded linear operator is compact. Proof: Easy exercise.

12.8. Hölder estimates. There is a second way to see the compactness of $H^{s+m}(\mathbb{T}^n) \hookrightarrow C^m(\mathbb{T}^n)$ that provides more information, while also yielding a practical interpretation of the motto that functions in $H^s(\mathbb{T}^n)$ are "$(s - n/2)$-times differentiable".

Assume Ω is a measurable subset of either $\mathbb{R}^n$ or $\mathbb{T}^n$, regarded in each case as a metric space with metric denoted by $\operatorname{dist}(x, y) = |x - y|$. Recall that a function $f : \Omega \to V$ is called **Lipschitz continuous** if there exists a constant $C > 0$ such that

$$|f(x) - f(y)| \leqslant C|x - y| \qquad \text{for all } x, y \in \Omega.$$

For example, a continuously differentiable function on an open domain $\mathcal{U} \subset \mathbb{R}^n$ is Lipschitz continuous on every subset $\Omega \subset \mathcal{U}$ on which the partial derivatives are bounded. Classic examples of non-Lipschitz continuous functions include $f(x) := |x|^\alpha$ for $0 < \alpha < 1$ on any neighborhood of $0 \in \mathbb{R}$. These instead satisfy the following condition, which is the same as Lipschitz contintuity for $\alpha = 1$, but weaker for $0 < \alpha < 1$:

Definition 12.34. A function f on $\Omega \subset \mathbb{R}^n$ is called **Hölder continuous** if there exists a number $\alpha \in (0, 1]$ and a constant $C > 0$ such that

$$|f(x) - f(y)| \leqslant C|x - y|^\alpha \qquad \text{for all } x, y \in \Omega.$$

The space of Hölder continuous functions on Ω with fixed **Hölder exponent** $\alpha \in (0, 1]$ is denoted by $C^{0,\alpha}(\Omega)$.

Hölder continuity can be quantified by the **Hölder seminorms**, defined for each $\alpha \in (0, 1]$ by

$$|f|_{C^{0,\alpha}} := \sup_{x \neq y \in \Omega} \frac{|f(x) - f(y)|}{|x - y|^\alpha},$$

thus a continuous function is α-Hölder continuous if and only if $|f|_{C^{0,\alpha}} < \infty$. This is a seminorm rather than a norm since it vanishes for constant functions, even if they are nonzero. A norm on the space $C^{0,\alpha}(\Omega)$ can then be defined by

$$\|f\|_{C^{0,\alpha}} := \|f\|_{C^0} + |f|_{C^{0,\alpha}}.$$

Exercise 12.35. Prove:

(a) $|\cdot|_{C^{0,\alpha}}$ is a seminorm.
(b) If $f_n \in C^{0,\alpha}(\Omega)$ converges uniformly to $f \in C^0(\Omega)$ and satisfies a uniform bound $|f_n|_{C^{0,\alpha}} \leqslant C$ for all n, then $f \in C^{0,\alpha}(\Omega)$.
(c) The norm $\|\cdot\|_{C^{0,\alpha}}$ on $C^{0,\alpha}(\Omega)$ is complete, i.e. $C^{0,\alpha}(\Omega)$ is a Banach space.
Hint: Show that if f_n is C^0-convergent to f and $|f_n - f_m|_{C^{0,\alpha}} < \epsilon$ holds for all $m, n \geqslant N$, then $|f - f_n|_{C^{0,\alpha}} \leqslant \epsilon$ holds for all $n \geqslant N$. Here is a start:

$$|(f - f_n)(x) - (f - f_n)(y)| \leqslant |(f - f_m)(x)| + |(f_m - f_n)(x) - (f_m - f_n)(y)| + |(f_m - f)(y)|.$$

Keep in mind that after fixing $n \geqslant N$ and $x \neq y$, m can be chosen arbitrarily large.

For functions that can be written down in simple formulas, it is typically easy to prove a $C^{0,1}$-bound by differentiating and bounding the derivative. As the example of the Weierstrass function in §11.8 shows, this trick cannot be relied upon for functions that arise as uniform limits of sequences. This is precisely the situation in which one often encounters functions that are Hölder but not necessarily Lipschitz continuous, and the following lemma provides a useful tool to recognize this.

Lemma 12.36. *Suppose f_k is a sequence of continuous functions on $\Omega \subset \mathbb{R}^n$ converging uniformly to f, and there exist constants $a > 1$, $b \geqslant 1$, $C > 0$ and $\beta \in (0, 1]$ such that*

$$\|f - f_k\|_{C^0} \leqslant \frac{C}{a^k} \qquad \text{and} \qquad |f_k|_{C^{0,\beta}} \leqslant C b^k.$$

Then $f \in C^{0,\alpha}(\Omega)$ *for* $\alpha := \dfrac{\beta}{1+\log_a b}$.

EXERCISE 12.37. Fill in the gaps in the following proof of Lemma 12.36. The estimate $|f(x) - f(y)| \leqslant C|x-y|^\alpha$ only needs to be proved for all $x, y \in \Omega$ with $0 < |x-y| \leqslant c$ for some constant $c > 0$. For any $k \in \mathbb{N}$, we have

$$|f(x) - f(y)| \leqslant |f(x) - f_k(x)| + |f_k(x) - f_k(y)| + |f_k(y) - f(y)| \leqslant \frac{2C}{a^k} + Cb^k|x-y|^\beta$$

for all $x, y \in \Omega$. Assuming $0 < |x-y| \leqslant c$ for some $c > 0$ sufficiently small, choose $k \in \mathbb{N}$ such that $\frac{1}{(ab)^{k+1}} \leqslant |x-y|^\beta \leqslant \frac{1}{(ab)^k}$. Use this to show $|f(x) - f(y)| \leqslant \frac{3aC}{a^{k+1}}$, and then use the identity $a^{1+\log_a b} = ab$.

EXERCISE 12.38. The Cantor function $f : [0,1] \to \mathbb{R}$ from Example 9.2 satisfies $f(1/3^n) = 1/2^n$ for every $n \in \mathbb{N}$. Use this to prove $f \notin C^{0,\alpha}([0,1])$ for $\alpha > \log_3 2$. Then show that the C^0-convergent sequence f_n in Example 9.2 satisfies $|f_n|_{C^{0,1}} = (3/2)^n$, and use it to prove $f \in C^{0,\alpha}([0,1])$ for all $\alpha \leqslant \log_3 2$.

EXERCISE 12.39. For any $\theta \in (0,1)$, there is a distinguished set $C_\theta \subset [0,1]$ of full measure such that $C_{1/3}$ is the usual Cantor ternary set: it is constructed by an inductive procedure in which at step $n \in \mathbb{N}$, one removes from the middle of each of 2^{n-1} intervals of identical lengths l_n an open interval of length θl_n. Follow this idea to its logical conclusion in order to prove the following statement: for every $\alpha_0 \in (0,1)$, there exists a surjective increasing function $f : [0,1] \to [0,1]$ such that $f \in C^{0,\alpha}([0,1])$ if and only if $\alpha \leqslant \alpha_0$, and f has vanishing derivative almost everywhere (and is therefore not absolutely continuous).

EXERCISE 12.40. Show that for $b \geqslant a > 1$, the Weierstrass function $f(x) = \sum_{k=0}^\infty \frac{1}{a^k} e^{2\pi i b^k x}$ belongs to $C^{0,\alpha}(\mathbb{R})$ for every $\alpha \in (0,1)$ with $\alpha \leqslant \log_b a$.
Remark: f is nowhere differentiable by Theorem 11.22, so it cannot be absolutely continuous and therefore (by Exercise 9.6) cannot be Lipschitz, even if $\log_b a = 1$.

EXERCISE 12.41. Suppose $g : [0,\infty) \to \mathbb{R}$ is a strictly increasing smooth function with $g^{(k)}(0) = 0$ for all $k \geqslant 0$, e.g. one can take $g(x) = e^{-1/x^2}$ for $x > 0$ and $g(0) = 0$. There is a unique extension of g to an odd function $\mathbb{R} \to \mathbb{R}$, which is also strictly increasing and continuous, so it admits a continuous inverse $f := g^{-1} : I \to \mathbb{R}$ on a sufficiently small interval $I = [-a,a]$, $a > 0$. Prove that f is absolutely continuous on I, but does not belong to $C^{0,\alpha}(I)$ for any $\alpha \in (0,1]$.
Hint: The vanishing of $g^{(k)}(0)$ implies an estimate of the form $|x|^{1/k} \leqslant c_k|f(x)|$ for some constant $c_k > 0$. For absolute continuity, prove directly that f satisfies the fundamental theorem of calculus, starting from the fact that this is clearly true on any interval not containing 0.

If the domain $\Omega \subset \mathbb{R}^n$ is open, then we can also discuss differentiability of functions on Ω and define for C^m-functions the norm

$$\|f\|_{C^{m,\alpha}} := \|f\|_{C^m} + \sum_{|\beta|=m} |\partial^\beta f|_{C^{0,\alpha}},$$

where the sum ranges over all multi-indices β of order m. This norm is finite if and only if f is of class C^m with bounded and α-Hölder continuous partial derivatives up to order m. (Note that the Hölder continuity of derivatives of order less than m follows already from the fact that derivatives of higher order are bounded, so the norm does not need to include any terms $|\partial^\beta f|_{C^{0,\alpha}}$ with $|\beta| < m$.) The space of functions satisfying this condition is denoted by

$$C^{m,\alpha}(\Omega) \subset C^m(\Omega).$$

EXERCISE 12.42. Prove that $C^{m,\alpha}(\Omega)$ is a Banach space for every integer $m \geqslant 0$ and $\alpha \in (0,1]$.

EXERCISE 12.43. Use the Arzelà-Ascoli theorem to prove that if Ω is an open subset of $\mathbb{R}^n$ or $\mathbb{T}^n$ with compact closure, then for every $\alpha \in (0,1]$, the obvious continuous inclusion

$$C^{0,\alpha}(\Omega) \hookrightarrow C^0(\Omega)$$

is compact. Then generalize by induction to the statement that for each integer $m \geqslant 0$ and $\alpha \in (0,1]$, the inclusion

$$C^{m,\alpha}(\Omega) \hookrightarrow C^m(\Omega)$$

is compact.

EXERCISE 12.44. Extend Exercise 12.43 to show that under the same assumption on Ω, for every integer $m \geqslant 0$ and $0 < \alpha < \beta \leqslant 1$, the obvious inclusion

$$C^{m,\beta}(\Omega) \hookrightarrow C^{m,\alpha}(\Omega)$$

is compact.
Hint: For $m = 0$, Exercise 12.43 guarantees for any $C^{0,\beta}$-bounded sequence a C^0-convergent subsequence, and Exercise 12.35 then implies that the limit is also of class $C^{0,\beta}$, though the convergence need not be in the $C^{0,\beta}$-topology. To show that the subsequence is $C^{0,\alpha}$-convergent for $\alpha < \beta$, the following relation can help:

$$\frac{|f(x) - f(y)|}{|x-y|^\alpha} = \left(\frac{|f(x) - f(y)|}{|x-y|^\beta}\right)^{\alpha/\beta} \cdot |f(x) - f(y)|^{1-\frac{\alpha}{\beta}}.$$

Exercise 12.43 holds in particular for fully periodic functions on $\mathbb{R}^n$ since $\mathbb{T}^n$ is compact. Thus Corollary 12.33 can now be seen as a consequence of the following enhancement of the Sobolev embedding theorem (Theorem 12.10):

THEOREM 12.45. *Assume $n \in \mathbb{N}$, $s > 0$ and $\alpha \in (0,1)$ satisfy $\alpha \leqslant s - \frac{n}{2}$. Then there exist continuous inclusions*

$$H^{s+m}(\mathbb{R}^n) \hookrightarrow C^{m,\alpha}(\mathbb{R}^n) \quad \textit{and} \quad H^{s+m}(\mathbb{T}^n) \hookrightarrow C^{m,\alpha}(\mathbb{T}^n)$$

for every integer $m \geqslant 0$.

REMARK 12.46. Note that Theorem 12.45 only gives us something new when $0 < s - n/2 \leqslant 1$, as the case $s - n/2 > 1$ is already handled by Theorem 12.10, which gives an inclusion $H^{s+m} \hookrightarrow C^{m+1}$ and therefore also into $C^{m,\alpha}$ for every $\alpha \in (0,1]$. In the case $s - n/2 = 1$, one should be careful to note that α is not allowed to equal 1, so we are not claiming anything about an inclusion $H^{s+m} \hookrightarrow C^{m,1}$. We will point out the specific step in the proof below that would fail if $\alpha = 1$, and an actual counterexample to the statement for this case may be found in Example 12.47.

PROOF OF THEOREM 12.45. We will establish the inclusion $H^s(\mathbb{R}^n) \hookrightarrow C^{0,\alpha}(\mathbb{R}^n)$ for $\alpha \in (0,1)$ with $\alpha \leqslant s - n/2$ and leave the remaining cases as exercises. In light of the inclusions $H^t \hookrightarrow H^s$ for $t > s$, we can assume

$$0 < s - n/2 = \alpha < 1$$

without loss of generality. Then Theorem 12.10 already implies that $f \in H^s(\mathbb{R}^n)$ is continuous and satisfies an estimate of the form $\|f\|_{C^0} \leqslant C\|f\|_{H^s}$, so our remaining task is to find a similar bound for its Hölder seminorm $|f|_{C^{0,\alpha}}$. In other words, we need to find a constant $C > 0$ independent of $f \in H^s(\mathbb{R}^n)$ such that

$$|f(x+y) - f(x)| \leqslant C\|f\|_{H^s} \cdot |y|^\alpha \quad \text{for all} \quad x, y \in \mathbb{R}^n \text{ with } y \neq 0.$$

The proof of Theorem 12.10 shows that $\hat{f} \in L^1(\mathbb{R}^n)$, thus we can write down the usual integral formula for f in terms of $\hat{f}$ and use the assumption $\|f\|_{H^s} = \|(1+|p|^2)^{s/2}\hat{f}\|_{L^2} < \infty$ to apply the Cauchy-Schwarz inequality:

(12.7)

$$\begin{aligned}|f(x+y)-f(x)| &= \left|\int_{\mathbb{R}^n} e^{2\pi i p\cdot(x+y)}\hat{f}(p)\,dp - \int_{\mathbb{R}^n} e^{2\pi i p\cdot x}\hat{f}(p)\,dp\right| \leqslant \int_{\mathbb{R}^n}\left|e^{2\pi i p\cdot y}-1\right|\cdot|\hat{f}(p)|\,dp \\ &= \int_{\mathbb{R}^n}\frac{\left|e^{2\pi i p\cdot y}-1\right|}{(1+|p|^2)^{s/2}}\cdot(1+|p|^2)^{s/2}|\hat{f}|\,dp \leqslant \left(\int_{\mathbb{R}^n}\frac{\left|e^{2\pi i p\cdot y}-1\right|^2}{(1+|p|^2)^s}\,dp\right)^{1/2}\cdot\|f\|_{H^s} \\ &\leqslant \left(\int_{\mathbb{R}^n}\frac{\left|e^{2\pi i p\cdot y}-1\right|^2}{|p|^{2s}}\,dp\right)^{1/2}\cdot\|f\|_{H^s}\end{aligned}$$

To estimate the integral in the last line, we first observe that the function $\mathbb{R}\to\mathbb{C}: t\mapsto e^{2\pi i t}$ has globally bounded derivative $2\pi i e^{2\pi i t}$ and thus satisifes $|e^{2\pi i t}-1| \leqslant 2\pi|t|$ for all $t\in\mathbb{R}$, implying

$$\left|e^{2\pi i p\cdot y}-1\right| \leqslant 2\pi|p\cdot y| \leqslant 2\pi|p|\cdot|y|. \tag{12.8}$$

Now partition $\mathbb{R}^n$ into the domains

$$E_0 := \left\{p\in\mathbb{R}^n \;\middle|\; |p| \leqslant 1/|y|\right\} \qquad\text{and}\qquad E_\infty := \left\{p\in\mathbb{R}^n \;\middle|\; |p| > 1/|y|\right\},$$

and let $\mathrm{Vol}(S^{n-1})$ denote the $(n-1)$-dimensional volume of the unit sphere in $\mathbb{R}^n$. Integrating in n-dimensional polar coordinates then gives

$$\begin{aligned}\int_{E_0}\frac{\left|e^{2\pi i p\cdot y}-1\right|^2}{|p|^{2s}}\,dp &\leqslant 4\pi^2|y|^2\int_{E_0}\frac{1}{|p|^{2s-2}}\,dp = 4\pi^2\,\mathrm{Vol}(S^{n-1})\cdot|y|^2\int_0^{1/|y|}\frac{r^{n-1}}{r^{2s-2}}\,dr \\ &= 4\pi^2\,\mathrm{Vol}(S^{n-1})\cdot|y|^2\int_0^{1/|y|} r^{n-2s+1}\,dr = \frac{4\pi^2\,\mathrm{Vol}(S^{n-1})}{n-2s+2}\cdot|y|^2\frac{1}{|y|^{n-2s+2}} \\ &= \frac{2\pi^2\,\mathrm{Vol}(S^{n-1})}{1-\alpha}\cdot|y|^{2\alpha},\end{aligned}$$

where the convergence of $\int_0^{1/|y|} r^{n-2s+1}\,dr$ relies on the assumption $n-2s+2 = 2(1-\alpha) > 0$. (This step in the proof would fail if we allowed $\alpha = 1$.) On E_∞, the estimate (12.8) is not useful since $|p|$ may be large, so instead we use the simpler estimate $|e^{2\pi i p\cdot y}-1| \leqslant 2$ arising from the triangle inequality, and the convergence of the integral will depend on the assumption $n-2s = -2\alpha < 0$:

$$\begin{aligned}\int_{E_\infty}\frac{\left|e^{2\pi i p\cdot y}-1\right|^2}{|p|^{2s}}\,dp &\leqslant 4\int_{E_\infty}\frac{1}{|p|^{2s}}\,dp = 4\,\mathrm{Vol}(S^{n-1})\int_{1/|y|}^\infty\frac{r^{n-1}}{r^{2s}}\,dr = 4\,\mathrm{Vol}(S^{n-1})\int_{1/|y|}^\infty r^{n-1-2s}\,dr \\ &= \frac{4\,\mathrm{Vol}(S^{n-1})}{n-2s}r^{n-2s}\bigg|_{r=1/|y|}^{r=\infty} = \frac{4\,\mathrm{Vol}(S^{n-1})}{2\alpha}\frac{1}{|y|^{-2\alpha}} = \frac{2\,\mathrm{Vol}(S^{n-1})}{\alpha}|y|^{2\alpha}.\end{aligned}$$

Putting both pieces of the integral together gives an estimate $\int_{\mathbb{R}^n}\frac{\left|e^{2\pi i p\cdot y}-1\right|^2}{|p|^{2s}}\,dp \leqslant c|y|^{2\alpha}$ for a suitable constant $c > 0$, so plugging this into (12.7) gives the result we were hoping for. □

EXAMPLE 12.47. Let $f(x) := \sum_{k=2}^\infty \frac{e^{2\pi i k x}}{k^2\ln k}$. Up to multiplication by a constant, differentiating this series term by term gives the Fourier series of Exercise 12.9, so $f \in H^{3/2}(S^1)$, and Theorem 12.45 thus implies $f \in C^{0,\alpha}(S^1)$ for every $\alpha\in(0,1)$. One can also show as in Exercise 12.8 that the differentiated series converges uniformly on compact subsets of $\{x \neq 0\}$, thus f is continuously differentiable on $S^1\backslash\{0\}$. But its derivative blows up at $x = 0$, showing that $f \notin C^{0,1}(S^1)$.

REMARK 12.48. One should not assume that the constants in Theorem 12.45 are always optimal. Consider for instance the Weierstrass function $f(x) = \sum_{k=0}^{\infty} \frac{1}{a^k} e^{2\pi i b^k x}$ for $b \in \mathbb{N}$ with $1 < \sqrt{b} < a \leqslant b$. According to Exercise 12.6, $f \in H^s(S^1)$ if and only if $s < \log_b a$. Since $a > \sqrt{b}$, this range includes values $s > 1/2$, so Theorem 12.45 implies $f \in C^{0,\alpha}(S^1)$ for all $\alpha < \log_b a - \frac{1}{2}$. But in fact, Exercise 12.40 shows that $f \in C^{0,\alpha}(S^1)$ for all $\alpha < \log_b a$.

12.9. Elliptic regularity. To demonstrate the power of the Fourier transform and Sobolev spaces, in this section we shall give a brief taste of the theory of elliptic PDEs.

To understand the goal, consider first a second-order ordinary differential equation of the form

$$\ddot{x}(t) = F(x(t), \dot{x}(t)) \tag{12.9}$$

for paths $x : (-\epsilon, \epsilon) \to \mathbb{R}^n$, where $F : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ is a function of class C^m, $1 \leqslant m \leqslant \infty$. A solution to this equation must by definition be twice differentiable at every point, but it is easy to see that it must in fact be better, i.e. more "regular," which for this discussion you can take to be a synonym for "smoother". Indeed, if $\ddot{x}$ always exists, then x and $\dot{x}$ are both continuous, and (12.9) thus implies that $\ddot{x}$ is continuous, hence x is of class C^2 and $\dot{x}$ is of class C^1. Since we also assumed $F \in C^1$, thus implies $F \circ (x, \dot{x}) \in C^1$ and therefore $\ddot{x} \in C^1$, so x is of class C^3. One can repeat this argument until F runs out of derivatives. The conclusion is that if the data in the equation is of class C^m, then any solution must be at least two steps more regular, namely of class C^{m+2}; in particular, if F is smooth, then so is x. This is true even though the equation itself makes sense for any function x that is everywhere twice differentiable.

The following example shows that *partial* differential equations, by contrast, do not always have this "regularizing" property.

EXAMPLE 12.49. The simplest version of the **wave equation** is the second-order PDE

$$\partial_t^2 u - \partial_x^2 u = 0$$

for a function $u : \mathbb{R}^2 \to \mathbb{R}$ of two variables $(t, x) \in \mathbb{R}^2$. For any C^2-function $f : \mathbb{R} \to \mathbb{R}$, the wave equation has solutions given by

$$u(t, x) := f(t \pm x).$$

Notice that although the wave equation is linear with constant (and thus smooth) coefficients, its solutions need not be smooth; the function $f \in C^2(\mathbb{R})$ can be chosen arbitrarily, and the solution u will then have only as many derivatives as f does.

There is a special class of PDEs, called *elliptic*, that do exhibit the same regularizing behavior as ODEs. For this discussion, we shall only consider the simplest and most popular example: the **Poisson equation**

$$\Delta f := \sum_{j=1}^{n} \partial_j^2 f = g,$$

where $g : \mathbb{R}^n \to \mathbb{R}$ is a given function, and the solution is meant to be a function $f : \mathbb{R}^n \to \mathbb{R}$. The second-order differential operator $-\Delta := -\sum_{j=1}^n \partial_j^2$ is called the **Laplacian**, and arises often in physics (e.g. in the study of electrostatic or gravitational potentials), as well as in differential geometry.[27] We shall consider the Poisson equation on the torus $\mathbb{T}^n$, that is, we assume $g : \mathbb{R}^n \to \mathbb{R}$ is a fully periodic function and consider solutions $f : \mathbb{R}^n \to \mathbb{R}$ that are also fully periodic.

THEOREM 12.50. *For any integer $m \geqslant 0$ and smooth function $g : \mathbb{T}^n \to \mathbb{R}$, all C^2-solutions $f : \mathbb{T}^n \to \mathbb{R}$ to the equation $\Delta f = g$ are also smooth.*

[27] The minus sign in the definition of the Laplace operator appears in some sources and not in others. It is appropriate if one wants to consider the *spectrum* of the operator: the minus sign ensures that all of its eigenvalues are positive. For our present discussion this makes no difference.

I would encourage the reader at this point to take out a piece of paper and consider whether Theorem 12.50 might be proved by some trick as simple as the ODE discussion at the top of this subsection. You will quickly run into difficulties, because the Laplace operator Δ gives us information about a particular linear combination of second partial derivatives of a solution f, but we cannot deduce from this anything about any individual partial derivative. From this perspective, Theorem 12.50 is a very surprising result. It follows from the next theorem, which is of a slightly more technical nature since it involves Sobolev spaces. To prepare the statement, observe that by Proposition 12.4, Δ defines a bounded linear operator

$$\Delta : H^2(\mathbb{T}^n) \to L^2(\mathbb{T}^n),$$

which is defined in the obvious way on the dense subspace $C^\infty(\mathbb{T}^n)$ but requires Fourier transforms in order to define Δf for $f \in H^2(\mathbb{T}^n)\backslash C^2(\mathbb{T}^n)$. Recall that functions in $H^2(\mathbb{T}^n)$ need not be twice differentiable in general; when $n > 3$, they need not even be continuous (cf. Theorem 12.10).

THEOREM 12.51. *If $m \in \mathbb{N}$ and $f \in H^2(\mathbb{T}^n)$ satisfies $\Delta f \in H^m(\mathbb{T}^n)$, then $f \in H^{m+2}(\mathbb{T}^n)$.*

To prove Theorem 12.50 from this statement, observe that if $g \in C^m(\mathbb{T}^n)$, then $g \in H^m(\mathbb{T}^n)$ since g has derivatives up to order m that are continuous, and therefore also in $L^2(\mathbb{T}^n)$. If $f \in C^2(\mathbb{T}^n)$ satisfies $\Delta f = g \in C^\infty(\mathbb{T}^n)$, it follows that $f \in H^2(\mathbb{T}^n)$ and $\Delta f \in H^m(\mathbb{T}^n)$ for every $m \in \mathbb{N}$. Theorem 12.51 then implies $f \in H^{m+2}(\mathbb{T}^n)$, thus f belongs to *all* of the Sobolev spaces $H^s(\mathbb{T}^n)$ for $s \geqslant 0$, and is therefore smooth according to the Sobolev embedding theorem (Theorem 12.10).

PROOF OF THEOREM 12.51. Suppose $f \in H^2(\mathbb{T}^n)$ and $\Delta f = g \in H^m(\mathbb{T}^n)$ for $m \in \mathbb{N}$. The Fourier coefficients of f and g are then related by

$$\widehat{\Delta f}_k = \sum_{j=1}^n \widehat{\partial_j^2 f}_k = \sum_{j=1}^n (2\pi i k_j)^2 \hat{f}_k = -4\pi^2 |k|^2 \hat{f}_k = \hat{g}_k$$

for all $k \in \mathbb{Z}^n$. For $s \leqslant m+2$, this implies

$$\begin{aligned}\sum_{k \in \mathbb{Z}^n\backslash\{0\}} |k|^{2s} |\hat{f}_k|^2 &= C \sum_{k \in \mathbb{Z}^n\backslash\{0\}} \frac{|k|^{2s}}{|k|^4} |\hat{g}_k|^2 = C \sum_{k \in \mathbb{Z}^n\backslash\{0\}} |k|^{2(s-2)} |\hat{g}_k|^2 \leqslant C \sum_{k \in \mathbb{Z}^n\backslash\{0\}} |k|^{2m} |\hat{g}_k|^2 \\ &\leqslant C \|g\|_{H^m}^2\end{aligned}$$

for a suitable constant $C > 0$, thus $f \in H^{m+2}(\mathbb{T}^n)$. □

13. Distributions

Throughout this section, we assume unless stated otherwise that

$$\Omega \subset \mathbb{R}^n$$

is an open subset, and we again consider functions on Ω with values in an arbitrary finite-dimensional inner product space $(V, \langle\ ,\ \rangle)$ over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$. At the beginning of §12, we heuristically sketched the definition of a Banach space $W^{m,p}(\Omega)$ consisting of functions in L^p that have derivatives up to order m also in L^p. Here we assume $m \geqslant 0$ is an integer and $1 \leqslant p \leqslant \infty$. The quickest rigorous definition of this space is as the closure with respect to the **$W^{m,p}$-norm**

$$\|f\|_{W^{m,p}} := \sum_{|\alpha| \leqslant m} \|\partial^\alpha f\|_{L^p}, \tag{13.1}$$

of the space of all smooth functions $f : \Omega \to V$ with $\|f\|_{W^{m,p}} < \infty$. There is nothing wrong with defining $W^{m,p}(\Omega)$ in this way, but it leaves open the question of precisely which functions actually belong to $W^{m,p}(\Omega)$. For $p = 2$ and $\Omega = \mathbb{R}^n$, we found an elegant solution to the this problem in §12 by using the Fourier transform to identify differentiation with the operation of multiplication by a polynomial, so that the space $H^m(\mathbb{R}^n) := W^{m,2}(\mathbb{R}^n)$ could be defined without having to explicitly differentiate its elements. We also saw that functions of class H^m really need not be m times differentiable, e.g. Example 12.47 describes a function in $H^1(S^1)$ that is of class C^1 on the complement of one point but has its derivative blowing up at that point. To understand this phenomenon properly in the cases $p \neq 2$ or $\Omega \subsetneq \mathbb{R}^n$ where the Fourier transform is not available, we need a new trick for talking about derivatives of functions that might not be classically differentiable.

13.1. Weak derivatives. The trick we have in mind arises from the following straightforward exercise combining Fubini's theorem with integration by parts:

Exercise 13.1. Show that if $f : \Omega \to V$ and $\varphi : \Omega \to \mathbb{K}$ are functions of class C^1 and φ has compact support in Ω, then for each $j = 1, \dots, n$,

$$\int_\Omega \varphi \cdot \partial_j f \, dm = -\int_\Omega \partial_j \varphi \cdot f \, dm.$$

Hint: The function φf has an obvious extension to a C^1-function on $\mathbb{R}^n$ that vanishes outside of Ω. Compute $\int_{\mathbb{R}^n} \partial_j(\varphi f) \, dm$.

In this exercise, requiring φ to have compact support ensures on the one hand that $\varphi \cdot \partial_j f$ and $\partial_j \varphi \cdot f$ are both Lebesgue-integrable functions, and it also eliminates the boundary terms that would otherwise appear when carrying out integration by parts. The resulting formula can be used to uniquely characterize the partial derivatives of f: namely, if $f : \Omega \to V$ and $g : \Omega \to V$ are functions of class C^1 and C^0 respectively such that

$$\int_\Omega \varphi g \, dm = -\int_\Omega \partial_j \varphi \cdot f \, dm \qquad \text{for all} \qquad \varphi \in C_0^\infty(\Omega), \tag{13.2}$$

then $g = \partial_j f$. Indeed, $h := g - \partial_j f$ is then a continuous function on Ω satisfying

$$\int_\Omega \varphi h = 0 \qquad \text{for all} \qquad \varphi \in C_0^\infty(\Omega),$$

and if $h(x) \neq 0$ for some $x \in \Omega$, then the latter relation is contradicted by any smooth bump function φ that satisfies $\varphi(x) = 1$ and vanishes outside a sufficiently small neighborhood of x.

Notice: the condition (13.2) does not explicitly mention any derivative of f. In fact, both sides of the relation are well defined as soon as f and g are locally integrable functions on Ω.

Definition 13.2. A function $f \in L^1_{\mathrm{loc}}(\Omega)$ is said to be **weakly differentiable** if there exist functions $g_1, \dots, g_n \in L^1_{\mathrm{loc}}(\Omega)$ such that for each $j = 1, \dots, n$,

$$\int_\Omega \varphi g_j \, dm = -\int_\Omega \partial_j \varphi \cdot f \, dm \qquad \text{for all} \qquad \varphi \in C_0^\infty(\Omega).^{28}$$

We then call g_j a **weak partial derivative** of f with respect to the variable x_j, and write $\partial_j f := g_j$.

Three important remarks should be understood immediately:

(1) If f is of class C^1, then its classical partial derivatives are also weak partial derivatives, thus for this class of functions there is no ambiguity in denoting weak derivatives by $\partial_j f$. (There will occasionally be ambiguity if we talk about functions that are differentiable almost everywhere—these sometimes also have weak derivatives, but sometimes they do not.)

[28] The function $\varphi \in C_0^\infty(\Omega)$ in this definition can be taken to have either real or complex values; it does not matter.

(2) In contrast with classical derivatives, weak derivatives are well defined only up to equality almost everywhere, i.e. if $g = \partial_j f$ weakly and $h = g$ almost everywhere, then h is also a weak derivative of f. Similarly, f can be changed on a set of measure zero without changing its weak derivatives.

(3) Related to the second point: weak differentiability is a property of the whole function $f \in L^1_{\text{loc}}(\Omega)$, and it is not purely *local*, i.e. it generally makes no sense to ask whether f is weakly differentiable at an individual point $x \in \Omega$, nor what the value of $\partial_j f(x)$ is, though one *can* ask what $\int_E \partial_j f \, dm$ is for any given measurable subset $E \subset \Omega$.

Since weak derivatives of locally integrable functions are also locally integrable functions, one can iterate the definition in obvious ways to define higher-order weak differentiability and weak derivatives $\partial^\alpha f$, which will be uniquely characterized by the relation

$$\int_\Omega \varphi \cdot \partial^\alpha f \, dm = (-1)^{|\alpha|} \int_\Omega \partial^\alpha \varphi \cdot f \, dm \qquad \text{for all} \qquad \varphi \in C_0^\infty(\Omega).$$

There is again no problem in making sense of this condition since φ is always assumed to be infinitely differentiable with compact support; we only need f and $\partial^\alpha f$ to be of class L^1_{loc}.

Let us clarify that a function may indeed have a weak derivative without being classically differentiable:

EXERCISE 13.3. Show that the function $f : \mathbb{R} \to \mathbb{R} : x \mapsto |x|$ has weak derivative $f'(x) := x/|x|$. (It is not necessary to specify a value for $f'(0)$ since $\{0\} \subset \mathbb{R}$ is a set of measure zero.) Then show that $f' \in L^1_{\text{loc}}(\mathbb{R})$ is not weakly differentiable.

For functions that are not of class C^1, we have not yet shown that weak derivatives are uniquely defined almost everywhere, but this is true, and follows from:

LEMMA 13.4. *If $f \in L^1_{\text{loc}}(\Omega)$ satisfies $\int \varphi f \, dm = 0$ for every smooth compactly supported function $\varphi : \Omega \to \mathbb{R}$, then $f = 0$ almost everywhere.*

PROOF. Given $x_0 \in \Omega$, choose $\epsilon > 0$ small enough so that the closed ϵ-ball $\bar{B}_\epsilon(x_0)$ about x_0 lies in Ω, and consider the function $g \in L^1(\mathbb{R}^n)$ defined by

$$g := \begin{cases} f & \text{on } \overline{B_\epsilon(x_0)}, \\ 0 & \text{everywhere else.} \end{cases}$$

Choose an approximate identity $\rho_j : \mathbb{R}^n \to [0, \infty)$ with shrinking support. For $x \in B_{\epsilon/2}(x_0)$ and j sufficiently large so that $\text{supp}(\rho_j) \subset B_{\epsilon/2}(0)$, the function $\rho_j(x - \cdot) : \mathbb{R}^n \to [0, \infty)$ then has compact support in $B_{\epsilon/2}(x) \subset B_\epsilon(x_0)$ and can therefore be regarded as an element of $C_0^\infty(\Omega)$, implying that the convolution

$$(\rho_j * g)(x) = \int_{\mathbb{R}^n} \rho_j(x - y) g(y) \, dy = \int_{B_{\epsilon/2}(x)} \rho_j(x - y) g(y) \, dy = \int_\Omega \rho_j(x - \cdot) f \, dm$$

vanishes for $x \in B_{\epsilon/2}(x_0)$. By Theorem 8.14, $\rho_j * g \to g$ in $L^1(\mathbb{R}^n)$ as $j \to \infty$, so we conclude that g (and therefore f) vanishes almost everywhere on $B_{\epsilon/2}(x_0)$. Since Ω can be covered by countably many subsets of the form $B_{\epsilon/2}(x_0)$ with $x_0 \in \Omega$ and $\epsilon > 0$, it follows that f vanishes almost everywhere in Ω. □

COROLLARY 13.5. *If $f \in L^1_{\text{loc}}(\Omega)$ is weakly differentiable, then its weak partial derivatives are unique up to equality almost everywhere.* □

EXERCISE 13.6. Consider the Cantor function f from Example 9.2 on the domain $\Omega := (0, 1) \subset \mathbb{R}$, which has classical derivative $f' = 0$ at almost every point. Show however that f is not weakly differentiable.

Hint: Show first that if a weak derivative existed, it would necessarily vanish almost everywhere on each of the intervals that are removed to form the Cantor set.

13.2. Test functions and the space of distributions. Let us fit the notion of weak derivatives into a larger context. We saw in Exercise 13.3 that locally integrable functions can be weakly differentiable without being classically differentiable, but also that not *all* functions in L^1_{loc} have weak derivatives. We will now see that if our notion of what a "function" can be is suitably enlarged, then *every* L^1_{loc} function can be understood to have a derivative in some sense.

The key observation is that for weak differentiation, what matters is not the values of a function $f : \Omega \to V$ at points in Ω, but rather the values of the linear map

$$\Lambda_f : C_0^\infty(\Omega) \to V : \varphi \mapsto \int_\Omega \varphi f.$$

This suggests that instead of talking about functions on Ω, we should talk about linear maps $C_0^\infty(\Omega) \to V$, e.g. in the case $V = \mathbb{K}$, we are talking about the dual space of $C_0^\infty(\Omega)$. To do this properly, we should consider only linear functionals that are continuous, which requires endowing $C_0^\infty(\Omega)$ with a topology.

Definition 13.7. A **test function** on Ω is defined to be a smooth function $\varphi : \Omega \to \mathbb{K}$ with compact support, and the vector space of all such functions is denoted by $\mathscr{D}(\Omega)$. A sequence $\varphi_j \in \mathscr{D}(\Omega)$ is said to **converge** to $\varphi_\infty \in \mathscr{D}(\Omega)$ if there exists a compact subset $K \subset \Omega$ such that φ_j has support contained in K for every $j \in \mathbb{N} \cup \{\infty\}$ and $\partial^\alpha \varphi_j$ converges uniformly to $\partial^\alpha \varphi_\infty$ for every multi-index α. A $\mathbb{K}$-linear map $\Lambda : \mathscr{D}(\Omega) \to V$ is then said to be **continuous** if and only if $\Lambda(\varphi_j) \to \Lambda(\varphi_\infty)$ for every convergent sequence $\varphi_j \to \varphi_\infty \in \mathscr{D}(\Omega)$.

Putting Definition 13.7 on firm mathematical footing requires the following result, whose proof is outsourced to §13.8 in order to avoid too much of a digression into abstract topology:

Proposition 13.8 (see §13.8). *The space of test functions $\mathscr{D}(\Omega)$ admits a natural topology that induces the notions of convergence and continuity described in Definition 13.7.*

Definition 13.9. A scalar-valued **distribution** on Ω is a continuous $\mathbb{K}$-linear functional $\Lambda : \mathscr{D}(\Omega) \to \mathbb{K}$. Similarly, a vector-valued distribution with values in the finite-dimensional vector space V over $\mathbb{K}$ is a continuous $\mathbb{K}$-linear map $\Lambda : \mathscr{D}(\Omega) \to V$. We shall generally assume that all our distributions take values in a fixed vector space V, and denote the the space of vector-valued distributions by

$$\mathscr{D}'(\Omega) = \left\{ \Lambda : \mathscr{D}(\Omega) \to V \;\middle|\; \Lambda \text{ is } \mathbb{K}\text{-linear and continuous} \right\}.$$

The space $\mathscr{D}'(\Omega)$ is endowed with the weak*-topology, i.e. the locally convex topology generated by the seminorms $\|\Lambda\|_\varphi := |\Lambda(\varphi)|$ for all $\varphi \in \mathscr{D}(\Omega)$. In particular, a sequence $\Lambda_j \in \mathscr{D}'(\Omega)$ converges to $\Lambda_\infty \in \mathscr{D}'(\Omega)$ if and only if $\Lambda_j(\varphi) \to \Lambda_\infty(\varphi)$ for every $\varphi \in \mathscr{D}(\Omega)$.

Remark 13.10. If V is a complex vector space, then one can regard it as a real vector space (of twice the dimension) and set $\mathbb{K} = \mathbb{R}$ without changing any result in the theory of distributions. The reason is that every real-linear map from the space of real-valued test functions to a complex vector space has a unique complex-linear extension to the space of complex-valued test functions. Thus for most purposes, it makes no difference whether we set $\mathbb{K}$ to be $\mathbb{R}$ or $\mathbb{C}$, and many books on distributions treat only the case $\mathbb{K} = \mathbb{R}$. We will need to set $\mathbb{K} = \mathbb{C}$ however when we discuss Fourier transforms in §13.6.

Note that choosing a basis of V identifies each vector-valued distribution with a finite tuple of scalar-valued distributions, just as for vector-valued functions. Since the choice of the space V almost never plays an important role in our discussion, we shall suppress it from the notation whenever possible.

EXAMPLE 13.11. There is a natural linear map

$$L^1_{\mathrm{loc}}(\Omega) \to \mathscr{D}'(\Omega) : f \mapsto \Lambda_f, \qquad \Lambda_f(\varphi) := \int_\Omega \varphi f \, dm,$$

and by Lemma 13.4, this map is injective. (Exercise: check that $\Lambda_f : \mathscr{D}(\Omega) \to V$ is continuous.) In this way, every locally integrable function determines a distribution, and we shall often abuse terminology by identifying one with the other, e.g. when we say "$\Lambda \in \mathscr{D}'(\Omega)$ is a function," we mean that there exists a (necessarily unique up to equality almost everywhere) function $f \in L^1_{\mathrm{loc}}(\Omega)$ such that $\Lambda = \Lambda_f$.

CONVENTION 13.12. We will sometimes also denote the action of a distribution $\Lambda \in \mathscr{D}'(\Omega)$ on test functions $\varphi \in \mathscr{D}(\Omega)$ by

$$(\Lambda, \varphi) := \Lambda(\varphi),$$

and abbreviate the case of a locally integrable function $f \in L^1_{\mathrm{loc}}(\Omega)$ by

$$(f, \varphi) := \Lambda_f(\varphi) := \int_\Omega \varphi f \, dm.$$

EXERCISE 13.13. Show that the map $L^1_{\mathrm{loc}}(\Omega) \to \mathscr{D}'(\Omega)$ in Example 13.11 is continuous, where $L^1_{\mathrm{loc}}(\Omega)$ is endowed with the Fréchet space topology defined in §4.3.

EXAMPLE 13.14. The most popular scalar-valued distribution that is not a function is what physicists call the **Dirac delta function**: for each $x \in \Omega$, we define $\delta_x \in \mathscr{D}'(\Omega)$ by

$$\delta_x(\varphi) := \varphi(x).$$

On $\Omega = \mathbb{R}^n$, one typically abbreviates $\delta := \delta_0$ for the δ-function centered at the origin, so that pretending δ is an actual function on $\mathbb{R}^n$ gives rise to the usual formula $\int_{\mathbb{R}^n} \varphi(x)\delta(x) \, dx = \varphi(0)$. A formal change of variables transforms this into $\delta_x(\varphi) = \varphi(x) = \int_{\mathbb{R}^n} \varphi(y+x)\delta(y) \, dy = \int_{\mathbb{R}^n} \varphi(u)\delta(u-x) \, du$, motivating the notation

$$\delta(\cdot - x) := \delta_x \in \mathscr{D}'(\Omega).$$

EXAMPLE 13.15. Suppose μ is a measure defined on the Borel subsets of $\Omega \subset \mathbb{R}^n$ such that $\mu(K) < \infty$ whenever $K \subset \Omega$ is compact. Then $\Lambda(\varphi) := \int_\Omega \varphi \, d\mu$ defines a real-valued distribution. The distributions Λ_f in Examples 13.11 (with $f : \Omega \to [0, \infty)$) and δ_x in Example 13.14 are both special cases of this, with measures defined by

$$\mu(E) := \int_E f \, dm \qquad \text{and} \qquad \mu(E) := \begin{cases} 1 & \text{if } x \in E, \\ 0 & \text{otherwise} \end{cases}$$

respectively. The latter is of course also known as the **Dirac measure** centered at x.

EXAMPLE 13.16. Here is a distribution that is not a special case of Example 13.15: for $k \in \mathbb{N}$ and $x \in \Omega \subset \mathbb{R}$, define

$$\Lambda(\varphi) := \varphi^{(k)}(x).$$

EXERCISE 13.17. Verify that the linear maps $\mathscr{D}(\Omega) \to V$ described in Examples 13.11, 13.14, 13.15 and 13.16 are all continuous.

The trick via integration by parts in the definition of weak differentiation now generalizes as follows.

DEFINITION 13.18. Given $\Lambda \in \mathscr{D}'(\Omega)$ and $j = 1, \ldots, n$, the **distributional derivative** (or derivative "in the sense of distributions") of Λ with respect to the variable x_j is a distribution $\partial_j \Lambda \in \mathscr{D}'(\Omega)$ defined by

$$(\partial_j \Lambda)(\varphi) := -\Lambda(\partial_j \varphi).$$

More generally, if α is any multi-index with order $|\alpha| \geqslant 0$, one defines $\partial^\alpha \Lambda \in \mathscr{D}'(\Omega)$ by

$$(\partial^\alpha \Lambda)(\varphi) := (-1)^{|\alpha|} \Lambda(\partial^\alpha \varphi).$$

It is easy to check that the distributions in Definition 13.18 are always well defined continuous linear maps, so every distribution is infinitely differentiable, and the operators $\partial^\alpha : \mathscr{D}'(\Omega) \to \mathscr{D}'(\Omega)$ are continuous linear maps. In this language, a function is weakly differentiable if and only if its derivative in the sense of distributions is also represented by a function. For functions of class C^1, the distributional derivatives can always be represented by the classical derivatives.

Exercise 13.19. The function $f(x) := x/|x|$ appeared in Exercise 13.3 as the weak derivative of the function $|x|$. Show that the derivative of f in the sense of distributions (meaning the derivative of the distribution Λ_f) is 2δ.

Example 13.20. Up to a sign, the distribution in Example 13.16 is the kth derivative of the δ-function: concretely, $\Lambda = (-1)^k \delta_x^{(k)}$.

When we talk about distributions represented by functions, we typically assume these functions to be locally integrable so that expressions like $\int_\Omega \varphi f \, dm$ make sense for all $\varphi \in \mathscr{D}(\Omega)$. This is not always strictly necessary, however: the next exercise exhibits a locally integrable function that is not weakly differentiable in the sense of Definition 13.2, but its distributional derivative can (with a little care) be represented by a function that is not of class L^1_{loc}.

Exercise 13.21. Show that the function $f(x) := \ln|x|$ is locally integrable on $\mathbb{R}$, and its derivative in $\mathscr{D}'(\mathbb{R})$ is given by

$$\Lambda_f'(\varphi) = \text{p. v.} \int_{\mathbb{R}} \frac{\varphi(x)}{x} \, dx := \lim_{\epsilon \to 0^+} \int_{|x| \geqslant \epsilon} \frac{\varphi(x)}{x} \, dx.$$

Here the notation p. v. stands for "Cauchy principal value" and is defined as the limit on the right. Check that this expression gives a well-defined distribution even though $1/x$ is not a locally integrable function on $\mathbb{R}$.

The product of a distribution $\Lambda \in \mathscr{D}'(\Omega)$ with a smooth scalar-valued function $f \in C^\infty(\Omega)$ defines a distribution $f\Lambda \in \mathscr{D}'(\Omega)$ via the obvious formula

$$(f\Lambda)(\varphi) := \Lambda(f\varphi).$$

This is well defined because $\varphi \mapsto f\varphi$ defines a continuous map $\mathscr{D}(\Omega) \to \mathscr{D}(\Omega)$; note that this depends on f having derivatives of all orders (though it does not need to have compact support), so the product of an arbitrary distribution Λ with a non-smooth function is not well defined in general. It is easy to check that for every $f \in C^\infty(\Omega)$, the linear map $\mathscr{D}'(\Omega) \to \mathscr{D}'(\Omega) : \Lambda \mapsto f\Lambda$ is also continuous.

Exercise 13.22. Show that for $f \in C^\infty(\Omega)$ and $\Lambda \in \mathscr{D}'(\Omega)$, distributional derivatives satisfy the Leibniz rule

$$\partial_j(f\Lambda) = (\partial_j f)\Lambda + f\partial_j \Lambda,$$

where on the right hand side, ∂_j denotes a classical derivative in the first term and a distributional derivative in the second.

13.3. Smoothness of distributions. For applications of distributions in the theory of PDEs, we need a more concrete understanding of the relationship between classical and distributional derivatives. This includes the answers to two questions:

- How well can an arbitrary distribution be approximated by a C^m-function? (see Corollary 13.33)

- How can one recognize whether a given distribution is representable by a C^m-function? (see Theorem 13.34)

The most useful tool toward these ends is a generalization of the convolution operator.

13.3.1. *The convolution.* Recall from §8.2 that for any locally integrable function $f : \mathbb{R}^n \to V$ and any test function $\varphi \in \mathscr{D}(\mathbb{R}^n)$, the convolution $\varphi * f : \mathbb{R}^n \to V$ is a well-defined function at every point $x \in \mathbb{R}^n$. It can be expressed in terms of the distribution $\Lambda_f \in \mathscr{D}'(\mathbb{R}^n)$ if we introduce two natural operations on the space of test functions: one is the translation operator

$$\tau_v : \mathscr{D}(\mathbb{R}^n) \to \mathscr{D}(\mathbb{R}^n), \qquad \tau_v\varphi(x) := \varphi(x+v) \quad \text{for} \quad v \in \mathbb{R}^n,$$

which we considered on L^p-spaces in §8.1. The other is the antipodal reflection operator

$$\sigma : \mathscr{D}(\mathbb{R}^n) \to \mathscr{D}(\mathbb{R}^n), \qquad \sigma\varphi(x) := \varphi(-x).$$

Both τ_v and σ extend naturally to operations on the space of distributions on $\mathbb{R}^n$. For $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ and $\varphi \in \mathscr{D}(\mathbb{R}^n)$, we have

$$\Lambda_{\tau_v f}(\varphi) = \int_{\mathbb{R}^n} \varphi(x) f(x+v)\, dx = \int_{\mathbb{R}^n} \varphi(x-v) f(x)\, dx = \Lambda_f(\tau_{-v}\varphi),$$

which motivates defining

$$\tau_v : \mathscr{D}'(\mathbb{R}^n) \to \mathscr{D}'(\mathbb{R}^n), \qquad \tau_v\Lambda := \Lambda \circ \tau_{-v}.$$

Similarly,

$$\Lambda_{\sigma f}(\varphi) = \int_{\mathbb{R}^n} \varphi(x) f(-x)\, dx = \int_{\mathbb{R}^n} \varphi(-x) f(x)\, dx = \Lambda_f(\sigma\varphi),$$

and we therefore define

$$\sigma : \mathscr{D}'(\mathbb{R}^n) \to \mathscr{D}'(\mathbb{R}^n), \qquad \sigma\Lambda := \Lambda \circ \sigma.$$

One verifies easily that τ_v and σ are each continuous linear maps on both $\mathscr{D}(\mathbb{R}^n)$ and $\mathscr{D}'(\mathbb{R}^n)$.

The convolution of $\varphi \in \mathscr{D}(\mathbb{R}^n)$ with $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ can now be expressed as

$$(\varphi * f)(x) = \int_{\mathbb{R}^n} \varphi(x-y) f(y)\, dy = \int_{\mathbb{R}^n} \sigma\varphi(y-x) f(y)\, dy = \int_{\mathbb{R}^n} \tau_{-x}\sigma\varphi(y) f(y)\, dy = \Lambda_f(\tau_{-x}\sigma\varphi).$$

It is therefore sensible to define the **convolution** of any distribution $\Lambda \in \mathscr{D}'(\mathbb{R}^n)$ with a test function $\varphi \in \mathscr{D}(\mathbb{R}^n)$ as the function $\varphi * \Lambda : \mathbb{R}^n \to V$ given by

$$(\varphi * \Lambda)(x) := \Lambda(\tau_{-x}\sigma\varphi) = \tau_x\Lambda(\sigma\varphi). \tag{13.3}$$

With a little care, this definition can be extended to include distributions that are defined only on an open subset $\Omega \subset \mathbb{R}^n$. Given subsets $A, B \subset \mathbb{R}^n$ and $v \in \mathbb{R}^n$, let us denote

$$\begin{aligned} A \pm v &:= \left\{ x \pm v \in \mathbb{R}^n \;\middle|\; x \in A \right\}, \\ A \pm B &:= \left\{ x \pm y \in \mathbb{R}^n \;\middle|\; x \in A \text{ and } y \in B \right\}, \\ -A &:= \left\{ -x \in \mathbb{R}^n \;\middle|\; x \in A \right\}. \end{aligned}$$

Then for any function $\varphi : \mathbb{R}^n \to \mathbb{K}$ with support in a subset $K \subset \mathbb{R}^n$, and for any $v \in \mathbb{R}^n$,

$$\operatorname{supp}(\varphi) \subset K \quad \Rightarrow \quad \operatorname{supp}(\tau_v\varphi) \subset K - v \quad \text{and} \quad \operatorname{supp}(\sigma\varphi) \subset -K.$$

It follows that for open sets $\Omega, \Omega' \subset \mathbb{R}^n$ and $v \in \mathbb{R}^n$, there is a well-defined continuous linear operator

$$\tau_v : \mathscr{D}(\Omega) \to \mathscr{D}(\Omega') \qquad \text{whenever} \qquad \Omega - v \subset \Omega',$$

and similarly

$$\sigma : \mathscr{D}(\Omega) \to \mathscr{D}(\Omega') \qquad \text{whenever} \qquad -\Omega \subset \Omega'.$$

Now if $\Lambda \in \mathscr{D}'(\Omega)$ and $K \subset \mathbb{R}^n$ is any compact set containing the support of $\varphi \in \mathscr{D}(\mathbb{R}^n)$, then (13.3) defines $\varphi * \Lambda$ as a function on the open set

$$\Omega' := \left\{ x \in \mathbb{R}^n \;\middle|\; -K + x \subset \Omega \right\}.$$

One must keep in mind that this set may be empty, but we will mostly be interested in situations where K is an arbitrarily small compact neighborhood of the origin, in which case Ω' is a nonempty subset of Ω. Since convolutions of functions are symmetric, we define

$$\Lambda * \varphi := \varphi * \Lambda.$$

Exercise 13.23. Prove that for $v \in \mathbb{R}^n$ and $k = 1, \ldots, n$, the operators τ_v, σ and ∂_k, acting on either the space of test functions or the space of distributions, are related to each other by

$$\tau_v \circ \partial_k = \partial_k \circ \tau_v, \qquad \sigma \circ \partial_k = -\partial_k \circ \sigma, \qquad \tau_v \circ \sigma = \sigma \circ \tau_{-v}.$$

We will see below that even in cases where Λ is not a function, the function $\varphi * \Lambda$ inherits the smoothness of φ. The proof of this rests on the smoothness of the translation operator τ_x as a function of $x \in \mathbb{R}^n$, i.e. the fact that for any fixed $\varphi \in \mathscr{D}(\mathbb{R}^n)$ and $\Lambda \in \mathscr{D}'(\Omega')$, the function $x \mapsto (\tau_x \Lambda)(\varphi) = \Lambda(\tau_{-x}\varphi)$ is smooth on a suitable open domain in $\mathbb{R}^n$. This follows in turn from a more general result related to differentiation under the integral sign.

The setting for the result we need is as follows. Assume $\mathcal{U} \subset \mathbb{R}^m$ and $\Omega \subset \mathbb{R}^n$ are open subsets, $\varphi : \mathcal{U} \times \Omega \to \mathbb{K}$ is a smooth function such that $\varphi_x := \varphi(x, \cdot) \in \mathscr{D}(\Omega)$ for every $x \in \mathcal{U}$, and $f \in L^1_{\text{loc}}(\Omega)$. One can then consider the function F on $\mathcal{U}$ defined via the parameter-dependent integral

$$F(x) := \int_\Omega \varphi(x, y) f(y)\, dy = \Lambda_f(\varphi_x).$$

If φ satisfies sufficient hypotheses for the application of Theorem 4.5, then one should expect this function to be smooth and satisfy

$$\frac{\partial^{|\alpha|} F}{\partial x^\alpha}(x) = \int_\Omega \frac{\partial^{|\alpha|}\varphi}{\partial x^\alpha}(x, y) f(y)\, dy = \Lambda_f \left(\frac{\partial^{|\alpha|}\varphi}{\partial x^\alpha}(x, \cdot) \right)$$

for every multi-index α in the variables $x = (x_1, \ldots, x_m) \in \mathcal{U} \subset \mathbb{R}^m$. It turns out that under a mild assumption about the support of φ, this also works when Λ_f is replaced by an arbitrary distribution:

Proposition 13.24. *Assume $\mathcal{U} \subset \mathbb{R}^m$ and $\Omega \subset \mathbb{R}^n$ are open subsets and $\varphi : \mathcal{U} \times \Omega \to \mathbb{K}$ is a smooth function such that for every compact set $K \subset \mathcal{U}$, $\varphi|_{K \times \Omega}$ has compact support. Then for any $\Lambda \in \mathscr{D}'(\Omega)$, the function*

$$F : \mathcal{U} \to V : x \mapsto \Lambda(\varphi(x, \cdot))$$

is smooth and satisfies

$$\frac{\partial^{|\alpha|} F}{\partial x^\alpha}(x) = \Lambda \left(\frac{\partial^{|\alpha|}\varphi}{\partial x^\alpha}(x, \cdot) \right)$$

for all multi-indices α in the variables $x = (x_1, \ldots, x_m) \in \mathcal{U} \subset \mathbb{R}^m$.

The proof requires two preparatory lemmas about the space of test functions.

Lemma 13.25. *Under the assumptions of Proposition 13.24, the map $\mathcal{U} \to \mathscr{D}(\Omega) : x \mapsto \varphi_x := \varphi(x, \cdot)$ is continuous.*

Proof. Given a convergent sequence $x_j \to x_\infty$ in $\mathcal{U}$, choose a compact set $C \subset \mathcal{U}$ containing an open neighborhood of x_∞. By assumption, there then exists a compact set $K \subset \Omega$ such that φ_x vanishes outside K for all $x \in C$, thus $\operatorname{supp}(\varphi_{x_j}) \subset K$ for all j sufficiently large. It thus remains only to prove C^∞-convergence of φ_{x_j} to φ_{x_∞}. Uniform convergence follows from the fact that since $C \times K$ is compact, φ is uniformly continuous on $C \times K$. The same argument proves uniform convergence $\partial^\alpha \varphi_{x_j} \to \partial^\alpha \varphi_{x_\infty}$ for all multi-indices α in the variables $y = (y_1, \dots, y_n) \in \Omega \subset \mathbb{R}^n$, since $\frac{\partial^{|\alpha|}\varphi}{\partial y^\alpha}$ is also continuous. $\square$

LEMMA 13.26. *Under the assumptions of Proposition 13.24, the functions $\varphi_x := \varphi(x, \cdot) : \Omega \to \mathbb{K}$ satisfy*

$$\lim_{h \to 0} \frac{\varphi_{x+he_k} - \varphi_x}{h} = \frac{\partial \varphi}{\partial x_k}(x, \cdot)$$

for every $x \in \mathcal{U}$ and $k = 1, \dots, m$, where $e_1, \dots, e_m \in \mathbb{R}^m$ denotes the standard Euclidean basis, and the convergence of the limit is in the topology of $\mathscr{D}(\Omega)$.

PROOF. Fix $x \in \mathcal{U}$ and $k \in \{1, \dots, m\}$. For all $h \in \mathbb{R}\backslash\{0\}$ close enough to 0, we can assume $x + he_k$ belongs to a compact subset in $\mathcal{U}$ such that all the functions $\varphi_{x+he_k} : \Omega \to \mathbb{R}$ have support contained in some fixed compact subset $K \subset \Omega$. Now use the fundamental theorem of calculus to write

$$\varphi_{x+he_k}(y) - \varphi_x(y) = h \int_0^1 \frac{\partial \varphi}{\partial x_k}(x + the_k, y)\, dt,$$

and note that for any multi-index α in the variables $y = (y_1, \dots, y_n) \in \Omega \subset \mathbb{R}^n$, the operator $\frac{\partial^{|\alpha|}}{\partial y^\alpha}$ can be passed under the integral sign on the right hand side since φ is smooth. We thus have

$$\left| \frac{\partial^{|\alpha|}}{\partial y^\alpha} \left(\frac{\varphi_{x+he_k}(y) - \varphi_x(y)}{h} \right) - \frac{\partial^{|\alpha|}}{\partial y^\alpha} \frac{\partial \varphi}{\partial x_k}(y) \right| \leqslant \int_0^1 \left| \frac{\partial^{|\alpha|}}{\partial y^\alpha} \frac{\partial \varphi}{\partial x_k}(x + the_k, y) - \frac{\partial^{|\alpha|}}{\partial y^\alpha} \frac{\partial \varphi}{\partial x_k}(x, y) \right| dt.$$

Since $\frac{\partial^{|\alpha|}}{\partial y^\alpha} \frac{\partial \varphi}{\partial x_k}(x + the_k, y)$ can be assumed to vanish for all $y \notin K$ and $|h|$ sufficiently small, uniform continuity implies that the integrand on the right hand side becomes arbitrarily small uniformly in $y \in \Omega$ as $h \to 0$. $\square$

PROOF OF PROPOSITION 13.24. The continuity of F follows immediately from Lemma 13.25 and the continuity of Λ. The main task is thus to prove that F has first partial derivatives given by

$$\frac{\partial F}{\partial x_k}(x) = \Lambda \left(\frac{\partial \varphi}{\partial x_k}(x, \cdot) \right),$$

since a similar application of Lemma 13.25 will then imply that these derivatives are also continuous, and the argument can be repeated inductively for all higher-order derivatives. For the computation of $\frac{\partial F}{\partial x_k}(x)$, one can again appeal to the continuity of Λ, together with Lemma 13.26, which gives

$$\frac{F(x + he_k) - F(x)}{h} = \frac{\Lambda(\varphi_{x+he_k}) - \Lambda(\varphi_x)}{h} = \Lambda \left(\frac{\varphi_{x+he_k} - \varphi_x}{h} \right) \to \Lambda \left(\frac{\partial \varphi}{\partial x_k}(x, \cdot) \right)$$

as $h \to 0$. $\square$

COROLLARY 13.27. *For an open set $\Omega \subset \mathbb{R}^n$ and a compact set $K \subset \mathbb{R}^n$, consider the open set*

$$\mathcal{U} := \left\{ x \in \mathbb{R}^n \;\middle|\; K + x \subset \Omega \right\}.$$

For any $\varphi \in \mathscr{D}(\mathbb{R}^n)$ with $\operatorname{supp}(\varphi) \subset K$, associate to each $\Lambda \in \mathscr{D}'(\Omega)$ the function F_Λ defined on $\mathcal{U}$ by

$$F_\Lambda(x) := (\tau_x \Lambda)(\varphi).$$

Then F_Λ is smooth and satisfies $\partial^\alpha F_\Lambda = F_{\partial^\alpha \Lambda}$ for every multi-index α.

PROOF. Apply Proposition 13.24 with the smooth function $\mathcal{U} \times \Omega \to \mathbb{R} : (x, y) \mapsto \varphi(y - x)$. $\square$

This is enough preparation to prove the first main result about the convolution.

THEOREM 13.28. *Suppose $\Lambda \in \mathscr{D}(\Omega)$ for an open set $\Omega \subset \mathbb{R}^n$, and $\varphi \in \mathscr{D}(\mathbb{R}^n)$ has support contained in the compact set $K \subset \mathbb{R}^n$. Then (13.3) defines a smooth function $\varphi * \Lambda$ on the open domain $\Omega' := \left\{ x \in \mathbb{R}^n \;\middle|\; -K + x \subset \Omega \right\}$, and it satisfies*

$$\partial^\alpha(\varphi * \Lambda) = (\partial^\alpha \varphi) * \Lambda = \varphi * (\partial^\alpha \Lambda)$$

for every multi-index α, where the operator ∂^α denotes a classical derivative in the first formula and a distributional derivative in the second.

PROOF. The second formula is immediate from Corollary 13.27 and the definition of the convolution. Since ∂_k commutes with translation operators and anticommutes with σ, we also have

$$\begin{aligned}(\varphi * \partial_k \Lambda)(x) &= \tau_x \partial_k \Lambda(\sigma\varphi) = \partial_k \Lambda(\tau_{-x}\sigma\varphi) = -\Lambda(\partial_k \tau_{-x}\sigma\varphi) = \Lambda(\tau_{-x}\sigma\partial_k\varphi) = \tau_x\Lambda(\sigma\partial_k\varphi) \\ &= (\partial_k\varphi * \Lambda)(x)\end{aligned}$$

for all $x \in \Omega'$ and $k = 1, \ldots, n$. The relation $\varphi * \partial^\alpha\Lambda = \partial^\alpha\varphi * \Lambda$ follows from this by induction on the order of differentiation. □

Since $\varphi * \Lambda$ is always a smooth function on Ω', it also defines an element of $\mathscr{D}'(\Omega')$. We would next like to give an alternative characterization of this distribution. For the case $\Lambda = \Lambda_f$ with $f \in L^1(\Omega)$, f can be extended to a function on $\mathbb{R}^n$ vanishing outside of Ω without changing the values of $\varphi * \Lambda_f = \varphi * f$ on Ω'. For any $\psi \in \mathscr{D}(\Omega')$, we can similarly extend ψ as 0 on $\mathbb{R}^n \backslash \Omega'$, and then use Fubini's theorem to show

$$\begin{aligned}(\varphi * f, \psi) &= \int_{\mathbb{R}^n} \psi(x)(\varphi * f)(x)\, dx = \int_{\mathbb{R}^n\times\mathbb{R}^n} \psi(x)\varphi(x-y)f(y)\, dx\, dy \\ &= \int_{\mathbb{R}^n\times\mathbb{R}^n} (\sigma\varphi)(y-x)\psi(x)f(y)\, dx\, dy = \int_{\mathbb{R}^n} (\sigma\varphi * \psi)(y)f(y)\, dy = (f, \sigma\varphi * \psi).\end{aligned}$$

It turns out that this formula remains valid when f is replaced by an arbitrary distribution. The proof requires a preparatory exercise.

EXERCISE 13.29. Show that for any $\varphi, \psi \in \mathscr{D}(\mathbb{R}^n)$ with $\operatorname{supp}(\varphi) \subset K \subset \mathbb{R}^n$ and $\operatorname{supp}(\psi) \subset K' \subset \mathbb{R}^n$, $\varphi * \psi$ is also in $\mathscr{D}(\mathbb{R}^n)$ and has $\operatorname{supp}(\varphi * \psi) \subset K + K'$. Moreover, if ψ_j is a sequence converging to ψ in $\mathscr{D}(\mathbb{R}^n)$, then $\varphi * \psi_j \to \varphi * \psi$ in $\mathscr{D}(\mathbb{R}^n)$.
Hint: Focus on proving uniform convergence of $\varphi * \psi_j$ to $\varphi * \psi$. Everything involving derivatives then follows easily from the formula $\partial^\alpha(\varphi * \psi) = \partial^\alpha\varphi * \psi = \varphi * \partial^\alpha\psi$.

PROPOSITION 13.30. *For any Λ and φ satisfying the assumptions of Theorem 13.28 and any $\psi \in \mathscr{D}(\Omega')$, the smooth function $\sigma\varphi * \psi$ has compact support in Ω, and*

$$(\varphi * \Lambda, \psi) = (\Lambda, \sigma\varphi * \psi).$$

PROOF. Since $\operatorname{supp}(\sigma\varphi) \subset -K$ and ψ has compact support in Ω', Exercise 13.29 together with the definition of Ω' in Theorem 13.28 imply $\sigma\varphi * \psi \in \mathscr{D}(\Omega)$.

To prove the stated formula, we shall exploit the linearity of Λ by approximating the integral defining $F(x) := (\sigma\varphi * \psi)(x) = \int_{\mathbb{R}^n} \sigma\varphi(x-y)\psi(y)\, dy$ with Riemann sums. For $\epsilon > 0$ and any given $x \in \mathbb{R}^n$, the compact support of ψ implies that the function $y \mapsto \sigma\varphi(x-y)\psi(y)$ is nonzero on at most finitely many points in the lattice $\epsilon\mathbb{Z}^n \subset \mathbb{R}^n$, thus we can define a function $F_\epsilon : \mathbb{R}^n \to \mathbb{K}$ by

$$F_\epsilon(x) := \epsilon^n \sum_{y\in\epsilon\mathbb{Z}^n} \sigma\varphi(x-y)\psi(y) = \epsilon^n \sum_{y\in\epsilon\mathbb{Z}^n} \tau_{-y}\sigma\varphi(x)\psi(y).$$

In fact, this is a finite linear combination of smooth functions with compact supports contained in $-K + \operatorname{supp}(\psi) \subset -K + \Omega' \subset \Omega$, thus it belongs to $\mathscr{D}(\Omega)$ and its support is contained in a compact subset of Ω independent of ϵ. The function $F_\epsilon(x)$ can also be written as $\int_{\mathbb{R}^n} f_{\epsilon,x}(y)\, dy$ for a step function $f_{\epsilon,x} : \mathbb{R}^n \to \mathbb{K}$ whose value at each y is the value of $f_{0,x}(y) := \sigma\varphi(x-y)\psi(y)$ at the nearest lattice point $y \in \epsilon\mathbb{Z}^n$. Since φ and ψ are both uniformly continuous, for every $\delta > 0$ there exists $\epsilon_0 > 0$ such that $\|f_{\epsilon,x} - f_{0,x}\|_{C^0} < \delta$ for all $x \in \mathbb{R}^n$ and $\epsilon < \epsilon_0$, thus F_ϵ converges uniformly to F as $\epsilon \to 0$. The same is then true for all derivatives: since $\partial^\alpha F(x) = (\partial^\alpha(\sigma\varphi) * \psi)(x)$ and $\partial^\alpha F_\epsilon(x) = \epsilon^n \sum_{y \in \epsilon\mathbb{Z}^n} \partial^\alpha(\sigma\varphi)(x-y)\psi(y)$ for all multi-indices α, the same arguments imply that all derivatives of F_ϵ converge uniformly to F, hence $F_\epsilon \to F$ in $\mathscr{D}(\Omega)$. The continuity and linearity of Λ then imply

$$(\Lambda, \sigma\varphi * \psi) = \Lambda(F) = \lim_{\epsilon\to 0^+} \Lambda(F_\epsilon) = \lim_{\epsilon\to 0^+} \epsilon^n \sum_{y\in\epsilon\mathbb{Z}^n} \psi(y)\Lambda(\tau_{-y}\sigma\varphi(x)) = \lim_{\epsilon\to 0^+} \epsilon^n \sum_{y\in\epsilon\mathbb{Z}^n} \psi(y)(\varphi * \Lambda)(y).$$

This last expression is a Riemann sum approximating the integral $\int_{\Omega'} \psi(y)(\varphi * \Lambda)(y)\,dy$, whose integrand is also a smooth function with compact support, so the sum converges to the integral as $\epsilon \to 0$. □

13.3.2. *Approximation of distributions by smooth functions.*

Example 13.31. The Dirac δ-function $\delta \in \mathscr{D}'(\mathbb{R}^n)$ satisfies $(\varphi*\delta)(x) = \delta(\tau_{-x}\sigma\varphi) = \tau_{-x}\sigma\varphi(0) = \sigma\varphi(-x) = \varphi(x)$, i.e. $\varphi * \delta = \delta * \varphi = \varphi$ for every $\varphi \in \mathscr{D}(\mathbb{R}^n)$.

The definition of the term *approximate identity* in §8.4 can now be restated as follows: a sequence of smooth functions $\rho_j : \mathbb{R}^n \to [0,\infty)$ is an approximate identity if and only if

$$\rho_j \to \delta \quad \text{in} \quad \mathscr{D}'(\mathbb{R}^n),$$

where we are of course identifying the functions ρ_j with the distributions $\Lambda_{\rho_j} \in \mathscr{D}'(\mathbb{R}^n)$ that they determine. If ρ_j also has shrinking support, then we can assume for any given open neighborhood $\Omega \subset \mathbb{R}^n$ of the origin that ρ_j belongs to $\mathscr{D}(\Omega)$ for large j.

Now suppose $\Lambda \in \mathscr{D}'(\Omega)$ is an arbitrary distribution on some open set $\Omega \subset \mathbb{R}^n$, and ρ_j is an approximate identity with $\mathrm{supp}(\rho_j) \subset B_{r_j}$ for some sequence $r_j \to 0$. The convolutions $\Lambda_j := \rho_j * \Lambda$ are then defined on the subsets

$$\Omega_j := \left\{ x \in \Omega \;\middle|\; \mathrm{dist}(x, \mathbb{R}^n \backslash \Omega) > r_j \right\}, \tag{13.4}$$

whose union for all j is Ω. It follows that any $\varphi \in \mathscr{D}(\Omega)$ has support contained in Ω_j for all j sufficiently large, so that the integrals $\int_\Omega \varphi\Lambda_j\,dm := \int_{\Omega_j} \varphi\Lambda_j\,dm$ can be defined for large j by regarding the integrand as 0 wherever φ vanishes. The statement of the following result should be understood in these terms.

Theorem 13.32. *Suppose $\rho_j : \mathbb{R}^n \to [0,\infty)$ is an approximate identity with shrinking support, $\Lambda \in \mathscr{D}'(\Omega)$ is a distribution defined on some open set $\Omega \subset \mathbb{R}^n$, and $\Lambda_j := \rho_j * \Lambda$. Then for every $\varphi \in \mathscr{D}(\Omega)$, $\int_\Omega \varphi\Lambda_j\,dm \to \Lambda(\varphi)$.*

Proof. Assume j is large enough for $\mathrm{supp}(\varphi)$ to be contained in the domain of Λ_j. Then according to Proposition 13.30,

$$\int_\Omega \varphi\Lambda_j\,dm = (\rho_j * \Lambda, \varphi) = \Lambda(\sigma\rho_j * \varphi).$$

The functions $\sigma\rho_j$ are also an approximate identity with shrinking support, so the result follows via the continuity of Λ and the following claim: for any approximate identity ρ_j with shrinking support and any $\varphi \in \mathscr{D}(\Omega)$, the functions $\rho_j * \varphi$ have compact support in Ω for all j sufficiently large and converge in $\mathscr{D}(\Omega)$ to φ as $j \to \infty$. Indeed, Exercise 13.29 implies that $\mathrm{supp}(\rho_j * \varphi)$ lives in an arbitrarily small compact neighborhood of $\mathrm{supp}(\varphi)$ for large j, and Theorem 8.17 gives convergence $\rho_j * \varphi \to \varphi$ in $C^\infty_{\mathrm{loc}}(\Omega)$. In light of the supports, C^∞_{loc}-convergence in this situation implies uniform convergence of all derivatives and thus convergence in $\mathscr{D}(\Omega)$. □

Corollary 13.33. *For every open set $\Omega \subset \mathbb{R}^n$, $C^\infty_0(\Omega)$ is dense in $\mathscr{D}'(\Omega)$.*

Proof. Given an approximate identity ρ_j with shrinking support, define $\Lambda_j := \rho_j * \Lambda$, a sequence of smooth functions defined on the nested sequence of open subsets $\Omega_1 \subset \Omega_2 \subset \ldots \subset \bigcup_{j\in\mathbb{N}} \Omega_j$ described in (13.4). Choose a corresponding sequence of smooth functions $\beta_j : \Omega \to [0,1]$ with $\operatorname{supp}(\beta_j) \subset \Omega_j$ and $\beta_j|_{\Omega_{j-1}} \equiv 1$. Then $\beta_j\Lambda_j$ can be extended to smooth functions on Ω that vanish outside of Ω_j, and since every $\varphi \in \mathscr{D}(\Omega)$ has support in Ω_j for j large, Theorem 13.32 implies $\beta_j\Lambda_j \to \Lambda$ in $\mathscr{D}'(\Omega)$. □

13.3.3. *Distributions of class C^m.*

Theorem 13.34. *For a distribution $\Lambda \in \mathscr{D}'(\Omega)$ on an open set $\Omega \subset \mathbb{R}^n$ and integers $m, k \geqslant 0$, the following conditions are equivalent:*

(1) Λ is represented by a function of class C^{k+m};
(2) $\partial^\alpha\Lambda$ is represented by a function of class C^k for each multi-index α of order m.

Proof. The main step is to prove the special case with $k = 0$ and $m = 1$, as the rest then follows by a straightforward inductive argument. Let us therefore assume $\Lambda \in \mathscr{D}'(\Omega)$ has the property that $\partial_k\Lambda = \Lambda_{g_k}$ for every $k = 1, \ldots, n$, with continuous functions $g_k \in C^0(\Omega)$. The goal is then to show that $\Lambda = \Lambda_f$ for some $f \in C^1(\Omega)$.

Choose an approximate identity ρ_j with shrinking support, and consider the sequence of smooth functions $f_j := \rho_j * \Lambda$, which are defined on a nested sequence of open subdomains $\Omega_j \subset \Omega$ whose union is Ω. By Theorem 13.28, $\partial_k f_j = \rho_j * g_k$ for each $k = 1, \ldots, n$, and since the g_k are continuous, it follows via Theorem 8.17 that $\partial_k f_j \to g_k$ in $C^0_{\mathrm{loc}}(\Omega)$. We claim that f_j also converges in $C^1_{\mathrm{loc}}(\Omega)$ to a function $f \in C^1(\Omega)$. Indeed, by the fundamental theorem of calculus, every $x_0 \in \Omega$ has a convex neighborhood $\mathcal{U}_{x_0} \subset \Omega$ in which for $x = x_0 + h \in \mathcal{U}_{x_0}$ with $h = (h_1, \ldots, h_n) \in \mathbb{R}^n$,

$$f_j(x) - f_j(x_0) = \sum_{k=1}^n h_k \int_0^1 \partial_k f_j(x_0 + th)\, dt, \tag{13.5}$$

and the right hand side converges uniformly in x to $\sum_{k=1}^n h_k \int_0^1 g_k(x_0 + th)\, dt$. If $f_j(x_0)$ converges, it follows that f_j converges uniformly on a neighborhood of x_0, and the limiting function will then satisfy

$$f(x) - f(x_0) = \sum_{k=1}^n h_k \int_0^1 g_k(x_0 + th)\, dt,$$

implying that f is of class C^1 on this neighborhood with $\partial_k f = g_k$. The claim will thus follow if we can prove that $f_j(x_0)$ converges. To this end, choose a test function $\varphi : \mathbb{R}^n \to [0,\infty)$ that is positive at x_0 and has support in a neighborhood $\mathcal{U}_{x_0}$ of x_0 which can be assumed to be arbitrarily small. By Theorem 13.32,

$$\lim_{j\to\infty} \int_{\mathcal{U}_{x_0}} \varphi f_j\, dm \to \Lambda(\varphi). \tag{13.6}$$

Now if $f_j(x_0)$ does not converge, then at least one of the following occurs after passing to a subsequence:

(1) $|f_j(x_0)| \to \infty$. Since (13.5) implies that $|f_j(x) - f_j(x_0)|$ is bounded independently of j for all $x \in \mathcal{U}_{x_0}$, it follows if $\operatorname{supp}(\varphi)$ is sufficiently concentrated around x_0 that $\left|\int_{\mathcal{U}_{x_0}} \varphi f_j\, dm\right| \to \infty$, contradicting (13.6).
(2) $f_{2j-1}(x_0)$ and $f_{2j}(x_0)$ each converge to different limits. A similar argument via (13.5) then implies that if φ has support sufficiently concentrated near x_0, then $\int_{\mathcal{U}_{x_0}} \varphi f_{2j-1}\, dm$ and $\int_{\mathcal{U}_{x_0}} \varphi f_{2j}\, dm$ each converge to different limits, giving another contradiction to (13.6).

These contradictions prove the claim.

We've now proved that f_j converges in $C^1_{\mathrm{loc}}(\Omega)$ to a function $f \in C^1(\Omega)$, and it follows that for every $\varphi \in \mathscr{D}(\Omega)$, $\int_\Omega \varphi f_j \, dm \to \int_\Omega \varphi f \, dm$. The latter equals $\Lambda(\varphi)$ according to Theorem 13.32, so $\Lambda = \Lambda_f$. $\square$

Here is a consequence that is much less obvious than it looks:

Corollary 13.35. *If f and g are two functions on a connected open set $\Omega \subset \mathbb{R}^n$ that have the same weak first-order partial derivatives almost everywhere, then $f - g$ is equal to a constant almost everywhere.*

Proof. The assumptions imply that $h := f - g$ satisfies $h' = 0$ in the sense of distributions. Since 0 is a continuous function, Theorem 13.34 then implies that h is equal almost everywhere to a C^1-function whose classical gradient is zero; since Ω is connected, that function is a constant. $\square$

Exercise 13.36. Consider a linear differential operator of the form $L = \sum_\alpha c_\alpha \partial^\alpha$ acting on scalar-valued functions on $\mathbb{R}^n$, where the coefficients c_α are scalars and the sum runs over finitely many multi-indices, which may be of various orders. A distribution $K \in \mathscr{D}'(\mathbb{R}^n)$ is called a **fundamental solution**[29] for the operator L if it satisfies $LK = \delta$.

(a) Show that if K is a fundamental solution for L, then for every smooth compactly supported function $f : \mathbb{R}^n \to \mathbb{K}$, $u := K * f$ is a smooth solution to the partial differential equation $Lu = f$.

(b) Find a locally integrable function $K : \mathbb{R} \to \mathbb{R}$ that is a fundamental solution for the operator ∂_x^2, and verify explicity that $u := K * f$ satisfies $u'' = f$ for any $f \in C_0^\infty(\mathbb{R})$.

Exercise 13.37. Show that the functions

$$K(x) := -\frac{1}{2\pi} \ln |x| \quad \text{for } n = 2, \qquad K(x) := \frac{1}{(n-2)\operatorname{Vol}(S^{n-1})|x|^{n-2}} \quad \text{for } n \geqslant 3,$$

where $\operatorname{Vol}(S^{n-1}) > 0$ denotes the volume of the unit sphere in $\mathbb{R}^n$, are in $L^1_{\mathrm{loc}}(\mathbb{R}^n)$ and are fundamental solutions for the Laplace operator $-\Delta := -\sum_{j=1}^n \partial_j^2$ on $\mathbb{R}^n$ with $n \geqslant 2$. In particular, they have (weak) first derivatives

$$K_j(x) := \partial_j K(x) = -\frac{1}{\operatorname{Vol}(S^{n-1})} \frac{x_j}{|x|^n},$$

and their second derivatives (in the sense of distributions) take the form

$$K_{jk}(x) := \partial_j \partial_k K(x) = \frac{1}{\operatorname{Vol}(S^{n-1})} \frac{x_j x_k}{|x|^{n+2}}, \qquad \text{for } j \neq k,$$

and $\partial_j^2 K = -\dfrac{1}{n}\delta + K_{jj}$, where

$$K_{jj}(x) := \frac{1}{\operatorname{Vol}(S^{n-1})} \sum_k \frac{x_j^2 - x_k^2}{|x|^{n+2}},$$

and the evaluation of $K_{jk} \in \mathscr{D}'(\mathbb{R}^n)$ on test functions is defined via principal value integrals as in Exercise 13.21, that is,

$$(K_{jk}, \varphi) := \lim_{\epsilon \to 0^+} \int_{\mathbb{R}^n \setminus B^n_\epsilon} K_{jk}(x) \varphi(x) \, dx.$$

[29]Fundamental solutions are also often called **Green's functions**.

13.4. Product distributions. In this subsection we assume for simplicity that all distributions are scalar valued, though the discussion can be generalized for vector-valued distributions with minor adjustments (see Remark 13.43).

Recall that for any two σ-finite measure spaces (X, μ) and (Y, ν), there is a *product measure* $\mu \otimes \nu$ on $X \otimes Y$, which is uniquely determined by the condition

$$(\mu \otimes \nu)(A \times B) = \mu(A)\nu(B)$$

for arbitrary measurable sets $A \subset X$ and $B \subset Y$. Fubini's theorem is essentially the statement that product measures exist and are unique, together with a useful recipe for computing integrals with respect to product measures. We would now like to establish a variation on Fubini's theorem for distributions.

DEFINITION 13.38. If $f : X \to \mathbb{K}$ and $g : Y \to \mathbb{K}$ are two scalar-valued functions on sets X and Y respectively, we define a scalar-valued function $f \otimes g : X \times Y \to \mathbb{K}$ by

$$(f \otimes g)(x, y) := f(x)g(y).$$

Given two open sets $\Omega_1 \subset \mathbb{R}^m$, $\Omega_2 \subset \mathbb{R}^n$ and distributions $\Lambda_1 \in \mathscr{D}'(\Omega_1)$ and $\Lambda_2 \in \mathscr{D}'(\Omega_2)$, a distribution on $\Omega_1 \times \Omega_2 \subset \mathbb{R}^{m+n}$ is called a **product distribution** for Λ_1 and Λ_2, and denoted by $\Lambda_1 \otimes \Lambda_2 \in \mathscr{D}'(\Omega_1 \times \Omega_2)$, if it satisfies

$$(\Lambda_1 \otimes \Lambda_2)(\varphi_1 \otimes \varphi_2) = \Lambda_1(\varphi_1)\Lambda_2(\varphi_2) \quad \text{for all} \quad \varphi_1 \in \mathscr{D}(\Omega_1) \text{ and } \varphi_2 \in \mathscr{D}(\Omega_2).$$

EXAMPLE 13.39. If Λ_1 and Λ_2 are given by measures as in Example 13.15, then the product measure defines a product distribution $\Lambda_1 \otimes \Lambda_2$. (Note that a measure satisfying the condition stated in Example 13.15 is always σ-finite.)

EXERCISE 13.40. Use Fubini's theorem to show that for any locally integrable scalar-valued functions $f \in L^1_{\text{loc}}(\Omega_1)$ and $g \in L^1_{\text{loc}}(\Omega_2)$, $f \otimes g$ belongs to $L^1_{\text{loc}}(\Omega_1 \times \Omega_2)$ and $\Lambda_{f \otimes g} = \Lambda_f \otimes \Lambda_g \in \mathscr{D}'(\Omega_1 \times \Omega_2)$.

In the setting of Exercise 13.40, Fubini's theorem provides the following recipe for evaluating $\Lambda_f \otimes \Lambda_g$ on an arbitrary test function $\varphi \in \mathscr{D}(\Omega_1 \otimes \Omega_2)$: extending f and g to functions on $\mathbb{R}^m$ and $\mathbb{R}^n$ that vanish outside Ω and Ω' respectively, the compact support of φ in $\Omega \times \Omega'$ makes $(x, y) \mapsto \varphi(x, y)f(x)g(y)$ a well-defined function in $L^1(\mathbb{R}^{m+n})$ and thus implies

$$\begin{aligned}
(\Lambda_f \otimes \Lambda_g, \varphi) &= \int_{\Omega \times \Omega'} \varphi(x, y)f(x)g(y)\, dx\, dy = \int_{\mathbb{R}^{m+n}} \varphi(x, y)f(x)g(y)\, dx\, dy \\
&= \int_{\mathbb{R}^m} \left(\int_{\mathbb{R}^n} \varphi(x, y)g(y)\, dy \right) f(x)\, dx = \Lambda_f\left(x \mapsto \Lambda_g(\varphi(x, \cdot))\right) \\
&= \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^m} \varphi(x, y)f(x)\, dx \right) g(y)\, dy = \Lambda_g\left(y \mapsto \Lambda_f(\varphi(\cdot, y))\right).
\end{aligned}$$

Implicit in our notation in the last two lines is that $x \mapsto \Lambda_g(\varphi(x, \cdot))$ and $y \mapsto \Lambda_f(\varphi(\cdot, y))$ define smooth compactly supported scalar-valued functions on Ω and Ω' respectively, so they can be regarded as test functions and fed into distributions for evaluation. As an easy consequence of Proposition 13.24, the same holds when Λ_f and Λ_g are replaced by arbitrary distributions:

EXERCISE 13.41 (cf. Proposition 13.24). Show that if $\Omega_1 \subset \mathbb{R}^m$ and $\Omega_2 \subset \mathbb{R}^n$ are open sets, $\varphi \in \mathscr{D}(\Omega_1 \times \Omega_2)$ and $\Lambda \in \mathscr{D}'(\Omega_1)$, then $\psi(y) := \Lambda(\varphi(\cdot, y))$ defines a smooth compactly supported function on Ω_2.

THEOREM 13.42 (Fubini's theorem for distributions). *In the setting of Definition 13.38, there exists a unique product distribution $\Lambda_1 \otimes \Lambda_2 \in \mathscr{D}'(\Omega_1 \times \Omega_2)$, and its evaluation on arbitrary test functions $\varphi \in \mathscr{D}(\Omega_1 \times \Omega_2)$ is given by*

$$(\Lambda_1 \otimes \Lambda_2)(\varphi) = \Lambda_1 \left(x \mapsto \Lambda_2(\varphi(x,\cdot))\right) = \Lambda_2 \left(y \mapsto \Lambda_1(\varphi(\cdot,y))\right). \tag{13.7}$$

PROOF. We first prove the uniqueness of $\Lambda_1 \otimes \Lambda_2$. Given two product distributions for Λ_1 and Λ_2, their difference is a distribution $\Lambda \in \mathscr{D}'(\Omega_1 \times \Omega_2)$ such that $\Lambda(\varphi \otimes \psi) = 0$ for all $\varphi \in \mathscr{D}(\Omega_1)$ and $\psi \in \mathscr{D}(\Omega_2)$. The idea is now to use an approximate identity to approximate Λ with smooth functions that vanish. For $k = 1, 2$, let $\rho_j^{(1)}$ and $\rho_j^{(2)} : \mathbb{R}^n \to [0,\infty)$ denote approximate identities on $\mathbb{R}^m$ and $\mathbb{R}^n$ respectively, both with shrinking support. The functions $\rho_j := \rho_j^{(1)} \otimes \rho_j^{(2)} : \mathbb{R}^{m+n} \to [0,\infty)$ then also have shrinking support, and by Fubini's theorem, they satisfy

$$\int_{\mathbb{R}^{m+n}} \rho_j \, dm = \left(\int_{\mathbb{R}^m} \rho_j^{(1)} \, dm\right) \left(\int_{\mathbb{R}^n} \rho_j^{(2)} \, dm\right) \to 1 \quad \text{as} \quad j \to \infty,$$

so by Lemma 8.12, ρ_j is an approximate identity on $\mathbb{R}^{m+n}$. Theorem 13.32 then implies that for any $\varphi \in \mathscr{D}(\Omega_1 \times \Omega_2)$, $\rho_j * \Lambda$ is a smooth function defined on a neighborhood of the support of φ for j sufficiently large and satisfying

$$\int_{\Omega_1 \times \Omega_2} \varphi(\rho_j * \Lambda) \to \Lambda(\varphi) \quad \text{as} \quad j \to \infty.$$

But the function $\rho_j * \Lambda$ is given by

$$(\rho_j * \Lambda)(x,y) = \tau_{(x,y)}\Lambda(\sigma\rho_j) = \Lambda(\tau_{-x}\sigma\rho_j^{(1)} \otimes \tau_{-y}\sigma\rho_j^{(2)})$$

for all $(x,y) \in \mathbb{R}^{m+n}$ in its domain of definition, so taking $(x,y) \in \Omega_1 \times \Omega_2$ and j large enough for $\tau_{-x}\sigma\rho_j^{(1)}$ and $\tau_{-y}\sigma\rho_j^{(2)}$ to have support in Ω_1 and Ω_2 respectively, the defining property of Λ implies that $\rho_j * \Lambda$ vanishes, proving $\Lambda(\varphi) = 0$.

It is easy to see that both of the expressions on the right hand side of (13.7) evaluate like a product distribution on test functions of the form $\varphi_1 \otimes \varphi_2 \in \mathscr{D}(\Omega_1 \times \Omega_2)$, thus with uniqueness established, the rest of the theorem will follow if we can show that both of these expressions really define distributions, i.e. they are continuous linear maps on $\mathscr{D}(\Omega_1 \times \Omega_2)$. The proof works the same for both expressions, so let us focus on the first one and consider the linear map $\Lambda : \mathscr{D}(\Omega_1 \times \Omega_2) \to \mathbb{K}$ defined by

$$\Lambda(\varphi) = \Lambda_1 \left(x \mapsto \Lambda_2(\varphi(x,\cdot))\right).$$

To show that this is continuous, suppose $\varphi_j \to \varphi_\infty$ in $\mathscr{D}(\Omega_1 \times \Omega_2)$, and pick compact subsets $K_1 \subset \Omega_1$ and $K_2 \subset \Omega_2$ such that $\operatorname{supp}(\varphi_j) \subset K := K_1 \times K_2$ for all j. Then the sequence φ_j is also convergent with respect to the C^∞-topology on

$$\mathscr{D}_K(\Omega_1 \times \Omega_2) := \left\{\varphi \in \mathscr{D}(\Omega_1 \times \Omega_2) \;\middle|\; \operatorname{supp}(\varphi) \subset K\right\},$$

which is a closed subspace of the Fréchet space of C^∞-functions with bounded derivatives of all orders on $\Omega_1 \times \Omega_2$. Since Λ restricts to a continuous linear functional on this subspace, a standard result on continuous linear operators (see Lemmas 13.94 and 13.95 in §13.8, or [RS80, §V.1]) implies that there exists a continuous seminorm $\|\cdot\|$ on $\mathscr{D}_K(\Omega_1 \times \Omega_2)$ such that $|\Lambda(\varphi)| \leqslant \|\varphi\|$ holds for every $\varphi \in \mathscr{D}_K(\Omega_1 \times \Omega_2)$. Since the topology on $\mathscr{D}_K(\Omega_1 \times \Omega_2)$ is generated by the increasing sequence of norms $\|\cdot\|_{C^m}$ for $m \in \mathbb{N}$, this actually means that for sufficiently large constants $C > 0$ and $m \in \mathbb{N}$,

$$|\Lambda(\varphi)| \leqslant C\|\varphi\|_{C^m} \quad \text{for all} \quad \varphi \in \mathscr{D}_K(\Omega_1 \times \Omega_2).$$

This estimate applies in particular to the sequence φ_j and its derivatives $\partial^\alpha \varphi_j$ for every multi-index α. Writing $\psi_j(x) := \Lambda_2(\varphi_j(x, \cdot))$, Proposition 13.24 gives

$$\partial^\alpha \psi_j(x) = \Lambda_2 \left(\frac{\partial^{|\alpha|} \varphi_j}{\partial x^\alpha}(x, \cdot) \right),$$

thus

$$|\partial^\alpha \psi_\infty(x) - \partial^\alpha \psi_j(x)| = \left| \Lambda_2 \left(\frac{\partial^{|\alpha|} \varphi_\infty}{\partial x^\alpha}(x, \cdot) - \frac{\partial^{|\alpha|} \varphi_j}{\partial x^\alpha}(x, \cdot) \right) \right| \leqslant C \left\| \frac{\partial^{|\alpha|} \varphi_\infty}{\partial x^\alpha}(x, \cdot) - \frac{\partial^{|\alpha|} \varphi_j}{\partial x^\alpha}(x, \cdot) \right\|_{C^m}$$
$$\leqslant C \left\| \varphi_\infty - \varphi_j \right\|_{C^{m+|\alpha|}} \to 0 \quad \text{as} \quad j \to \infty,$$

giving C^∞-convergence $\psi_j \to \psi_\infty$. Since $\operatorname{supp}(\varphi_j) \subset K_1 \times K_2$, we also have $\operatorname{supp}(\psi_j) \subset K_1$ for all j, thus $\psi_j \to \psi_\infty$ in $\mathscr{D}(\Omega_1)$, and the continuity of Λ_1 now implies $\Lambda(\varphi_j) = \Lambda_1(\psi_j) \to \Lambda_1(\psi_\infty) = \Lambda(\varphi_\infty)$. □

REMARK 13.43. One can also define the notion of a product distribution $\Lambda_1 \otimes \Lambda_2$ if Λ_1 is scalar valued and Λ_2 is vector valued (or the other way around), but in this case an extra definition is needed before one can make sense of (13.7), as $x \mapsto \Lambda_2(\varphi(x, \cdot))$ is now a vector-valued function and thus does not belong to $\mathscr{D}(\Omega_1)$. The quickest way to rectify this is to choose a basis $e_1, \ldots, e_k$ of V and extend $\Lambda_1 : \mathscr{D}(\Omega_1) \to \mathbb{K}$ to a linear map from the space of compactly supported smooth functions $\Omega_1 \to V$ to V by $\Lambda_1(\sum_j \varphi_j e_j) := \sum_j \Lambda_1(\varphi_j) e_j$ for $\varphi_1, \ldots, \varphi_k \in \mathscr{D}(\Omega_1)$. It is easy to check that this definition is independent of the choice of basis, and Theorem 13.42 then becomes valid for the product of a scalar-valued and a vector-valued distribution.

13.5. The Sobolev spaces $W^{m,p}(\Omega)$. Let us now explain how to generalize the Sobolev spaces $H^m(\mathbb{R}^n)$ to arbitrary open domains $\Omega \subset \mathbb{R}^n$ and $p \neq 2$. The theory of distributions is not strictly needed for this discussion, but it makes some aspects of it seem easier and more natural.

DEFINITION 13.44. For an open set $\Omega \subset \mathbb{R}^n$, an integer $m \geqslant 0$ and a real number $p \in [1, \infty]$, the space $W^{m,p}(\Omega)$ is defined to consist of all $f \in L^p(\Omega)$ such that for every multi-index α with $|\alpha| \leqslant m$, the weak derivative $\partial^\alpha f$ exists and is also in $L^p(\Omega)$. The norm on $W^{m,p}(\Omega)$ is defined by

$$\|f\|_{W^{m,p}} := \sum_{|\alpha| \leqslant m} \|\partial^\alpha f\|_{L^p}.$$

REMARK 13.45. In contrast to §12, we are not considering non-integer values of m in our definition of $W^{m,p}(\Omega)$. Such a notion does exist but is much more complicated to define; details may be found in [**AF03**].

It is not hard to show that $W^{m,p}(\Omega)$ is a Banach space, as it admits a natural continuous linear inclusion

$$W^{m,p}(\Omega) \hookrightarrow \bigoplus_{|\alpha| \leqslant m} L^p(\Omega)$$

sending each $f \in W^{m,p}(\Omega)$ to a finite tuple of L^p-functions whose "α-coordinate" is $\partial^\alpha f$, and Exercise 13.46 below shows that the image of this inclusion is a closed subspace. More generally, one defines

$$W^{m,p}_{\mathrm{loc}}(\Omega) := \left\{ f \in L^p_{\mathrm{loc}}(\Omega) \;\middle|\; f \text{ has weak derivatives } \partial^\alpha f \in L^p_{\mathrm{loc}}(\Omega) \text{ for all } |\alpha| \leqslant m \right\},$$

which is equivalently the space of functions on Ω (up to equality almost everywhere) whose restrictions to every open subset with compact closure are of class $W^{m,p}$. As with L^p_{loc} (cf. §4.3), one can use the $W^{m,p}$-norms over an exhausting nested sequence of open subsets with compact closures $\Omega_1 \subset \bar{\Omega}_1 \subset \Omega_2 \subset \bar{\Omega}_2 \subset \ldots \subset \bigcup_{j \in \mathbb{N}} \Omega_j = \Omega$ to endow $W^{m,p}_{\mathrm{loc}}(\Omega)$ with the structure of a Fréchet space.

Exercise 13.46. Suppose $f_j \in W^{m,p}(\Omega)$ is a sequence such that for every multi-index α of order at most m, $\partial^\alpha f_j$ is L^p-convergent to some $g_\alpha \in L^p(\Omega)$. Show that the function $f := \lim_{j\to\infty} f_j \in L^p(\Omega)$ is then in $W^{m,p}(\Omega)$ and satisfies $\partial^\alpha f = g_\alpha$ for all $|\alpha| \leqslant m$.
Hint: For any test function $\varphi \in \mathscr{D}(\mathbb{R}^n)$, the L^p-convergence $\partial^\alpha f_j \to g_\alpha$ implies L^1-convergence on the support of φ.

Example 13.47. As shown in Exercise 13.3, the function $f(x) := |x|$ on $\mathbb{R}$ has a bounded weak derivative, thus $f \in W^{1,p}(\Omega)$ for every bounded open interval $\Omega \subset \mathbb{R}$ and $1 \leqslant p \leqslant \infty$. This shows that there is no value of p for which functions of class $W^{1,p}$ must be everywhere differentiable in the classical sense.

One can use approximate identities to show that the subspace

$$W^{m,p}(\Omega) \cap C^\infty(\Omega) \subset W^{m,p}(\Omega)$$

is dense for all $p < \infty$, thus an equivalent definition of $W^{m,p}(\Omega)$ for these cases would be as the closure of the space of smooth functions on Ω with respect to the $W^{m,p}$-norm. The next exercise proves a slightly stronger variant of this result in the case $\Omega = \mathbb{R}^n$.

Exercise 13.48. Prove via the following steps that $C_0^\infty(\mathbb{R}^n)$ is dense in $W^{m,p}(\mathbb{R}^n)$ for every $m \geqslant 0$ and $p < \infty$:

(a) If $f \in W^{m,p}(\mathbb{R}^n)$ and $\rho_j : \mathbb{R}^n \to [0,\infty)$ is an approximate identity with shrinking support, use Theorems 8.14 and 13.28 to show that $f_j := \rho_j * f$ is in $W^{m,p}(\mathbb{R}^n) \cap C^\infty(\mathbb{R}^n)$ and converges in $W^{m,p}$ to f as $j \to \infty$.

(b) Fix a smooth function $\psi : \mathbb{R}^n \to [0,1]$ that equals 1 on the unit ball and has compact support in the ball of radius 2, and let $\psi_\epsilon(x) := \psi(\epsilon x)$ for $\epsilon > 0$. Show that for any $f \in W^{m,p}(\mathbb{R}^n) \cap C^\infty(\mathbb{R}^n)$, $\psi_\epsilon f \to f$ in $W^{m,p}$ as $\epsilon \to 0$.
Hint: You need to estimate $\|\partial^\alpha[(1-\psi_\epsilon)f]\|_{L^p}$ for every multi-index α with $|\alpha| \leqslant m$. Consider separately the terms that either do or do not involve derivatives of ψ_ϵ.

Remark 13.49. While $C^\infty(\Omega) \cap W^{m,p}(\Omega)$ is always dense in $W^{m,p}(\Omega)$, it is not true for arbitrary open domains $\Omega \subset \mathbb{R}^n$ that $C_0^\infty(\Omega)$ is dense in $W^{m,p}(\Omega)$. In general, the $W^{m,p}$-closure of $C_0^\infty(\Omega)$ defines a closed subspace $W_0^{m,p}(\Omega) \subset W^{m,p}(\Omega)$ that is often useful in applications to boundary value problems, as it can be regarded as the space of $W^{m,p}$-functions on Ω that "vanish at the boundary". The proof in Exercise 13.48 that $C_0^\infty(\mathbb{R}^n)$ is dense in $W^{m,p}(\mathbb{R}^n)$ implicitly makes use of the fact that one has an infinite amount of room in $\mathbb{R}^n$ to "stretch out" the cutoff functions ψ_ϵ without losing control of their derivatives. This trick does not work more generally, e.g. when $\Omega \subset \mathbb{R}^n$ is bounded.

We can now clarify the relationship of $W^{m,p}(\Omega)$ to the Sobolev spaces we defined earlier via the Fourier transform.

Proposition 13.50. *For every integer $m \geqslant 0$, $W^{m,2}(\mathbb{R}^n) = H^m(\mathbb{R}^n)$.*

Proof. Both spaces are linear subspaces of $L^2(\mathbb{R}^n)$, and by Theorem 12.12 and Exercise 13.48, both contain the Schwartz space $\mathscr{S}(\mathbb{R}^n)$ as a dense subspace. One can easily show that the $W^{m,2}$-norm and H^m-norm are equivalent on $\mathscr{S}(\mathbb{R}^n)$, thus the two spaces are the closures of $\mathscr{S}(\mathbb{R}^n)$ with respect to equivalent norms, and are therefore identical. $\square$

Exercise 13.51. Prove:

(a) If f is an absolutely continuous function on an interval $[a,b]$, then its classical derivative f' (defined almost everywhere according to Theorem 9.7) is also its weak derivative on the domain (a,b), hence $f \in W^{1,1}((a,b))$.

Hint: For any $\varphi \in \mathscr{D}((a,b))$, φf defines an absolutely continuous function on $[a,b]$ that vanishes at the end points.

(b) If $f \in W^{1,1}_{\mathrm{loc}}(\Omega)$ for an open subset $\Omega \subset \mathbb{R}$, then f is equal almost everywhere to a function that is absolutely continuous on every compact subinterval of Ω.
Hint: On $[a,b] \subset \Omega$, define $g(x) := \int_a^x f'(t)\,dt$ and apply Corollary 13.35.

(c) For any open interval $\Omega \subset \mathbb{R}$, there exists a constant $c > 0$ such that

$$\|f\|_{C^0} \leqslant c\|f\|_{W^{1,1}} \quad \text{for all} \quad f \in W^{1,1}(\Omega).$$

Hint: The fundamental theorem of calculus implies $|f(x) - f(y)| \leqslant \|f'\|_{L^1}$ for all $x, y \in \Omega$, and thus $|f(x)| \geqslant \|f\|_{C^0} - \|f'\|_{L^1}$ for all $x \in \Omega$.

Exercise 13.52. Consider the function $f(x) := \ln\big|\ln|x|\big|$ on the r-ball $B_r \subset \mathbb{R}^n$ about the origin for some $r \in (0,1)$.

(a) Show that the classical first derivatives $\partial_j f$, defined on $B_r \backslash \{0\}$, are also weak derivatives of f on B_r.
Hint: Since f and $\partial_j f$ are both in $L^1(B_r)$, for any $\varphi \in \mathscr{D}(B_r)$ supported in some cube $Q \subset B_r$ around 0, you can approximate $\int_Q \partial_j(\varphi f)\,dm$ by integrating over $Q\backslash\{|x_j| < \epsilon\}$ for small $\epsilon > 0$, and then use integration by parts. There will be a boundary term; you need to show that the singularity of f at 0 is not bad enough to make the boundary term matter as $\epsilon \to 0$.

(b) Show that for $n = 1$, $f \notin W^{1,p}(B_r)$ for any $p \geqslant 1$, but for $n \geqslant 2$, $f \in W^{1,p}(B_r)$ if and only if $p \leqslant n$.

We saw in §12 that in general, functions of class $W^{1,p}$ need not be anywhere differentiable, and on higher-dimensional domains, Exercise 13.52 shows that they need not even be continuous—the continuity result in Exercise 13.51 is special to one-dimensional domains. The Sobolev embedding theorem gives sharp criteria saying to what extent the functions in any given Sobolev space must be classically differentiable. The proof of this important result, which generalizes Theorems 12.10 and 12.45 beyond the case $p = 2$ and $\Omega = \mathbb{R}^n$, belongs more properly to a course on PDEs, so we will not include it, but here is the statement:

Theorem 13.53 (Sobolev embedding theorem). *Suppose $k \in \mathbb{N}$ and $p \in [1,\infty)$ satisfy the relation*

$$0 < k - n/p \leqslant 1,$$

and $\Omega \subset \mathbb{R}^n$ is either $\mathbb{R}^n$ or an open subset whose closure is a compact C^1-smooth manifold with boundary.[30] *Then for every integer $m \geqslant 0$ and every $\alpha \in (0,1)$ with $\alpha \leqslant k - n/p$, there exists a continuous inclusion*

$$W^{k+m,p}(\Omega) \hookrightarrow C^{m,\alpha}(\Omega).$$

Exercise 13.54. Show that in the situation of Theorem 13.53, whenever Ω is bounded and the strict inequality $\alpha < k - n/p$ is satisfied, the inclusion $W^{k+m,p}(\Omega) \hookrightarrow C^{m,\alpha}(\Omega)$ is compact. In particular, there is a continuous inclusion $W^{k+m,p}(\Omega) \hookrightarrow C^m(\bar{\Omega})$ whenever $kp > n$, and it is compact if $\Omega \subset \mathbb{R}^n$ is bounded. (See §4.3 for the definition of the Banach space $C^m(\bar{\Omega})$.)

Theorem 13.53 motivates thinking of functions in $W^{k,p}(\Omega)$ as functions that have "$k - n/p$ continuous derivatives" whenever $kp > n$, where the number $k - n/p$ need not be an integer. This intuition is further supported by the following generalization of the obvious inclusion $H^t(\mathbb{R}^n) \hookrightarrow$

[30] The hypothesis on Ω can be generalized considerably; here we are only stating a version that can be understood without too many extra definitions. The theorem as stated remains true for any (bounded or unbounded) open domain $\Omega \subset \mathbb{R}^n$ whose boundary satisfies something called the "strong local Lipschitz condition"; see [**AF03**, §4.12] for details.

$H^s(\mathbb{R}^n)$ for $t > s$. The case with Ω bounded is known as the *Rellich-Kondrachov compactness theorem* (cf. Theorem 12.31):

THEOREM 13.55. *Under the same assumptions on Ω as in Theorem 13.53, suppose $1 \leqslant p, q < \infty$ and $k, m \geqslant 0$ are integers satisfying*

$$k \geqslant m, \quad p \leqslant q, \qquad \text{and} \qquad k - \frac{n}{p} \geqslant m - \frac{n}{q}.$$

Then there exists a continuous inclusion $W^{k,p}(\Omega) \hookrightarrow W^{m,q}(\Omega)$, and this inclusion is compact if the inequality $k - \frac{n}{p} \geqslant m - \frac{n}{q}$ is strict and Ω is bounded.

EXERCISE 13.56. When Ω is a bounded interval $(a, b) \subset \mathbb{R}$, Theorem 13.53 says that for all integers $m \geqslant 0$, there are continuous inclusions

$$\begin{aligned} W^{1+m,p}((a,b)) &\hookrightarrow C^{m,\alpha}((a,b)) \quad \text{if} \quad 0 < \alpha < 1,\ 1 < p \leqslant \infty \text{ and } \alpha \leqslant 1 - \frac{1}{p} \\ W^{2+m,1}((a,b)) &\hookrightarrow C^{m,\alpha}((a,b)) \quad \text{if} \quad 0 < \alpha < 1. \end{aligned}$$

Prove this as follows:

(a) Deduce the inclusions $W^{2,1} \hookrightarrow C^{0,\alpha}$ for $\alpha \in (0, 1]$ from a continuous inclusion $W^{2,1} \hookrightarrow C^1$ using Exercise 13.51.
(b) Deduce the inclusion $W^{1,p} \hookrightarrow C^0$ for every $p \geqslant 1$ from Exercise 13.51.
(c) For $a \leqslant x < y \leqslant b$, the fundamental theorem of calculus implies $|f(x) - f(y)| \leqslant \|f'\|_{L^1([x,y])}$ for $f \in W^{1,p}((a, b))$ since (by Exercise 13.51) f is absolutely continuous. Use Hölder's inequality to deduce a Hölder-type estimate $|f(x)-f(y)| \leqslant c\|f'\|_{L^p} \cdot |x-y|^\alpha$ for $0 < \alpha \leqslant 1 - 1/p$ whenever $p > 1$. The proof for $m = 0$ is thus complete.
(d) Extend the result to all $m \in \mathbb{N}$ by induction.

According to Theorem 13.53, the condition $kp > n$ guarantees continuity for functions of class $W^{k,p}$ on n-dimensional domains. We saw in Exercise 13.51 that the situation is slightly better when $n = 1$: here the condition $kp = n$ already suffices for continuity, but the function in Exercise 13.52 demonstrates that this is false in dimensions $n \geqslant 2$. The situation with $kp = n$ is often called the *Sobolev borderline case*. Even in dimension one, the borderline case has the disadvantage that functions of class $W^{1,1}$ need not be Hölder continuous, and so in contrast to Exercise 13.54, the inclusion $W^{1,1}(\Omega) \hookrightarrow C^0(\Omega)$ for bounded intervals $\Omega \subset \mathbb{R}$ is not compact.

EXERCISE 13.57. Find a sequence of smooth functions $f_j : (-1, 1) \to \mathbb{R}$ such that $\|f_j\|_{L^1}$ and $\|f_j'\|_{L^1}$ are bounded but f_j has no C^0-convergent subsequence.
Hint: Construct f_j so that it converges in L^1 to a (discontinuous) characteristic function.

REMARK 13.58. The Sobolev embedding theorem furnishes one major reason why it is useful to study the properties of *all* the L^p-spaces for $1 \leqslant p \leqslant \infty$, rather than just L^2, which might otherwise be easier since the latter is a Hilbert space. As a concrete example, suppose you are studying a first-order PDE for functions on 2-dimensional domains. If you want to work only with Hilbert spaces but also want all your functions to be continuous, then Theorem 13.53 requires you to take functions of class $H^m = W^{m,2}$ with $m \geqslant 2$, which involves at least one more (weak) derivative than the PDE itself actually needs. In such a situation, it may be easier to work with functions of class $W^{1,p}$ for some $p > 2$, as these are continuous, and one only needs to compute first-order derivatives in order to verify whether a given function belongs to this space.

13.6. Tempered distributions and Fourier transforms. Since we are going to talk about Fourier transforms in this subsection, we need to assume $\mathbb{K} = \mathbb{C}$.

We would now like to define Fourier transforms of functions for which the usual integral formula cannot even approximately make sense, e.g. functions that are not in $L^2(\mathbb{R}^n)$, and ideally, distributions. One can *almost* deduce the correct definition by considering the distribution $\Lambda_f \in \mathscr{D}'(\mathbb{R}^n)$ corresponding to a function $f \in L^1(\mathbb{R}^n)$: by Fubini's theorem, we have

$$(\mathscr{F}f, \varphi) = \int_{\mathbb{R}^n} \varphi(p) \left(\int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} f(x)\, dx \right) dp = \int_{\mathbb{R}^n} \left(\int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} \varphi(p)\, dp \right) f(x)\, dx = (f, \mathscr{F}\varphi)$$

for all $\varphi \in \mathscr{D}(\mathbb{R}^n)$. This suggests defining $\mathscr{F}\Lambda \in \mathscr{D}'(\mathbb{R}^n)$ for arbitrary $\Lambda \in \mathscr{D}'(\mathbb{R}^n)$ by $(\mathscr{F}\Lambda)(\varphi) := \Lambda(\mathscr{F}\varphi)$, but this definition as it stands does not quite make sense: $\mathscr{F}\varphi$ might not have compact support, in which case it is not a test function and $\Lambda(\mathscr{F}\varphi)$ will not make sense for arbitrary distributions Λ. The solution is to replace the usual space of test functions with the Schwartz space $\mathscr{S}(\mathbb{R}^n)$, since the latter is closed under the Fourier transform.

Before defining what a continuous linear functional on $\mathscr{S}(\mathbb{R}^n)$ is, we need to define a topology on $\mathscr{S}(\mathbb{R}^n)$. As with $\mathscr{D}(\mathbb{R}^n)$, we would like this topology to be relatively strong, so that as many functionals as possible are continuous, but also to have the property that continuity can be characterized purely in terms of convergent sequences (cf. Proposition 13.8). This turns out to be easier for $\mathscr{S}(\mathbb{R}^n)$ than for $\mathscr{D}(\mathbb{R}^n)$: the natural choice is to endow $\mathscr{S}(\mathbb{R}^n)$ with the topology generated by the countable family of seminorms

$$\|\varphi\|_{\alpha,\beta} := \|x^\alpha \partial^\beta \varphi\|_{C^0}$$

for all multi-indices α, β, so convergence $\varphi_k \to \varphi$ in $\mathscr{S}(\mathbb{R}^n)$ will mean that for every polynomial function $P : \mathbb{R}^n \to \mathbb{R}$ and every multi-index β, the functions $P\partial^\beta \varphi_k$ converge uniformly on $\mathbb{R}^n$ to $P\partial^\beta\varphi$. It follows easily from the completeness of the C^0-norm that sequences that are Cauchy with respect to all of these seminorms must also converge, hence $\mathscr{S}(\mathbb{R}^n)$ is now a Fréchet space. In particular, the topology we have defined on $\mathscr{S}(\mathbb{R}^n)$ is metrizable, thus continuity and sequential continuity of functions defined on $\mathscr{S}(\mathbb{R}^n)$ are equivalent notions.

Exercise 13.59. Show that the natural inclusions $\mathscr{D}(\mathbb{R}^n) \hookrightarrow \mathscr{S}(\mathbb{R}^n)$ and $\mathscr{S}(\mathbb{R}^n) \hookrightarrow W^{m,p}(\mathbb{R}^n)$ for all $m \geqslant 0$ and $p \in [1, \infty]$ are continuous.

Exercise 13.60. Show that the following linear operators $\mathscr{S}(\mathbb{R}^n) \to \mathscr{S}(\mathbb{R}^n)$ are continuous:

(a) ∂^α and $\varphi \mapsto x^\alpha\varphi$ for every multi-index α;
(b) $\mathscr{F}$ and $\mathscr{F}^*$.

Definition 13.61. A complex-valued **tempered distribution** on $\mathbb{R}^n$ is a continuous complex-linear functional $\Lambda : \mathscr{S}(\mathbb{R}^n) \to \mathbb{C}$. Similarly, a vector-valued tempered distribution with values in the finite-dimensional complex vector space V is a continuous complex-linear map $\Lambda : \mathscr{S}(\mathbb{R}^n) \to V$. We shall generally assume that all tempered distributions take values in a fixed vector space V, and denote the vector space of vector-valued tempered distributions by

$$\mathscr{S}'(\mathbb{R}^n) = \left\{ \Lambda : \mathscr{S}(\mathbb{R}^n) \to V \;\middle|\; \Lambda \text{ is complex linear and continuous} \right\}.$$

The space $\mathscr{S}'(\mathbb{R}^n)$ is endowed with the weak*-topology, i.e. the locally convex topology generated by the seminorms $\|\Lambda\|_\varphi := |\Lambda(\varphi)|$ for all $\varphi \in \mathscr{S}(\mathbb{R}^n)$, hence a sequence $\Lambda_j \in \mathscr{S}'(\mathbb{R}^n)$ converges to $\Lambda_\infty \in \mathscr{S}'(\mathbb{R}^n)$ if and only if $\Lambda_j(\varphi) \to \Lambda_\infty(\varphi)$ for every $\varphi \in \mathscr{S}(\mathbb{R}^n)$.

The inclusion $\mathscr{D}(\mathbb{R}^n) \hookrightarrow \mathscr{S}(\mathbb{R}^n)$ in Exercise 13.59 gives rise to a natural continuous inclusion $\mathscr{S}'(\mathbb{R}^n) \hookrightarrow \mathscr{D}'(\mathbb{R}^n)$, i.e. every tempered distribution is also a distribution in the usual sense. The converse is false, and in fact $\mathscr{S}'(\mathbb{R}^n)$ does not even contain all locally integrable functions, e.g. $f(x) := e^{x^2}$ does not define an element $\Lambda_f \in \mathscr{S}'(\mathbb{R})$ since there exist functions $\varphi \in \mathscr{S}(\mathbb{R})$

for which $\int_{\mathbb{R}} \varphi f \, dm$ is not defined. However, most important examples of distributions are also tempered distributions: these include large classes of functions as in the following two exercises,[31] as well as standard singular examples like the Dirac δ-function and its derivatives. By a slight abuse of notation, we shall write $L^1_{\mathrm{loc}}(\mathbb{R}^n) \cap \mathscr{S}'(\mathbb{R}^n)$ for the space of all locally integrable functions f on $\mathbb{R}^n$ such that $\varphi f \in L^1(\mathbb{R}^n)$ for every $\varphi \in \mathscr{S}(\mathbb{R}^n)$ and the formula $\Lambda_f(\varphi) := \int_{\mathbb{R}^n} \varphi f \, dm$ defines a tempered distribution $\Lambda_f \in \mathscr{S}'(\mathbb{R}^n)$.

Exercise 13.62. A function $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ is said to have **polynomial growth** if it satisfies $|f| \leqslant |P|$ for some polynomial function $P : \mathbb{R}^n \to \mathbb{R}$; equivalently, this is true if and only if there exist constants $C > 0$ and $k \in \mathbb{N}$ such that

$$|f(x)| \leqslant C(1 + |x|^k) \quad \text{for all} \quad x \in \mathbb{R}^n.$$

Show that any function with this property is in $\mathscr{S}'(\mathbb{R}^n)$.

Exercise 13.63. Show that $L^p(\mathbb{R}^n) \subset \mathscr{S}'(\mathbb{R}^n)$ for every $p \in [1, \infty]$, and the inclusions $L^p(\mathbb{R}^n) \hookrightarrow \mathscr{S}'(\mathbb{R}^n)$ are continuous.
Hint: Use the continuity of the inclusions $\mathscr{S}(\mathbb{R}^n) \hookrightarrow L^q(\mathbb{R}^n)$ and the natural injection $L^p \hookrightarrow (L^q)^*$ for $\frac{1}{p} + \frac{1}{q} = 1$.

Partial derivative operators are defined as continuous linear maps on $\mathscr{S}'(\mathbb{R}^n)$ in the same way as $\mathscr{D}'(\mathbb{R}^n)$; continuity in this case follows from the continuity of ∂^α on $\mathscr{S}(\mathbb{R}^n)$ (Exercise 13.60). The product of a smooth function $f \in C^\infty(\mathbb{R}^n)$ with a tempered distribution $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$ is not well defined unless $\varphi \mapsto f\varphi$ is a continuous map $\mathscr{S}(\mathbb{R}^n) \to \mathscr{S}(\mathbb{R}^n)$, which is not true e.g. for $f(x) := e^{x^2}$ on $\mathbb{R}$, but is true if f and its derivatives of all orders have polynomial growth as in Exercise 13.62. Under this assumption, it is straightforward to show that the Leibniz rule in Exercise 13.22 also holds for tempered distributions.

Remark 13.64. For a function $f \in L^1_{\mathrm{loc}}(\mathbb{R}^n)$ that defines a tempered distribution, we now have two potentially inequivalent definitions for the notion of weak derivatives $\partial_j f$, depending whether we want $\partial_j f$ to define an element of $\mathscr{D}'(\mathbb{R}^n)$ or $\mathscr{S}'(\mathbb{R}^n)$. In the latter case, it needs to satisfy a stronger condition involving integration against test functions in $\mathscr{S}(\mathbb{R}^n)$, a larger space than $\mathscr{D}(\mathbb{R}^n)$; it could happen for instance that f has a locally integrable weak derivative $\partial_j f$ that grows too fast at infinity to define a tempered distribution, in which case the stronger condition fails. However, if a weak derivative $\partial_j f$ does define a tempered distribution—which is always the case for instance if $\partial_j f$ is of class L^p for some p, and notably if f belongs to a suitable Sobolev space—then it also satisfies the stronger condition, i.e. it is also a derivative of f in the sense of *tempered* distributions. The reason is that, by Exercise 13.65 below, $\mathscr{D}(\mathbb{R}^n)$ is dense in $\mathscr{S}(\mathbb{R}^n)$, so any two tempered distributions that evaluate the same on $\mathscr{D}(\mathbb{R}^n)$ are identical.

Exercise 13.65. Show that for any $\varphi \in \mathscr{S}(\mathbb{R}^n)$ and the family of compactly supported smooth cutoff functions $\psi_\epsilon : \mathbb{R}^n \to [0, 1]$ in Exercise 13.48, $\psi_\epsilon \varphi \to \varphi$ in $\mathscr{S}(\mathbb{R}^n)$ as $\epsilon \to 0$. In particular, $\mathscr{D}(\mathbb{R}^n)$ is dense in $\mathscr{S}(\mathbb{R}^n)$.
Hint: For any multi-indices α and β, the condition $\varphi \in \mathscr{S}(\mathbb{R}^n)$ implies $\|x^\alpha \partial^\beta \varphi\|_{C^0(\mathbb{R}^n \setminus B_R)} \to 0$ as $R \to \infty$. (Why?)

The next set of exercises generalizes the convolution operator and its main properties from §13.3 to the context of tempered distributions.

Exercise 13.66. Recall the translation operator τ_v for functions f on $\mathbb{R}^n$ and $v \in \mathbb{R}^n$, defined by $(\tau_v f)(x) := f(x + v)$.

[31] The word "tempered" refers to conditions as in Exercise 13.62 and 13.63 that rule out functions like e^{x^2}, which grow too fast at infinity.

(a) Show that for every pair of multi-indices α and β, there exists a constant $C > 0$ and a finite set of pairs of multi-indices $\{(\alpha_i, \beta_i)\}_{i=1}^N$ such that

$$\|\tau_v \varphi\|_{\alpha,\beta} \leqslant C \left(\sum_{i=1}^N \|\varphi\|_{\alpha_i,\beta_i} \right) (1 + |v|^{|\alpha|}) \quad \text{for all} \quad \varphi \in \mathscr{S}(\mathbb{R}^n),\ v \in \mathbb{R}^n.$$

In particular, $\tau_v \varphi$ is also in $\mathscr{S}(\mathbb{R}^n)$ for every $\varphi \in \mathscr{S}(\mathbb{R}^n)$ and $v \in \mathbb{R}^n$.
Hint: By Exercise 11.2, you can assume $|\partial^\beta \varphi(x)| \leqslant \frac{c}{1+|x|^k}$ for some $k \in \mathbb{N}$ arbitrarily large and a constant $c > 0$ determined by k and finitely many seminorms of φ. Estimate $|x^\alpha \partial^\beta \varphi(x + v)|$ by looking separately at the cases $|x| \leqslant 2|v|$ and $|x| \geqslant 2|v|$.

(b) Show that Lemmas 13.25 and 13.26 remain valid with the space of test functions replaced by the Schwartz space (ignoring all conditions that involve compact support).

Exercise 13.67. Reread the proof of Corollary 13.27 and verify that, in light of Exercise 13.66, the function $F_\Lambda(x) := (\tau_x \Lambda)(\varphi)$ defined on $\mathbb{R}^n$ for any $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$ and $\varphi \in \mathscr{S}(\mathbb{R}^n)$ is smooth and satisfies $\partial^\alpha F_\Lambda = F_{\partial^\alpha \Lambda}$ for all multi-indices α.

Exercise 13.68. Show that for any $\varphi \in \mathscr{S}(\mathbb{R}^n)$ and $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$, the formula (13.3) defines a smooth function $\varphi * \Lambda$ on $\mathbb{R}^n$ satisfying $\partial^\alpha(\varphi * \Lambda) = (\partial^\alpha \varphi) * \Lambda = \varphi * (\partial^\alpha \Lambda)$ for all multi-indices α.

Exercise 13.69. Consider the convolution of two Schwartz functions $\varphi, \psi \in \mathscr{S}(\mathbb{R}^n)$.

(a) Show that $\varphi * \psi$ is continuous and bounded on $\mathbb{R}^n$.
(b) Show that if ψ_j is a sequence converging in $\mathscr{S}(\mathbb{R}^n)$ to ψ, then $\varphi * \psi_j$ converges uniformly to $\varphi * \psi$.
(c) For $k = 1, \dots, n$ and a function f on $\mathbb{R}^n$, let $P_k f$ denote the function on $\mathbb{R}^n$ defined by $(P_k f)(x) := x_k f(x)$, so e.g. by Exercise 13.60, P_k defines a continuous linear operator $\mathscr{S}(\mathbb{R}^n) \to \mathscr{S}(\mathbb{R}^n)$. Show that $P_k(\varphi * \psi) = (P_k \varphi) * \psi + \varphi * P_k \psi$, and deduce that $P_k(\varphi * \psi)$ is continuous and bounded.
(d) Deduce that $\varphi * \psi \in \mathscr{S}(\mathbb{R}^n)$, and for any sequence $\psi_j \to \psi$ in $\mathscr{S}(\mathbb{R}^n)$, $\varphi * \psi_j \to \varphi * \psi$ in $\mathscr{S}(\mathbb{R}^n)$.

Proposition 13.70. *For $\varphi \in \mathscr{S}(\mathbb{R}^n)$ and $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$, $\varphi * \Lambda$ is a polynomially bounded function and thus defines an element of $\mathscr{S}'(\mathbb{R}^n)$. Moreover, if $\varphi_j \in \mathscr{S}(\mathbb{R}^n)$ converges to φ in $\mathscr{S}(\mathbb{R}^n)$, then the tempered distributions $\varphi_j * \Lambda$ converge to $\varphi * \Lambda$ in $\mathscr{S}'(\mathbb{R}^n)$.*

Proof. We start by proving that $\varphi * \Lambda$ has polynomial growth. By one of the standard characterizations of continuity for linear operators on locally convex spaces (see Lemmas 13.94 and 13.95 in §13.8, or [**RS80**, §V.1]), the continuity of $\Lambda : \mathscr{S}(\mathbb{R}^n) \to V$ means that there exists a finite set of pairs of multi-indices $\{(\alpha_i, \beta_i)\}_{i=1}^N$ and a constant $C > 0$ such that

$$|\Lambda(\varphi)| \leqslant C \sum_{i=1}^N \|\varphi\|_{\alpha_i,\beta_i} \quad \text{for all} \quad \varphi \in \mathscr{S}(\mathbb{R}^n).$$

Using Exercise 13.66, the convolution $\varphi * \Lambda$ thus satisfies

$$|(\varphi * \Lambda)(x)| = |\Lambda(\tau_{-x} \sigma \varphi)| \leqslant C \sum_{i=1}^N \|\tau_{-x}(\sigma \varphi)\|_{\alpha_i,\beta_i} \leqslant C \sum_{i=1}^N c_i (1 + |x|^{|\alpha_i|}) \leqslant C'(1 + |x|^k)$$

for suitable constants $c_i > 0$, $C' > 0$ and $k \in \mathbb{N}$ sufficiently large. In this expression, the constant $C > 0$ is determined entirely by Λ, while only $c_1, \dots, c_N$ (and therefore also C') depend on φ; looking more closely at Exercise 13.66, we see moreover that they can be bounded linearly in terms of finitely many of the seminorms $\|\varphi\|_{\alpha,\beta}$. For this reason, if $\varphi_j \to \varphi$ is a convergent sequence in $\mathscr{S}(\mathbb{R}^n)$, the same argument gives

$$|(\varphi * \Lambda)(x) - (\varphi_j * \Lambda)(x)| = |\Lambda(\tau_{-x}\sigma(\varphi - \varphi_j))| \leqslant C_j(1 + |x|^k)$$

for constants $C_j > 0$ that converge to 0 as $j \to \infty$, thus for any $\psi \in \mathscr{S}(\mathbb{R}^n)$,

$$|(\varphi * \Lambda, \psi) - (\varphi_j * \Lambda, \psi)| \leqslant \int_{\mathbb{R}^n} |\psi| \cdot |\varphi * \Lambda - \varphi_j * \Lambda| \, dm \leqslant C_j \int_{\mathbb{R}^n} |\psi(x)|(1 + |x|^k) \, dx \to 0.$$

□

EXERCISE 13.71. Use Proposition 13.30 and the density of $\mathscr{D}(\mathbb{R}^n)$ in $\mathscr{S}(\mathbb{R}^n)$ to deduce that the relation $(\varphi * \Lambda, \psi) = \Lambda(\sigma\varphi * \psi)$ also holds for all $\varphi, \psi \in \mathscr{S}(\mathbb{R}^n)$ and $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$.

EXERCISE 13.72. Suppose $\rho_j : \mathbb{R}^n \to [0, \infty)$ is an approximate identity with shrinking support. Prove:

(a) For any $\varphi \in \mathscr{S}(\mathbb{R}^n)$, $\rho_j * \varphi \to \varphi$ in $\mathscr{S}(\mathbb{R}^n)$ as $j \to \infty$.
(b) For any $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$, $\rho_j * \Lambda \to \Lambda$ in $\mathscr{S}'(\mathbb{R}^n)$ as $j \to \infty$. (This proves that $C^\infty(\mathbb{R}^n) \cap \mathscr{S}'(\mathbb{R}^n)$ is dense in $\mathscr{S}'(\mathbb{R}^n)$.)
(c) For any $f \in C^\infty(\mathbb{R}^n) \cap \mathscr{S}'(\mathbb{R}^n)$ and the family of compactly supported smooth cutoff functions $\psi_\epsilon : \mathbb{R}^n \to [0, 1]$ in Exercise 13.48, $\psi_\epsilon f \to f$ in $\mathscr{S}'(\mathbb{R}^n)$ as $\epsilon \to 0$. (This proves that $C_0^\infty(\mathbb{R}^n)$ is dense in $\mathscr{S}'(\mathbb{R}^n)$.)

While distributions are easier to work with than tempered distributions for many purposes, the major advantage of the latter is that they admit natural definitions of the Fourier transform and Fourier inverse operators.

DEFINITION 13.73. We define $\mathscr{F}, \mathscr{F}^* : \mathscr{S}'(\mathbb{R}^n) \to \mathscr{S}'(\mathbb{R}^n)$ by

$$(\mathscr{F}\Lambda)(\varphi) := \widehat{\Lambda}(\varphi) := \Lambda(\widehat{\varphi}) \qquad \text{and} \qquad (\mathscr{F}^*\Lambda)(\varphi) := \check{\Lambda}(\varphi) := \Lambda(\check{\varphi}).$$

The continuity of $\mathscr{F}$ and $\mathscr{F}^*$ on $\mathscr{S}(\mathbb{R}^n)$ (Exercise 13.60) implies that they are also continuous on $\mathscr{S}'(\mathbb{R}^n)$, and the relations $\mathscr{F}\mathscr{F}^* = \mathscr{F}^*\mathscr{F} = \mathbb{1}$ extend immediately from $\mathscr{S}(\mathbb{R}^n)$ to $\mathscr{S}'(\mathbb{R}^n)$. The calculation via Fubini's theorem at the beginning of this subsection shows that our definition of $\mathscr{F}\Lambda$ and $\mathscr{F}^*\Lambda$ for any $\Lambda = \Lambda_f$ with $f \in L^1(\mathbb{R}^n)$ matches the result of the usual integral formula.

EXERCISE 13.74. For $f \in L^2(\mathbb{R}^n)$, use approximation by L^1-functions to show that $\mathscr{F}\Lambda_f = \Lambda_{\mathscr{F}f}$ and $\mathscr{F}^*\Lambda_f = \Lambda_{\mathscr{F}^*f}$, where $\mathscr{F}f$ and $\mathscr{F}^*f$ are defined as in §11.

REMARK 13.75. Recall from Lemma 13.4 that two locally integrable functions are equal almost everywhere if and only if they define the same distribution. The same is true for tempered distributions since $\mathscr{D}(\mathbb{R}^n) \subset \mathscr{S}(\mathbb{R}^n)$. Exercise 13.74 thus shows that the most general possible definition of the Fourier transform, given by Definition 13.73, matches the definition we previously had for functions in $L^2(\mathbb{R}^n)$.

We can now make rigorous sense of formal relations such as $\int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} \, dx = \delta(x)$ that appeared in §11, for instance:

EXERCISE 13.76. Regarding the Dirac δ-function and the constant function 1 as tempered distributions on $\mathbb{R}^n$, show that $\mathscr{F}(\delta) = \mathscr{F}^*\delta = 1$, hence $\mathscr{F}^*(1) = \mathscr{F}(1) = \delta$.

EXERCISE 13.77. Show that the relations in (11.4) between the operators $\mathscr{F}$, $\mathscr{F}^*$ and ∂^α remain valid when $f \in \mathscr{S}(\mathbb{R}^n)$ is replaced by a tempered distribution $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$.

EXERCISE 13.78. Show that the relations $\mathscr{F}(\varphi * \Lambda) = \widehat{\varphi}\widehat{\Lambda}$ and $\mathscr{F}^*(\varphi * \Lambda) = \check{\varphi}\check{\Lambda}$ hold for all $\varphi \in \mathscr{S}(\mathbb{R}^n)$ and $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$.

13.7. Distributions with compact support. We saw in §11 that the Fourier transform exchanges regularity properties of a function with decay conditions at infinity, e.g. one can see this in the relations (11.4) that transform differentiation into multiplication by polynomials, and the fact that Lebesgue-integrable functions have continuous Fourier transforms. We would now like to explain a beautiful extension of this phenomenon into the realm of distributions.

DEFINITION 13.79. The **support** $\operatorname{supp}(\Lambda) \subset \Omega$ of a distribution $\Lambda \in \mathscr{D}'(\Omega)$ is the complement of the union of all open subsets $\mathcal{U} \subset \Omega$ such that $\Lambda(\varphi) = 0$ for all $\varphi \in \mathscr{D}(\Omega)$ with $\operatorname{supp}(\varphi) \subset \mathcal{U}$. Equivalently, $\operatorname{supp}(\Lambda)$ is the intersection of all closed subsets $\mathcal{V} \subset \Omega$ such that $\Lambda(\varphi) = 0$ for all $\varphi \in \mathscr{D}(\Omega)$ with $\operatorname{supp}(\varphi) \cap \mathcal{V} = \varnothing$.

REMARK 13.80. The support of $\Lambda \in \mathscr{D}'(\Omega)$ is in fact the smallest closed subset such that Λ vanishes on all test functions with support disjoint from $\operatorname{supp}(\Lambda)$, or equivalently, its complement is the largest open subset $\mathcal{U} \subset \Omega$ such that $\Lambda(\varphi)$ vanishes whenever $\operatorname{supp}(\varphi) \subset \mathcal{U}$. To see that $\Omega \backslash \operatorname{supp}(\Lambda)$ has the latter property, observe that for any $\varphi \in \mathscr{D}(\Omega)$ with $\operatorname{supp}(\varphi) \cap \operatorname{supp}(\Lambda) = \varnothing$, the compactness of $\operatorname{supp}(\varphi)$ implies that it is contained in the union of a finite collection of open subsets $\mathcal{U}_1, \ldots, \mathcal{U}_N$ such that $\operatorname{supp}(\psi) \subset \mathcal{U}_i$ implies $\Lambda(\psi) = 0$ for any $\psi \in \mathscr{D}(\Omega)$. One can then use a partition of unity to write φ as $\sum_{i=1}^N \varphi_i$ for some $\varphi_i \in \mathscr{D}(\Omega)$ with $\operatorname{supp}(\varphi_i) \subset \mathcal{U}_i$, implying $\Lambda(\varphi_i) = 0$ for all i and thus $\Lambda(\varphi) = 0$.

EXAMPLE 13.81. If $f \in L^1_{\mathrm{loc}}(\Omega)$ vanishes outside of some closed subset $\mathcal{V} \subset \Omega$, then $\operatorname{supp}(\Lambda_f) \subset \mathcal{V}$.

EXAMPLE 13.82. For any $\Lambda \in \mathscr{D}'(\Omega)$ and $f \in C^\infty(\Omega)$, $f\Lambda \in \mathscr{D}'(\Omega)$ has $\operatorname{supp}(f\Lambda) \subset \operatorname{supp}(f)$ since $f\varphi \equiv 0$ whenever $\varphi \in \mathscr{D}(\Omega)$ has support disjoint from that of f.

LEMMA 13.83. *A distribution $\Lambda \in \mathscr{D}'(\Omega)$ has compact support if and only if there exists a distribution $\Lambda' \in \mathscr{D}'(\Omega)$ and a smooth compactly supported function $f : \Omega \to \mathbb{K}$ such that $\Lambda = f\Lambda'$.*

PROOF. The statement is obvious in one direction since $\operatorname{supp}(f\Lambda') \subset \operatorname{supp}(f)$. Conversely, suppose there exists a compact subset $K \subset \Omega$ such that $\Lambda(\varphi) = 0$ whenever $\operatorname{supp}(\varphi) \cap K = \varnothing$. Choose an open neighborhood $\mathcal{U} \subset \Omega$ of K with compact closure and a compactly supported function $f : \Omega \to [0,1]$ such that $f|_{\mathcal{U}} \equiv 1$. We claim that $f\Lambda = \Lambda$. Indeed, for any $\varphi \in \mathscr{D}(\Omega)$, we can write $\varphi = f\varphi + (1-f)\varphi$, where $(1-f)\varphi$ vanishes on $\mathcal{U}$, thus its support is disjoint from K, implying $\Lambda(\varphi) = \Lambda(f\varphi) = (f\Lambda)(\varphi)$. □

PROPOSITION 13.84. *If $\Lambda \in \mathscr{D}'(\Omega)$ has compact support, then Λ extends to a continuous linear map on the space $C^\infty(\Omega)$ of all scalar-valued smooth functions with the C^∞_{loc}-topology.*

PROOF. Suppose $\varphi_j \in \mathscr{D}(\Omega)$ is a sequence converging in the C^∞_{loc}-topology to $\varphi_\infty \in \mathscr{D}(\Omega)$. By Lemma 13.83, we can write $\Lambda = f\Lambda'$ for some $\Lambda' \in \mathscr{D}'(\Omega)$ and a smooth function $f : \Omega \to \mathbb{K}$ with support in a compact set $K \subset \Omega$. Since $\varphi_j \to \varphi_\infty$ in the C^∞-topology over K, it follows that $f\varphi_j$ is C^∞-convergent to $f\varphi_\infty$, thus $f\varphi_j \to f\varphi_\infty$ in $\mathscr{D}(\Omega)$, so that the continuity of Λ' implies

$$\Lambda(\varphi_j) = \Lambda'(f\varphi_j) \to \Lambda'(f\varphi_\infty) = \Lambda(\varphi_j).$$

This proves that $\Lambda : \mathscr{D}(\Omega) \to V$ is continuous with respect to C^∞_{loc}-convergence. Since $\mathscr{D}(\Omega)$ is dense in $C^\infty(\Omega)$ with respect to this topology, it follows that Λ has a unique continuous extension to the larger space. □

REMARK 13.85. Proposition 13.84 also has a converse; see Proposition 13.105.

In light of the obvious continuous inclusion $\mathscr{S}(\mathbb{R}^n) \hookrightarrow C^\infty(\mathbb{R}^n)$, in which $C^\infty(\mathbb{R}^n)$ carries the C^∞_{loc}-topology, we also have:

COROLLARY 13.86. *Every compactly supported distribution $\Lambda \in \mathscr{D}'(\mathbb{R}^n)$ is also a tempered distribution, i.e. it has a unique extension to a continuous linear map on $\mathscr{S}(\mathbb{R}^n)$.* □

If $f \in L^1_{\text{loc}}(\mathbb{R}^n)$ has compact support, then f also belongs to $L^1(\mathbb{R}^n)$, so its Fourier transform is given by

$$\hat{f}(p) = \int_{\mathbb{R}^n} e^{-2\pi i p\cdot x} f(x)\, dx = \Lambda_f(e^{-2\pi i p\cdot x}),$$

where we have used Proposition 13.84 to extend the domain of Λ_f to smooth functions such as $x \mapsto e^{-2\pi i p\cdot x}$ that need not have compact support. As we saw in §11, the fact that f is of class L^1 implies that $\hat{f}$ is continuous, but we can now say more: since the product of f with any polynomial is also a compactly supported L^1_{loc}-function and therefore belongs to $L^1(\mathbb{R}^n)$, $\hat{f}$ also has continuous derivatives of all orders, i.e. it is smooth. The remarkable fact is that this result still holds when f is replaced by an arbitrary compactly supported distribution, which may have very badly behaved local singularities but still satisfies the best possible "decay" condition at infinity:

THEOREM 13.87. *For any compactly supported distribution Λ on $\mathbb{R}^n$, $\mathscr{F}\Lambda$ and $\mathscr{F}^*\Lambda$ are smooth functions on $\mathbb{R}^n$ given by*

$$\mathscr{F}\Lambda(p) = \Lambda(e^{-2\pi i p\cdot x}), \qquad \mathscr{F}^*\Lambda(p) = \Lambda(e^{2\pi i p\cdot x}),$$

where Proposition 13.84 is used for evaluating Λ on smooth functions with noncompact support.

PROOF. By Exercise 13.77, smoothness will follow immediately once we have proved that the stated formulas for $\mathscr{F}\Lambda$ and $\mathscr{F}^*\Lambda$ are correct, as multiplying Λ by any polynomial preserves the condition of compact support. We shall focus on the formula for $\mathscr{F}\Lambda$, since the parallel statement for $\mathscr{F}^*\Lambda$ has an almost identical proof. By Lemma 13.83, it would be equivalent to prove that for every $\Lambda \in \mathscr{D}'(\mathbb{R}^n)$ and every compactly supported smooth function $\psi : \mathbb{R}^n \to \mathbb{K}$, $\widehat{\psi\Lambda} \in \mathscr{S}'(\mathbb{R}^n)$ is given by the function $p \mapsto \Lambda(\psi e^{-2\pi i p\cdot x})$ on $\mathbb{R}^n$. The latter encapsulates two claims:

1. The function $g(p) := \Lambda(\psi e^{-2\pi i p\cdot x})$ has sufficiently tame behavior at infinity to define a tempered distribution;
2. For all $\varphi \in \mathscr{S}(\mathbb{R}^n)$,

$$(\psi\Lambda, \hat{\varphi}) = \int_{\mathbb{R}^n} \varphi(p) g(p)\, dp. \tag{13.8}$$

For the first claim, let us show that g has polynomial growth. Indeed, a straightforward change-of-variable calculation gives

$$\psi(x) e^{-2\pi i p\cdot x} = \mathscr{F}^*(\tau_p \hat{\psi})(x),$$

thus

$$g(p) = \Lambda(\mathscr{F}^* \tau_p \hat{\psi}) = (\mathscr{F}^*\Lambda)(\tau_p\hat{\psi}) = \tau_{-p}\check{\Lambda}(\hat{\psi}) = \tau_{-p}\check{\Lambda}(\sigma(\sigma\hat{\psi})) = (\sigma\hat{\psi} * \check{\Lambda})(-p),$$

and the claim follows from Proposition 13.70 since $\sigma\hat{\psi} \in \mathscr{S}(\mathbb{R}^n)$ and $\check{\Lambda} \in \mathscr{S}'(\mathbb{R}^n)$.

In light of this result, both sides of (13.8) now clearly define continuous linear functions of $\varphi \in \mathscr{S}(\mathbb{R}^n)$, so to prove that they are identical, it will suffice to show this for all φ in the dense subspace $\mathscr{D}(\mathbb{R}^n)$. The goal is thus to prove that

$$\Lambda\left(x \mapsto \psi(x) \int_{\mathbb{R}^n} e^{-2\pi i p\cdot x} \varphi(p)\, dp\right) = \int_{\mathbb{R}^n} \varphi(p) \Lambda(\psi e^{-2\pi i p\cdot x})\, dp$$

holds for all $\Lambda \in \mathscr{D}'(\mathbb{R}^n)$ and $\varphi, \psi \in \mathscr{D}(\mathbb{R}^n)$. Writing $\mathbb{1} \in \mathscr{D}'(\mathbb{R}^n)$ for the scalar-valued distribution $\mathbb{1}(\varphi) := \int_{\mathbb{R}^n} \varphi\, dm$, Theorem 13.42 identifies both sides of this equation with $(\Lambda \otimes \mathbb{1})(F)$ for the test function $F \in \mathscr{D}(\mathbb{R}^n \times \mathbb{R}^n)$ given by $F(x,p) := \psi(x)\varphi(p)e^{-2\pi i p\cdot x}$. □

13.8. Appendix: The topology of the space of test functions. For a working knowledge of the theory of distributions, it is usually not necessary to understand the topology of the space $\mathscr{D}(\Omega)$ beyond the notions described in Definition 13.7 of convergent sequences and continuity of linear maps on $\mathscr{D}(\Omega)$. Nonetheless, the further development of the theory requires knowing that $\mathscr{D}(\Omega)$ can also be viewed as a topological vector space, in which convergence and continuity are determined by the topology. You may have noticed in Definition 13.7 that the notion of convergence in $\mathscr{D}(\Omega)$ is extremely strict, i.e. it is very hard for a sequence of test functions to converge. This strictness is an advantage, because it means that it is that much easier for a linear functional on $\mathscr{D}(\Omega)$ to be continuous; in other words, having fewer convergent sequences in $\mathscr{D}(\Omega)$ makes the space of distributions $\mathscr{D}'(\Omega)$ larger. This will mean that the topology of $\mathscr{D}(\Omega)$ needs to be quite strong,[32] e.g. it needs to involve conditions on derivatives of arbitrarily high orders, and therefore cannot be described merely in terms of a norm, so $\mathscr{D}(\Omega)$ will not be a Banach space. One might reasonably hope for it to be a Fréchet space, like the Schwartz space $\mathscr{S}(\mathbb{R}^n)$ (see §13.6), but this will also turn out to be too ambitious (see Remark 13.90). The next best thing would be a locally convex space, and this is not hard to achieve.

13.8.1. *Definition and properties of the topology.* For a compact subset $K \subset \Omega$, consider the linear subspace

$$\mathscr{D}_K(\Omega) := \left\{\varphi \in \mathscr{D}(\Omega) \;\middle|\; \operatorname{supp}(\varphi) \subset K\right\}.$$

The countable family of norms $\|\cdot\|_{C^m}$ for integers $m \geqslant 0$ endows $\mathscr{D}_K(\Omega)$ with the structure of a Fréchet space such that convergence of a sequence $\varphi_j \to \varphi_\infty$ in $\mathscr{D}_K(\Omega)$ means uniform convergence $\partial^\alpha \varphi_j \to \partial^\alpha \varphi_\infty$ for every multi-index α. We shall assume $\mathscr{D}_K(\Omega)$ to be endowed with this Fréchet space topology from now on. It will frequently be useful to observe that since $\mathscr{D}_K(\Omega)$ is metrizable, a function defined on $\mathscr{D}_K(\Omega)$ is continuous if and only if it is sequentially continuous.

According to Definition 13.7, a convergent sequence in $\mathscr{D}_K(\Omega)$ is also convergent in $\mathscr{D}(\Omega)$, so the topology we define on $\mathscr{D}(\Omega)$ should have the property that the obvious inclusion

$$\mathscr{D}_K(\Omega) \hookrightarrow \mathscr{D}(\Omega) \tag{13.9}$$

is sequentially continuous for every compact set $K \subset \Omega$. If these inclusions are continuous, and $\|\cdot\|$ is a seminorm on $\mathscr{D}(\Omega)$ that is continuous with respect to its topology, then $\|\cdot\|$ will also restrict to a continuous seminorm on $\mathscr{D}_K(\Omega)$. The following definition therefore produces the strongest locally convex topology on $\mathscr{D}(\Omega)$ for which the inclusions (13.9) are all continuous.

Definition 13.88. A **good seminorm** on $\mathscr{D}(\Omega)$ is a seminorm whose restriction to the subspace $\mathscr{D}_K(\Omega) \subset \mathscr{D}(\Omega)$ is continuous for every $K \subset \Omega$ compact. We endow $\mathscr{D}(\Omega)$ with the locally convex topology generated by the family of all good seminorms, i.e. a set $\mathcal{U} \subset \mathscr{D}(\Omega)$ is open if and only if for every $\varphi \in \mathcal{U}$, there exists a seminorm $\|\cdot\|$ on $\mathscr{D}(\Omega)$ such that

$$\left\{\psi \in \mathscr{D}(\Omega) \;\middle|\; \|\psi - \varphi\| < 1\right\} \subset \mathcal{U}$$

and $\|\cdot\|$ is continuous on $\mathscr{D}_K(\Omega)$ for all $K \subset \Omega$ compact.[33]

---

[32] Given two topologies $\mathcal{T}_1$ and $\mathcal{T}_2$ on the same set X, one says that $\mathcal{T}_1$ is **stronger** (or **finer**, or **larger**) than $\mathcal{T}_2$ if every set in $\mathcal{T}_2$ also belongs to $\mathcal{T}_1$. One also says in this case that $\mathcal{T}_2$ is **weaker** (or **coarser** or **smaller**) than $\mathcal{T}_1$. Making a topology on X stronger makes it harder for sequences in X to converge and harder for maps from other spaces into X to be continuous, but easier for maps from X to other spaces to be continuous.

[33] In describing the topology of $\mathscr{D}(\Omega)$ in this way, we are using the easily verifiable fact that the maximum of any finite collection of good seminorms is also a good seminorm, and so is any positive multiple of a good seminorm. This implies that for any collection of good seminorms $\|\cdot\|_i$ and any $\epsilon_i > 0$ with $i = 1, \ldots, N$, the finite intersection of the open neighborhoods $\{\psi \in \mathscr{D}(\Omega) \mid \|\psi - \varphi\|_i < \epsilon_i\}$ for $i = 1, \ldots, N$ can equally well be described as $\{\psi \in \mathscr{D}(\Omega) \mid \|\psi - \varphi\| < 1\}$ where $\|f\| := \max\left\{\frac{\|f\|_1}{\epsilon_1}, \ldots, \frac{\|f\|_N}{\epsilon_N}\right\}$ defines another good seminorm.

The next exercise shows that good seminorms on $\mathscr{D}(\Omega)$ exist in abundance, thus the topology we have defined on $\mathscr{D}(\Omega)$ is quite large.

EXERCISE 13.89. Show that each of the following defines a good seminorm on $\mathscr{D}(\Omega)$:

(a) $\|\varphi\|_\alpha := \max_{x\in\Omega} |\partial^\alpha \varphi(x)|$ for any multi-index α. (The C^m-norm for any $m \geqslant 0$ is a finite sum of seminorms of this type, thus it is also a good seminorm.)
(b) $\|\varphi\|_f := \|f\varphi\|$ where $\|\cdot\|$ is any good seminorm and $f : \Omega \to \mathbb{R}$ is any smooth function.
(c) $\|\varphi\|_{f,\alpha} := \|f\varphi\|_\alpha$ for any multi-index α and continuous function $f : \Omega \to \mathbb{R}$. For this example, the open set $\{\|\varphi\|_{f,\alpha} < 1\}$ describes all $\varphi \in \mathscr{D}(\Omega)$ that satisfy $|\partial^\alpha\varphi| < 1/|f|$ everywhere on Ω, where we adopt the convention $1/0 := \infty$ so that the condition is vacuous wherever $f = 0$.

REMARK 13.90. The following observations show that the topology we've defined on $\mathscr{D}(\Omega)$ cannot be metrizable, so $\mathscr{D}(\Omega)$ is not a Fréchet space. If d were a metric defining the topology of $\mathscr{D}(\Omega)$, then for every $\varphi \in \mathscr{D}(\Omega)$, the sets $\mathcal{U}_j := \{\psi \in \mathscr{D}(\Omega) \mid d(\varphi,\psi) < 1/j\}$ for $j \in \mathbb{N}$ would define a countable sequence of neighborhoods of φ with the property that *every* neighborhood of φ contains $\mathcal{U}_k$ for some $k \in \mathbb{N}$.[34] Since the topology is determined by good seminorms, this would equivalently mean that there exists a sequence of good seminorms $\|\cdot\|_j$ for $j \in \mathbb{N}$ such that for every good seminorm $\|\cdot\|$, the set $\{\psi \in \mathscr{D}(\Omega) \mid \|\psi\| < 1\}$ contains $\{\psi \in \mathscr{D}(\Omega) \mid \|\psi\|_k < 1\}$ for some $k \in \mathbb{N}$; in other words,

$$\|\cdot\| \leqslant \|\cdot\|_k \quad \text{for some} \quad k \in \mathbb{N}.$$

By Exercise 13.89, it would follow that for every continuous function $f : \Omega \to \mathbb{R}$, there exists $k \in \mathbb{N}$ such that

$$\|f\varphi\|_{C^0} \leqslant \|\varphi\|_k \quad \text{for all} \quad \varphi \in (\Omega).$$

To see that this is impossible, pick a sequence of nontrivial functions $\varphi_1, \varphi_2, \varphi_3, \ldots \in \mathscr{D}(\Omega)$ whose supports are pairwise disjoint compact sets $K_1, K_2, K_3, \ldots \subset \Omega$, and choose $f : \Omega \to \mathbb{R}$ to be a continuous function that satisfies

$$f > \frac{\|\varphi_j\|_j}{\|\varphi_j\|_{C^0}} \quad \text{on } K_j \quad \text{for all } j \in \mathbb{N}.$$

Then $\|f\varphi_j\|_{C^0} > \|\varphi_j\|_j$ for every $j \in \mathbb{N}$, giving a contradiction.

LEMMA 13.91. *If $\varphi_j \to \varphi_\infty$ in the topology of $\mathscr{D}(\Omega)$, then there exists a compact subset $K \subset \Omega$ such that $\varphi_j \in \mathscr{D}_K(\Omega)$ for every $j \in \mathbb{N} \cup \{\infty\}$.*

PROOF. If not, then after replacing φ_j with a subsequence, we can find a sequence of points $x_j \in \Omega$ that lie outside the support of φ_∞, have no accumulation point, and satisfy $\varphi_j(x_j) \neq 0$ for every j. Choose a continuous function $f : \Omega \to (0,\infty)$ such that $f(x_j) \leqslant |\varphi_j(x_j)|$ for every j. Then by Exercise 13.89, $\mathcal{U} := \{\varphi \in \mathscr{D}(\Omega) \mid |\varphi - \varphi_\infty| < f\}$ is an open neighborhood of φ_∞ in $\mathscr{D}(\Omega)$, but $\varphi_j \notin \mathcal{U}$ for every j, so φ_j cannot converge to φ_∞. □

COROLLARY 13.92. *A sequence $\varphi_j \in \mathscr{D}(\Omega)$ converges to $\varphi_\infty \in \mathscr{D}(\Omega)$ if and only if there exists a compact set $K \subset \Omega$ such that $\varphi_j \in \mathscr{D}_K(\Omega)$ for all $j \in \mathbb{N} \cup \{\infty\}$ and $\varphi_j \to \varphi_\infty$ in the topology of $\mathscr{D}_K(\Omega)$.* □

As preparation for the next result, we need some general facts about continuity for linear maps between locally convex spaces. A preliminary remark about locally convex topologies is in order. If X carries the locally convex topology generated by a given family of seminorms

[34] A collection of neighborhoods with this property is called a **countable neighborhood base** of φ. A topological space in which every point admits a countable neighborhood base is called **first countable**. What Remark 13.90 shows in effect is that every metrizable space is first countable, but $\mathscr{D}(\Omega)$ is not.

$\{\|x\|_\alpha\}_{\alpha\in I}$,then by definition, every open set in X is a union of finite intersections of sets of the form $\{x \in X \,|\|x - x_0\|_\alpha < \epsilon\}$ for arbitrary $x_0 \in X$, $\alpha \in I$ and $\epsilon > 0$. Equivalently, a set $\mathcal{U} \subset Y$ is open if and only if for every $x_0 \in \mathcal{U}$, there exists a nonempty finite subset $I_0 \subset I$ and numbers $\epsilon_\alpha > 0$ for $\alpha \in I_0$ such that

$$x \in X \text{ with } \|x - x_0\|_\alpha < \epsilon_\alpha \text{ for all } \alpha \in I_0 \quad \Rightarrow \quad x \in \mathcal{U}.$$

The seminorms $\|\cdot\|_\alpha : X \to [0,\infty)$ are each continuous functions, and in the situation above, $\|x\| := \sum_{\alpha\in I_0} \frac{\|x\|_\alpha}{\epsilon_\alpha}$ also defines a continuous seminorm; the aforementioned condition can then equally well be described as

$$x \in X \text{ with } \|x - x_0\| < 1 \quad \Rightarrow \quad x \in \mathcal{U}.$$

This provides a briefer way of characterizing open sets: $\mathcal{U} \subset X$ is open if and only if for every $x_0 \in \mathcal{U}$, there exists a continuous seminorm $\|\cdot\|$ such that every $x \in X$ with $\|x - x_0\| < 1$ belongs to $\mathcal{U}$. The sufficiency of this condition is clear since continuity of $\|\cdot\|$ implies that every set of the form $\{x \in X \mid \|x - x_0\| < 1\}$ is open.

Lemma 13.93. *On a topological vector space X, a seminorm $\|\cdot\| : X \to \mathbb{R}$ is continuous if and only if the set $\{x \in X \mid \|x\| < 1\} \subset X$ is open.*

Proof. In one direction, the implication is an immediate consequence of the definition of continuity and the fact that $(-1,1) \subset \mathbb{R}$ is open. For the converse, we use the fact that for every $x_0 \in X$ and $\epsilon > 0$, the invertible affine map $\Phi : X \to X : x \mapsto x_0 + \epsilon x$ is a homeomorphism, thus if $B := \{x \in X \,|\|x\| < 1\}$ is open, then so is $\Phi(B) = \{x \in X \mid \|x - x_0\| < \epsilon\}$. Given this, if $\mathcal{V} \subset [0,\infty)$ is any open subset and $x_0 \in X$ satisfies $\|x_0\| \in \mathcal{V}$, then choosing any $\epsilon > 0$ such that $(\|x_0\| - \epsilon, \|x_0\| + \epsilon) \subset \mathcal{V}$, the triangle inequality implies that every $x \in X$ in the open set $\{x \in X \mid \|x - x_0\| < \epsilon\}$ satisfies $\|x\| \leqslant \|x_0\| + \|x - x_0\| < \|x_0\| + \epsilon$ and $\|x\| \geqslant \|x_0\| - \|x_0 - x\| > \|x_0\| - \epsilon$, so this open subset belongs to the preimage of $\mathcal{V}$ under $\|\cdot\|$, proving that this preimage is open. $\square$

Lemma 13.94. *Suppose X is a locally convex space whose topology is determined by the family of seminorms $\{\|\cdot\|_\alpha\}_{\alpha\in I}$. Then a seminorm $\|\cdot\|$ on X is continuous if and only if there exists a nonempty finite subset $I_0 \subset I$ and a constant $C > 0$ such that*

$$\|x\| \leqslant C \sum_{\alpha\in I_0} \|x\|_\alpha \quad \textit{for all} \quad x \in X.$$

Proof. We claim first that if $\|\cdot\|_1$ is a continuous seminorm and $\|\cdot\| \leqslant \|\cdot\|_1$, then $\|\cdot\|$ is also continuous. Indeed, consider $B := \{x \in X \mid \|x\| < 1\}$, and for any $x_0 \in B$, choose $\epsilon > 0$ such that $\|x_0\| + \epsilon < 1$. If $\|\cdot\|_1$ is continuous, then the set $\mathcal{U} := \{x \in X \mid \|x - x_0\|_1 < \epsilon\}$ is an open neighborhood of x_0, and if $\|\cdot\| \leqslant \|\cdot\|_1$, then every $x \in \mathcal{U}$ satisfies

$$\|x\| \leqslant \|x_0\| + \|x - x_0\| \leqslant \|x_0\| + \|x - x_0\|_1 < \|x_0\| + \epsilon < 1,$$

implying $\mathcal{U} \subset B$. This proves that $B \subset X$ is open, so by Lemma 13.93, $\|\cdot\|$ is continuous.

By the assumptions of the lemma, the seminorms $\|\cdot\|_\alpha$ are continuous for all $\alpha \in I$, thus $C \sum_{\alpha\in I_0} \|\cdot\|_\alpha$ is also a continuous seminorm for any $C > 0$ and any finite set $I_0 \subset I$. The claim in the previous paragraph thus implies one direction of the lemma.

For the other direction, assume $\|\cdot\|$ is continuous, so $B := \{x \in X \mid \|x\| < 1\}$ is an open set. Since the family of seminorms $\{\|\cdot\|_\alpha\}_{\alpha\in I}$ generates the topology of X, it follows that B contains a neighborhood of $0 \in X$ in the form

$$\mathcal{U} := \{x \in X \mid \|x\|_\alpha < \epsilon_\alpha \text{ for every } \alpha \in I_0\}$$

for some nonempty finite subset $I_0 \subset I$ and real numbers $\{\epsilon_\alpha > 0\}_{\alpha\in I_0}$. In other words,

$$\|x\|_\alpha < \epsilon_\alpha \text{ for every } \alpha \in I_0 \quad \Rightarrow \quad \|x\| < 1. \tag{13.10}$$

We claim that $\|x\| \leqslant C \sum_{\alpha\in I_0} \|x\|_\alpha$ holds for every $x \in X$, where $C > 0$ is a constant independent of x. There is nothing to prove if $\|x\| = 0$, so consider $x \in X$ with $\|x\| > 0$. At least one of the $\|x\|_\alpha$ for $\alpha \in I_0$ must then also be positive, as otherwise multiplying x by a sufficiently large positive scalar would produce a contradiction to (13.10). The quotient

$$Q(x) := \frac{\|x\|}{\sum_{\alpha\in I_0} \|x\|_\alpha}$$

is therefore well defined whenever $\|x\| > 0$, and we claim that on this subset of X, it is bounded. If not, then there exists a sequence $x_j \in X$ with $\|x_j\| > 0$ and $Q(x_j) \to \infty$. But each x_j can be multiplied by a positive scalar without changing the value of $Q(x_j)$, thus we are free to assume without loss of generality that the denominator in the definition of $Q(x_j)$ some fixed constant less than $\min_{\alpha\in I_0} \epsilon_\alpha$ for every j. In this case, (13.10) implies that the numerator is less than 1 and thus gives a bound on $Q(x_j)$, which is a contradiction. □

LEMMA 13.95. *For two locally convex spaces X and Y, a linear map $\Lambda : X \to Y$ is continuous if and only if for every continuous seminorm $\|\cdot\|_Y$ on Y, there exists a continuous seminorm $\|\cdot\|_X$ on X such that $\|\Lambda(x)\|_Y \leqslant \|x\|_X$.*

PROOF. Assume the second condition holds, $\mathcal{V} \subset Y$ is an open set, and $x_0 \in X$ is a point with $y_0 := \Lambda(x_0) \in \mathcal{V}$. The openness of $\mathcal{V}$ implies that for some continuous seminorm which we will denote by $\|\cdot\|_Y$, $\{y \in Y \mid \|y - y_0\|_Y < 1\}$ defines an open neighborhood of y that is contained in $\mathcal{V}$. If $\|\cdot\|_X$ is a continuous seminorm on X satisfying $\|\Lambda(x)\|_Y \leqslant \|x\|_X$ for all X, it follows that $\{x \in X \mid \|x - x_0\|_X < 1\}$ is an open neighborhood of x_0 in X such that for all x in this neighborhood, $\|\Lambda(x) - y_0\|_Y = \|\Lambda(x - x_0)\|_Y \leqslant \|x - x_0\|_X < 1$, implying $x \in \Lambda^{-1}(\mathcal{V})$. This proves that $\Lambda^{-1}(\mathcal{V}) \subset X$ is open and thus that Λ is continuous.

Conversely, suppose Λ is continuous and $\|\cdot\|_Y$ is an arbitrary continuous seminorm on Y. Then $B := \{y \in Y \mid \|y\|_Y < 1\}$ is open, hence $\Lambda^{-1}(B) \subset X$ is an open neighborhood of 0 and therefore contains $\mathcal{U} := \{x \in X \mid \|x\|_X < 1\}$ for some continuous seminorm $\|\cdot\|_X$ on X. In other words, we have

$$\|x\|_X < 1 \qquad \Rightarrow \qquad \|\Lambda(x)\|_Y < 1 \tag{13.11}$$

for all $x \in X$. We claim that $\|\Lambda(x)\|_Y \leqslant \|x\|_X$ holds for all $x \in X$. If $\|\Lambda(x)\|_Y = 0$ there is nothing to prove, so assume $\|\Lambda(x)\| > 0$. Then $\|x\|_X$ must also be positive, as otherwise multiplying x by a sufficiently large positive scalar produces a contradiction to (13.11). It follows that the quotient $Q(x) := \|\Lambda(x)\|_Y / \|x\|_X > 0$ is well defined whenever its numerator is nonzero, and clearly it does not change if x is multiplied by any positive scalar, thus we are free to assume $\|x\|_X = 1 - \epsilon$ for any $\epsilon > 0$ arbitrarily small. Under this assumption, (13.11) implies $|\Lambda(x)\|_Y < 1$ and thus $Q(x) < 1/(1 - \epsilon)$; since $\epsilon > 0$ was arbitrary, it follows that $Q(x) \leqslant 1$. □

PROPOSITION 13.96. *For any locally convex space X, a linear map $\Lambda : \mathscr{D}(\Omega) \to X$ is continuous if and only if its restrictions $\Lambda|_{\mathscr{D}_K(\Omega)} : \mathscr{D}_K(\Omega) \to X$ are continuous for all compact $K \subset \Omega$.*

PROOF. Since the inclusions $\mathscr{D}_K(\Omega) \hookrightarrow \mathscr{D}(\Omega)$ are continuous, the statement is obvious in one direction. We need to show that if Λ has a continuous restriction to every $\mathscr{D}_K(\Omega)$, then it is continuous on $\mathscr{D}(\Omega)$. For this, it will be convenient to choose an open covering of Ω by countably many subsets $\{\Omega_j\}_{j\in\mathbb{N}}$ with the following properties:

(1) The covering is **locally finite**, i.e. every point in Ω has a neighborhood that intersects at most finitely many of the Ω_j;
(2) $K_j := \bar{\Omega}_j$ is compact for every j.

For a concrete construction of $\{\Omega_j\}_{j\in\mathbb{N}}$, choose a strictly increasing sequence $r_j > 0$ with $\lim_{j\to\infty} r_j = \sup\{|x| \mid x \in \Omega\}$, another sequence $\epsilon_j > 0$ such that $r_1 - \epsilon_1 > 0$ and $r_j - \epsilon_j > r_{j-1}$ for every $j \geqslant 2$, and define

$$\Omega_j := \left\{ x \in \Omega \;\middle|\; r_{j-1} - \epsilon_{j-1} < |x| < r_j \right\}$$

where for $j = 1$ we interpret the lower bound on $|x|$ as a vacuous condition. With this construction, it is clear that one can also find a sequence of smooth functions $\rho_j : \Omega \to [0,1]$ such that each ρ_j is supported in Ω_j and $\sum_{j=1}^{\infty} \rho_j \equiv 1$, where the sum is finite at every point due to the local finiteness of the open covering.[35] Any $\varphi \in \mathscr{D}(\Omega)$ can now be decomposed as

$$\varphi = \sum_{j=1}^{\infty} \varphi_j, \qquad \text{where} \qquad \varphi_j := \rho_j \varphi \text{ has support in } K_j.$$

Observe that for every $\varphi \in \mathscr{D}(\Omega)$, only finitely many of the functions φ_j can be nonzero: indeed, the local finiteness of the covering $\{\Omega_j\}$ implies that at most finitely many of the sets Ω_j can intersect the compact set $\operatorname{supp}(\varphi)$.

To show that $\Lambda : \mathscr{D}(\Omega) \to X$ is continuous, it suffices by Lemma 13.95 to show that for any continuous seminorm $\|\cdot\|_X$ on X, there exists a good seminorm $\|\cdot\|$ on $\mathscr{D}(\Omega)$ such that $\|\Lambda(\varphi)\|_X \leqslant \|\varphi\|$ for all $\varphi \in \mathscr{D}(\Omega)$. The topology of $\mathscr{D}_{K_j}(\Omega)$ for each $j \in \mathbb{N}$ is generated by the monotone sequence of norms $\|\cdot\|_{C^m}$ for $m = 0, 1, 2, \ldots$, thus continuity of Λ on $\mathscr{D}_{K_j}(\Omega)$ implies that there exists an integer $m_j \geqslant 0$ and a positive number c_j such that

$$\|\Lambda(\psi)\|_X \leqslant c_j \|\psi\|_{C^{m_j}} \quad \text{for all} \quad \psi \in \mathscr{D}_{K_j}(\Omega).$$

Since the sum $\varphi = \sum_j \varphi_j$ is finite for each $\varphi \in \mathscr{D}(\Omega)$ and $\varphi_j \in \mathscr{D}_{K_j}(\Omega)$ for $j = 1, 2, 3, \ldots$, we can apply the triangle inequality and write

$$\|\Lambda(\varphi)\|_X \leqslant \sum_j \|\Lambda(\varphi_j)\|_X \leqslant \sum_j c_j \|\varphi_j\|_{C^{m_j}} \leqslant c_j' \|\varphi\|_{C^{m_j}},$$

where each of the modified constants $c_j' > 0$ depends on the C^{m_j}-norm of ρ_j but not on φ. With these constants fixed, it is easy to check that

$$\|\varphi\| := \sum_{j=1}^{\infty} c_j' \|\varphi\|_{C^{m_j}(\Omega_j)}$$

defines a good seminorm on $\mathscr{D}(\Omega)$, as for any compact $K \subset \Omega$, the restriction of this seminorm to $\mathscr{D}_K(\Omega)$ has only finitely many nonzero terms, and C^∞-convergence in $\mathscr{D}_K(\Omega)$ implies that each individual term converges. Since $\|\Lambda\varphi\|_X \leqslant \|\varphi\|$ by construction, this establishes that $\Lambda : \mathscr{D}(\Omega) \to Y$ is continuous. □

The following easy consequence completes the proof of Proposition 13.8:

COROLLARY 13.97. *For any locally convex space X, a linear map $\Lambda : \mathscr{D}(\Omega) \to X$ is continuous if and only if it is sequentially continuous, i.e. for every convergent sequence $\varphi_j \to \varphi_\infty$ in $\mathscr{D}(\Omega)$, $\Lambda(\varphi_j) \to \Lambda(\varphi_\infty)$.*

PROOF. By a standard result in point-set topology, continuous maps are always sequentially continuous. Conversely, if $\Lambda : \mathscr{D}(\Omega) \to X$ is sequentially continuous, then its restriction to $\mathscr{D}_K(\Omega)$ for each compact set $K \subset \Omega$ is sequentially continuous, and since $\mathscr{D}_K(\Omega)$ is metrizable, it follows that the restriction of Λ to $\mathscr{D}_K(\Omega)$ is also continuous. By Proposition 13.96, Λ itself is therefore continuous. □

[35]A collection of functions ρ_j with these properties is called a **partition of unity** on Ω.

REMARK 13.98. By another standard result in point-set topology (see e.g. [**Wen23**, §4]), a sequentially continuous map $f : X \to Y$ between two topological spaces is continuous whenever X is first countable. We did not claim this to be true for $\mathscr{D}(\Omega)$, which is *not* first countable according to Remark 13.90, but Corollary 13.97 says that it is nonetheless true specifically for *linear* maps to other locally convex spaces. Philosophically, the reason this works is that $\mathscr{D}(\Omega)$ can be viewed—in a sense to be made precise in §13.8.2 below—as a *limit* of a family of spaces in which sequential continuity does imply continuity, namely the metrizable spaces $\mathscr{D}_K(\Omega)$ for $K \subset \Omega$ compact.

13.8.2. *Inductive limits.* The topology we've defined on $\mathscr{D}(\Omega)$ is often referred to as an *inductive limit* topology. While one can understand all of its properties without knowing what this term means, let us take a moment to discuss the wider context in which it arises.

We need to introduce a few notions from abstract category theory. For the particular application relevant here, the "category" we have in mind is the class of locally convex spaces (these are the *objects* of the category), and the natural class of maps between two such spaces consists of all continuous linear maps (these are the *morphisms* of the category). We shall formulate the definitions below in terms of this particular category just for concreteness, but they would still make sense in any other category, e.g. topological spaces and continuous maps, vector spaces and linear maps, groups and group homomorphisms, and so forth.

Suppose I is a set with a pre-order $\prec$, i.e. $\prec$ is reflexive ($\alpha \prec \alpha$) and transitive ($\alpha \prec \beta$ and $\beta \prec \gamma$ implies $\alpha \prec \gamma$), but the relations $\alpha \prec \beta$ and $\beta \prec \alpha$ need not imply $\alpha = \beta$, so $\prec$ need not be a partial order. The pair $(I, \prec)$ is called a **directed set** if for every pair $\alpha, \beta \in I$, there exists $\gamma \in I$ with $\gamma \succ \alpha$ and $\gamma \succ \beta$. An obvious example is $\mathbb{N}$ with its usual total order $\prec := \leqslant$. A more interesting and relevant example for our purposes is to define I as the set of all compact subsets in a fixed open set $\Omega \subset \mathbb{R}^n$, with $K \prec K'$ defined to mean $K \subset K'$. Notice that the ordering relation in this example is a partial order, but not a total order since for any two compact subsets, it need not be true that either is contained in the other. It forms a directed set because whenever $K, K' \subset \Omega$ are both compact, $K \cup K'$ is another compact subset of Ω that contains both of them.

DEFINITION 13.99. A **direct system** (or **inductive system**) of locally convex spaces consists of a directed set $(I, \prec)$ and a family of locally convex spaces $\{X_\alpha\}_{\alpha \in I}$ together with continuous linear maps $\varphi_{\beta\alpha} : X_\alpha \to X_\beta$ defined for each $\alpha, \beta \in I$ with $\alpha \prec \beta$, such that

$$\varphi_{\alpha\alpha} = \mathrm{Id}_{X_\alpha}$$

and the diagram

$$\begin{tikzcd} X_\alpha \arrow[r, "\varphi_{\beta\alpha}"] \arrow[rr, bend right, "\varphi_{\gamma\alpha}"'] & X_\beta \arrow[r, "\varphi_{\gamma\beta}"] & X_\gamma \end{tikzcd}$$

commutes for every triple $\alpha, \beta, \gamma \in I$ with $\alpha \prec \beta \prec \gamma$.

The notion of "convergence" for a direct system must necessarily look somewhat different from what we've seen before for sequences, as there is no meaningful topology to be defined on the "set" of all locally convex spaces.[36] The idea is instead to measure the convergence of a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ in terms of the continuous linear maps from each X_α to other fixed spaces.

DEFINITION 13.100. For a direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ of locally convex spaces over the directed set $(I, \prec)$, a **target** $\{Y, f_\alpha\}$ of the system consists of a locally convex space Y together with

[36] And strictly speaking, the collection of all locally convex spaces is far too large to be called a *set*; it is instead a *proper class*. This remark is included only for the sake of readers who truly care about abstract set theory.

associated continuous linear maps $f_\alpha : X_\alpha \to Y$ for each $\alpha \in I$ such that the diagram

$$\begin{tikzcd} X_\alpha \arrow[rr, "\varphi_{\beta\alpha}"] \arrow[dr, "f_\alpha"'] & & X_\beta \arrow[dl, "f_\beta"] \\ & Y & \end{tikzcd}$$

commutes for every pair $\alpha, \beta \in I$ with $\alpha \prec \beta$.

Definition 13.101. A target $\{X_\infty, \varphi_\alpha\}$ of the direct system $\{X_\alpha, \varphi_{\beta\alpha}\}$ is called a **direct limit** (or **inductive limit** or **colimit**) of the system and written as

$$X_\infty = \varinjlim\{X_\alpha\}$$

if it satisfies the following "universal" property: for all targets $\{Y, f_\alpha\}$ of $\{X_\alpha, \varphi_{\beta\alpha}\}$, there exists a unique continuous linear map $f_\infty : X_\infty \to Y$ such that the diagram

$$\begin{tikzcd} X_\alpha \arrow[r, "\varphi_\alpha"] \arrow[dr, "f_\alpha"'] & X_\infty \arrow[d, "f_\infty"] \\ & Y \end{tikzcd}$$

commutes for every $\alpha \in I$.

The essential meaning of a direct limit can be encoded in the diagram

$$\begin{tikzcd} X_\alpha \arrow[r, "\varphi_{\beta\alpha}"] \arrow[drrrr] & X_\beta \arrow[r, "\varphi_{\gamma\beta}"] \arrow[drrr] & X_\gamma \arrow[r] \arrow[drr] & \dots \arrow[r] & \varinjlim\{X_\alpha\} \arrow[d, dashed] \\ & & & & Y \end{tikzcd}$$

where we assume $\alpha \prec \beta \prec \gamma \prec \ldots \in I$. The key feature of the space $\varinjlim\{X_\alpha\}$ is that whenever a space Y and continuous linear maps $X_\alpha \to Y$ in a commuting diagram of this type are given, the "limit" map from $\varinjlim\{X_\alpha\}$ to Y indicated by the dashed arrow must also exist (as a continuous linear map) and be unique.

Note that these definitions on their own give no guarantee for any given direct system that a direct limit must exist, and if it exists, then it is generally not unique. Indeed:

Exercise 13.102. If $\{X, f_\alpha\}$ is a direct limit of $\{X_\alpha, \varphi_{\beta\alpha}\}$ and Y is another locally convex space such that there exists a continuous linear isomorphism $\psi : X \to Y$ with a continuous inverse, show that $\{Y, \psi \circ f_\alpha\}$ is also a direct limit of $\{X_\alpha, \varphi_{\beta\alpha}\}$.
Remark: The invertibility of ψ is needed only for showing that $\{Y, \psi \circ f_\alpha\}$ satisfies the universal property; it is already a target without this.

The non-uniqueness exhibited by the exercise above is however the worst thing that can happen: if $\{X, f_\alpha\}$ and $\{Y, g_\alpha\}$ are any two direct limits of the same system $\{X_\alpha, \varphi_{\beta\alpha}\}$, then the universal property provides unique continuous linear maps $g_\infty : X \to Y$ and $f_\infty : Y \to X$ satisfying $g_\infty \circ f_\alpha = g_\alpha$ and $f_\infty \circ g_\alpha = f_\alpha$ for every $\alpha \in I$. It follows that $f_\infty \circ g_\infty$ is the unique continuous linear map $X \to X$ satisfying $(f_\infty \circ g_\infty) \circ f_\alpha = f_\alpha$ for every $\alpha \in I$, which implies $f_\infty \circ g_\infty = \mathrm{Id}_X$. A similar argument shows $g_\infty \circ f_\infty = \mathrm{Id}_Y$, thus X and Y are isomorphic, and there is a distinguished isomorphism relating them. For this reason, we typically refer to "the" (rather than "a") direct limit of any system for which a limit exists.

Example 13.103. Given an open set $\Omega \subset \mathbb{R}^n$, take $(I, \prec)$ to be the set of all compact subsets $K \subset \Omega$ with $K \prec K'$ defined to mean $K \subset K'$. There is then a direct system $\{X_K, \varphi_{K',K}\}$ over $(I, \prec)$ such that $X_K = \mathscr{D}_K(\Omega)$ and $\varphi_{K',K}$ is the obvious inclusion map $\mathscr{D}_K(\Omega) \hookrightarrow \mathscr{D}_{K'}(\Omega)$, defined whenever $K \subset K'$. Define $\varphi_K : \mathscr{D}_K(\Omega) \hookrightarrow \mathscr{D}(\Omega)$ also as the natural inclusion for each $K \in I$. Proposition 13.96 can then be reinterpreted as the statement that $\{\mathscr{D}(\Omega), \varphi_K\}$ is a universal target for the direct system $\{\mathscr{D}_K(\Omega), \varphi_{K',K}\}$, in other words,

$$\mathscr{D}(\Omega) = \varinjlim\{\mathscr{D}_K(\Omega)\}.$$

This is why the topology of $\mathscr{D}(\Omega)$ is often called the **inductive limit topology** determined by the natural Fréchet space topologies of $\mathscr{D}_K(\Omega)$ for all compact $K \subset \Omega$.

REMARK 13.104. One really should call the topology on $\mathscr{D}(\Omega)$ a **locally convex inductive limit topology**, as omitting the words "locally convex" can potentially cause confusion. A topologist would interpret the words "inductive limit topology" to mean a universal target in the sense of Definition 13.101, but with X_∞ and Y allowed in general to be arbitrary topological spaces (not necessarily topological vector spaces), and all maps required to be continuous but not necessarily linear. It is not hard to show that the direct limit in this sense of the system $\{\mathscr{D}_K(\Omega), \varphi_{K',K}\}$ can be identified again with the vector space $\mathscr{D}(\Omega)$, but endowed with an even stronger topology, for which a set $\mathcal{U} \subset \mathscr{D}(\Omega)$ is open if and only if $\mathcal{U} \cap \mathscr{D}_K(\Omega) \subset \mathscr{D}_K(\Omega)$ is open for every compact $K \subset \Omega$. This topology has the same notion of convergent sequences as the locally convex topology we defined, and it has the nice property that for any topological space X, a (not necessarily linear) map $f : \mathscr{D}(\Omega) \to X$ is continuous if and only if its restriction to $\mathscr{D}_K(\Omega)$ is continuous for every compact $K \subset \Omega$. However, since this topology contains sets that are not open in the locally convex inductive limit topology, it cannot be locally convex—in fact there is no good reason to expect $\mathscr{D}(\Omega)$ with this topology to be a topological vector space.

13.8.3. *Comparison with other topologies.* There are other natural topologies one could imagine defining on the space of smooth functions with compact support, and it is natural to wonder why the inductive limit topology defined in §13.8.1 is a better choice. The obvious answer is that since we defined the topology on $\mathscr{D}(\Omega)$ to be as strong as possible while still being locally convex, this makes the space of distributions $\mathscr{D}'(\Omega)$ as large as possible. But let us briefly discuss some alternatives. In order to avoid confusion, we will refer to the space of smooth compactly supported functions $\Omega \to \mathbb{R}$ in this subsection as

$$C_0^\infty(\Omega),$$

reserving the notation $\mathscr{D}(\Omega)$ for the case where this space is endowed with the specific topology from §13.8.1.

Alternative 1: *The C^∞_{loc}-topology.*
The space $C^\infty(\Omega)$ of all (not necessarily compactly supported) smooth functions $\Omega \to \mathbb{R}$ admits a natural Fréchet space topology for which convergence means uniform convergence of derivatives of all orders on compact subsets. This is often called C^∞_{loc}-convergence. A countable family of seminorms for the C^∞_{loc}-topology, also sometimes called the **weak** or **compact-open C^∞-topology**, is given by

$$\|\varphi\|_{m,j} := \|\varphi\|_{C^m(K_j)} \qquad \text{for integers } m \geqslant 0,\ j \geqslant 1,$$

where $K_1 \subset K_2 \subset K_3 \subset \ldots \bigcup_{j\in\mathbb{N}} K_j = \Omega$ is any exhausting sequence of compact subsets such that K_j is contained in the interior of K_{j+1} for every j. This gives the right notion of convergence because every compact set is contained in K_j for j sufficiently large, and it defines a metrizable topology since the family of seminorms is countable (see e.g. [RS80, Theorem V.5]). Continuity on $C^\infty(\Omega)$ is thus equivalent to sequential continuity, and since the notion of C^∞_{loc}-convergence can be expressed without referring to the specific choice of exhaustion $K_1 \subset K_2 \subset K_3 \subset \ldots$, the C^∞_{loc}-topology is also independent of that choice.

As a subspace of $C^\infty(\Omega)$, $C_0^\infty(\Omega)$ inherits a metrizable topology for which convergence means C^∞_{loc}-convergence. The first thing to notice, however, is that $C_0^\infty(\Omega)$ is not a Fréchet space with this topology, i.e. it is not complete, because it is not a *closed* subspace of $C^\infty(\Omega)$. In fact, by choosing any sequence of smooth cutoff functions $\rho_j \in C_0^\infty(\Omega)$ with $\rho_j|_{K_j} \equiv 1$ for every j, it is easy to check that for every $\varphi \in C^\infty(\Omega)$, $\varphi_j := \rho_j\varphi \to \varphi$ in C^∞_{loc}, thus $C_0^\infty(\Omega)$ is dense in $C^\infty(\Omega)$ with respect to the C^∞_{loc}-topology. This has an immediate consequence for the space of continuous linear functionals on $C_0^\infty(\Omega)$: any linear functional $\Lambda : C_0^\infty(\Omega) \to \mathbb{R}$ that is continuous in the C^∞_{loc}-topology must admit a continuous extension to a linear functional on $C^\infty(\Omega)$. Most distributions clearly do not have this property; since functions $\varphi \in C^\infty(\Omega)$ can grow arbitrarily large near infinity or near the boundary of Ω, even globally integrable functions $f : \Omega \to \mathbb{R}$ do not generally define continuous functionals of $\varphi \in C^\infty(\Omega)$ under the pairing $\Lambda_f(\varphi) := \int_\Omega \varphi f\, dm$. On the other hand, it is possible to give a precise characterization of the distributions for which this works.

Proposition 13.105. *A distribution $\Lambda \in \mathscr{D}'(\Omega)$ is continuous with respect to the C^∞_{loc}-topology on $C_0^\infty(\Omega)$ if and only if it has compact support.*

Proof. In one direction, this statement follows from Proposition 13.84. For the converse, continuity of $\Lambda \in \mathscr{D}'(\Omega)$ with respect to C^∞_{loc}-convergence implies since $C_0^\infty(\Omega)$ is dense in $C^\infty(\Omega)$ that Λ extends to a C^∞_{loc}-continuous linear functional on $C^\infty(\Omega)$. If $\mathrm{supp}(\Lambda)$ is not compact, then for every compact set $K \subset \Omega$, there exists a test function $\varphi \in \mathscr{D}(\Omega)$ with $\mathrm{supp}(\varphi) \cap K = \varnothing$ and $\Lambda(\varphi) \neq 0$. We can therefore find an exhausting sequence of compact subsets $K_1 \subset K_2 \subset K_3 \subset \ldots \subset \bigcup_{j\in\mathbb{N}} K_j = \Omega$ and associated test functions $\varphi_1, \varphi_2, \varphi_3, \ldots \in \mathscr{D}(\Omega)$ such that $\mathrm{supp}(\varphi_j) \subset K_j \backslash K_{j-1}$ and $\Lambda(\varphi_j) \neq 0$ for all j. For any choice of constnts $c_j \in \mathbb{R}$, the sequence $\psi_k := \sum_{j=1}^k \varphi_j \in \mathscr{D}(\Omega)$ is then C^∞_{loc}-convergent to a smooth function $\psi_\infty \in C^\infty(\Omega)$, but the constants c_j can easily be chosen to make sure that $\Lambda(\psi_k) = \sum_{j=1}^k c_j\Lambda(\varphi_j)$ diverges as $k \to \infty$, giving a contradiction. □

Alternative 2: *The C^∞-topology.*
The countable family of norms $\|\cdot\|_{C^m}$ for integers $m \geqslant 0$ determines a Fréchet space topology on the subspace

$$C_b^\infty(\Omega) := \left\{\varphi \in C^\infty(\Omega) \;\middle|\; \partial^\alpha\varphi \text{ is bounded for every multi-index } \alpha\right\} \subset C^\infty(\Omega).$$

The associated notion of *C^∞-convergence* means uniform convergence for derivatives of all orders, not just on compact subsets but globally on Ω, thus C^∞-convergence implies (but is not implied by) C^∞_{loc}-convergence, and the C^∞-topology on the subspace $C_0^\infty(\Omega)$ is strictly stronger than the C^∞_{loc}-topology. This sounds like good news for the theory of distributions, as it means that the space of C^∞-continuous linear functionals is larger than the space of compactly supported distributions considered in Proposition 13.105. But the next exercise shows that it is still not large enough to contain all locally integrable functions.

Exercise 13.106. Find a sequence $\varphi_j \in C_0^\infty(\mathbb{R})$ such that $\varphi_j \to 0$ in the C^∞-topology, but $\int_{\mathbb{R}} \varphi_j\, dm = 1$ for all j. This implies that the distribution $\Lambda_f : \mathscr{D}(\Omega) \to \mathbb{R}$ defined via the locally integrable function $f := 1$ on $\mathbb{R}$ is not continuous with respect to the C^∞-topology.

A further hint that the C^∞-topology is not an ideal choice for $\mathscr{D}(\Omega)$ arises from the observation that $C_0^\infty(\Omega)$ is not a C^∞-closed subspace of $C_b^\infty(\Omega)$; one can easily find C^∞-convergent sequences of compactly supported functions whose limits do not have compact support. It follows that every C^∞-continuous linear functional on $C_0^\infty(\Omega)$ must admit a continuous extension to a subspace of $C_b^\infty(\Omega)$ that is strictly larger than $C_0^\infty(\Omega)$. Exercise 13.106 shows that even relatively tame functions like $f \equiv 1$ on $\mathbb{R}$ need not define distributions that are extendable in this sense.

Remark 13.107. One other theoretical drawback of the C^∞-topology is worth mentioning. All the other topologies discussed in this section can be defined in more general contexts, e.g. for the space of compactly supported smooth functions on a finite-dimensional manifold, and the weak and strong C^∞-topologies can even be defined for spaces of smooth maps from one finite-dimensional manifold to another. Generalizations of this type are essential for certain fundamental perturbation results in differential topology (see e.g. [**Hir94**, Chapter 2]). The definitions in this setting become more complicated, as they necessarily involve choices of local coordinate charts, and one must then verify that the topologies defined in this way are independent of choices. For the weak C^∞-topology and its strong variants to be discussed below, this is not difficult, because while the C^m-norm of a function can certainly change if one composes the function with a smooth coordinate transformation, this change can be bounded as long as it is only being considered over a compact subset. The ordinary C^∞-topology for functions $\Omega \to \mathbb{R}$ is simpler to define, but since it involves C^m-norms over noncompact sets, it does not have such coordinate-invariant properties and thus cannot be defined in a meaningful way for functions on a noncompact manifold.

Alternative 3: *The (strong) Whitney C^∞-topology.*
To describe the Whitney C^∞-topology, one should first describe the Whitney C^m-topology for $0 \leqslant m < \infty$ on $C^\infty(\Omega)$. We give two definitions: first, it is the smallest topology containing all sets of the form

$$\mathcal{U}(\varphi, \alpha, f) := \left\{ \psi \in C^\infty(\Omega) \;\middle|\; |\partial^\alpha(\psi - \varphi)| < f \right\}$$

for arbitrary choices of $\varphi \in C^\infty(\Omega)$, multi-indices α of order at most m and continuous functions $f : \Omega \to (0, \infty)$. Equivalently, one can generate this topology with sets of the form

$$\mathcal{V}(\varphi, \{\Omega_i\}, \{k_i\}, \{\epsilon_i\}) := \left\{ \psi \in C^\infty(\Omega) \;\middle|\; \|\psi - \varphi\|_{C^{k_i}(\Omega_i)} < \epsilon_i \right\}, \tag{13.12}$$

for all positive choices of $\varphi \in C^\infty(\Omega)$, locally finite open coverings $\{\Omega_i\}_{i \in I}$ of Ω, and collections of numbers $\{k_i \in \{0, \dots, m\}\}_{i \in I}$ and $\{\epsilon_i > 0\}_{i \in I}$.

Exercise 13.108. Show that the two definitions of the Whitney C^m-topology given above are equivalent.

The **Whitney C^∞-topology** is now defined to be the smallest topology on $C^\infty(\Omega)$ that contains the Whitney C^m-topology for every $m \geqslant 0$, i.e. it is generated by the sets $\mathcal{U}(\varphi, \alpha, f)$ without any bound on the order of the multi-index α, or by $\mathcal{V}(\varphi, \{\Omega_i\}, \{k_i\}, \{\epsilon_i\})$, in which the set of integers $\{k_i \geqslant 0\}_{i \in I}$ is always required to be bounded, but no fixed bound is imposed.

It is straightforward to transform the definitions of $\mathcal{U}(\varphi, \alpha, f) \subset C^\infty(\Omega)$ and $\mathcal{V}(\varphi, \{\Omega_i\}, \{k_i\}, \{\epsilon_i\})$ into conditions of the form $\{\|\psi - \varphi\| < 1\}$ for suitable seminorms $\|\cdot\|$, thus the Whitney C^∞-topology is locally convex. One can however use the argument of Remark 13.90 to show that it is not first countable, and thus not metrizable. Here is a clear advantage of the Whitney topology in comparison with the C^∞ and C^∞_{loc}-topologies:

Proposition 13.109. *In the Whitney C^∞-topology, $C_0^\infty(\Omega)$ is a closed subspace of $C^\infty(\Omega)$.*

Proof. One needs to show that $C^\infty(\Omega) \backslash C_0^\infty(\Omega)$ is open. If $\varphi \in C^\infty(\Omega)$ does not have compact support, then there exists a continuous function $f : \Omega \to (0, \infty)$ and a sequence $x_j \in \Omega$ with no accumulation point such that $f(x_j) < |\varphi(x_j)|$ for all j. The set $\left\{ \psi \in C^\infty(\Omega) \;\middle|\; |\psi - \varphi| < f \right\}$ is then a Whitney-open neighborhood of φ consisting of functions ψ that satisfy $\psi(x_j) \neq 0$ for all j and thus never have compact support. □

A similar argument to Lemma 13.91 and Corollary 13.92 also shows:

Proposition 13.110. *A sequence in $C_0^\infty(\Omega)$ converges in the Whitney C^∞-topology if and only if it converges in $\mathscr{D}(\Omega)$.* □

One can easily show that the seminorms one uses to define the open sets $\mathcal{U}(\varphi, \alpha, f)$ or $\mathcal{V}(\varphi, \{\Omega_i\}, \{k_i\}, \{\epsilon_i\})$ are also *good* seminorms in the sense of Definition 13.88, thus the topology of $\mathscr{D}(\Omega)$ contains the Whitney C^∞-topology, and Proposition 13.110 reveals that the two topologies are evidently quite similar. In fact, the Whitney C^0-topology is already strong enough to make all functionals of the form $\Lambda_f(\varphi) = \int_\Omega \varphi f\, dm$ for $f \in L^1_{\mathrm{loc}}(\Omega)$ continuous, thus the Whitney C^∞-topology also has this clearly desirable property. The next example shows however that the topology of $\mathscr{D}(\Omega)$ is *strictly* stronger, so that the space of distributions is still strictly larger than the space of linear functionals on $C_0^\infty(\Omega)$ that are continuous in the Whitney topology. In light of Corollary 13.97, this also reveals that for the Whitney topology on $C_0^\infty(\Omega)$, sequential continuity of a linear functional does not imply continuity.

EXAMPLE 13.111. Consider the real-valued distribution $\Lambda : \mathscr{D}(\mathbb{R}) \to \mathbb{R}$ defined by

$$\Lambda(\varphi) := \sum_{k=0}^{\infty} \varphi^{(k)}(k).$$

This is well defined on any individual test function $\varphi \in \mathscr{D}(\mathbb{R})$ since only finitely many terms in the sum are nonzero, and the same is true for any convergent sequence of test functions, thus Λ is sequentially continuous and therefore continuous on $\mathscr{D}(\mathbb{R})$. But it is not continuous with respect to the Whitney C^∞-topology on $C_0^\infty(\mathbb{R})$. To see this, consider $\Lambda^{-1}((-1,1))$. If this were open in the Whitney topology, then there would need to exist a finite collection of multi-indices $\alpha_1, \dots, \alpha_N$ and continuous functions $f_1, \dots, f_N : \mathbb{R} \to (0, \infty)$ such that $\bigcap_{j=1}^N \mathcal{U}(0, \alpha_j, f_j) \subset \Lambda^{-1}((-1,1))$, meaning

$$|\partial^{\alpha_j}\varphi| < f_j \text{ for all } j = 1, \dots, N \qquad \Rightarrow \qquad |\Lambda(\varphi)| < 1.$$

But this condition constrains only finitely many derivatives of φ, thus one can always find a function that satisfies it but has $|\Lambda(\varphi)| \geqslant 1$ due to the behavior of some derivative of even higher order.

Alternative 4: *The (very) strong C^∞-topology.* A minor modification to the definition of the Whitney C^∞-topology gives rise to an even stronger topology which we shall refer to as the *strong C^∞-topology.*[37] It is generated by all sets of the form $\mathcal{V}(\varphi, \{\Omega_i\}, \{k_i\}, \{\epsilon_i\})$ as in (13.12), for arbitrary locally finite open coverings $\{\Omega_i\}_{i\in I}$, sets of nonnegative integers $\{k_i\}_{i\in I}$ and positive numbers $\{\epsilon_i\}_{i\in I}$. The crucial difference is that in our definition of the Whitney C^∞-topology, the set of integers $\{k_i\}_{i\in I}$ was always required to be bounded, and this is no longer required. Note that since the open covering $\{\Omega_i\}_{i\in I}$ is locally finite, only finitely many of the sets can intersect any given compact subset of Ω, but there still may be infinitely many sets in the covering. The result is that neighborhoods generating the strong topology are required to satisfy conditions on only finitely many derivatives over each individual compact subset, but globally on Ω, there may be conditions on derivatives of all orders.

If one only considers convergence of sequences, then there is no difference between the strong and Whitney C^∞-topologies: Proposition 13.110 admits the same proof for the strong topology and shows that it also has the same notion of convergence as $\mathscr{D}(\Omega)$. That it is nonetheless *strictly* stronger than the Whitney topology follows from Example 13.111 and the following:

EXERCISE 13.112. Show that the strong C^∞-topology on $C_0^\infty(\Omega)$ is equivalent to the topology of $\mathscr{D}(\Omega)$.

The strong C^∞-topology is thus merely a different perspective on the locally convex inductive limit topology, one that does not require talking about the Fréchet subspaces $\mathscr{D}_K(\Omega)$ with $K \subset \Omega$ compact. This approach to the topology of $\mathscr{D}(\Omega)$ is discussed in more detail in [**Hor66**, §2.12], which gives in particular an explicit family of good seminorms generating the topology.

[37] The literature is not unanimous on the terminology for these topologies: different sources may use the words "strong topology" or "Whitney topology" to refer to either of alternatives 3 and 4, and one occasionally even finds an authoritative source that fails to distinguish between them. (I am thinking especially of [**Hir94**], which defines the Whitney topology in §2.1 and the strong topology in §2.4 but states erroneously that they are equivalent.) Alternative 4 is occasionally also called the *very strong C^∞*-topology, e.g. in [**Ill03**].

Part 3: Abstract Banach spaces

The next portion of this course concerns a number of important theorems about the general properties of Banach spaces and bounded linear operators.

14. Hahn-Banach theorem

14.1. Motivation. It has already been mentioned that the following question is often nontrivial: Given a topological vector space E over the field $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$, are there any nontrivial bounded linear functionals $\Lambda : E \to \mathbb{K}$? Exercise 2.38 shows that the answer is sometimes no: concretely, for $0 < p < 1$, the space $L^p([0,1])$ of functions $f : [0,1] \to \mathbb{R}$ satisfying $\|f\|_{L^p}^p := \int_0^1 |f(t)|^p \, dt < \infty$, with topology defined via the metric $d(f,g) := \|f - g\|_{L^p}^p$, has a trivial dual space. We've also seen however that Hilbert spaces and L^p spaces for $p \geqslant 1$ always have nontrivial dual spaces. So do the C^k-spaces, e.g. integration with respect to any finite measure defines a bounded linear functional on the space of bounded C^0-functions.

A more quantitative version of the question is the following: For which spaces E does there exist for every $x \in E$ a bounded linear functional $\Lambda : E \to \mathbb{K}$ such that $\Lambda(x) \neq 0$?

The Hahn-Banach theorem will imply that this is true whenever E is a locally convex space, so in particular, it is true for all Banach spaces.

For a bit of motivation, notice that constructing interesting functionals $\Lambda \in E^*$ would be easy if we could always answer the following question: Given a subspace $W \subset E$ and a bounded linear functional $\lambda : W \to \mathbb{K}$, can λ be extended to a bounded linear functional $\Lambda : E \to \mathbb{K}$ matching λ on W? Indeed, if this "bounded extension" problem can always be solved, then for any given $x \neq 0 \in E$, one could take any nontrivial functional λ on the 1-dimensional subspace spanned by x and extend it to a bounded linear functional $\Lambda \in E^*$ satisfying $\Lambda(x) \neq 0$. We already know from Lemma 2.31 that if W is not closed, then λ admits a unique continuous extension over the closure $\overline{W} \subset E$, so we lose no generality if we assume that W in the first place is closed. But if $W \subset E$ is a proper closed subspace of infinite codimension, it is not immediately obvious whether λ admits a bounded extension from W to E. In finite dimensions, you would know what to do: You would choose an inner product, so that you can define the orthogonal complement $W^\perp \subset E$, and then define the extension $\Lambda : E \to \mathbb{K}$ uniquely by the conditions $\Lambda|_W = \lambda$ and $\Lambda|_{W^\perp} = 0$. This trick also works if E is a general Hilbert space. If E is an infinite-dimensional Banach space, then we do not have an inner product in general, but we could try doing the same trick with $W^\perp$ replaced by a more general *complementary* subspace $W' \subset E$, i.e. a subspace such that

$$E = W \oplus W',$$

so that every $x \in E$ is uniquely the sum of two vectors $w \in W$ and $w' \in W'$. On a purely algebraic level, one can always find such a complement: we can define it using a Hamel basis of the quotient space E/W. We will not go into details on this, because analytically, it is not useful: without more information about the complement $W' \subset E$ of W, there is no guarantee that the unique linear functional $\Lambda : E \to \mathbb{K}$ determined by $\Lambda|_W = \lambda$ and $\Lambda|_{W'} = 0$ will be bounded. It is not difficult to show that this works whenever the complementary subspace $W' \subset E$ is also *closed*, so if every closed subspace of E admits a complementary closed subspace, then we're in business, and the bounded extension problem can be solved. There's just one problem: Not every closed subspace of a Banach space admits a complementary closed subspace—the existence of such complements is yet another "obvious" fact from finite-dimensional linear algebra that turns out to be false in more general infinite-dimensional settings. We'll come back to this topic in §15.3. The good news is that the bounded extension problem can nonetheless be solved, even when closed complements do not exist, and this is what the Hahn-Banach theorem does.

14.2. Statement and proof. The key idea behind the Hahn-Banach theorem is convexity. Recall from §2.3 that seminorms $\|\cdot\| : E \to \mathbb{R}$ on a vector space E are convex functions, meaning that for all $x, y \in E$ and $\tau \in [0,1]$, we have

$$\|\tau x + (1-\tau)y\| \leqslant \tau\|x\| + (1-\tau)\|y\|.$$

Moreover, by Lemma 13.95, a linear functional $\Lambda : E \to \mathbb{K}$ on a locally convex space E is continuous if and only if there exists a continuous seminorm $\|\cdot\|$ on E such that $|\Lambda(x)| \leqslant \|x\|$ for all $x \in E$. Note that if $\mathbb{K} = \mathbb{R}$, then here it suffices to assume $\Lambda(x) \leqslant \|x\|$, since $-\Lambda(x) = \Lambda(-x) \leqslant \|-x\| = \|x\|$ then implies $|\Lambda(x)| \leqslant \|x\|$. I recommend keeping these observations in mind when you read the following somewhat technical statement.

Theorem 14.1 (Hahn-Banach). *Assume E is a real vector space, $p : E \to \mathbb{R}$ is a convex function, $W \subset E$ is a linear subspace and $\lambda : W \to \mathbb{R}$ is a linear map satisfying $\lambda(x) \leqslant p(x)$ for all $x \in W$. Then there exists an extension of λ to a linear functional $\Lambda : E \to \mathbb{R}$ satisfying $\Lambda|_W = \lambda$ and $\Lambda(x) \leqslant p(x)$ for all $x \in E$.*

Several interesting corollaries will be discussed in the next subsection, but here is one to start with:

Corollary 14.2. *For E a normed vector space over $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ and $W \subset E$ a linear subspace, every bounded linear functional $\lambda \in W^*$ can be extended to a bounded linear functional $\Lambda \in E^*$ such that $\Lambda|_W = \lambda$ and $\|\Lambda\| = \|\lambda\|$.*

Proof. In the case $\mathbb{K} = \mathbb{R}$, $\lambda \in W^*$ is a linear map $\lambda : W \to \mathbb{R}$ satisfying $\lambda(x) \leqslant c\|x\|$ for all $x \in W$ and some constant $c = \|\lambda\| \geqslant 0$ independent of x. Defining the convex function $p : E \to \mathbb{R} : x \mapsto c\|x\|$, the Hahn-Banach theorem implies the existence of a functional $\Lambda : E \to \mathbb{R}$ such that $\Lambda|_W = \lambda$ and $\Lambda(x) \leqslant p(x) = c\|x\|$ for all $x \in E$; replacing x with $-x$ in cases where $\Lambda(x) < 0$, this implies $|\Lambda(x)| \leqslant c\|x\|$ and thus $\|\Lambda\| = c$.

The case $\mathbb{K} = \mathbb{C}$ can be reduced to the case $\mathbb{K} = \mathbb{R}$ by considering for any complex-linear functional $\lambda : W \to \mathbb{C}$ the associated real-linear functional $\operatorname{Re}\lambda : W \to \mathbb{R}$, which satisfies $|\operatorname{Re}\lambda(x)| \leqslant |\lambda(x)| \leqslant c\|x\|$ for all $x \in W$. Applying the result in the real case then produces a functional $\Phi : E \to \mathbb{R}$ such that $\Phi|_W = \operatorname{Re}\lambda|_W$ and $\|\Phi(x)\| \leqslant c\|x\|$ for all $x \in E$, and one checks that $\Lambda(x) := \Phi(x) - i\Phi(ix)$ defines a suitable complex-linear extension of λ. □

Example 14.3. For any $x \neq 0$ in a normed vector space E, there is a unique bounded linear functional λ on the subspace spanned by x satisfying $\lambda(x) = \|x\|$ and thus $\|\lambda\| = 1$. Applying Corollary 14.2 then produces a bounded linear functional $\Lambda \in E^*$ such that $\Lambda(x) = \|x\|$ and $\|\Lambda\| = 1$. Recall from Exercise 5.3 that for most interesting measure spaces X and the Banach space $E := L^\infty(X)$, there exist functions $f \in E$ such that all nontrivial bounded linear functionals $\Lambda_g \in E^*$ represented by functions $g \in L^1(X)$ satisfy the *strict* inequality $|\Lambda_g(f)| < \|g\|_{L^1} \cdot \|f\|_{L^\infty} = \|\Lambda_g\| \cdot \|f\|_{L^\infty}$. For $\Lambda \in E^*$ and $x \in E$ as obtained above from the Hahn-Banach theorem, this strict inequality is not satisfied, so it follows that Λ is not representable by an L^1-function, i.e. the Hahn-Banach theorem fills in the gap previously left open in our proof that the Riesz representation theorem fails for $L^\infty(X)$.

The proof of the Hahn-Banach theorem is non-obvious but relatively elementary when $\dim E < \infty$, and extending to the case $\dim E = \infty$ is also straightforward with the aid of Zorn's lemma. The need to use Zorn's lemma, however, makes the theorem slightly unpopular in certain quarters of the mathematical community. In practice, for most of the spaces that one is typically interested in, there are ways of proving the important corollaries of the Hahn-Banach theorem without using the theorem itself, and thus without relying on the axiom of choice.

Remark 14.4. While the standard applications of the Hahn-Banach theorem are analytical results, the theorem itself is really more a result of algebra than analysis: neither the statement nor the proof involves any notions of convergence or completeness.

Proof of Theorem 14.1. Step 1 is the elementary portion of the proof: the goal is namely to show that the theorem holds in the special case where $W \subset E$ has codimension 1. Equivalently, we assume there exists a point $y \in E \backslash W$ and extend $\lambda : W \to \mathbb{R}$ to a suitable functional Λ defined on the larger subspace $W \oplus \mathbb{R}y \subset E$. Such an extension will be uniquely determined by the number $a := \Lambda(y) \in \mathbb{R}$, so the question becomes: for what values of $a \in \mathbb{R}$ does Λ satisfy the inequality

$$\Lambda(x + cy) \leqslant p(x + cy) \qquad \text{for all} \qquad x \in W,\ c \in \mathbb{R}? \tag{14.1}$$

Assuming $\Lambda(x) = \lambda(x)$ for $x \in W$, the relation already holds by hypothesis when $c = 0$. For $c > 0$, we will have $\Lambda(x + cy) = \Lambda(x) + ca \leqslant p(x + cy)$ if and only if the relation

$$a \leqslant \frac{1}{c}\left[p(x + cy) - \Lambda(x)\right] \qquad \text{for all} \qquad x \in W,\ c > 0 \tag{14.2}$$

is satisfied, and similarly, $\Lambda(x - cy) = \Lambda(x) - ca \leqslant p(x - cy)$ holds if and only if

$$a \geqslant \frac{1}{c}\left[\Lambda(x) - p(x - cy)\right] \qquad \text{for all} \qquad x \in W,\ c > 0. \tag{14.3}$$

It is possible to find a number $a \in \mathbb{R}$ satisfying both inequalities if and only if the right hand side of (14.2) is at least as large as the right hand side of (14.3), and we claim that the latter is always true: in fact, for all $x, x' \in W$ and $\alpha, \beta > 0$, we claim

$$\frac{1}{\alpha}\left[\Lambda(x) - p(x - \alpha y)\right] \leqslant \frac{1}{\beta}\left[p(x' + \beta y) - \Lambda(x')\right]. \tag{14.4}$$

Indeed, this is equivalent to the relation

$$\Lambda(\beta x + \alpha x') \leqslant \beta p(x - \alpha y) + \alpha p(x' + \beta y),$$

and the latter holds due to the following sequence of relations, the last of which follows from the convexity of p:

$$\begin{aligned}
\Lambda(\beta x + \alpha x') &= (\alpha + \beta)\Lambda\left(\frac{\beta}{\alpha + \beta}x + \frac{\alpha}{\alpha + \beta}x'\right) \leqslant (\alpha + \beta)p\left(\frac{\beta}{\alpha + \beta}x + \frac{\alpha}{\alpha + \beta}x'\right) \\
&= (\alpha + \beta)p\left(\frac{\beta}{\alpha + \beta}(x - \alpha y) + \frac{\alpha}{\alpha + \beta}(x' + \beta y)\right) \\
&\leqslant (\alpha + \beta)\left[\frac{\beta}{\alpha + \beta}p(x - \alpha y) + \frac{\alpha}{\alpha + \beta}p(x' + \beta y)\right] = \beta p(x - \alpha y) + \alpha p(x' + \beta y).
\end{aligned}$$

This proves the claim and thus completes Step 1.

Step 2 is to apply Zorn's lemma (see §2.7). Define a partially ordered set $(S, \prec)$, where S consists of all pairs (W', Λ') such that W' is a linear subspace of E containing W and $\Lambda' : W' \to \mathbb{R}$ is a linear functional satisfying $\Lambda'|_W = \lambda$ and $\Lambda' \leqslant p$, and

$$(X_1, \Lambda_1) \prec (X_2, \Lambda_2) \qquad \Leftrightarrow \qquad X_1 \subset X_2 \text{ and } \Lambda_2|_{X_1} = \Lambda_1.$$

The set S is nonempty since $(W, \lambda) \in S$, and if $S_0 \subset S$ is a totally ordered subset, then one finds an upper bound $(W', \Lambda') \in S$ for S_0 by setting

$$W' := \bigcup_{(W_0, \Lambda_0) \in S_0} W_0,$$

so that Λ' becomes the unique extension of each Λ_0. Zorn's lemma then provides a maximal element $(W_\infty, \Lambda_\infty) \in S$. If $W_\infty \neq E$, then Step 1 can be applied to extend Λ_∞ suitably to an even larger subspace than W_∞, contradicting the condition that $(W_\infty, \Lambda_\infty)$ is maximal. We conclude that Λ_∞ is the desired extension of λ to E. □

14.3. Applications. We've already seen the most popular application of the Hahn-Banach theorem, namely Corollary 14.2 on the ability to extend bounded linear functionals from subspaces without changing their norms. A few other applications have been mentioned in previous sections, and we can now establish them rigorously. For instance, the fact that locally convex spaces (and Banach spaces in particular) *always* have nontrivial dual spaces:

THEOREM 14.5. *For any locally convex space X and any $x \neq 0 \in X$, there exists a continuous linear functional $\Lambda \in X^*$ such that $\Lambda(x) \neq 0$.*

PROOF. Given $x \neq 0 \in X$, choose any continuous seminorm $\|\cdot\|$ on X with $\|x\| \neq 0$, let $W \subset X$ denote the 1-dimensional subspace spanned by x, and choose any linear functional λ on W such that $|\lambda(x)| = \|x\|$. If X is a real vector space, the Hahn-Banach theorem extends λ to linear functional $\Lambda : X \to \mathbb{R}$ satisfying $\Lambda(x) \leqslant \|x\|$ and therefore also $|\Lambda(x)| \leqslant \|x\|$ for all $x \in X$; by Lemma 13.95, the latter implies that Λ is continuous. If X is a complex vector space, one can instead apply the Hahn-Banach theorem to $\operatorname{Re} \lambda : W \to \mathbb{R}$ and produce from the resulting real-linear extension a complex-linear extension as in the proof of Corollary 14.2. □

The following consequence of Theorem 14.5 looks obvious at first, but it is both nontrivial and useful. In topological language, a similar argument shows that the weak topology on a locally convex space is Hausdorff.

COROLLARY 14.6. *In any locally convex space X, if $x_n \in X$ is a weakly convergent sequence with $x_n \rightharpoonup x$ and $x_n \rightharpoonup y$, then $x = y$.*

PROOF. Assuming $x_n \rightharpoonup x$ and $y \neq x$, Theorem 14.5 provides a functional $\Lambda \in X^*$ with $\Lambda(x - y) \neq 0$, and thus $\Lambda(x) \neq \Lambda(y)$. Weak convergence then implies $\Lambda(x_n) \to \Lambda(x)$, and it follows that $\Lambda(x_n)$ cannot converge to $\Lambda(y)$, thus x_n is not weakly convergent to y. □

The next result was stated previously as Corollary 5.7, and it follows easily from Corollary 14.2 and Example 14.3.

THEOREM 14.7. *For every Banach space E, the canonical map $\Phi : E \to E^{**}$ is an injective isometry, i.e. it satisfies $\|\Phi(x)\| = \|x\|$ for every $x \in E$.* □

We can now also extend to all Banach spaces a result that, for Hilbert spaces, follows from Corollary 3.14 and the Riesz representation theorem:

THEOREM 14.8. *For any Banach space E, a linear subspace $W \subset E$ is dense if and only if the subspace $\{\Lambda \in E^* \mid \Lambda|_W = 0\} \subset E^*$ is trivial.*

PROOF. If $W \subset E$ is dense and $\Lambda \in E^*$ vanishes on W, then it must also vanish on E due to the uniqueness of continuous linear extensions from dense subspaces (Lemma 2.31).

Conversely, suppose W is not dense, so its closure $\overline{W} \subset E$ is a proper closed subspace. For any $y \in E \backslash \overline{W}$, define $\lambda : \overline{W} \oplus \mathbb{K} y \to \mathbb{K}$ as the unique linear functional that vanishes on $\overline{W}$ and satisfies $\lambda(y) = 1$.

We claim that the distance between y and $\overline{W}$ is positive, i.e. there exists a number $\delta > 0$ such that $\|v - y\| \geqslant \delta$ for all $v \in \overline{W}$. Indeed, there must otherwise exist a sequence $v_k \in \overline{W}$ such that $\|v_k - y\| \to 0$, meaning $v_k \to y$, which would imply $y \in \overline{W}$ since $\overline{W}$ is closed, and is thus a contradiction.

Next, we claim that the functional $\lambda : \overline{W} \oplus \mathbb{K}y \to \mathbb{K}$ is bounded. Indeed, since $\operatorname{dist}(\overline{W}, y) =: \delta > 0$, we have for every $c \neq 0 \in \mathbb{K}$ and $v \in \overline{W}$,

$$\|v + cy\| = |c| \cdot \left\| y - \left(-\frac{v}{c} \right) \right\| \geqslant \delta |c|,$$

and thus

$$|\lambda(v + cy)| = |c| \leqslant \frac{1}{\delta} \|v + cy\|,$$

proving the claim.

By Corollary 14.2, λ now extends to a bounded linear functional $\Lambda : E \to \mathbb{K}$, which vanishes on $W \subset E$ but not at $y \in E$. □

Exercise 14.9. Use uniform convexity to give a second proof of Theorem 14.8 for the special case $E := L^p(X)$ with $1 < p < \infty$, without appealing to the Hahn-Banach theorem.
Hint: You may find some useful inspiration in the proofs of Theorem 3.12 and Corollary 3.14.

Since we now know that the dual of L^∞ is larger than L^1, implying in particular that L^1 is not reflexive, it follows from the next result that L^∞ is also not reflexive:

Theorem 14.10. *A Banach space is reflexive if and only if its dual space is reflexive.*

Proof. Assume E is a reflexive Banach space, so the canonical injective isometry $\Phi : E \to E^{**}$ is surjective. Since it is an isometry, its inverse is then also an isometry and thus a bounded linear operator, and one easily checks that the transpose of that inverse

$$E^* \xrightarrow{(\Phi^{-1})^\top} E^{***}$$

is the canonical injective isometry from E^* to $(E^*)^{**}$, which is therefore also invertible, proving that E^* is also reflexive.

The converse requires the Hahn-Banach theorem; more precisely, we shall deduce it from Theorem 14.8. Suppose namely that E is not reflexive, so $\Phi : E \to E^{**}$ is not surjective. Since Φ is an isometry, Cauchy sequences in the subspace $\Phi(E) \subset E^{**}$ correspond to Cauchy sequences in E and therefore converge, proving that the image $\Phi(E) \subset E^{**}$ is a closed subspace, so in particular, it is not dense in E^{**}. It then follows from Theorem 14.8 that there exists a nontrivial bounded linear functional $\Psi : E^{**} \to \mathbb{K}$ that vanishes on $\Phi(E)$. One easily checks however that for any $\lambda \in E^*$, any functional in the image of the canonical injective isometry $E^* \to E^{***}$ that vanishes on $\Phi(E)$ must be trivial, proving that Ψ is not in the image of $E^* \to E^{***}$, so E^* is not reflexive. □

15. Sums, quotients, and complements

One of the most fundamental results about Hilbert spaces is the theorem on orthogonal complements: For every Hilbert space $\mathcal{H}$ and closed subspace $W \subset \mathcal{H}$, one has $\mathcal{H} = W \oplus W^\perp$, thus presenting $\mathcal{H}$ as the direct sum of two closed subspaces. Recall that for a vector space Z and two linear subspaces $X, Y \subset Z$, the algebraic notation

$$Z = X \oplus Y$$

means

$$X + Y = Z \qquad \text{and} \qquad X \cap Y = \{0\},$$

or in other words, every $x \in X$ can be presented as $x = v + w$ for unique elements $v \in V$ and $w \in W$. The absence of an inner product makes the question of splitting Banach spaces into sums of closed subspaces considerably subtler, but we are now in a position to prove a few results about this.

15.1. Direct sums of Banach spaces. We start by defining a second meaning for the symbol "$X \oplus Y$", meant for a slightly different context. The **direct sum** of two Banach spaces X and Y, denoted by $X \oplus Y$, is the Banach space

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

equipped with the norm

$$\|(x, y)\| := \|x\| + \|y\|. \tag{15.1}$$

It is an easy exercise to check that this definition satisfies the axioms of a norm on $X \oplus Y$, and that it is complete whenever the respective norms on X and Y are complete, thus making $X \oplus Y$ into a Banach space.

Exercise 15.1. Show that

$$\|(x, y)\|_2 := \sqrt{\|x\|^2 + \|y\|^2}$$

also defines a norm on $X \oplus Y$, and it is equivalent to the norm in (15.1).

Remark 15.2. If Z is a Banach space and $X, Y \subset Z$ are closed subspaces, then X and Y are also naturally Banach spaces, giving rise to a Banach space $X \oplus Y$ as defined above, and the definition of the norm on $X \oplus Y$ makes

$$X \oplus Y \to Z : (x, y) \mapsto x + y$$

into a bounded linear map. The splitting condition $Z = X \oplus Y$ then holds if and only if this bounded linear map is invertible.

15.2. Quotients by closed subspaces. Given a normed vector space X and linear subspace $W \subset X$, we endow the **quotient space**

$$X/W = \{[x] \mid x \in X, \text{ where } [x] = [y] \text{ if and only if } x - y \in W\}$$

with the seminorm

$$\|[x]\| := \inf_{w \in W} \|x + w\|.$$

It is straightforward to check that $\|\cdot\|$ as defined on X/W satisfies the axioms of a seminorm.

Proposition 15.3. *If X is a Banach space and $W \subset X$ is a closed subspace, then the induced seminorm on X/W is a norm and makes X/W into a Banach space.*

Proof. If $\|[x]\| = 0$, then there exists a sequence $w_n \in W$ such that $\|x + w_n\| \to 0$, meaning $w_n \to -x$. Since W is closed, this implies $-x \in W$ and thus $x \in W$, so $[x] = 0 \in X/W$, proving that the seminorm on X/W is a norm.

To prove completeness, we recall the following fact from first-year analysis: A normed vector space is complete if and only if every absolutely convergent series in that space converges. To make use of this, suppose $[x_n] \in X/W$ is a sequence such that $\sum_n \|[x_n]\| < \infty$. We can then choose $w_n \in W$ for each n so that $\|x_n + w_n\| \leqslant 2\|[x_n]\|$, in which case $\sum_n \|x_n + w_n\| < \infty$, and since X is complete, it follows that $\sum_n (x_n + w_n)$ converges to some element $y \in X$. We then have

$$\left\| \sum_{n=1}^{N} [x_n] - [y] \right\| = \left\| \left[\sum_{n=1}^{N} (x_n + w_n) - y \right] \right\| \leqslant \left\| \sum_{n=1}^{N} (x_n + w_n) - y \right\| \to 0$$

as $N \to \infty$, proving $\sum_{n=1}^{\infty} [x_n] = [y] \in X/W$. □

We observe that the definition of the norm on X/W makes the quotient projection

$$X \xrightarrow{\pi} X/W : x \mapsto [x]$$

into a bounded linear operator.

15.3. Complemented subspaces. In a Banach space X, a closed subspace $W \subset X$ is said to be **complemented** if there exists another closed subspace $W' \subset X$ such that $X = W \oplus W'$.

Example 15.4. If X is a Hilbert space, then Theorem 3.12 implies that every closed subspace $W \subset X$ is complemented, since its orthogonal complement $W^\perp \subset X$ is another closed subspace and $X = W \oplus W^\perp$.

Example 15.5. If X is a Banach space and $W \subset X$ is a closed subspace of finite codimension

$$\operatorname{codim}(W) := \dim(X/W) < \infty,$$

then W is complemented. Indeed, a closed complement of W can be defined as the span of any collection of elements $w_1, \ldots, w_n \in X$ whose equivalence classes $[w_1], \ldots, [w_n]$ form a basis of X/W.

As was discussed in §14.1, the Hahn-Banach theorem would be unnecessary if every closed subspace of a Banach space were complemented. Counterexamples are not so easy to find, but they exist: For instance, in the Banach space $\ell^\infty := L^\infty(\mathbb{N}, \nu)$ of bounded sequences $\{x_n\}_{n\in\mathbb{N}}$, the closed subspace consisting of sequences $\{x_n\}$ with $\lim_{n\to\infty} x_n = 0$ is known to admit no closed complement. A short proof of this fact (using Corollary 15.9 below) appears in [**Whi66**].

Aside from their relevance to the bounded extension problem, the importance of complemented subspaces has much to do with the continuity of linear projection maps.

Definition 15.6. For a vector space X with two subspaces $W, W' \subset X$ such that $X = W \oplus W'$, the **projection to W along W'** is the linear map

$$X \xrightarrow{\Pi} X : w + w' \mapsto w \qquad \text{for } w \in W,\ w' \in W';$$

equivalently, it is the unique linear map $\Pi : X \to X$ such that $\operatorname{im}(\Pi) = W$, $\ker(\Pi) = W'$, and $\Pi^2 = \Pi$. More generally, a linear map $\Pi : X \to X$ is called a **projection** whenever $\Pi^2 = \Pi$.

Remark 15.7. It is often useful to observe that if $X = W \oplus W'$ and Π is the projection to W along W', then $\mathbb{1} - \Pi$ is the projection to W' along W.

Theorem 15.8. *For a Banach space X with linear subspaces $W, W' \subset X$ such that $X = W \oplus W'$, the projection $\Pi : X \to X$ to W along W' is continuous if and only if both W and W' are closed.*

Corollary 15.9. *A closed subspace $W \subset X$ of a Banach space is complemented if and only if there exists a continuous linear projection map $\Pi : X \to X$ with $\operatorname{im}(\Pi) = W$.*

Proof. If a continuous projection $\Pi : X \to X$ with image W exists, then $W' := \ker(\Pi)$ is a closed complement of W. Conversely, if W has a closed complement W', then Theorem 15.8 implies that the projection to W along W' is continuous. □

We are doing things slightly out of order, because the complete proof of Theorem 15.8 depends on the following important result, to be proved in §16 below:

Theorem (Inverse mapping theorem). *Every bounded linear bijection between two Banach spaces has a bounded inverse.*

You may view the following proof in part as motivation for the effort that will be required in proving the inverse mapping theorem.

Proof of Theorem 15.8. If the projection $\Pi : X \to X$ is continuous, then $W' = \ker(\Pi)$ is closed, and the complementary projection $\mathbb{1} - \Pi : X \to X$ is also continuous, implying that $W = \ker(\mathbb{1} - \Pi)$ is also closed.

Conversely, suppose $W, W' \subset X$ are both closed, in which case both are naturally Banach spaces, and we can consider their direct sum $W \oplus W'$ as a Banach space, for which the natural bounded linear map

$$W \oplus W' \xrightarrow{\Phi} X : (w, w') \mapsto w + w'$$

is invertible, as mentioned in Remark 15.2. The inverse of Φ can be written in terms of the projection Π as

$$\Phi^{-1} = (\Pi, \mathbb{1} - \Pi).$$

Since $W \oplus W'$ and X are both Banach spaces and Φ is bounded, the inverse mapping theorem implies that Φ^{-1} is also bounded, and so therefore are Π and $\mathbb{1} - \Pi$. □

With this characterization of complemented subspaces in hand, we conclude by extracting from the Hahn-Banach theorem a result dual to Example 15.5:

Theorem 15.10. *In a Banach space X, every finite-dimensional linear subspace $W \subset X$ is complemented.*

Proof. Choose a basis $e_1, \dots, e_n \in W$; using Corollary 14.2, we can then associate to this basis a set of bounded linear functionals $\Lambda_1, \dots, \Lambda_n \in X^*$ such that

$$\Lambda_i(e_j) = \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j. \end{cases}$$

A bounded linear projection map $\Pi : X \to X$ with $\operatorname{im}(\Pi) = W$ is then defined by $\Pi(x) := \sum_{i=1}^n \Lambda_i(x) e_i$, and the result now follows from Corollary 15.9. □

Exercise 15.11. For a Banach space X with two subspaces $W, W' \subset X$ such that $X = W \oplus W'$ and W is closed, let $\pi : X \to X/W$ denote the quotient projection, so that $\pi|_{W'} : W' \to X/W$ is a bijection. Show that the inverse of this bijection is continuous if and only if the subspace $W' \subset X$ is also closed.

Note: For one direction of the implication, you may use the inverse mapping theorem.

16. The Baire category theorem and its consequences

There was some unfinished business hidden in the proof of Theorem 15.8: we cited the inverse mapping theorem, which states that invertible bounded linear operators between Banach spaces *always* have bounded inverses. Among other things, we will derive that result in this section from the closely related *open* mapping theorem, whose proof requires the Baire category theorem.

16.1. Baire category theorem. One way of motivating the Baire category theorem is via the following question: In a metric space M, what kinds of subsets should be considered "insignificant," and what should it mean for a property to hold "almost everywhere"?

These questions have standard answers whenever M is endowed with a measure: the insignificant sets are those which have measure zero, and a property holds almost everywhere if and only if it holds on the complement of such a set. We will be interested however in cases where M is an infinite-dimensional Banach or Fréchet space, on which there is no standard choice of measure. The idea is then to focus on the topological properties of sets in M, and a special role is played by subsets that are both *open* and *dense*; intuitively, we would like to say that a property is true "almost everywhere" if (but not only if) it holds over an open and dense set. Note that in measure theory, this is not quite right, i.e. an open and dense subset of $\mathbb{R}^n$ need not have full measure (see Example 16.5 below); we will see nonetheless that for qualitative purposes, open and dense sets can reasonably play the role of sets with "full measure" in settings where there is a metric but no natural measure. In keeping with the usual principles of measure theory, we would also like to say that a countable union of insignificant sets is always insignificant, whereas no set that contains an open ball of any size should be considered insignificant. The Baire category theorem tells us, essentially, that in any *complete* metric space, the items on this wish-list for the notion of an "insignificant" set do not contradict each other.

DEFINITION 16.1. In a metric space M, a set $S \subset M$ is called **nowhere dense** if its closure $\bar{S}$ has empty interior, i.e. $\bar{S}$ does not contain any open ball. Equivalently, S is nowhere dense if it is contained in the complement of an open and dense subset of M.

A subset of M is called **meager** (or **of the first category**) if it is a countable union of nowhere dense sets, and **comeager** (or **residual**) if it is the complement of a meager set. Equivalently, $S \subset M$ is comeager if it contains a countable intersection of open and dense sets.

THEOREM 16.2 (Baire category theorem). *A nonempty complete metric space is never a union of countably many nowhere dense subsets.*

REMARK 16.3. Closed balls in complete metric spaces are also complete metric spaces, and open balls always contain smaller closed balls, thus an equivalent formulation of the theorem would be that a meager subset of a nonempty complete metric space can never contain an open or closed ball of any size. Rephrasing in terms of the complements of these sets, another equivalent statement is that in a complete metric space, *comeager* subsets are always dense. The latter makes it reasonable to define "almost everywhere" to mean "on a comeager subset". By convention, the term "almost everywhere" is reserved for measure-theoretic contexts, but it is common to use the word "generic" for the concept under discussion here, i.e. a property dependent on points x in a complete metric space M is said to be true **generically** if it holds for all x in some comeager subset of M.

PROOF OF THEOREM 16.2. Assuming M is a complete metric space and $M_1, M_2, M_3, \ldots \subset M$ is a sequence of nowhere dense subsets, we shall construct a Cauchy sequence $x_n \in M$ whose limit cannot belong to $\bigcup_{n\in\mathbb{N}} M_n$. Since $\overline{M}_1$ does not contain any open ball, $M\backslash\overline{M}_1$ is nonempty, and we can therefore choose $x_1 \in M\backslash\overline{M}_1$ and $\epsilon_1 > 0$ such that $B_{\epsilon_1}(x_1) \cap \overline{M}_1 = \varnothing$, where $\epsilon_1 < 1$ without loss of generality. We can then similarly choose $x_2 \in B_{\epsilon_1}(x_1)\backslash\overline{M}_2$ and $\epsilon_2 \in (0, 1/2)$ such that $\overline{B_{\epsilon_2}(x_2)} \subset B_{\epsilon_1}(x_1)$ but $B_{\epsilon_2}(x_2) \cap \overline{M}_2 = \varnothing$. Continuing inductively in this way, one finds sequences $x_n \in M$ and $\epsilon_n > 0$ such that for each $n \geqslant 2$,

$$\epsilon_n < \frac{1}{2^n}, \qquad \overline{B_{\epsilon_n}(x_n)} \subset B_{\epsilon_{n-1}}(x_{n-1}), \quad \text{and} \quad B_{\epsilon_n}(x_n) \cap \overline{M}_n = \varnothing.$$

This makes x_n a Cauchy sequence, so $x_n \to x$ for some $x \in \bigcap_{n\in\mathbb{N}} \overline{B_{\epsilon_n}(x_n)}$, but the latter implies $x \notin M_n$ for every $n \in \mathbb{N}$. □

REMARK 16.4. For the types of sets that one typically encounters in nature, those which have measure zero are often also meager, and vice versa: every countable subset of $\mathbb{R}^n$, for example, is both. Certain standard "genericity" theorems are equally valid in terms of either concept, e.g. Sard's theorem implies that the set of critical values of a smooth function $\mathbb{R}^n \to \mathbb{R}^m$ has measure zero and is also meager. The next two examples show, however, that neither of these conditions actually implies the other in general.

EXAMPLE 16.5. Since $\mathbb{Q} \subset \mathbb{R}$ is countable, one can choose any enumeration $\{x_1, x_2, x_3, \ldots\} = \mathbb{Q}$ of the rational numbers and, for any $\epsilon > 0$, construct an open subset

$$S_\epsilon := \bigcup_{n\in\mathbb{N}} \left(x_n - \frac{\epsilon}{2^{n+1}}, x_n + \frac{\epsilon}{2^{n+1}} \right) \subset \mathbb{R},$$

which is also dense since $\mathbb{Q} \subset \mathbb{R}$ is dense, and its measure satisfies $m(S_\epsilon) \leqslant \sum_{n=1}^\infty \frac{\epsilon}{2^n} = \epsilon$. It follows that $\mathbb{R}\backslash S_\epsilon$ is nowhere dense but has nonzero (and in fact infinite) measure.

EXAMPLE 16.6. Using the sets $S_\epsilon \subset \mathbb{R}$ from Example 16.5, one obtains a nested sequence of open and dense subsets

$$\mathbb{R} \supset S_1 \supset S_{1/2} \supset S_{1/3} \supset \ldots \supset \bigcap_{N\in\mathbb{N}} S_{1/N} =: S,$$

whose intersection $S \subset \mathbb{R}$ is comeager and has measure $m(S) = \lim_{N\to\infty} m(S_{1/N}) \leqslant \lim_{N\to\infty} \frac{1}{N} = 0$. This proves the intuitively surprising fact that $\mathbb{R}$ is the union of a meager set with a disjoint set of measure zero.

16.2. Open and inverse mapping theorems. Here is a first application of the Baire category theorem.

THEOREM 16.7 (Open mapping theorem). *If X and Y are Banach spaces and $T \in \mathscr{L}(X,Y)$ is a surjective bounded linear operator, then T is also an* ***open map****, i.e. for every open subset $\mathcal{U} \subset X$, the image $T(\mathcal{U}) \subset Y$ is also open.*

Before proving this, let us fill in the gap that was left open in §15.3:

COROLLARY 16.8 (Inverse mapping theorem). *If X and Y are Banach spaces and $T \in \mathscr{L}(X,Y)$ is a bijective bounded linear operator, then $T^{-1} : Y \to X$ is also bounded.*

PROOF. Recall that a map between two topological spaces is continuous if and only if the preimage of every open set is open. In the case at hand, T^{-1} is thus continuous if and only if $T(\mathcal{U}) \subset Y$ is open for every open set $\mathcal{U} \subset X$, and this follows from the open mapping theorem. □

REMARK 16.9. A bounded linear operator $A : X \to Y$ between two Banach spaces is sometimes called a **Banach space isomorphism** if it is invertible. Corollary 16.8 justifies this terminology by guaranteeing that not only A but also its inverse A^{-1} is continuous.[38]

We can now remove an unnecessary hypothesis from a result stated at the very beginning of this course, Corollary 1.12:

COROLLARY 16.10. *For two Banach spaces X, Y, the invertible bounded linear maps form an open subset of $\mathscr{L}(X,Y)$.* □

PROOF OF THEOREM 16.7. As a preliminary remark, notice that it will suffice to prove that for any surjective bounded linear operator $T : X \to Y$ between two Banach spaces,

$$T(B_r(0)) \supset B_R(0) \qquad \text{for some } r, R > 0.$$

Indeed, for any $x_0 \in X$ and $\epsilon > 0$, one then has

$$T(B_\epsilon(x_0)) = Tx_0 + \frac{\epsilon}{r} T(B_r(0)) \supset Tx_0 + B_{\epsilon R/r}(0) = B_{\epsilon R/r}(Tx_0),$$

proving that the image of any open ball contains an open ball about the image of its center, hence the image of any open set is open.

[38] If you like category-theoretic terminology, the proper way to say this is that A^{-1} is a valid *morphism* in the category of Banach spaces with bounded linear maps, and the correct notion of *isomorphisms* in that category is therefore the invertible bounded linear maps.

Claim 1: For some $r > 0$, $\overline{T(B_r(0))}$ has nonempty interior.
Indeed, since T is surjective,

$$Y = \bigcup_{n=1}^{\infty} T(B_n(0))$$

implies via the Baire category theorem that not all of the sets $T(B_n(0))$ are nowhere dense, so for some $n \in \mathbb{N}$, $\overline{T(B_n(0))}$ contains a ball $B_\epsilon(y) \subset Y$.

Claim 2: $\overline{T(B_1(0))}$ contains a ball about $0 \in Y$.
To see this, we start with the ball $B_\epsilon(y) \subset \overline{T(B_n(0))}$ from Claim 1, which implies that for every $y' \in B_\epsilon(0)$, both y and $y + y' \in B_\epsilon(y)$ are limits of sequences of points in the image of $B_n(0)$, i.e. there exist sequences $x_k, x_k' \in B_n(0)$ such that

$$Tx_k \to y \qquad \text{and} \qquad Tx_k' \to y + y'.$$

It follows that $T(x_k' - x_k) \to y'$, and since $\|x_k - x_k'\| \leqslant \|x_k\| + \|x_k'\| < 2n$, this proves

$$B_\epsilon(0) \subset \overline{T(B_{2n}(0))},$$

and therefore also $B_{\epsilon/2n}(0) \subset \overline{T(B_1(0))}$. Let us henceforth relabel $\epsilon/2n$ as ϵ, so that without loss of generality,

$$B_\epsilon(0) \subset \overline{T(B_1(0))}.$$

Claim 3: $\overline{T(B_1(0))} \subset T(B_2(0))$.
Note that the proof will then be complete, because in combination with Claim 2, this implies that the image under T of $B_2(0)$ contains $B_\epsilon(0)$. To establish the claim, suppose $y \in \overline{T(B_1(0))}$ is given, so there exist points in the image of $B_1(0)$ that are arbitrarily close to y, and in particular, we can choose $x_1 \in B_1(0)$ such that

$$y - Tx_1 \in B_{\epsilon/2}(0) \subset \overline{T(B_{1/2}(0))}.$$

Applying the same reasoning to the point $y - Tx_1$, we can then choose a point $x_2 \in B_{1/2}(0)$ such that

$$y - Tx_1 - Tx_2 \in B_{\epsilon/4}(0) \subset \overline{T(B_{1/4}(0))}.$$

Continuing inductively, this produces a sequence

$$x_n \in B_{1/2^{n-1}}(0) \qquad \text{such that} \qquad y - Tx_1 - Tx_2 - \ldots - Tx_n \in B_{\epsilon/2^n}(0).$$

Since $\|x_n\| < 1/2^{n-1}$, we have $\sum_{n=1}^\infty \|x_n\| < \sum_{n=1}^\infty \frac{1}{2^{n-1}} = 2 < \infty$ and can thus appeal to the fact that absolutely convergent series in Banach spaces converge: we obtain a limit

$$x := \sum_{n=1}^{\infty} x_n \in B_2(0),$$

which satisfies $Tx = y$. □

16.3. Closed graph theorem. The **graph** of a map $A : X \to Y$ is defined to be the set

$$\Gamma_A := \big\{(x, Ax) \in X \times Y \;\big|\; x \in X\big\} \subset X \times Y.$$

If A is a linear map, then its graph is a linear subspace of $X \oplus Y$, and we recall that the latter is naturally a Banach space if X and Y are both Banach spaces. The next result will provide some useful inspiration when we study the spectral theory of *unbounded* linear operators at the end of this course.

THEOREM 16.11 (Closed graph theorem). *For two Banach spaces X, Y, a linear map $A : X \to Y$ is bounded if and only if its graph Γ_A is a closed subspace of $X \oplus Y$.*

Proof. Suppose A is continuous and $(x_n, Ax_n) \in \Gamma_A$ is a convergent sequence of points in its graph, with $(x_n, Ax_n) \to (x, y) \in X \oplus Y$. This means $x_n \to x$ and $Ax_n \to y$, so the continuity of A and the uniqueness of limits of sequences in metric spaces implies $Ax = y$, hence $(x, y) = (x, Ax) \in \Gamma_A$, proving that Γ_A is closed.

For the converse, suppose $A : X \to Y$ is a linear map for which $\Gamma_A \subset X \oplus Y$ is a closed subspace, thus making Γ_A into a Banach space with norm

$$\|(x, Ax)\| = \|x\| + \|Ax\|.$$

The projection map $\pi_1 : \Gamma_A \to X : (x, Ax) \mapsto x$ is then a linear bijection and satisfies

$$\|\pi_1(x, Ax)\| = \|x\| \leqslant \|x\| + \|Ax\| = \|(x, Ax)\|,$$

thus it is bounded, and the inverse mapping theorem then implies that its inverse $\pi_1^{-1} : X \to \Gamma_A$ is also bounded. The other projection map $\pi_2 : \Gamma_A \to Y : (x, Ax) \mapsto Ax$ is similarly bounded, and it follows that

$$A = \pi_2 \circ \pi_1^{-1}$$

is also bounded. □

We deduce from the closed graph theorem that self-adjoint linear operators on Hilbert spaces are *automatically* bounded.[39]

Corollary 16.12 (Hellinger-Toeplitz theorem). *For a Hilbert space $\mathcal{H}$, every linear operator $A : \mathcal{H} \to \mathcal{H}$ satisfying*

$$\langle x, Ay \rangle = \langle Ax, y \rangle \qquad \textit{for all} \qquad x, y \in \mathcal{H}$$

is bounded.

Proof. Suppose $(x_n, Ax_n) \in \Gamma_A$ is a convergent sequence of points in the graph of A, with $(x_n, Ax_n) \to (x, y) \in \mathcal{H} \times \mathcal{H}$, so $x_n \to x$ and $Ax_n \to y$. Then for any $z \in \mathcal{H}$, we have

$$\langle x_n, Az \rangle \to \langle x, Az \rangle = \langle Ax, z \rangle,$$

but the left hand side is also $\langle Ax_n, z \rangle$, which converges to $\langle y, z \rangle$, implying

$$\langle y, z \rangle = \langle Ax, z \rangle \qquad \text{for all} \qquad z \in \mathcal{H}.$$

This proves $Ax = y$, so $(x, y) = (x, Ax) \in \Gamma_A$, and the graph is therefore closed. □

16.4. Principle of uniform boundedness. As another application of the Baire category theorem, we can now answer this:

Question 16.13. *Can a weakly convergent sequence be unbounded?*

An affirmative answer would seem surprising, but it does not follow easily from the definition of weak convergence. What we will show, in fact, is that if $x_n \in X$ is an unbounded sequence in a Banach space, then $\Lambda(x_n)$ is also unbounded and therefore cannot converge for *generic* bounded linear functionals $\Lambda \in X^*$ (see Remark 16.3 on the use of the term "generic"). The Baire category theorem then guarantees that this holds in particular for all Λ in a dense subset of X^*.

In §16.5 below, we shall also address the following more concrete question, whose answer may seem legitimately surprising:

Question 16.14. *For a continuous periodic function $f \in C^0(S^1) \subset L^2(S^1)$ of one variable, does the L^2-convergent Fourier series $\sum_{k \in \mathbb{Z}} e^{2\pi i k x} \hat{f}_k$ also converge* pointwise *to f?*

[39]There is nonetheless a well-developed spectral theory for "unbounded self-adjoint operators" which we plan to discuss at the end of this course. The Hellinger-Toeplitz theorem is the reason why such operators need to be defined on dense subspaces, rather than the entire Hilbert space.

THEOREM 16.15 (Principle of condensation of singularities). *Suppose X is a Banach space, I is a set, and $\{T_\alpha : X \to Y_\alpha\}_{\alpha \in I}$ is a collection of bounded linear maps to normed vector spaces Y_α. If the set $\{\|T_\alpha\| \mid \alpha \in I\} \subset \mathbb{R}$ is unbounded, then there exists a comeager subset $R \subset X$ such that the set $\{\|T_\alpha x\| \mid \alpha \in I\} \subset \mathbb{R}$ is unbounded for every $x \in R$. In other words, if the operator norms $\|T_\alpha\|$ are unbounded, then $\|T_\alpha x\|$ is also unbounded for "generic" $x \in X$.*

PROOF. For $n \in \mathbb{N}$, let $R_n := \{x \in X \mid \sup_{\alpha \in I} \|T_\alpha x\| > n\}$. It will suffice to show that $R_n \subset X$ is open and dense for each $n \in \mathbb{N}$; the desired comeager subset is then $R := \bigcap_{n\in\mathbb{N}} R_n$.

To show that R_n is open, suppose $x_0 \in R_n$, so there exists $\alpha \in I$ such that $\|T_\alpha x_0\| > n$. Then since T_α and the norm are both continuous, one also has $\|T_\alpha x\| > n$ for all $x \in X$ sufficiently close to x_0, proving $x \in R_n$.

We now prove that R_n is dense. If not, then there exist $x_0 \in X$ and $\epsilon > 0$ such that $B_\epsilon(x_0) \cap R_n = \varnothing$, implying that for all $x \in B_\epsilon(x_0)$, $\sup_{\alpha\in I} \|T_\alpha x\| \leqslant n$. It then follows that for any $x \in X$ with $\|x\| = \epsilon/2$,

$$\|T_\alpha x\| = \|T_\alpha x_0 - T_\alpha(x_0 - x)\| \leqslant \|T_\alpha x_0\| + \|T_\alpha(x_0 - x)\| \leqslant n + n = 2n,$$

and thus $\|T_\alpha\| \leqslant \frac{2n}{\epsilon/2}$ for all $\alpha \in I$, contradicting the assumption that $\{\|T_\alpha\| \mid \alpha \in I\}$ is unbounded. □

Note that Theorem 16.15 on its own does not use the Baire category theorem, but without Baire, the result would be useless because we could not guarantee that the comeager subset $R \subset X$ is nonempty. This knowledge is used in the next two corollaries:

COROLLARY 16.16 (Principle of uniform boundedness). *For any collection of bounded linear operators $\{T_\alpha : X \to Y_\alpha\}_{\alpha\in I}$ between Banach spaces such that $\{\|T_\alpha x\| \mid \alpha \in I\}$ is bounded for every $x \in X$, $\{\|T_\alpha\| \mid \alpha \in I\}$ is also bounded.* □

COROLLARY 16.17. *Weakly convergent sequences in Banach spaces are bounded.*

PROOF. Given $x_n \rightharpoonup x$ in X, the convergence $\Lambda(x_n) \to \Lambda(x)$ for every $\Lambda \in X^*$ implies that the set $\{|\Lambda(x_n)| \mid n \in \mathbb{N}\} \subset \mathbb{R}$ is bounded. Using the natural injective isometry $\Phi : X \to X^{**}$, we obtain a collection of bounded linear functionals $\Phi(x_n) \in X^{**}$ on X^* such that $\{|\Phi(x_n)\Lambda| \mid n \in \mathbb{N}\} \subset \mathbb{R}$ is bounded for every $\Lambda \in X^*$, so by the principle of uniform boundedness, it follows that the norms $\|\Phi(x_n)\| = \|x_n\|$ satisfy a bound independent of n. □

16.5. Pointwise convergence of Fourier series. As a final application, we turn our attention to Fourier series and the surprising answer to Question 16.14.

THEOREM 16.18. *Given any countable set $I \subset S^1$, there exists a comeager set $R_I \subset C^0(S^1)$ of continuous periodic functions such that for every $x \in I$ and $f \in R_I$, the sum $\sum_{k\in\mathbb{Z}} e^{2\pi ikx} \widehat{f}_k$ diverges.*

PROOF. It suffices to prove that the result holds when I is a single point $\{x_0\} \in S^1$; the general case then follows because the countable intersection $R_I := \bigcap_{x\in I} R_{\{x\}}$ is comeager if each of the sets $R_{\{x\}}$ is comeager. With this in mind, we fix in the following a point $x_0 \in S^1$. We shall assume as usual that our functions on S^1 take values in a finite-dimensional complex inner product space V.

Define for each $N \in \mathbb{N}$ a bounded linear operator $T_N \in \mathscr{L}(C^0(S^1), V)$ that evaluates the corresponding *partial* Fourier series of $f \in C^0(S^1)$ at the point x_0:

$$T_N f := \sum_{k=-N}^{N} e^{2\pi ikx_0} \widehat{f}_k \in V.$$

Our main task is to show that

$$\|T_N\| \to \infty \qquad \text{as} \qquad N \to \infty.$$

Indeed, plugging $\{T_N : C^0(S^1) \to V\}_{N \in \mathbb{N}}$ into Theorem 16.15 will then give a comeager subset $R \subset C^0(S^1)$ such that for every $f \in R$, the sequence of partial Fourier series of f evaluated at x_0 is unbounded, and therefore divergent.

Using the characteristic function of the interval $[-N, N]$, we can write $T_N f$ as

$$T_N f = \sum_{k \in \mathbb{Z}} e^{2\pi i k x_0} \chi_{[-N,N]}(k) \hat{f}_k = (\rho_N * f)(x_0),$$

where $\rho_N \in C^\infty(S^1)$ is the unique complex-valued function with Fourier coefficients $(\hat{\rho}_N)_k = \chi_{[-N,N]}(k)$, and we've used Exercise 11.20 to rewrite $T_N f$ as a convolution. The function ρ_N is known as the *Nth Dirichlet kernel*, and we will take a closer look at it in the next paragraph. For now, we observe that ρ_N is a finite linear combination of complex exponential functions and is thus analytic, so it has isolated zeroes; in particular, it is nonzero almost everywhere on S^1. In order to compute the operator norm $\|T_N\|$, we then observe that for every $f \in C^0(S^1)$,

$$|T_N f| = |(\rho_N * f)(x_0)| = \left| \int_{S^1} \rho_N(x_0 - x) f(x)\, dx \right| \leqslant \|\rho_N\|_{L^1} \cdot \|f\|_{C^0},$$

and this inequality can be made arbitrarily close to an equality if we choose $f \in C^0(S^1)$ to be any continuous function sufficiently L^1-close to the L^1-function $x \mapsto \rho_N(x_0 - x)/|\rho_N(x_0 - x)|$, which is defined almost everywhere. This proves

$$\|T_N\| = \|\rho_N\|_{L^1},$$

and the remaining task is thus to prove

$$\|\rho_N\|_{L^1} = \int_{S^1} |\rho_N(x)|\, dx \to \infty \qquad \text{as} \qquad N \to \infty.$$

In light of its Fourier coefficients, a precise formula for the Dirichlet kernel ρ_N is given by

$$\rho_N(x) = \sum_{k=-N}^{N} e^{2\pi i k x}.$$

To rewrite this in a more useful form, we recall first the standard formula for the partial sums of a geometric series,

$$\sum_{k=0}^{N} a^k = \frac{1 - a^{N+1}}{1 - a} \qquad \text{for} \qquad a \neq 1,$$

from which one derives the related formula

$$\begin{aligned}
\sum_{k=-N}^{N} a^k &= \sum_{k=0}^{N} a^k + \sum_{k=0}^{N} (1/a)^k - 1 = \frac{1-a^{N+1}}{1-a} + \frac{1-a^{-N-1}}{1-a^{-1}} - 1 \\
&= \frac{(1-a^{N+1})(1-a^{-1}) + (1-a^{-N-1})(1-a)}{(1-a)(1-a^{-1})} - \frac{(1-a)(1-a^{-1})}{(1-a)(1-a^{-1})} \\
&= \frac{1-a^{N+1}-a^{-1}+a^N+1-a^{-N-1}-a+a^{-N}-(1-a-a^{-1}+1)}{(1-a)(1-a^{-1})} \\
&= \frac{-a^{N+1}+a^N-a^{-N-1}+a^{-N}}{(1-a)(1-a^{-1})} \\
&= \frac{a^{-N-1}}{a^{-1}} \cdot \frac{-a^{2N+2}+a^{2N+1}-1+a}{(1-a)(a-1)} = a^{-N}\frac{(a-1)(1-a^{2N+1})}{(1-a)(a-1)} \\
&= a^{-N}\frac{1-a^{2N+1}}{1-a} \qquad \text{for} \qquad a \ne 1.
\end{aligned}$$

To put the latter in a more symmetric form, we multiply the numerator and denominator of $\frac{a^{-N}-a^{N+1}}{1-a}$ by $a^{-1/2}$, producing

$$\sum_{k=-N}^{N} a^k = \frac{a^{-N-\frac{1}{2}} - a^{N+\frac{1}{2}}}{a^{-\frac{1}{2}} - a^{\frac{1}{2}}} = \frac{a^{N+\frac{1}{2}} - a^{-N-\frac{1}{2}}}{a^{\frac{1}{2}} - a^{-\frac{1}{2}}} \qquad \text{for} \qquad a \ne 1.$$

Since $e^{2\pi i x} \ne 1$ for almost every $x \in S^1 = \mathbb{R}/\mathbb{Z}$, we can plug $a := e^{2\pi i x}$ into this formula and apply the identity $e^{i\theta} - e^{-i\theta} = 2i\sin(\theta)$, obtaining

$$\rho_N(x) = \frac{\sin\left[2\pi\left(N+\frac{1}{2}\right)x\right]}{\sin(\pi x)} \qquad \text{for almost every } x \in S^1.$$

We now use this formula to establish a lower bound for $\|\rho_N\|_{L^1}$. Note that the denominator is positive for $x \in (0,1)$ and satisfies

$$\sin(\pi x) \leqslant \pi x, \tag{16.1}$$

implying

$$|\rho_N(x)| \geqslant \frac{\left|\sin\left[2\pi\left(N+\frac{1}{2}\right)x\right]\right|}{\pi x} \qquad \text{for} \qquad 0 < x < 1,$$

while the numerator has fixed signs on intervals of length $\frac{1}{2N+1}$. We deduce

$$\begin{aligned}
\|\rho_N\|_{L^1} &\geqslant \frac{1}{\pi}\sum_{j=1}^{2N+1} \int_{\frac{j-1}{2N+1}}^{\frac{j}{2N+1}} \frac{\left|\sin\left[2\pi\left(N+\frac{1}{2}\right)x\right]\right|}{x}\,dx = \frac{1}{\pi}\sum_{j=1}^{2N+1}\int_{\pi(j-1)}^{\pi j} \frac{|\sin u|}{u}\,du \\
&\geqslant \frac{1}{\pi}\sum_{j=1}^{2N+1} \frac{1}{\pi j}\int_{\pi(j-1)}^{\pi j} |\sin u|\,du = \frac{2}{\pi^2}\sum_{j=1}^{2N+1}\frac{1}{j} \to \infty \qquad \text{as} \qquad N \to \infty.
\end{aligned}$$

□

17. Fredholm operators

Recall from §12.6 that a bounded linear operator $T : X \to Y$ is called *compact* if it maps bounded sets to sets with compact closure. This is a property that *all* linear operators in finite dimensions have, and in infinite-dimensional settings, the compact operators furnish an especially nice class of operators for which many familiar results of finite-dimensional linear algebra can be generalized; this will be apparent in particular when we study their spectral theory in §18.5. Fredholm operators—which are related to compact operators but are not the same thing in general—form another class of operators for which some important things from the finite-dimensional setting can be rescued. They play an important role in the study of elliptic PDEs, and especially in differential geometry, e.g. the famous Atiyah-Singer index theorem is a formula that relates topological data to the Fredholm indices of certain differential operators defined in the setting of a closed differentiable manifold.

17.1. Definitions and examples. A bounded linear operator $T : X \to Y$ is called a **Fredholm operator** if

$$\dim \ker T < \infty \qquad \text{and} \qquad \dim \operatorname{coker} T < \infty,$$

where by definition the **cokernel** of T is

$$\operatorname{coker} T := Y \big/ \operatorname{im} T,$$

so the second condition means that the image of T has finite codimension. The **Fredholm index** of T is then the integer

$$\operatorname{ind}(T) := \dim \ker T - \dim \operatorname{coker} T \in \mathbb{Z}.$$

Fredholm operators arise naturally in the study of linear PDEs, in particular as certain types of differential operators for functions on compact domains (often with suitable boundary conditions imposed).

Example 17.1. For periodic functions of one variable $x \in S^1 = \mathbb{R}/\mathbb{Z}$ with values in a finite-dimensional vector space V, the derivative $\partial_x : C^k(S^1) \to C^{k-1}(S^1)$ is a Fredholm operator with index 0 for any $k \in \mathbb{N}$. Indeed,

$$\ker \partial_x = \left\{\text{constant functions } S^1 \to V\right\} \subset C^k(S^1),$$

and

$$\operatorname{im} \partial_x = \left\{ g \in C^{k-1}(S^1) \;\middle|\; \int_{S^1} g(x)\, dx = 0 \right\};$$

the latter follows from the fundamental theorem of calculus since the condition $\int_{S^1} g(x)\, dx = 0$ ensures that the function $f(x) := \int_0^x g(t)\, dt$ on $\mathbb{R}$ is periodic. The surjective linear map

$$C^{k-1}(S^1) \to V : g \mapsto \int_{S^1} g(x)\, dx$$

thus has $\operatorname{im} \partial_x$ as its kernel, so it descends to an isomorphism $\operatorname{coker} \partial_x \to V$, implying $\operatorname{ind}(\partial_x) = \dim V - \dim V = 0$.

Example 17.2. For the same reasons as explained in Example 17.1, $\partial_x : C^{k,\alpha}(S^1) \to C^{k-1,\alpha}(S^1)$ is Fredholm with index 0 for every $k \in \mathbb{N}$ and $\alpha \in (0, 1]$.

Exercise 17.3. Use Fourier series to show that the unique extension of $\partial_x : C^\infty(S^1) \to C^\infty(S^1)$ to a bounded linear operator $H^{s+1}(S^1) \to H^s(S^1)$ is also Fredholm with index 0 for every $s \geqslant 0$.

Exercise 17.4. Show that for functions taking values in a vector space V of dimension n, the derivative $\partial_x : C^k([0, 1]) \to C^{k-1}([0, 1])$ is a surjective Fredholm operator with index n, but imposing the boundary condition $f(0) = f(1) = 0$ produces an injective Fredholm operator

$$\left\{ f \in C^k([0, 1]) \;\middle|\; f(0) = f(1) = 0 \right\} \xrightarrow{\partial_x} C^{k-1}([0, 1])$$

with index $-n$.

Exercise 17.5. Show that for $n \geqslant 2$ and each $j = 1, \ldots, n$, the bounded linear operators $\partial_j : C^k(\mathbb{T}^n) \to C^{k-1}(\mathbb{T}^n)$, $\partial_j : C^{k,\alpha}(\mathbb{T}^n) \to C^{k-1,\alpha}(\mathbb{T}^n)$ and $\partial_j : H^{s+1}(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$ have infinite-dimensional kernels and are thus *not* Fredholm.

EXAMPLE 17.6. The Laplacian $\Delta := \sum_{j=1}^n \partial_j^2$ on fully periodic functions of n variables valued in a finite-dimensional vector space V defines a Fredholm operator

$$\Delta : H^{s+2}(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$$

with index 0 for each $s \geqslant 0$. Indeed, if $u \in H^{s+2}(\mathbb{T}^n)$ and $f = \Delta u$, then f has Fourier coefficients

$$\hat{f}_k = \sum_{j=1}^n \widehat{\partial_j^2 u}_k = -4\pi^2|k|^2\hat{u}_k \in V,$$

thus

$$\begin{aligned}\|f\|_{H^s}^2 &= \sum_{k\in\mathbb{Z}^n}(1+|k|^2)^s|\hat{f}_k|^2 = 16\pi^4\sum_{k\in\mathbb{Z}^n}(1+|k|^2)^s|k|^4|\hat{u}_k|^2 \leqslant 16\pi^4\sum_{k\in\mathbb{Z}^n}(1+|k|^2)^{s+2}|\hat{u}_k|^2 \\ &= 16\pi^4\|u\|_{H^{s+2}}^2,\end{aligned}$$

proving that Δ is bounded from H^{s+2} to H^s. If $\Delta u = 0$, then $\hat{f}_k = 0$ for all $k \in \mathbb{Z}^n$, implying $\hat{u}_k = 0$ for all $k \in \mathbb{Z}^n\backslash\{0\}$, but there is no condition on the coefficient $\hat{u}_0 \in V$, thus $\ker\Delta$ is the space of functions in $H^{s+2}(\mathbb{T}^n)$ whose only nonvanishing Fourier coefficient is $\hat{u}_0$, also known as the constant functions $\mathbb{T}^n \to V$. Similarly, the equation $\Delta u = f$ can be solved for a given $f \in H^s(\mathbb{T}^n)$ by writing $\hat{u}_k = -\frac{1}{4\pi^2|k|^2}\hat{f}_k$ for all $k \in \mathbb{Z}^n\backslash\{0\}$, but this is only possible if $\hat{f}_0 = 0$, thus $\operatorname{im}\Delta = \left\{f \in H^s(\mathbb{T}^n) \;\middle|\; \hat{f}_0 = 0\right\}$, and the surjective linear map $H^s(\mathbb{T}^n) \to V : f \mapsto \hat{f}_0$ therefore descends to an isomorphism $\operatorname{coker}\Delta \to V$. We conclude $\operatorname{ind}\Delta = \dim V - \dim V = 0$.

EXERCISE 17.7. Show that the wave operator $\partial_t^2 - \partial_x^2 : H^{s+2}(\mathbb{T}^2) \to H^s(\mathbb{T}^2)$ for fully periodic functions of two variables $(t, x) \in \mathbb{R}^2$ has infinite-dimensional kernel, so it is not Fredholm.
Hint: Consider functions of the form $(t, x) \mapsto f(t \pm x)$.

EXAMPLE 17.8. On any bounded open domain $\Omega \subset \mathbb{R}^n$, the Laplacian defines bounded linear operators $C^{k+2}(\bar{\Omega}) \to C^k(\bar{\Omega})$, $C^{k+2,\alpha}(\Omega) \to C^{k,\alpha}(\Omega)$ for each $k \geqslant 0$ and $\alpha \in (0, 1]$, as well as $W^{k+2,p}(\Omega) \to W^{k,p}(\Omega)$ for each $p \in [1, \infty]$, but none of these operators are Fredholm. The reason is that all smooth solutions to the equation $\Delta u = 0$ on $\mathbb{R}^n$ (these are called **harmonic** functions) belong to the kernels of these operators, and there is an infinite-dimensional space of such solutions. This is especially easy to see in the case $n = 2$, where one can identify $\mathbb{R}^2 = \mathbb{C}$ and extract harmonic functions from the real parts of holomorphic functions $\mathbb{C} \to \mathbb{C}$.

17.2. The Sobolev spaces $H_0^1(\Omega)$ and $H^{-1}(\Omega)$. The following is an extended digression from the general theory of Fredholm operators, and may be considered inessential. Its purpose is to present a more interesting example of a Fredholm operator than the ones that we've seen so far—in fact, the example discussed below is of fundamental importance in the theory of *elliptic* PDEs.

The Laplacian Δ is the most popular example of an elliptic operator; in contrast to the wave operator of Exercise 17.7, it has the right properties to produce a Fredholm operator in suitable functional-analytic settings, as demonstrated by Example 17.6. The problem with Example 17.8 turns out to be not the operator Δ itself, but the fact that it is being considered on a bounded domain without imposing any boundary condition.[40]

[40]The reader might protest at this point that in Exercise 17.4, we saw an example of a differential operator for functions on a bounded domain that was Fredholm despite no mention of any boundary condition. The domain in that example, however, had dimension 1, and since boundaries of 1-dimensional domains are isolated points at which functions take values in a finite-dimensional space, the difference between having a boundary condition or not in this situation is finite dimensional. Moreover, differential equations for functions of one variable are governed by the Picard-Lindelöf theorem on local existence and uniqueness of solutions. The situation for PDEs in more than one variable is radically different.

To discuss the Laplacian with boundary conditions, it is useful to introduce a few new variations on the usual Sobolev spaces $H^s(\mathbb{R}^n)$. We shall assume in the following that all functions take values in a fixed finite-dimensional complex inner product space $(V, \langle\ ,\ \rangle)$ unless otherwise noted. Recall that $H^s(\mathbb{R}^n)$ is defined for each $s \geqslant 0$ as the space of functions $f \in L^2(\mathbb{R}^n)$ with the property that the product of the Fourier transform $\widehat{f} : \mathbb{R}^n \to V$ with the function $\mathbb{R}^n \to \mathbb{R} : p \mapsto (1 + |p|^2)^{s/2}$ is also in $L^2(\mathbb{R}^n)$. The same definition does not quite make sense for $s < 0$ since, in that case, $(1 + |p|^2)^{s/2}\widehat{f}$ could very well be of class L^2 without $\widehat{f}$ itself being of class L^2, in which case one should not require $H^s(\mathbb{R}^n)$ to be a subspace of $L^2(\mathbb{R}^n)$. The remedy is to define $H^s(\mathbb{R}^n)$ for $s < 0$ as a space of *tempered distributions* rather than functions. In fact, the resulting definition also makes sense for $s \geqslant 0$, but reduces then to the previous definition since tempered distributions whose Fourier transforms are L^2-functions can always be represented by L^2-functions.

Definition 17.9. For any $s \in \mathbb{R}$, we define $H^s(\mathbb{R}^n) \subset \mathscr{S}'(\mathbb{R}^n)$ as the space of all tempered distributions Λ whose Fourier transforms are represented by functions of the form $\widehat{\Lambda}(p) = (1 + |p|^2)^{-s/2} f(p)$ for some $f \in L^2(\mathbb{R}^n)$. The H^s-norm is then defined via the inner product

$$\langle \Lambda, \Lambda' \rangle_{H^s} := \left\langle (1 + |p|^2)^{s/2}\widehat{\Lambda}, (1 + |p|^2)^{s/2}\widehat{\Lambda}' \right\rangle_{L^2} = \int_{\mathbb{R}^n} (1 + |p|^2)^s \langle \widehat{\Lambda}(p), \widehat{\Lambda}'(p) \rangle \, dp.$$

It is easy to see that $H^s(\mathbb{R}^n)$ is a Hilbert space, as it admits a natural unitary isomorphism to $L^2(\mathbb{R}^n)$, defined by taking Fourier transforms and multiplying by $(1 + |p|^2)^{s/2}$.

Exercise 17.10. Assuming $s \in \mathbb{R}$, prove:

(a) A distribution $\Lambda \in \mathscr{S}'(\mathbb{R}^n)$ is in $H^{-s}(\mathbb{R}^n)$ if and only if it satisfies a bound $|\Lambda(\varphi)| \leqslant c\|\varphi\|_{H^s}$ for all test functions $\varphi \in \mathscr{S}(\mathbb{R}^n)$.
(b) The space of vector-valued Schwartz-class functions $\mathscr{S}(\mathbb{R}^n)$ is dense in $H^s(\mathbb{R}^n)$.
(c) The pairing $\mathscr{S}(\mathbb{R}^n) \times \mathscr{S}(\mathbb{R}^n) \to \mathbb{C} : (\varphi, \psi) \mapsto \langle \varphi, \psi \rangle_{L^2}$ extends to a continuous real-bilinear pairing

$$\langle\ ,\ \rangle_s : H^{-s}(\mathbb{R}^n) \times H^s(\mathbb{R}^n) \to \mathbb{C},$$

$$\langle \Lambda, f \rangle_s := \left\langle (1 + |p|^2)^{-s/2}\widehat{\Lambda}, (1 + |p|^2)^{s/2}\widehat{f} \right\rangle_{L^2} = \int_{\mathbb{R}^n} \langle \widehat{\Lambda}(p), \widehat{f}(p) \rangle \, dp,$$

such that the real-linear map $\Lambda \mapsto \langle \Lambda, \cdot \rangle_s$ sends $H^{-s}(\mathbb{R}^n)$ isomorphically to the dual space of $H^s(\mathbb{R}^n)$.

Definition 17.11. For each $s \in \mathbb{R}$ and an open subset $\Omega \subset \mathbb{R}^n$, we identify $C_0^\infty(\Omega)$ with the space of smooth functions $\mathbb{R}^n \to V$ that have compact support in Ω, and define the closed subspace

$$\widetilde{H}^s(\Omega) \subset H^s(\mathbb{R}^n)$$

as the closure of $C_0^\infty(\Omega)$ in the H^s-norm. We also define the closed subspace

$$H^s_{\Omega^c}(\mathbb{R}^n) := \left\{ \Lambda \in H^s(\mathbb{R}^n) \;\middle|\; \Lambda(\varphi) = 0 \text{ for all } \varphi \in \mathscr{D}(\Omega) \right\} \subset H^s(\mathbb{R}^n),$$

i.e. $H^s_{\Omega^c}(\mathbb{R}^n)$ is the space of distributions in $H^s(\mathbb{R}^n)$ whose supports are contained in $\Omega^c := \mathbb{R}^n \backslash \Omega$. Finally, we define the quotient Banach space

$$H^s(\Omega) := H^s(\mathbb{R}^n) / H^s_{\Omega^c}(\mathbb{R}^n).$$

Exercise 17.12. Given $s \in \mathbb{R}$, let $I : H^{-s}(\mathbb{R}^n) \to (H^s(\mathbb{R}^n))^*$ denote the natural real-linear isomorphism in Exercise 17.10, and assume $\Omega \subset \mathbb{R}^n$ is an open set. Prove:

(a) I maps $H^{-s}_{\Omega^c}(\mathbb{R}^n) \subset H^{-s}(\mathbb{R}^n)$ onto the annihilator of $\widetilde{H}^s(\Omega)$, i.e. the space of bounded linear functionals on $H^s(\mathbb{R}^n)$ that vanish on $\widetilde{H}^s(\Omega)$.

(b) The map $H^{-s}(\mathbb{R}^n) \to \left(\widetilde{H}^s(\Omega)\right)^* : \Lambda \mapsto I(\Lambda)|_{\widetilde{H}^s(\Omega)}$ descends to the quotient $H^{-s}(\Omega) = H^{-s}(\mathbb{R}^n)/H^{-s}_{\Omega^c}(\mathbb{R}^n)$ to define a Banach space isomorphism $H^{-s}(\Omega) \to \left(\widetilde{H}^s(\Omega)\right)^*$.

It is straightforward to show that the obvious inclusion

$$H^s(\mathbb{R}^n) \hookrightarrow \mathscr{S}'(\mathbb{R}^n)$$

is continuous, as the topology defined on $H^s(\mathbb{R}^n)$ by the H^s-norm is stronger than the usual weak*-topology on the space of tempered distributions. Given an open set $\Omega \subset \mathbb{R}^n$, the continuous inclusion $\mathscr{D}(\Omega) \hookrightarrow \mathscr{S}(\mathbb{R}^n)$ dualizes to define a continuous restriction map $\mathscr{S}'(\mathbb{R}^n) \to \mathscr{D}'(\Omega) : \Lambda \mapsto \Lambda|_{\mathscr{D}(\Omega)}$, so that composing this with the inclusion above yields a natural continuous linear map

$$H^s(\mathbb{R}^n) \to \mathscr{D}'(\Omega).$$

The kernel of this map is $H^s_{\Omega^c}(\mathbb{R}^n)$, thus it descends to the quotient $H^s(\Omega) = H^s(\mathbb{R}^n)/H^s_{\Omega^c}(\mathbb{R}^n)$ as a natural continuous linear injection

$$H^s(\Omega) \hookrightarrow \mathscr{D}'(\Omega) : [\Lambda] \mapsto \Lambda|_{\mathscr{D}(\Omega)}. \tag{17.1}$$

It is useful to keep this injection in mind and regard elements of $H^s(\Omega)$ as distributions on Ω: from this perspective, $H^s(\Omega)$ is precisely the space of distributions on Ω that arise as restrictions to Ω of distributions in $H^s(\mathbb{R}^n)$. For elements $[\Lambda] \in H^s(\Omega)$ such that Λ can be represented by a locally integrable function $f : \mathbb{R}^n \to V$ (as is for instance always possible when $s \geqslant 0$), the corresponding distribution on Ω is represented by $f|_\Omega$, which is uniquely determined up to equality almost everywhere on Ω. In particular, $H^s(\Omega)$ for each $s \geqslant 0$ is identified in this way with a linear subspace of $L^2(\Omega)$.

Exercise 17.13. For any integer $k \geqslant 0$ and open subset $\Omega \subset \mathbb{R}^n$, show that the the map $[f] \mapsto f|_\Omega$ defines an injective bounded linear operator $H^k(\Omega) \to W^{k,2}(\Omega)$ with norm at most 1. *Hint: Recall that $H^k(\mathbb{R}^n) = W^{k,2}(\mathbb{R}^n)$, so the H^k-norm for functions on $\mathbb{R}^n$ is equivalent to a norm written in terms of weak derivatives instead of Fourier transforms.*

Remark 17.14. The injection $H^k(\Omega) \hookrightarrow W^{k,2}(\Omega)$ is also surjective, and thus a Banach space isomorphism, whenever it can be shown that every $f \in W^{k,2}(\Omega)$ admits an extension over $\mathbb{R}^n$ that belongs to $W^{k,2}(\mathbb{R}^n) = H^k(\mathbb{R}^n)$. This is not true in general, but it is trivially true for $k = 0$, thus giving a natural isomorphism $H^0(\Omega) = L^2(\Omega)$. By a standard extension result in the theory of Sobolev spaces, it is also true for every $k \in \mathbb{N}$ if the boundary of Ω satisfies certain regularity assumptions, e.g. it is true whenever Ω is bounded and its closure in $\mathbb{R}^n$ is a smooth submanifold with boundary.

As a closed subspace of a Hilbert space, $\widetilde{H}^s(\Omega)$ is also a Hilbert space, and a Hilbert space structure can also be assigned to $H^s(\Omega)$ by identifying it with the H^s-orthogonal complement of $H^s_{\Omega^c}(\mathbb{R}^n) \subset H^s(\mathbb{R}^n)$, though for our purposes, it will usually suffice to regard $H^s(\Omega)$ as a Banach space with the natural quotient norm. Notice that since $C_0^\infty(\Omega) \subset H^s(\mathbb{R}^n)$ for every $s \in \mathbb{R}^n$, there is a natural inclusion

$$C_0^\infty(\Omega) \hookrightarrow H^s(\Omega) : f \mapsto [f].$$

The following definition and subsequent proposition are not strictly necessary for our exposition in these notes, but we include them in order to make our notation consistent with what is found in most textbooks.

Definition 17.15. For $s > 0$ and an open subset $\Omega \subset \mathbb{R}^n$, the closed subspace

$$H_0^s(\Omega) \subset H^s(\Omega)$$

is defined as the H^s-closure of $C_0^\infty(\Omega) \subset H^s(\Omega)$.

Remark 17.16. The reason to restrict to $s > 0$ in Definition 17.15 is that for $s \leqslant 0$, $C_0^\infty(\Omega)$ is already dense in $H^s(\Omega)$, so the definition would in those cases give nothing new. The density of $C_0^\infty(\Omega) \subset H^s(\Omega)$ for $s \leqslant 0$ is an easy consequence of the case $s = 0$, for which Exercise 17.13 and Remark 17.14 give a natural identification $H^0(\Omega) = L^2(\Omega)$ and we can appeal to the fact that $C_0^\infty(\Omega)$ is dense in $L^2(\Omega)$. One can then deduce from the density of $\mathscr{S}(\mathbb{R}^n) \subset H^s(\mathbb{R}^n)$ that $L^2(\Omega)$ is dense in $H^s(\Omega)$ for every $s < 0$.

Proposition 17.17. *For any open set $\Omega \subset \mathbb{R}^n$ and each $k \in \mathbb{N}$, there is a natural isomorphism $\widetilde{H}^k(\Omega) \to H_0^k(\Omega) : f \mapsto [f]$.*

Proof. The quotient projection $H^k(\mathbb{R}^n) \to H^k(\Omega) : f \mapsto [f]$ restricts to an injection on the subspace $C_0^\infty(\Omega) \subset H^k(\mathbb{R}^n)$, and since it is bounded with respect to the H^k-norm, it has a unique extension to a bounded linear map $\Phi : \widetilde{H}^k(\Omega) \to H^k(\Omega)$. The closure of the image of Φ is $H_0^k(\Omega)$ by definition, thus our task is to prove that Φ is injective with closed image. To see this, we use the $W^{k,2}$-inner product as a substitute for the H^k-inner product and claim that every $f \in C_0^\infty(\Omega)$ is $W^{k,2}$-orthogonal to the subspace $H^k_{\Omega^c}(\mathbb{R}^n)$. Indeed, elements of $H^k_{\Omega^c}(\mathbb{R}^n)$ can be regarded as functions $g \in L^2(\mathbb{R}^n)$ that vanish almost everywhere on Ω and have weak derivatives $\partial^\alpha g \in L^2(\mathbb{R}^n)$ for every multi-index α of order $|\alpha| \leqslant k$, and it follows that these weak derivatives also vanish almost everywhere on Ω. We thus have

$$\langle f, g\rangle_{W^{k,2}} = \sum_{|\alpha| \leqslant k} \langle \partial^\alpha f, \partial^\alpha g\rangle_{L^2} = 0,$$

as the integrand in each of these L^2-inner products vanishes almost everywhere. Since the $W^{k,2}$- and H^k-norms are equivalent, it follows that the closure $\widetilde{H}^k(\Omega)$ is contained in the $W^{k,2}$-orthogonal complement of $H^k_{\Omega^c}(\mathbb{R}^n) \subset H^k(\mathbb{R}^n)$. The quotient projection $H^k(\mathbb{R}^n) \to H^k(\Omega)$ restricts to the latter space as a Banach space isomorphism, and the restriction of that isomorphism to the smaller closed subspace $\widetilde{H}^k(\Omega)$ is Φ, whose injective image is therefore also a closed subspace. $\square$

For any multi-index α of order $|\alpha| = m$, it is straightforward to show that the differential operator ∂^α defines bounded linear maps

$$\partial^\alpha : H^{s+m}(\mathbb{R}^n) \to H^s(\mathbb{R}^n)$$

for each $s \in \mathbb{R}$, and since ∂^α preserves the space of distributions with support disjoint from any given open set $\Omega \subset \mathbb{R}^n$, ∂^α descends to a bounded linear map of the quotients

$$\partial^\alpha : H^{s+m}(\Omega) \to H^s(\Omega).$$

For $f \in H^{s+m}(\Omega)$ and $g \in H^s(\Omega)$, each represented via (17.1) as distributions on Ω, the meaning of the relation $\partial^\alpha f = g$ can now be understood as follows. By definition, f and g each admit extensions $\widetilde{f} \in H^{s+m}(\mathbb{R}^n)$ and $\widetilde{g} \in H^s(\mathbb{R}^n)$ whose restrictions to Ω are f and g respectively, and $\partial^\alpha f = g$ then holds in $H^s(\Omega) = H^s(\mathbb{R}^n)/H^s_{\Omega^c}(\mathbb{R}^n)$ if and only if $\partial^\alpha \widetilde{f} - \widetilde{g}$ is in $H^s_{\Omega^c}(\mathbb{R}^n)$, which projects trivially to the quotient. This means that $\partial^\alpha \widetilde{f} - \widetilde{g}$ vanishes on all test functions supported in Ω, which is the same as saying that $\partial^\alpha f$ and g are identical distributions in $\mathscr{D}'(\Omega)$. If f and g are locally integrable functions, this means exactly what one would expect: g is equal to a weak derivative $\partial^\alpha f$ over the domain Ω. (Note that the notion of weak differentiation depends on the choice of domain, and no claim is being made here about the relation $\partial^\alpha \widetilde{f} = \widetilde{g}$ holding on any domain larger than Ω.)

Allowing linear combinations of such differential operators, the remarks of the previous paragraph apply in particular to the Laplace operator Δ with $m = 2$. The following can then be regarded as an existence and uniqueness result for distributional solutions of the Poisson equation

$\Delta u = f$ with boundary condition $u|_{\partial\bar{\Omega}} \equiv 0$; in particular, it provides for every $f \in L^2(\Omega) \subset H^{-1}(\Omega)$ a unique weak solution u in the space $H_0^1(\Omega)$.

THEOREM 17.18. *For any bounded open set $\Omega \subset \mathbb{R}^n$, the Laplacian defines an injective Fredholm operator*

$$H_0^1(\Omega) \xrightarrow{\Delta} H^{-1}(\Omega)$$

of index 0, i.e. it is a Banach space isomorphism.

This result is the first step in the study of the *Dirichlet problem*, which seeks solutions to the Laplace equation $\Delta u = 0$ on a bounded domain $\Omega \subset \mathbb{R}^n$ with prescribed boundary values $u|_{\partial\bar{\Omega}}$ (see for instance [**Tay96**, §5.1]). We will prove the theorem via a series of exercises in §17.7, by showing that Δ can be written as the sum of an isomorphism with a compact operator; as a "compact perturbation" of an index 0 Fredholm operator, it is therefore an index 0 Fredholm operator. A straightforward integration by parts argument then shows that the operator is injective, and since its index is 0, surjectivity follows immediately.

17.3. Main theorems. The simplest examples of Fredholm operators come from finite-dimensional linear algebra: every linear map $A : \mathbb{K}^n \to \mathbb{K}^m$ is Fredholm, and the fact that A descends to an isomorphism $\mathbb{K}^n/\ker A \to \operatorname{im} A$ reveals that the index of A is

$$\dim\ker A - (m - \dim\operatorname{im} A) = n - (n - \ker A) + \dim\operatorname{im} A - m = n - m.$$

Notice that this result depends only on the dimensions of the domain and target of A, not on A itself. For a Fredholm operator $T : X \to Y$ in infinite dimensions, one cannot so readily extract information from the isomorphism $X/\ker T \cong \operatorname{im} T$ since both sides are infinite dimensional. The remarkable fact is that the *index* of T, while dependent on more data than merely the spaces X and Y, still does not change under small perturbations or continuous deformations of T through families of Fredholm operators.

THEOREM 17.19. *For any Banach spaces X, Y, the set $\operatorname{Fred}(X, Y) \subset \mathscr{L}(X, Y)$ of Fredholm operators from X to Y is open with respect to the operator norm, and the function*

$$\operatorname{ind} : \operatorname{Fred}(X, Y) \to \mathbb{Z}$$

is continuous, i.e. it is locally constant.

COROLLARY 17.20. *For any Banach spaces X, Y and any continuous map $[0, 1] \to \operatorname{Fred}(X, Y)$: $s \mapsto T_s$, $\operatorname{ind}(T_s)$ is independent of s.* □

Our second main result is the following theorem on "compact perturbations," proved in §17.6.

THEOREM 17.21. *If X, Y are Banach spaces, $T \in \operatorname{Fred}(X, Y)$, and $K \in \mathscr{L}(X, Y)$ is a compact operator, then $T + K$ is also Fredholm.*

Notice that in the setting of Theorem 17.21, the operators $tK : X \to Y$ are also compact for every $t \in [0, 1]$, giving rise to a continuous family of Fredholm operators $T_s := T + sK$. Corollary 17.20 thus implies

$$\operatorname{ind}(T) = \operatorname{ind}(T + K) \qquad \text{whenever } T \text{ is Fredholm and } K \text{ is compact.}$$

This applies in particular to all operators of the form $\mathbb{1} - K$ for compact $K \in \mathscr{L}(X)$: they are Fredholm with index 0 since the isomorphism $\mathbb{1} : X \to X$ also is. The following consequence of this observation is known as the *Fredholm alternative*, and should remind you of finite-dimensional linear algebra.

COROLLARY 17.22 (the Fredholm alternative). *For any compact operator $K : X \to X$ on a Banach space X, exactly one of the following holds:*

(i) *The linear homogeneous equation* $x - Kx = 0$ *has a nontrivial finite-dimensional space of solutions;*

(ii) *The linear inhomogeneous equation* $x - Kx = y$ *has a unique solution for every* $y \in X$.

PROOF. Since $\operatorname{ind}(\mathbb{1} - K) = 0$, it satisfies $\dim \ker(\mathbb{1} - K) = \dim \operatorname{coker}(\mathbb{1} - K) < \infty$. The first alternative occurs when these dimensions are positive, and the second when they are zero. □

As a concrete example of Corollary 17.22, consider the Banach space isomorphism $\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$ for a bounded open subset $\Omega \subset \mathbb{R}^n$. We will see in Exercise 17.36 that the natural inclusion $j : H_0^1(\Omega) \hookrightarrow H^{-1}(\Omega)$ is compact, thus for any scalar $\lambda \in \mathbb{C}$, the operator $\lambda j \Delta^{-1} : H^{-1}(\Omega) \to H^{-1}(\Omega)$ is the composition of a bounded operator with a compact operator, and is therefore compact. The equation

$$\left(\mathbb{1} - \lambda j \Delta^{-1}\right) u = 0$$

is then satisfied if and only if u is a function in $H_0^1(\Omega)$ with $\Delta u = \lambda u$, i.e. it is an eigenfunction of Δ over the domain Ω, with eigenvalue λ. Corollary 17.22 thus guarantees among other things that for any eigenvalue λ of Δ on a bounded domain Ω with the boundary condition $u|_{\partial\bar{\Omega}} \equiv 0$, the corresponding eigenspace is finite dimensional.

EXERCISE 17.23. Show that for any finite-dimensional vector space V and any continuous periodic functions $g : S^1 \to \mathbb{R}$ and $h : S^1 \to V$ satisfying $\int_{S^1} g(x)\, dx \neq 0$, the linear inhomogeneous differential equation

$$f' + gf = h$$

has a unique periodic solution $f : S^1 \to V$.
Hint: Show that $T : C^1(S^1) \to C^0(S^1) : f \mapsto f' + gf$ *is a compact perturbation of the Fredholm operator in Example 17.1.*

As preparation for the proof of Theorem 17.19, the following result is a minor variation on the observation in Example 15.5 that finite-codimensional closed subspaces are always complemented.

LEMMA 17.24. *If* X, Y *are Banach spaces and* $T \in \mathscr{L}(X, Y)$ *has finite-dimensional cokernel, then* $\operatorname{im} T \subset Y$ *is closed.*

PROOF. Choose $w_1, \ldots, w_n \in Y$ such that the equivalence classes $[w_1], \ldots, [w_n]$ form a basis of $\operatorname{coker} T = Y / \operatorname{im} T$, and define the linear injection

$$\Phi : \mathbb{R}^n \hookrightarrow Y : (\lambda_1, \ldots, \lambda_n) \mapsto \sum_{j=1}^n \lambda_j w_j.$$

We can use this to define a surjective bounded linear operator

$$\Psi : X \oplus \mathbb{R}^n \to Y : (x, z) \mapsto Tx + \Phi(z),$$

whose kernel is $\ker T \oplus \{0\} \subset X \oplus \mathbb{R}^n$. The surjectivity of this operator implies that it has closed image, so by Exercise 17.25 below, there exists a constant $c > 0$ such that

$$\|\Psi(x, z)\| \geqslant c \cdot \inf_{v \in \ker T} \|(x + v, z)\|$$

for all $(x, z) \in X \oplus \mathbb{R}^n$. In particular, setting $z = 0$ in this estimate yields

$$\|Tx\| \geqslant c \cdot \inf_{v \in \ker T} \|x + v\|,$$

which by Exercise 17.25 implies that $\operatorname{im} T \subset Y$ is closed. □

EXERCISE 17.25. For Banach spaces X, Y, prove:

(a) An injective operator $T \in \mathscr{L}(X, Y)$ has closed image if and only if it satisfies a bound of the form $\|Tx\| \geqslant c\|x\|$ for some constant $c > 0$ independent of $x \in X$.

(b) Every $T \in \mathscr{L}(X,Y)$ descends to a bounded linear operator $X/\ker T \to Y : [x] \mapsto Tx$.
(c) An operator $T \in \mathscr{L}(X,Y)$ has closed image if and only if it satisfies a bound of the form

$$\|Tx\| \geqslant c \cdot \inf_{v \in \ker T} \|x+v\|$$

for some constant $c > 0$ independent of $x \in X$.

Remark 17.26. Some sources explicitly include the condition that $\operatorname{im} T \subset Y$ is closed as part of the definition of a Fredholm operator $T : X \to Y$. Lemma 17.24 shows that this is unnecesary, but it makes little difference in practice, as the standard ways of proving that T is a Fredholm operator (e.g. Lemma 17.34 below) typically include an explicit proof that T has closed image.

Recall from §15.3 that finite-dimensional and finite-codimensional closed subspaces always admit closed complements. Thanks to Lemma 17.24, we therefore obtain the following useful picture of an arbitrary Fredholm operator $T_0 : X \to Y$. Let us abbreviate

$$K := \ker T_0 \subset X, \qquad W := \operatorname{im} T_0 \subset Y,$$

and choose closed subspaces $V \subset X$ and $C \subset Y$ such that

$$X = V \oplus K \qquad \text{and} \qquad Y = W \oplus C, \tag{17.2}$$

keeping in mind that $\dim K < \infty$ and $\dim C = \dim(Y/W) < \infty$. The restriction $A_0 := T_0|_V : V \to \operatorname{im} T_0 = W$ is now a bounded linear bijection, so by the inverse mapping theorem, it is a Banach space isomorphism, meaning its inverse $A_0^{-1} : W \to V$ is also bounded. In block form with respect to the splittings (17.2), T_0 now takes the form

$$T_0 = \begin{pmatrix} A_0 & 0 \\ 0 & 0 \end{pmatrix} : V \oplus K \to W \oplus C. \tag{17.3}$$

We can of course use the same splittings to write any other operator $T \in \mathscr{L}(X,Y)$ in a similar block form

$$T = \begin{pmatrix} A & B \\ C & D \end{pmatrix} : V \oplus K \to W \oplus C \tag{17.4}$$

for bounded linear operators $A : V \to W$, $B : K \to W$, $C : V \to C$ and $D : K \to C$, e.g. A is the composition $\Pi_W T \iota_V$ where $\iota_V : V \hookrightarrow X$ is the continuous inclusion and $\Pi_W : Y \to W$ is the continuous projection along C, and so forth. Note that the fact that the subspaces in these splittings are *closed* is crucial here: by Theorem 15.8, this is the reason why e.g. the projection Π_W is continuous, so that each of the blocks A, B, C, D depend continuously on $T \in \mathscr{L}(X,Y)$. The most useful observation will be that since the space of Banach space isomorphisms $V \to W$ is an open subset of $\mathscr{L}(V,W)$, the term $A : V \to W$ will remain invertible whenever T is sufficiently close to T_0.

Proof of Theorem 17.19. Given $T_0 \in \operatorname{Fred}(X,Y)$, choose splittings as in (17.2) with $K = \ker T_0$ and $W = \operatorname{im} T_0$ in order to write each $T \in \mathscr{L}(X,Y)$ in block from as in (17.4). Since the block $A \in \mathscr{L}(V,W)$ depends continuously on T and the set of invertible bounded linear maps is open, we can define an open neighborhood $\mathcal{U} \subset \mathscr{L}(X,Y)$ of T_0 by

$$\mathcal{U} := \left\{ T \in \mathscr{L}(X,Y) \;\middle|\; A : V \to W \text{ is invertible} \right\}.$$

We claim that every $T \in \mathcal{U}$ is Fredholm, with

$$\dim \ker T \leqslant \dim \ker T_0, \quad \dim \operatorname{coker} T \leqslant \dim \operatorname{coker} T_0, \quad \text{and} \quad \operatorname{ind} T = \operatorname{ind} T_0.$$

To see this, we can associate to each $T \in \mathcal{U}$ a pair of Banach space isomorphisms $\Phi \in \mathscr{L}(X)$ and $\Psi \in \mathscr{L}(Y)$, expressed in block form with respect to the splittings $X = V \oplus K$ and $Y = W \oplus C$ as

$$\Phi := \begin{pmatrix} \mathbb{1} & -A^{-1}B \\ 0 & \mathbb{1} \end{pmatrix}, \qquad \Psi := \begin{pmatrix} \mathbb{1} & 0 \\ -CA^{-1} & \mathbb{1} \end{pmatrix}.$$

That these are both Banach space isomorphisms is straightforward to check: their inverses are namely

$$\Phi^{-1} := \begin{pmatrix} \mathbb{1} & A^{-1}B \\ 0 & \mathbb{1} \end{pmatrix}, \qquad \Psi^{-1} := \begin{pmatrix} \mathbb{1} & 0 \\ CA^{-1} & \mathbb{1} \end{pmatrix}.$$

The linear map $T : X \to Y$ is thus conjugate to

$$T' := \Psi T \Phi = \begin{pmatrix} A & 0 \\ 0 & T^{\mathrm{red}} \end{pmatrix},$$

where we define the "reduced" operator

$$T^{\mathrm{red}} := D - CA^{-1}B \in \mathscr{L}(K, C).$$

There are two crucial things to observe about the block-diagonal operator T': its top left block is invertible, and its bottom right block is a linear map between *finite-dimensional* vector spaces. We thus have

$$\ker T = \Phi(\ker T') = \Phi(\{0\} \oplus \ker T^{\mathrm{red}}),$$

implying $\dim \ker T = \dim \ker T^{\mathrm{red}} \leqslant \dim K = \dim \ker T_0$. Similarly, Ψ^{-1} maps $\operatorname{im} T' = W \oplus \operatorname{im} T^{\mathrm{red}}$ isomorphically to $\operatorname{im} T$ and thus descends to an isomorphism $\Psi^{-1} : \operatorname{coker} T' \to \operatorname{coker} T$, where

$$\operatorname{coker} T' = (W \oplus C)\big/(W \oplus \operatorname{im} T^{\mathrm{red}}) \cong C\big/\operatorname{im} T^{\mathrm{red}} = \operatorname{coker} T^{\mathrm{red}},$$

which gives $\dim \operatorname{coker} T = \dim \operatorname{coker} T^{\mathrm{red}} \leqslant \dim C = \dim \operatorname{coker} T_0$. Observe finally that as an operator between finite-dimensional spaces, the index of $T^{\mathrm{red}} : K \to C$ depends only on the spaces themselves, so it is the same as the index of the zero map $K \to C$, giving

$$\operatorname{ind} T = \operatorname{ind} T^{\mathrm{red}} = \operatorname{ind}\left(K \xrightarrow{0} C\right) = \dim K - \dim C = \operatorname{ind} T_0.$$

□

17.4. Some preparatory results. The results of this and the next section will serve as preparation for the proof of Theorem 17.21 on compact perturbations.

PROPOSITION 17.27. *A normed vector space is finite dimensional if and only if the closed unit ball about the origin is compact.*

PROOF. One direction of the statement follows from first-year analysis, since all closed and bounded subsets of finite-dimensional vector spaces are compact. For the converse, assume X is a normed vector space and $\overline{B_1(0)} \subset X$ is compact. We will give an argument that, with minor modifications,[41] also applies to arbitrary topological vector spaces, proving that the finite-dimensional vector spaces are the only *locally compact* topological vector spaces.

Let $\mathcal{U} := B_1(0) \subset X$ and assume its closure $\bar{\mathcal{U}}$ is compact. Observe that for each $x \in X$, the set $x + \frac{1}{2}\mathcal{U} \subset X$ is a neighborhood of x, so compactness implies

$$\mathcal{U} \subset \bar{\mathcal{U}} \subset \bigcup_{i=1}^{n} \left(x_i + \frac{1}{2}\mathcal{U}\right) \tag{17.5}$$

[41] See in particular https://terrytao.wordpress.com/2011/05/24/locally-compact-topological-vector-spaces/

for some finite set $x_1, \ldots, x_n \in X$. We will show that the finite-dimensional subspace $V \subset X$ spanned by $x_1, \ldots, x_n$ is in fact X. Indeed, (17.5) implies $\mathcal{U} \subset V + \frac{1}{2}\mathcal{U}$, and rescaling then implies $\frac{1}{2}\mathcal{U} \subset V + \frac{1}{4}\mathcal{U}$ since V is a linear subspace, and thus

$$\mathcal{U} \subset V + \frac{1}{2}\mathcal{U} \subset V + \frac{1}{4}\mathcal{U}.$$

Repeating this argument finitely many times produces

$$\mathcal{U} \subset V + \frac{1}{2^n}\mathcal{U}$$

for every $n \in \mathbb{N}$. It follows that every $x \in \mathcal{U}$ belongs for each $n \in \mathbb{N}$ to the ball of radius $1/2^n$ about some point in V, and is therefore in the closure of V. Since $\dim V < \infty$, V is already closed, so this implies $\mathcal{U} \subset V$. For an arbitrary $x \in X$, we can now choose $\epsilon > 0$ so that $\epsilon x \in \mathcal{U}$, and it follows that $x = \frac{1}{\epsilon}\epsilon x \in V$. ☐

Remark 17.28. Another popular proof of Proposition 17.27 (which however does not generalize to topological vector spaces) uses a basic geometric result called the *Riesz lemma*, which states that for any closed proper subspace V in a normed vector space X,

$$\sup_{x \in X,\ \|x\|=1} \operatorname{dist}(x, V) = 1.$$

If $\dim X = \infty$, one can use this to construct for any $\delta \in (0,1)$ a sequence $x_n \in X$ that satisfies $\|x_n\| = 1$ for all n but $\|x_n - x_m\| \geqslant \delta$ for all $m \neq n$, so that no subsequence can be Cauchy. (See e.g. [**BS18**, §2.2].) If X is an inner product space, then one can do better and achieve $\delta = \sqrt{2}$ by constructing x_n to be orthonormal.

Proposition 17.29. *If $K \in \mathscr{L}(X, Y)$ is compact, then so is $K^\top \in \mathscr{L}(Y^*, X^*)$.*

Proof. Assume $K : X \to Y$ is compact and $\Lambda_n \in Y^*$ is a sequence satisfying $\|\Lambda_n\| \leqslant C$ for some constant $C > 0$. Letting $\overline{B_1(0)} \subset X$ denote the closed unit ball in X, the set

$$M := \overline{K(\overline{B_1(0)})} \subset Y$$

is then compact. The functions $\Lambda_n|_M : M \to \mathbb{K}$ now satisfy

$$|\Lambda_n(y)| \leqslant C \cdot \max_{z \in M} \|z\|$$

and are thus uniformly bounded; they also satisfy the Lipschitz condition

$$|\Lambda_n(y) - \Lambda_n(y')| \leqslant C \cdot \|y - y'\|$$

for $y, y' \in M$, so they are equicontinuous. It now follows from the Arzelà-Ascoli theorem that after replacing Λ_n with a subsequence, the sequence $\Lambda_n|_M : M \to \mathbb{K}$ is uniformly convergent. (Note that in applying the Arzelà-Ascoli theorem, we are using the fact that $M \subset Y$ is compact, which follows from the compactness of K.) This implies that the sequence $K^\top \Lambda_n|_{\overline{B_1(0)}} = \Lambda_n \circ K|_{\overline{B_1(0)}} : \overline{B_1(0)} \to \mathbb{K}$ also converges uniformly, hence it is uniformly Cauchy, implying that $K^\top \Lambda_n$ is also a Cauchy sequence and therefore convergent in X^*. ☐

Remark 17.30. We will not need to use this, but the converse of Proposition 17.29 is also true; see [**BS18**, Theorem 4.28(iii)].

The **annihilator** of a subset $V \subset X$ is defined by

$$V^\perp := \left\{ \Lambda \in X^* \;\middle|\; \Lambda|_V = 0 \right\} \subset X^*,$$

and similarly, the **pre-annihilator** of a set of dual vectors $V \subset X^*$ is

$${}^\perp V := \left\{ x \in X \;\middle|\; \Lambda(x) = 0 \text{ for all } \Lambda \in V \right\} \subset X.$$

In other words, ${}^{\perp}V = \Phi^{-1}(V^{\perp}) \subset X$ for the canonical inclusion $\Phi : X \to X^{**}$. It is easy to check that whenever V is a linear subspace of X or X^*, $V^{\perp}$ or ${}^{\perp}V$ respectively is a closed linear subspace.

EXERCISE 17.31. For a closed subspace $V \subset X$ with inclusion map $i : V \hookrightarrow X$ and quotient projection $\pi : X \to X/V$, prove:

(a) The transpose $i^{\top} : X^* \to V^*$ of i descends to a Banach space isomorphism $X^*/V^{\perp} \to V^*$.
(b) The transpose $\pi^{\top} : (X/V)^* \to X^*$ of π defines a Banach space isomorphism to the subspace $V^{\perp} \subset X^*$.

PROPOSITION 17.32. *For any $T \in \mathscr{L}(X, Y)$,*

$$(\operatorname{im} T)^{\perp} = \ker T^{\top} \quad \text{and} \quad {}^{\perp}(\operatorname{im} T^{\top}) = \ker T.$$

If additionally $\operatorname{im} T \subset Y$ is closed, then $\operatorname{im} T^{\top} \subset X^*$ is also closed, and

$$\operatorname{im} T = {}^{\perp}(\ker T^{\top}) \quad \text{and} \quad \operatorname{im} T^{\top} = (\ker T)^{\perp}.$$

PROOF. The first two equalities are readily verified from the definitions, where in the second case, one needs to use the fact that $y \in Y$ vanishes if and only if $\Lambda(y) = 0$ for every $\Lambda \in Y^*$, which follows from the Hahn-Banach theorem. It is similarly straightforward to verify the inclusions $\operatorname{im} T \subset {}^{\perp}(\ker T^{\top})$ and $\operatorname{im} T^{\top} \subset (\ker T)^{\perp}$.

We claim moreover that $\operatorname{im} T$ is always dense in ${}^{\perp}(\ker T^{\top})$. Indeed, consider a bounded linear functional $\Lambda : {}^{\perp}(\ker T^{\top}) \to \mathbb{K}$ such that $\Lambda|_{\operatorname{im} T} = 0$, and use the Hahn-Banach theorem to extend Λ to a bounded linear functional on Y. Then $\Lambda \in (\operatorname{im} T)^{\perp} = \ker T^{\top}$, thus $\Lambda(y) = 0$ for all $y \in {}^{\perp}(\ker T^{\top})$ by definition, implying that the original unextended functional was trivial. The density of $\operatorname{im} T$ now follows from Theorem 14.8 (a corollary of the Hahn-Banach theorem). This proves that $\operatorname{im} T = {}^{\perp}(\ker T^{\top})$ if and only if $\operatorname{im} T$ is closed.

It remains to prove that if $\operatorname{im} T$ is closed, then $(\ker T)^{\perp} \subset \operatorname{im} T^{\top}$; this will imply that $\operatorname{im} T^{\top}$ is also closed since $(\ker T)^{\perp}$ always is. As a first step, Exercise 17.25 gives an estimate

$$\|Tx\| \geqslant c \cdot \inf_{v \in \ker T} \|x + v\| \tag{17.6}$$

for some constant $c > 0$. Now suppose $\Lambda \in (\ker T)^{\perp} \subset X^*$, so $\Lambda(v) = 0$ for all $v \in \ker T$ and thus

$$|\Lambda(x)| = |\Lambda(x + v)| \leqslant \|\Lambda\| \cdot \|x + v\|$$

for all $x \in X$ and $v \in \ker T$. Taking the infimum over $v \in \ker T$ and combining this with (17.6) gives

$$|\Lambda(x)| \leqslant \frac{1}{c}\|\Lambda\| \cdot \|Tx\|. \tag{17.7}$$

To show that $\Lambda \in \operatorname{im} T^{\top}$, observe that there exists a unique bounded linear functional $\lambda_0 : \operatorname{im} T \to \mathbb{K}$ such that

$$\lambda_0(Tx) = \Lambda(x) \quad \text{for all} \quad x \in X;$$

indeed, the value of $\Lambda(x)$ is independent of the choice of $x \in T^{-1}(Tx)$ since $\Lambda(x) = 0$ whenever $Tx = 0$, and the estimate (17.7) implies that this functional is bounded. Extending λ_0 to $\lambda \in Y^*$ via the Hahn-Banach theorem, we now have $\Lambda = \lambda_0 \circ T = \lambda \circ T = T^{\top}\lambda$. □

17.5. The semi-Fredholm property. A bounded linear map $T : X \to Y$ is said to be **semi-Fredholm** if

$$\dim \ker T < \infty \qquad \text{and} \qquad \operatorname{im} T \subset Y \text{ is closed.}$$

This condition often turns out to be a convenient stepping stone toward proving the Fredholm property.

LEMMA 17.33. *The following conditions on an operator $T \in \mathscr{L}(X, Y)$ are equivalent:*

(1) T and $T^\top$ are both semi-Fredholm.
(2) T is Fredholm.
(3) $T^\top$ is Fredholm.

Moreover, when these conditions hold,

$$\dim \ker T^\top = \dim \operatorname{coker} T \quad \text{and} \quad \dim \operatorname{coker} T^\top = \dim \ker T,$$

hence $\operatorname{ind} T^\top = -\operatorname{ind} T$.

PROOF. Assume T and $T^\top$ are both semi-Fredholm, so $\ker T$ and $\ker T^\top$ are both finite dimensional and $\operatorname{im} T$ and $\operatorname{im} T^\top$ are both closed. Using the isomorphisms from Exercise 17.31 and Proposition 17.32, we have

$$(\ker T)^* \cong X^*/(\ker T)^\perp = X^*/\operatorname{im} T^\top = \operatorname{coker} T^\top,$$

and

$$(\operatorname{coker} T)^* = (Y/\operatorname{im} T)^* \cong (\operatorname{im} T)^\perp = \ker T^\top,$$

so the finite-dimensionality of $\ker T$ and $\ker T^\top$ implies that $\operatorname{coker} T$ and $\operatorname{coker} T^\top$ are also both finite dimensional. This proves (1) $\Rightarrow$ (2) and (1) $\Rightarrow$ (3), along with the stated relations between dimensions.

If we instead assume T is Fredholm, then Lemma 17.24 and Proposition 17.32 imply that $\operatorname{im} T$ and $\operatorname{im} T^\top$ are both closed, so the isomorphisms above still hold, and the finite-dimensionality of $\ker T$ and $\operatorname{coker} T$ implies the same for $\ker T^\top$ and $\operatorname{coker} T^\top$, proving (2) $\Rightarrow$ (3). The proof that (3) implies (1) or (2) requires an additional argument to show that $\operatorname{im} T$ is closed whenever $\operatorname{im} T^\top$ is closed, but we shall omit this since it is not needed in the sequel. (A proof may be found in [**BS18**, Theorem 4.16].) □

LEMMA 17.34. *An operator $T \in \mathscr{L}(X, Y)$ is semi-Fredholm if and only if there exists a Banach space Z and compact operator $K \in \mathscr{L}(X, Z)$ satisfying*

$$\|x\| \leqslant c\left(\|Tx\| + \|Kx\|\right) \tag{17.8}$$

for all $x \in X$ and some constant $c > 0$ independent of x.

PROOF. Assume $T : X \to Y$ is semi-Fredholm. Its finite-dimensional kernel $Z := \ker T \subset X$ then admits a closed complement, so there is a continuous linear projection map $K : X \to Z$. This operator is clearly compact since it has finite rank. The linear map

$$X \to Y \oplus Z : x \mapsto (Tx, Kx)$$

is then bounded and injective, with image the closed subspace $\operatorname{im} T \oplus Z \subset Y \oplus Z$, so Exercise 17.25 gives the estimate

$$\|(Tx, Kx)\| = \|Tx\| + \|Kx\| \geqslant c\|x\|$$

for some constant $c > 0$ independent of $x \in X$.

Conversely, suppose a compact operator $K : X \to Z$ is given such that the estimate $\|x\| \leqslant c\|Tx\| + c\|Kx\|$ is satisfied. We will show that the closed unit ball in $\ker T$ is compact, implying via Proposition 17.27 that $\dim \ker T < \infty$. Indeed, if $x_n \in \ker T$ is a sequence satisfying $\|x_n\| \leqslant 1$, then after reducing to a subsequence, we can assume Kx_n converges in Z, due to the compactness of K. In particular, Kx_n is a Cauchy sequence, and since $Tx_n = 0$ for all n, applying the estimate (17.8) to $x_n - x_m$ yields

$$\|x_n - x_m\| \leqslant c\|Kx_n - Kx_m\|.$$

This proves that x_n is a Cauchy sequence in X, so x_n converges, and $\ker T$ is therefore finite dimensional.

To prove that $\operatorname{im} T$ is closed, we first simplify the situation by restricting T to a closed subspace $V \subset X$ that is complementary to $\ker T$; such a subspace necessarily exists since finite-dimensional subspaces are always complemented, and the restricted operator $T|_V : V \to Y$ is now injective but has the same image as T. Now if $x_n \in V$ is a sequence such that $Tx_n \to y \in Y$, we claim that x_n must be bounded. If not, then after restricting to a subsequence, we can assume $\|x_n\| \to \infty$ and thus $T(x_n/\|x_n\|) \to 0$, while the boundedness of $x_n/\|x_n\|$ implies without loss of generality that $K(x_n/\|x_n\|)$ converges. Arguing as in the previous paragraph via Cauchy sequences, we now conclude from (17.8) that $x_n/\|x_n\|$ converges to some $x_\infty \in V$ with $\|x_\infty\| = 1$ but $Tx_\infty = 0$, and that is impossible since $\ker T \cap V = \{0\}$. But now that we know x_n is bounded, Kx_n must in turn have a convergent subsequence, while Tx_n converges by assumption, so another application of (17.8) to Cauchy sequences proves that x_n has a subsequence convergent to some element $x \in V$, which must then satisfy $Tx = y$, proving that $\operatorname{im} T$ is closed. □

17.6. Compact perturbations. Let us now restate and prove Theorem 17.21.

THEOREM 17.35. *If $T \in \mathscr{L}(X, Y)$ is Fredholm, then so is $T + K$ for every compact operator $K \in \mathscr{L}(X, Y)$.*

PROOF. We claim first that if T is semi-Fredholm, then so is $T+K$. We use the characterization of the semi-Fredholm condition in Lemma 17.34: assume T satisfies an estimate of the form $\|x\| \leqslant c\|Tx\| + c\|K_0 x\|$ for some compact operator $K_0 : X \to Z$. The perturbed operator $T + K : X \to Y$ then satisfies

$$\|x\| \leqslant c\|Tx\| + c\|K_0 x\| \leqslant c\|(T+K)x\| + c\|Kx\| + c\|K_0 x\| = c\left(\|(T+K)x\| + \|K_1 x\|\right),$$

where we define the operator $K_1 : X \to Y \oplus Z : x \mapsto (Kx, K_0 x)$, and K_1 is compact since both K_0 and K are compact. Lemma 17.34 then implies that $T + K$ is semi-Fredholm.

Now if T is Fredholm, Lemma 17.33 implies that $T^\top : Y^* \to X^*$ is also Fredholm, while $K^\top : Y^* \to X^*$ is also compact due to Proposition 17.29. The result of the previous paragraph thus implies that $T^\top + K^\top = (T+K)^\top$ is also semi-Fredholm, so by Lemma 17.33, $T + K$ is also Fredholm. □

17.7. The Dirichlet problem. The next four exercises furnish the proof of Theorem 17.18, which stated that the Laplace operator defines a Banach space isomorphism

$$\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$$

for any bounded open subset $\Omega \subset \mathbb{R}^n$. It will be convenient to assume in the following that

$$\bar{\Omega} \subset (0,1)^n, \tag{17.9}$$

though it should be clear via scaling and translation that if the result is true for this special case, then it is true in general.

EXERCISE 17.36. Assume (17.9) holds.

(a) Associate to each $f \in C_0^\infty(\Omega)$ the unique function $F \in C^\infty(\mathbb{T}^n)$ such that $f(x) = F(x)$ for $x \in (0,1)^n$. Show that the map $C_0^\infty(\Omega) \to C^\infty(\mathbb{T}^n) : f \mapsto F$ extends to bounded linear injections

$$L^2(\Omega) \hookrightarrow L^2(\mathbb{T}^n) \quad \text{and} \quad \widetilde{H}^1(\Omega) \hookrightarrow H^1(\mathbb{T}^n)$$

whose images are closed.

Hint: Avoid Fourier analysis here by replacing the usual H^1-norm with the equivalent norm $\|u\|_{W^{1,2}} := \sum_{|\alpha| \leqslant 1} \|\partial^\alpha u\|_{L^2}$. This works equally well on $\mathbb{R}^n$ or $\mathbb{T}^n$.

(b) Deduce via the natural isomorphism $\widetilde{H}^1(\Omega) \to H_0^1(\Omega)$ from Proposition 17.17 and the compactness of the inclusion $H^1(\mathbb{T}^n) \hookrightarrow L^2(\mathbb{T}^n)$ that the linear injection $H_0^1(\Omega) \hookrightarrow L^2(\Omega) : [f] \mapsto f|_\Omega$ is also compact.

(c) Deduce that the natural inclusion $H_0^1(\Omega) \hookrightarrow H^{-1}(\Omega)$ is compact by presenting it as a composition of bounded linear operators in which at least one is compact.

Remark: This result is a special case of the Rellich-Kondrachov compactness theorem. Notice that the boundedness of Ω plays an essential role in the proof; by contrast, the inclusion $H^1(\mathbb{R}^n) \hookrightarrow H^{-1}(\mathbb{R}^n)$ for instance is not compact.

Exercise 17.37. Consider the bounded linear operator

$$\widetilde{\Phi} : H^1(\mathbb{R}^n) \to H^{-1}(\mathbb{R}^n) : u \mapsto u - \frac{1}{4\pi^2}\Delta u,$$

which descends to quotients to define a bounded linear operator $\Phi : H^1(\Omega) \to H^{-1}(\Omega)$. Prove:

(a) $\widetilde{\Phi} : H^1(\mathbb{R}^n) \to H^{-1}(\mathbb{R}^n)$ is a unitary isomorphism.

(b) Let $I : H^{-1}(\Omega) \to \left(\widetilde{H}^1(\Omega)\right)^*$ denote the natural real-linear isomorphism from Exercise 17.12, and denote by $Q : H^{-1}(\mathbb{R}^n) \to H^{-1}(\Omega) = H^{-1}(\mathbb{R}^n)/H^{-1}_{\Omega^c}(\mathbb{R}^n)$ the quotient projection. Then the map

$$I \circ Q \circ \widetilde{\Phi}|_{\widetilde{H}^1(\Omega)} : \widetilde{H}^1(\Omega) \to \left(\widetilde{H}^1(\Omega)\right)^*$$

is an isometric real-linear isomorphism. Deduce from this that $\Phi : H^1(\Omega) \to H^{-1}(\Omega)$ restricts to $H_0^1(\Omega)$ as an isomorphism $H_0^1(\Omega) \to H^{-1}(\Omega)$.
Hint: Write down an explicit formula for $IQ\widetilde{\Phi}(u)f$ for $u, f \in \widetilde{H}^1(\Omega)$.

Exercise 17.38. Writing $j : H_0^1(\Omega) \hookrightarrow H^{-1}(\Omega)$ for the natural inclusion, deduce from the formula $\Delta = 4\pi^2(j - \Phi) : H_0^1(\Omega) \to H^{-1}(\Omega)$ and the previous exercises that Δ is a Fredholm operator with index 0.

Exercise 17.39. Prove:

(a) Every $u \in C_0^\infty(\mathbb{R}^n)$ satisfies $-\int_{\mathbb{R}^n} \langle u, \Delta u\rangle\, dm = \int_{\mathbb{R}^n} |\nabla u|^2\, dm.$

(b) For any open subset $\Omega \subset \mathbb{R}^n$, the relation in part (a) extends to $-\langle \Delta u, u\rangle_1 = \|\nabla u\|_{L^2}^2$ for every $u \in \widetilde{H}^1(\Omega)$, where $\langle\ ,\ \rangle_1 : H^{-1}(\mathbb{R}^n) \times H^1(\mathbb{R}^n) \to \mathbb{C}$ denotes the duality pairing in Exercise 17.10.

(c) The operator $\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$ is injective.
Caution: This is not difficult, but since $H^{-1}(\Omega)$ is a quotient, it is slightly more complicated than just assuming $u \in \widetilde{H}^1(\Omega)$ satisfies $\Delta u = 0$ and applying part (b).

(d) The operator $\Delta : H_0^1(\Omega) \to H^{-1}(\Omega)$ is also surjective if $\Omega \subset \mathbb{R}^n$ is bounded.

Part 4: Spectral theory

18. Bounded operators

18.1. Introduction. Spectral theory is often considered to be one of the main selling points of functional analysis. Its development was motivated in the early 20th century by forces external to mathematics: the laws of quantum mechanics were formulated in terms of self-adjoint operators on complex Hilbert spaces, and those laws required assuming as a "postulate" that such operators can always be diagonalized in terms of orthonormal bases of eigenvectors. Strictly speaking, that is false, e.g. by the usual mathematical definition of the word "eigenvector," a self-adjoint operator on an infinite-dimensional Hilbert space need not have any eigenvectors at all. The following example illustrates the problem.

Example 18.1. For complex-valued functions on a nontrivial compact interval $I := [a, b]$, the linear operator

$$L^2(I) \xrightarrow{Q} L^2(I), \qquad (Qf)(x) := xf(x)$$

is easily shown to be bounded and self-adjoint. If $f \in L^2(I)$ is an eigenvector (also called an **eigenfunction** since $L^2(I)$ is a function space) of Q with eigenvalue $\lambda \in \mathbb{C}$, then the equation $xf(x) = \lambda f(x)$ must be satisfied for almost every $x \in I$, implying $(x - \lambda)f(x) = 0$ and thus $f \equiv 0$ almost everywhere. This proves that Q does not have any eigenvalues or eigenvectors.

Incidentally, the example above is not just something artificially cooked up to prove a point. The operator in that example is important in quantum mechanics: it is the self-adjoint operator corresponding to the measurable quantity known as "position" for a particle that we imagine constrained to move inside a 1-dimensional box, and the lack of eigenvectors has a physical interpretation, namely the impossibility of ever knowing the position of such a particle with absolute precision. However, the conventional physicists' perspective does not see any diagonalizability problem in Example 18.1: one can simply define for each $\lambda \in I$ the function

$$e_\lambda(x) := \begin{cases} 1 & \text{if } x = \lambda, \\ 0 & \text{if } x \neq \lambda, \end{cases}$$

then notice that this function satisfies

$$Qe_\lambda = \lambda e_\lambda,$$

and also that arbitrary functions $f : I \to \mathbb{C}$ can be written in the form

$$f = \sum_{\lambda \in I} c_\lambda e_\lambda \tag{18.1}$$

for suitable coefficients $c_\lambda \in \mathbb{C}$, and thus regard $\{e_\lambda\}_{\lambda \in I}$ as an "orthonormal" eigenbasis of the operator Q. Mathematically, this is nonsense: there is no obvious way to define the uncountable sum in (18.1), and the uncountable collection of functions $\{e_\lambda\}_{\lambda \in I}$ definitely does not form an orthonormal basis of the separable Hilbert space $L^2(I)$—in fact, these functions do not even define nontrivial elements of $L^2(I)$, since they each vanish almost everywhere.

It is, however, rarely a good idea for mathematicians to ignore physicists' intuition. There is a way to make rigorous sense of all of this, and from the correct perspective, the self-adjoint operator $Q : L^2(I) \to L^2(I)$ described above is not just diagonalizable, but is *already* diagonalized. Making sense of this will require a slightly non-obvious way of generalizing the word "spectrum" to infinite-dimensional settings, so that the spectrum of a bounded linear operator can, in general, contain more than just its eigenvalues.

CONVENTION 18.2. From now on, the field $\mathbb{K}$ is $\mathbb{C}$ rather than $\mathbb{R}$, so all our vector spaces are *complex* vector spaces. In the finite-dimensional setting, you know the reason for this convention already: by the fundamental theorem of algebra, working over $\mathbb{C}$ ensures that every linear transformation $\mathbb{C}^n \to \mathbb{C}^n$ has a nonempty spectrum, which would not be true if we worked over $\mathbb{R}$. The same thing will turn out to be true in the infinite-dimensional case.

18.2. Resolvent and spectrum. We start with the observation that $\lambda \in \mathbb{C}$ is an eigenvalue of a linear transformation $A : \mathbb{C}^n \to \mathbb{C}^n$ if and only if any of the following conditions hold:

(1) There exists a nonzero vector $v \in \mathbb{C}^n$ such that $Av = \lambda v$;
(2) The linear transformation $\lambda\mathbb{1} - A$ on $\mathbb{C}^n$ has nontrivial kernel;
(3) The linear transformation $\lambda\mathbb{1} - A$ on $\mathbb{C}^n$ is not invertible.

These conditions are all equivalent, but in infinite-dimensional settings, the third condition is in general strictly weaker than the other two.

NOTATION. For a complex Banach space X and a scalar $\lambda \in \mathbb{C}$, we write $\mathbb{1} : X \to X$ for the identity map and abbreviate the bounded linear map $\lambda\mathbb{1} : X \to X$ in the following simply as λ.

DEFINITION 18.3. The **resolvent set** of a bounded linear operator $T \in \mathscr{L}(X)$ on a complex Banach space X is defined as

$$\rho(T) := \left\{ \lambda \in \mathbb{C} \;\middle|\; \lambda - T \in \mathscr{L}(X) \text{ is invertible} \right\}.$$

The **spectrum** of T is the complement in $\mathbb{C}$ of the resolvent set, denoted by

$$\sigma(T) := \mathbb{C}\backslash\rho(T).$$

For $\lambda \in \rho(T)$, the bounded linear operator

$$R_\lambda(T) := (\lambda - T)^{-1} \in \mathscr{L}(X)$$

is called the **resolvent** of T at λ. A number $\lambda \in \sigma(T)$ is called an **eigenvalue** of T if

$$\ker(\lambda - T) \neq \{0\},$$

and this nontrivial kernel is then called an **eigenspace** of T, and its elements **eigenvectors**.

One sometimes also hears the word **eigenfunctions** in place of "eigenvectors" when X is a space of functions.

PROPOSITION 18.4. *The resolvent set $\rho(T) \subset \mathbb{C}$ of a bounded linear operator $T \in \mathscr{L}(X)$ is open, and the spectrum $\sigma(T) \subset \mathbb{C}$ is closed.*

PROOF. The second statement is synonymous with the first, and the first holds because invertibility is an open condition on $\mathscr{L}(X)$, by Corollary 16.10. □

The next example illustrates two phenomena that may conflict with your finite-dimensional intuition:

(1) The set of eigenvalues need not be discrete;
(2) Not every $\lambda \in \sigma(T)$ must be an eigenvalue.

EXAMPLE 18.5. Let ℓ^1 denote the Banach space of absolutely summable sequences $\mathbf{x} = (x_1, x_2, x_3, \dots)$ in $\mathbb{C}$, with $\|\mathbf{x}\|_{\ell^1} := \sum_{n=1}^{\infty} |x_n|$. The surjective operator

$$\ell^1 \xrightarrow{T} \ell^1 : (x_1, x_2, x_3, \dots) \mapsto (x_2, x_3, x_4, \dots)$$

then satisfies $T\mathbf{x} = \lambda\mathbf{x}$ if and only if $x_{n+1} = \lambda x_n$ for every $n \in \mathbb{N}$, hence $x_n = \lambda^{n-1} x_1$, implying

$$\|\mathbf{x}\|_{\ell^1} = |x_1| \sum_{n=1}^{\infty} |\lambda|^{n-1} < \infty \qquad \Leftrightarrow \qquad |\lambda| < 1,$$

and every complex number in the open unit disk is therefore an eigenvalue of T. Since the spectrum $\sigma(T)$ is closed, it must also contain the unit circle, though the calculation above shows that elements on the unit circle are not eigenvalues.

Now identify the dual space of ℓ^1 with the space of bounded sequences ℓ^∞ and consider the transpose of T, which is the injective map

$$\ell^\infty \xrightarrow{T^\top} \ell^\infty : (x_1, x_2, x_3, \dots) \mapsto (0, x_1, x_2, \dots).$$

Injectivity implies that 0 is not an eigenvalue of $T^\top$. If $T^\top\mathbf{x} = \lambda\mathbf{x}$ for some $\lambda \in \mathbb{C}\backslash\{0\}$, then $0 = \lambda x_1$ implies $x_1 = 0$, and $x_1 = \lambda x_2$ then implies $x_2 = 0$, and so forth, leading to the conclusion that $\mathbf{x} = 0$, so there are no eigenvalues at all. However, if $|\lambda| < 1$, we know that there is an eigenvector $\mathbf{x}_\lambda \in \ell^1\backslash\{0\}$ of T, satisfying $T\mathbf{x}_\lambda = \lambda\mathbf{x}_\lambda$. Writing

$$(\mathbf{y}, \mathbf{x}) := \sum_{n=1}^{\infty} y_n x_n$$

for the natural duality pairing of $\mathbf{y} \in \ell^\infty$ with $\mathbf{x} \in \ell^1$, it follows that for every $\mathbf{y} \in \ell^\infty$,

$$\big((\lambda - T^\top)\mathbf{y}, \mathbf{x}_\lambda\big) = (\mathbf{y}, (\lambda - T)\mathbf{x}_\lambda) = 0,$$

thus $\mathbf{x}_\lambda$ defines a nontrivial element in the dual space of ℓ^∞ that annihilates $\operatorname{im}(\lambda - T^\top) \subset \ell^\infty$, implying that $\lambda - T^\top \in \mathscr{L}(\ell^\infty)$ is not surjective. We deduce that in spite of the lack of eigenvalues, $\sigma(T^\top)$ also contains the entire unit disk in $\mathbb{C}$.

18.3. Point and residual spectrum. Qualitatively, there are various things that the condition $\lambda \in \sigma(T)$ can mean.

DEFINITION 18.6. The **point spectrum** of $T \in \mathscr{L}(X)$ is the set of eigenvalues of T, i.e. the set of all $\lambda \in \mathbb{C}$ such that $\ker(\lambda - T)$ is nontrivial. The **residual spectrum** consists of all $\lambda \in \mathbb{C}$ such that $\ker(\lambda - T)$ is trivial but $\operatorname{im}(\lambda - T)$ is not dense in X.

Note that $\sigma(T)$ can also contain elements $\lambda \in \mathbb{C}$ that are in neither the point spectrum nor the residual spectrum: this happens whenever $\lambda - T$ is injective and has dense image but fails to be surjective. The next exercise shows that λ can in this case be understood as an "approximate eigenvalue" of T.

EXERCISE 18.7. Show that if $T \in \mathscr{L}(X)$ and $\lambda \in \sigma(T)$ belongs to neither the point spectrum nor the residual spectrum, then there exists a sequence $x_n \in X$ with $\|x_n\| = 1$ for all n such that $(\lambda - T)x_n \to 0$ as $n \to \infty$.

EXERCISE 18.8. For a Banach space X and $T \in \mathscr{L}(X)$ with transpose $T^\top \in \mathscr{L}(X^*)$, prove:

(a) If λ belongs to the residual spectrum of T, then it is an eigenvalue of $T^\top$.
(b) If λ is an eigenvalue of T, then it is either an eigenvalue or belongs to the residual spectrum of $T^\top$.

We recall that for a Hilbert space $\mathcal{H}$, the **adjoint** $T^* \in \mathscr{L}(\mathcal{H})$ of an operator $T \in \mathscr{L}(\mathcal{H})$ is uniquely determined by the condition $\langle x, Ty \rangle = \langle T^* x, y \rangle$ for all $x, y \in \mathcal{H}$. We call T **self-adjoint** if $T^* = T$.

REMARK 18.9. The Riesz representation theorem identifies a Hilbert space $\mathcal{H}$ with its own dual space $\mathcal{H}^*$ such that for any bounded linear operator $A : \mathcal{H} \to \mathcal{H}$, the transpose $A^\top : \mathcal{H}^* \to \mathcal{H}^*$ gets identified with the adjoint $A^* : \mathcal{H} \to \mathcal{H}$. Exercise 18.8 therefore implies relations between the spectrum of any bounded linear operator on a Hilbert space and that of its adjoint.

EXERCISE 18.10. For a Hilbert space $\mathcal{H}$ and $T \in \mathscr{L}(\mathcal{H})$, show that

$$\sigma(T^*) = \left\{ \bar{\lambda} \in \mathbb{C} \;\middle|\; \lambda \in \sigma(T) \right\}.$$

PROPOSITION 18.11. *For a Hilbert space $\mathcal{H}$, if $T \in \mathscr{L}(\mathcal{H})$ is self-adjoint, then $\sigma(T) \subset \mathbb{R}$, the residual spectrum of T is empty, and any two eigenvectors for distinct eigenvalues are orthogonal.*

PROOF. Suppose $\lambda, \mu \in \mathbb{R}$. Then for $x \in \mathcal{H}$,

$$\begin{aligned} \left\| \left[T - (\lambda + i\mu) \right] x \right\|^2 &= \langle (T - \lambda)x - i\mu x, (T - \lambda)x - i\mu x \rangle \\ &= \|(T - \lambda)x\|^2 + \mu^2 \|x\|^2 - i\mu \langle (T - \lambda)x, x \rangle + i\mu \langle x, (T - \lambda)x \rangle \\ &= \|(T - \lambda)x\|^2 + \mu^2 \|x\|^2, \end{aligned}$$

implying in particular the estimate $\left\| \left[T - (\lambda + i\mu) \right] x \right\| \geqslant \mu \|x\|$. If $\mu \neq 0$, it now follows from Exercise 17.25 that $T - (\lambda + i\mu)$ is injective with closed image. If $\lambda + i\mu \in \sigma(T)$, this means that $\lambda + i\mu$ cannot be an eigenvalue, but must instead belong to the residual spectrum of T. By Exercise 18.8 and Remark 18.9, it follows that $\lambda - i\mu$ is an eigenvalue of $T^* = T$, which is a contradiction since $\lambda - i\mu$ also has a nontrivial imaginary part. We conclude $\sigma(T) \subset \mathbb{R}$, and if $\lambda \in \sigma(T)$ belongs to the residual spectrum, then the argument above implies that λ is also an eigenvalue, which is likewise a contradiction.

What remains to be proved should be familiar from finite-dimensional linear algebra: if $Tv = \lambda v$ and $Tw = \mu w$ with $\lambda \neq \mu$, then

$$(\lambda - \mu)\langle v, w \rangle = \langle \lambda v, w \rangle - \langle v, \mu w \rangle = \langle Tv, w \rangle - \langle v, Tw \rangle = 0,$$

implying $\langle v, w \rangle = 0$. $\square$

18.4. Spectral radius. A word of preparation is required before stating the next result. A large portion of the standard theory of complex-analytic functions can be generalized to the setting of functions

$$\mathbb{C} \supset \Omega \xrightarrow{f} X$$

defined on an open subset $\Omega \subset \mathbb{C}$ and taking values in an arbitrary complex Banach space X. Analyticity, in this setting, can be taken to mean that for every $z_0 \in \Omega$, there exists a radius $r > 0$ such that f is expressible on the open disk $B_r(z_0) \subset \Omega$ about z_0 as an absolutely convergent power series

$$f(z_0 + z) = f(z_0) + \sum_{k=1}^{\infty} z^k a_k, \qquad z \in B_r(0) \subset \mathbb{C},$$

for some $a_k \in X$. It follows from this condition that for any bounded linear functional $\Lambda : X \to \mathbb{C}$, $\Lambda \circ f : \Omega \to \mathbb{C}$ is a complex-analytic function in the usual sense. The next two results apply this with functions valued in the Banach space $\mathscr{L}(X)$ of bounded linear operators.

THEOREM 18.12. *For any complex Banach space X and any $T \in \mathscr{L}(X)$, the function $\mathbb{C} \supset \rho(T) \to \mathscr{L}(X) : \lambda \mapsto R_\lambda(T)$ is complex analytic.*

PROOF. Suppose $\lambda_0 \in \rho(T)$, so $\lambda_0 - T$ is invertible. We recall that for any $A \in \mathscr{L}(X)$ with $\|A\| < 1$, the absolutely convergent series $\sum_{k=0}^{\infty}(-1)^k A^k$ converges to an inverse for $\mathbb{1} + A$. It follows that for $A_0, B \in \mathscr{L}(X)$ with A_0 invertible and $\|B\| < 1/\|A_0^{-1}\|$, $A_0 + B$ is also invertible, with

$$(A_0 + B)^{-1} = (\mathbb{1} + A_0^{-1}B)^{-1}A_0^{-1} = \sum_{k=0}^{\infty}(-1)^k(A_0^{-1}B)^k A_0^{-1}.$$

Applying this with $\mu \in \mathbb{C}$ sufficiently small gives

$$R_{\lambda_0+\mu}(T) = [(\lambda_0 - T) + \mu]^{-1} = \sum_{k=0}^{\infty}(-1)^k \left[R_{\lambda_0}(T)\mu\right]^k R_{\lambda_0}(T) = \sum_{k=0}^{\infty}(-1)^k\mu^k R_{\lambda_0}(T)^{k+1},$$

valid in particular whenever $|\mu| < \frac{1}{\|R_{\lambda_0}(T)\|}$. □

COROLLARY 18.13. *For every bounded linear operator $T \in \mathscr{L}(X)$ on a complex Banach space, $\sigma(T) \neq \varnothing$.*

PROOF. If $\sigma(T) = \varnothing$, then $\mathbb{C} \to \mathscr{L}(X) : \lambda \mapsto R_\lambda(T)$ is a globally-defined analytic function by Theorem 18.12, and we observe

$$R_\lambda(T) = (\lambda - T)^{-1} = \frac{1}{\lambda}\left(\mathbb{1} - \frac{1}{\lambda}T\right)^{-1} \to 0 \qquad \text{as } |\lambda| \to \infty.$$

It follows that for every bounded linear functional $\Lambda : \mathscr{L}(X) \to \mathbb{C}$, $\lambda \mapsto \Lambda(R_\lambda(T))$ is a globally-defined analytic function $\mathbb{C} \to \mathbb{C}$ that decays to zero at infinity; by Liouville's theorem, this function is therefore identically zero. We deduce that $\Lambda(R_\lambda(T)) = 0$ for every $\Lambda \in \mathscr{L}(X)^*$ and $\lambda \in \mathbb{C}$, and thus $R_\lambda(T) = 0$, which is a contradiction since $R_\lambda(T)$ should be invertible. □

REMARK 18.14. The Hahn-Banach theorem implies that if $A \neq 0 \in \mathscr{L}(X)$, then there exists a functional $\Lambda \in \mathscr{L}(X)^*$ with $\Lambda(A) \neq 0$, as was used at the end of the proof above. If A is surjective, however—as was also the case with $R_\lambda(T)$ in the proof above—then we don't really need the full strength of the Hahn-Banach theorem, we only need to know that the dual space X^* is nontrivial. Indeed, assuming the latter, pick a nontrivial functional $\Psi \in X^*$, along with an element $x \in X$ such that $\Psi(x) \neq 0$, and use the surjectivity of A to extract an element $y \in X$ with $Ay = x$. We can then define a functional $\Lambda \in \mathscr{L}(X)^*$ by

$$\Lambda(B) := \Psi(By),$$

which satisfies $\|\Lambda\| \leqslant \|\Psi\| \cdot \|y\|$, and we observe that $\Lambda(A) = \Psi(Ay) = \Psi(x) \neq 0$ by construction.

DEFINITION 18.15. The **spectral radius** of an operator $T \in \mathscr{L}(X)$ is defined as

$$r(T) := \sup_{\lambda \in \sigma(T)} |\lambda| \geqslant 0.$$

We observe: $\lambda - T = \lambda\left(\mathbb{1} - \frac{T}{\lambda}\right)$ is invertible whenever $|\lambda| > \|T\|$, thus:

PROPOSITION 18.16. *The spectral radius of any bounded operator $T \in \mathscr{L}(X)$ satisfies*

$$r(T) \leqslant \|T\|.$$

□

REMARK 18.17. As a closed and bounded subset of $\mathbb{C}$, the spectrum $\sigma(T)$ of a bounded linear operator is always compact, so the supremum in Definition 18.15 can freely be replaced by a maximum.

EXAMPLE 18.18. The matrix $\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ has 0 as its only eigenvalue, demonstrating that the inequality in Proposition 18.16 is sometimes strict.

On the other hand, the matrix in the example above is not self-adjoint, and if you recall that self-adjoint transformations in finite-dimensions always admit orthonormal bases of eigenvectors, you will easily see that $r(T) = \|T\|$ for such transformations. Let us show that the latter also holds in infinite dimensions:

THEOREM 18.19. *For a Hilbert space $\mathcal{H}$ and self-adjoint operator $T \in \mathscr{L}(\mathcal{H})$, $r(T) = \|T\|$.*

The proof of this theorem requires a bit of preparation. We will need the following input from complex analysis (see e.g. the discussion of the Cauchy integral formula in [**BS18**, §5.1.3]): For a complex Banach space X and an analytic function $\mathbb{C} \supset \Omega \xrightarrow{f} X$, f can be written as a convergent power series $f(z) = \sum_{n=0}^{\infty}(z - z_0)^n a_n$ with $a_n \in X$ for all z in the open ball $B_R(z_0) \subset \Omega$ about any given point $z_0 \in \Omega$, where the radius of convergence $R > 0$ is given by

$$R = \frac{1}{\limsup_{n\to\infty} \|a_n\|^{1/n}} = \sup \left\{ r > 0 \;\middle|\; B_r(z_0) \subset \Omega \right\}. \tag{18.2}$$

For $T \in \mathscr{L}(X)$, the function $z \mapsto R_{1/z}(T)$ is well defined and analytic for all $z \in \mathbb{C}$ with $0 < |z| < 1/r(T)$, and for $|z| < 1/\|T\|$, setting $\lambda := 1/z$ and writing

$$R_\lambda(T) = (\lambda - T)^{-1} = \left[\lambda\left(\mathbb{1} - \frac{1}{\lambda}T\right)\right]^{-1} = \frac{1}{\lambda}\sum_{k=0}^{\infty}\frac{1}{\lambda^k}T^k$$

gives rise to the power series formula

$$R_{1/z}(T) = z\sum_{k=0}^{\infty} z^k T^k. \tag{18.3}$$

This function thus extends analytically to $z = 0$ by setting $R_\infty(T) := 0$, but it does not extend to $\left\{|z| = 1/r(T)\right\}$ since the circle of radius $r(T)$ necessarily contains points in $\sigma(T)$. The radius of convergence of (18.3) is therefore precisely $1/r(T)$, so that (18.2) implies the formula

$$r(T) = \limsup_{n\to\infty} \|T^n\|^{1/n}. \tag{18.4}$$

LEMMA 18.20. *For a Hilbert space $\mathcal{H}$ and operator $T \in \mathscr{L}(\mathcal{H})$, $\|T^*T\| = \|T\|^2$.*

PROOF. The duality isomorphism $\mathcal{H} \to \mathcal{H}^* : x \mapsto \langle x, \rangle$ coming from the Riesz representation theorem is an isometry, so we have

$$\|x\| = \sup_{\|y\|=1} |\langle x, y\rangle|$$

for all $x \in \mathcal{H}$, and thus

$$\|T\| = \sup_{\|x\|=1}\|Tx\| = \sup_{\|x\|=\|y\|=1} |\langle Tx, y\rangle|.$$

Using the Cauchy-Schwarz inequality, this implies

$$\|T^*Tx\| = \sup_{\|x\|=\|y\|=1} |\langle T^*Tx, y\rangle| = \sup_{\|x\|=\|y\|=1} |\langle Tx, Ty\rangle| \leqslant \sup_{\|x\|=\|y\|=1} \|Tx\| \cdot \|Ty\| = \|T\|^2.$$

To see that this is actually an equality, choose a sequence $x_n \in \mathcal{H}$ with $\|x_n\| = 1$ such that $\|Tx_n\| \to \|T\|$, and observe that $|\langle Tx_n, Tx_n\rangle| = \|Tx_n\|^2 \to \|T\|^2$. □

COROLLARY 18.21. *For any self-adjoint operator $T \in \mathscr{L}(\mathcal{H})$ on a Hilbert space, $\|T^2\| = \|T\|^2$.* □

Proof of Theorem 18.19. We already know $r(T) \leqslant \|T\|$, so by (18.4), it will suffice to find a sequence $k_n \in \mathbb{N}$ with $k_n \to \infty$ such that $\|T^{k_n}\|^{1/k_n} \to \|T\|$. Take $k_n := 2^n$, so $T^{k_n} = (T^{k_{n-1}})^2$, and applying Corollary 18.21 inductively gives $\|T^{k_n}\| = \|T\|^{k_n}$. □

18.5. Spectra of compact operators. The spectral theory of compact operators on Banach spaces is not so different from the finite-dimensional case. We will see in §19.6 that this fact has useful consequences for the spectral theory of unbounded operators as well.

Theorem 18.22 (Riesz-Schauder theorem). *For a compact operator $K : X \to X$ on a complex Banach space, the spectrum $\sigma(K) \subset \mathbb{C}$ has no limit points except possibly 0, and every $\lambda \in \sigma(K)\backslash\{0\}$ is an eigenvalue of finite multiplicity, meaning $\dim \ker(\lambda - K) < \infty$.*

Proof. For $\lambda \in \mathbb{C}\backslash\{0\}$, the operator $\lambda\mathbb{1} : X \to X$ is an isomorphism and is thus Fredholm with index 0. By Theorem 17.35, it follows that $\lambda - K$ is also Fredholm with index 0, and is thus injective if and only if it is surjective, proving that $\lambda \in \sigma(K)$ if and only if $0 < \dim \ker(\lambda - K) < \infty$.

It remains to prove that if $\lambda_0 \in \sigma(K)\backslash\{0\}$, then there exists a neighborhood $\mathcal{U} \subset \mathbb{C}$ of λ_0 such that $\sigma(K) \cap \mathcal{U} = \{\lambda_0\}$. We will deduce this from the following claim: If $T : X \to X$ is a complex-linear Fredholm operator of index 0, then there exists a neighborhood $\mathcal{U} \subset \mathbb{C}$ of 0 such that one of the following is true:

(1) $T + \mu$ is invertible for all $\mu \in \mathcal{U}\backslash\{0\}$;
(2) $T + \mu$ is not invertible for all $\mu \in \mathcal{U}$.

To see this, assume without loss of generality that T has nontrivial kernel, since otherwise it must be invertible and the claim follows immediately from the fact that invertibility is an open condition. We now borrow the setup from the proof of Theorem 17.19 and write

$$X = V \oplus K = W \oplus C,$$

where $K = \ker(T)$ and $W = \operatorname{im}(T)$, so K and C are finite dimensional, and T restricts to V as a Banach space isomorphism $A : V \to W$. Our two splittings of X into sums of closed subspaces produce a block matrix representation of any bounded linear operator $X \to X$, so in particular, the identity map can be written as

$$\mathbb{1} = \begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} : V \oplus K \to W \oplus C$$

for suitable operators $\alpha, \beta, \gamma, \delta$, and the corresponding representation of $T + \mu$ is then

$$T + \mu = \begin{pmatrix} A + \mu\alpha & \mu\beta \\ \mu\gamma & \mu\delta \end{pmatrix} : V \oplus K \to W \oplus C.$$

Applying the same linear coordinate transformations as in the proof of Theorem 17.19 now makes $T + \mu$ conjugate to the operator

$$\begin{pmatrix} A + \mu\alpha & 0 \\ 0 & \mu\delta - \mu\gamma(A + \mu\alpha)^{-1}\mu\beta \end{pmatrix}.$$

Since $A : V \to W$ is invertible, we can assume $A + \mu\alpha : V \to W$ is also invertible whenever $|\mu|$ is sufficiently small, and in this case, $T + \mu$ will be invertible if and only if

$$\Phi(\mu) := \mu\delta - \mu\gamma(A + \mu\alpha)^{-1}\mu\beta = \mu\delta - \mu^2\gamma(A + \mu\alpha)^{-1}\beta \in \mathscr{L}(K, C)$$

is invertible. The detail to notice now is that for $\mu \in \mathbb{C}$ in a sufficiently small neighborhood of 0, $\Phi(\mu)$ is an analytic function of μ, taking values in the finite-dimensional complex vector space $\mathscr{L}(K, C)$, and it satisfies $\Phi(0) = 0$. Since $\operatorname{ind}(T) = 0$, the spaces K and C have the same dimension $q \in \mathbb{N}$, and we can therefore choose isomorphisms of both with $\mathbb{C}^q$, which identifies Φ with an analytic function valued in the space of matrices $\mathbb{C}^{q \times q}$, and

$$\mu \mapsto \det(\Phi(\mu))$$

is then a complex-valued analytic function that vanishes at $\mu = 0$. It follows that $\det(\Phi(\mu))$ either vanishes identically for μ near 0 or the zero at $\mu = 0$ is isolated, and this proves the claim.

Applying the claim to the Fredholm operator $T := \lambda_0 - K$, we deduce that if the desired result is not true, then *every* other λ in some neighborhood of λ_0 is also an eigenvalue of T. This leads to a contradiction as follows: choose a continuous path $c : [0, \infty) \to \mathbb{C}$ from $c(0) = \lambda_0$ out to infinity, avoiding $0 \in \mathbb{C}$ along the way, and let $t_0 := \sup \{ t \in [0, \infty) \mid c(t) \in \sigma(T) \}$. Since T is bounded and therefore has finite spectral radius, $t_0 < \infty$, and since the spectrum is a closed set, we therefore find an eigenvalue $\lambda := c(t_0) \in \sigma(T)$ that is not isolated in $\sigma(T)$, implying in light of the claim above that $c(t)$ must also be an eigenvalue for $t > t_0$ sufficiently close to t_0. This is a contradiction. $\square$

We now come to the simplest infinite-dimensional generalization of the fact that self-adjoint operators in finite dimensions can always be orthogonally diagonalized. The result for compact self-adjoint operators in infinite dimensions is not so different: the size of the resulting diagonal "matrix" may now be infinite, though the nonzero portion of it will be only *countably* infinite (see Remark 18.24).

THEOREM 18.23 (Hilbert-Schmidt theorem). *Assume $K : \mathcal{H} \to \mathcal{H}$ is a compact self-adjoint operator on a complex Hilbert space $\mathcal{H}$. Then $\mathcal{H}$ admits an orthonormal basis $\{e_\alpha\}_{\alpha \in I}$ such that $K e_\alpha = \lambda_\alpha e_\alpha$ for a collection of real eigenvalues $\lambda_\alpha \in \mathbb{R}$.*

PROOF. Each eigenspace $E_\lambda := \ker(\lambda - K) \subset \mathcal{H}$ is a closed subspace and therefore also a Hilbert space. Since eigenvectors with distinct eigenvalues are orthogonal by Proposition 18.11, we can choose an orthonormal basis for each eigenspace and, in this way, form an orthonormal set spanning some closed subspace $E \subset \mathcal{H}$ with $K(E) \subset E$. We need to show that $E = \mathbb{H}$. Note that by construction, E contains $E_0 = \ker K$.

If $v \in E^\perp \subset \mathcal{H}$, then for all $x \in E$, we have $\langle x, Kv \rangle = \langle Kx, v \rangle = 0$ since $Kx \in E$, implying that Kv also lies in $E^\perp$. This makes $K|_{E^\perp} : E^\perp \to E^\perp$ a compact self-adjoint operator on the Hilbert space $E^\perp$, but with no eigenvalues, so if $E^\perp$ is nontrivial, then by Theorem 18.22, the spectrum of this restricted operator can only be $\{0\}$. Applying Theorem 18.19 then gives

$$\|K|_{E^\perp}\| = r(K|_{E^\perp}) = 0,$$

so $K|_{E^\perp} = 0$, meaning $E^\perp \subset \ker K$, which is a contradiction since $\ker K \subset E$. $\square$

REMARK 18.24. The Riesz-Schauder theorem tells us that $\sigma(K) \backslash \{0\}$ is a discrete set of eigenvalues, each with finite multiplicity, from which it follows that the set $\{ \alpha \in I \mid \lambda_\alpha \neq 0 \}$ is at most countable. If this set is infinite, it can therefore be rewritten as a sequence $\{\alpha_n \in I\}_{n \in \mathbb{N}}$, and it follows that $\lambda_{\alpha_n} \to 0$ as $n \to \infty$. Countability also implies that $(\ker K)^\perp \subset \mathcal{H}$ is separable, and the orthonormal basis of $\ker K \subset \mathcal{H}$ will then be countable if and only if $\mathcal{H}$ is separable.

EXAMPLE 18.25. Given a function $F : \mathbb{T}^n \to \mathbb{R}$ of class L^2 with $F(x) = F(-x)$ for all $x \in \mathbb{T}^n$, we can use convolution to define the operator

$$T : L^2(\mathbb{T}^n) \to L^2(\mathbb{T}^n) : u \mapsto F * u,$$

which can also be characterized in terms of Fourier series via the condition

$$\widehat{Tu}_k = \widehat{F}_k \widehat{u}_k \qquad \text{for } k \in \mathbb{Z}^n.$$

The fact that $F \in L^2(\mathbb{T}^n)$ means $\widehat{F} \in \ell^2$, and thus $|\widehat{F}_k| \to 0$ as $|k| \to \infty$, while the condition $F(-x) = F(x)$ implies

$$\widehat{F}_k = \int_{\mathbb{T}^n} e^{-2\pi i k\cdot x} F(x)\, dx = \int_{\mathbb{T}^n} e^{2\pi i k\cdot x} F(-x)\, dx = \overline{\widehat{F}}_k,$$

and thus $\widehat{F}_k \in \mathbb{R}$. Using Parseval's identity, we deduce

$$\langle u, Tv\rangle_{L^2} = \sum_{k\in\mathbb{Z}^n} \langle \hat{u}_k, \widehat{F}_k \hat{v}_k\rangle = \sum_{k\in\mathbb{Z}^n} \langle \widehat{F}_k \hat{u}_k, \hat{v}_k\rangle = \langle Tu, v\rangle_{L^2}$$

for all $u, v \in L^2(\mathbb{T}^n)$, so T is self-adjoint. Exercise 18.26 below shows that T is also compact. One easily checks that $u \in L^2(\mathbb{T}^n)$ is an eigenfunction of T whenever it has only one nontrivial Fourier coefficient, so in particular, the complex exponential functions

$$\{x \mapsto e^{2\pi k\cdot x}\}_{k\in\mathbb{Z}^n}$$

are all eigenfunctions, with $T(e^{2\pi k\cdot x}) = \widehat{F}_k e^{2\pi i k\cdot x}$, and they form an orthonormal basis of $L^2(\mathbb{T}^n)$. Since the eigenvalues are the values of the Fourier coefficients, one sees explicitly from the condition $\lim_{|k|\to\infty} |\widehat{F}_k| = 0$ that every eigenvalue other than 0 has finite multiplicity.

Exercise 18.26. Deduce from the condition $\lim_{|k|\to\infty} |\widehat{F}_k| = 0$ that the convolution operator $u \mapsto F * u$ in Example 18.25 is a compact operator $L^2(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$.
Hint: For inspiration, look again at the proof that the inclusions $H^s(\mathbb{T}^n) \hookrightarrow H^t(\mathbb{T}^n)$ for $s > t$ are compact.

Example 18.27. Recall from §17.7 that the Laplace operator $\Delta : H^1_0(\Omega) \to H^{-1}(\Omega)$ is an isomorphism for any bounded open subset $\Omega \subset \mathbb{R}^n$, and the inclusion $j : H^1_0(\Omega) \hookrightarrow L^2(\Omega)$ is compact. The operator

$$T := j\Delta^{-1}|_{L^2(\Omega)} : L^2(\Omega) \to L^2(\Omega)$$

is then compact due to the compactness of j. It is also injective, meaning it does not have 0 as an eigenvalue, and a routine exercise in integration by parts shows that it is self-adjoint. It follows that $L^2(\Omega)$ has an orthonormal basis $\{f_k \in L^2(\Omega)\}_{k\in\mathbb{N}}$ with $Tf_k = \mu_k f_k$ for a sequence of real eigenvalues $\mu_k \in \mathbb{R}\backslash\{0\}$ with $\mu_k \to 0$. Now observe: $Tf_k = j\Delta^{-1} f_k = \mu_k f_k$ is equivalent to the condition that $\mu_k f_k \in H^1_0(\Omega)$ and $\Delta(\mu_k f_k) = \mu_k \Delta f_k = f_k$, or in other words,

$$\Delta f_k = \lambda_k f_k \qquad \text{where} \qquad \lambda_k := \frac{1}{\mu_k} \in \mathbb{R}\backslash\{0\} \text{ satisfies } |\lambda_k| \to \infty.$$

Corollary 18.28. *For any bounded open set $\Omega \subset \mathbb{R}^n$, $L^2(\Omega)$ admits an orthonormal basis consisting of functions in $H^1_0(\Omega)$ that are eigenfunctions of the Laplace operator, i.e. they satisfy $\Delta f_k = \lambda_k f_k$, with nonzero real eigenvalues λ_k that accumulate only at ∞.* □

This result about the Laplace operator gives a foretaste of the spectral theory of *unbounded* self-adjoint operators, in particular those with "compact resolvent," which we'll study in §19.6.

The diagonalization of compact self-adjoint operators also has the following application to more general compact operators on Hilbert spaces, beyond those that are self-adjoint:

Theorem 18.29 (Singular value decomposition). *Assume $\mathcal{H}$ is an infinite-dimensional Hilbert space and $K : \mathcal{H} \to \mathcal{H}$ is a compact operator. Then there exist (not necessarily complete) countable orthonormal sets $\{\varphi_n \in \mathcal{H}\}_{n\in\mathbb{N}}$ and $\{\psi_n \in \mathcal{H}\}_{n\in\mathbb{N}}$, and nonnegative numbers $\{\lambda_n \geqslant 0\}_{n\in\mathbb{N}}$ with $\lambda_n \to 0$ such that*

$$Kf = \sum_{n=1}^{\infty} \lambda_n \langle \psi_n, f\rangle \varphi_n$$

for every $f \in \mathcal{H}$.

Proof. The operator $K^*K : \mathcal{H} \to \mathcal{H}$ is compact and self-adjoint, so by Theorem 18.23 and Remark 18.24, there exists a countable orthonormal set $\{\psi_n \in \mathcal{H}\}_{n\in\mathbb{N}}$ spanning the closed subspace $(\ker K)^\perp = (\ker K^*K)^\perp \subset \mathcal{H}$ such that $K^*K\psi_n = \mu_n\psi_n$ for each $n \in \mathbb{N}$, with eigenvalues μ_n given explicitly by

$$\mu_n = \langle \psi_n, K^*K\psi_n\rangle = \|K\psi_n\|^2 > 0.$$

Define $\lambda_n := \sqrt{\mu_n}$ and let $\varphi_n := \frac{1}{\lambda_n}K\psi_n$. Then for $m, n \in \mathbb{N}$,

$$\langle \varphi_n, \varphi_m\rangle = \frac{1}{\lambda_m\lambda_n}\langle K\psi_m, K\psi_n\rangle = \frac{1}{\lambda_m\lambda_n}\langle \psi_m, K^*K\psi_n\rangle = \frac{\mu_n}{\lambda_m\lambda_n}\langle \psi_m, \psi_n\rangle = \begin{cases} 0 & \text{if } m \neq n, \\ 1 & \text{if } m = n, \end{cases}$$

so $\{\varphi_n \in \mathcal{H}\}_{n\in\mathbb{N}}$ is another orthonormal set, and we observe that since $\{\psi_n \in \mathcal{H}\}_{n\in\mathbb{N}}$ spans $(\ker K)^\perp$, $\{\varphi_n \in \mathcal{H}\}_{n\in\mathbb{N}}$ consequently spans $\operatorname{im} K \subset \mathcal{H}$. One now checks that for any $f \in \mathcal{H}$, $\langle \varphi_n, Kf\rangle = \lambda_n\langle \psi_n, f\rangle$, and the stated formula for K follows. □

18.6. Spectral representations. We now consider self-adjoint operators that are not compact. We saw a revealing example already at the beginning of this section: On a compact interval $I := [a, b]$, the quantum-mechanical "position" operator $Qf(x) := xf(x)$ is a self-adjoint bounded linear operator $Q : L^2(I) \to L^2(I)$ with no eigenvalues. Now that we have the correct definition of the spectrum $\sigma(Q)$, however, it is easy to show that every $\lambda \in I$ is in it, as every such number is an "approximate" eigenvalue of Q: indeed, using cutoff functions, one finds a sequence $f_n : I \to [0, \infty)$ with $\|f_n\|_{L^2} = 1$ and $\operatorname{supp}(f_n) \subset (\lambda - 1/n, \lambda + 1/n)$, which then satisfies

$$\|(Q-\lambda)f_n\|_{L^2}^2 = \int_a^b (x-\lambda)^2|f_n(x)|^2\,dx = \int_{\lambda-1/n}^{\lambda+1/n} (x-\lambda)^2|f_n(x)|^2\,dx \leqslant \frac{1}{n^2}\|f_n\|_{L^2}^2 = \frac{1}{n^2} \to 0$$

as $n \to \infty$. It follows that there is no estimate of the form $\|(Q-\lambda)f\|_{L^2} \geqslant c\|f\|_{L^2}$, so by Exercise 17.25, $\operatorname{im}(Q-\lambda) \subset L^2(I)$ is not closed, and λ therefore belongs to $\sigma(Q)$.

So, how can an operator like Q, with no eigenvalues, be diagonalized? This is a trick question, in fact, because Q is of a very special class, defined via multiplication with a scalar-valued function; the definition below makes operators of this type the natural infinite-dimensional generalization of the notion of a diagonal matrix.

Remark 18.30. For technical reasons (see Exercise 18.34 below), the measure space (X, μ) appearing in the following definition needs to be **semifinite**, meaning that every measurable set $A' \subset X$ with $\mu(A') = \infty$ contains another measurable set $A \subset A'$ with $0 < \mu(A) < \infty$. A measure space is easily shown to have this property if it is σ-finite, which covers most cases that are of interest in practice. An example that is semifinite but not σ-finite is the disjoint union of an uncountable collection of finite measure spaces; this type of example can arise in the spectral theory of operators on nonseparable Hilbert spaces (see Remark 18.54).

Definition 18.31. A **spectral representation** of an operator $A \in \mathscr{L}(\mathcal{H})$ on a complex Hilbert space $\mathcal{H}$ consists of a semifinite measure space (X, μ), a bounded measurable function $F : X \to \mathbb{C}$, and a unitary isomorphism

$$\mathcal{H} \xrightarrow{U} L^2(X, \mu)$$

such that $UAU^{-1} =: T_F$ is the "multiplication" operator

$$L^2(X, \mu) \xrightarrow{T_F} L^2(X, \mu) : u \mapsto Fu.$$

As in finite dimensions, an isomorphism $U : \mathcal{H}_0 \to \mathcal{H}_1$ between two inner product spaces is called **unitary** if it satisfies

$$\langle Ux, Uy\rangle = \langle x, y\rangle \qquad \text{for all} \qquad x, y \in \mathcal{H}_0.$$

EXERCISE 18.32. Show that an isomorphism $U : \mathcal{H}_0 \to \mathcal{H}_1$ between inner product spaces is unitary if and only if it satisfies $\|Ux\| = \|x\|$ for all $x \in \mathcal{H}_0$.
Hint: Consider the norms of $x + y$ and $x + iy$ and their images under U for arbitrary $x, y \in \mathcal{H}$.

EXAMPLE 18.33. Suppose $A : \mathbb{C}^n \to \mathbb{C}^n$ is a linear map admitting an orthonormal basis of eigenvectors $v_1, \ldots, v_n \in \mathbb{C}^n$ with $Av_j = \lambda_j v_j$ for some $\lambda_j \in \mathbb{C}$. Then there exists a unitary transformation $U : \mathbb{C}^n \to \mathbb{C}^n$ such that UAU^{-1} is a diagonal matrix with the eigenvalues λ_j on the diagonal. Define the measure space (X, μ) where $X := \{1, \ldots, n\}$ and μ is the counting measure; then the space of complex-valued L^2-functions on (X, μ) is unitarily equivalent to $\mathbb{C}^n$ via an isomorphism that identifies $j \in X$ with the jth standard basis vector of $\mathbb{C}^n$ for each $j = 1, \ldots, n$. Under this identification, $UAU^{-1} : L^2(X, \mu) \to L^2(X, \mu)$ becomes the multiplication operator T_F with $F : X \to \mathbb{C}$ defined by $F(j) := \lambda_j$.

Notice that in the setting of a spectral representation, we have $U(\lambda - A)U^{-1} = \lambda - T_F$ for every $\lambda \in \mathbb{C}$, thus

$$\sigma(A) = \sigma(T_F).$$

EXERCISE 18.34. Show that for the multiplication operator $T_F : L^2(X, \mu) \to L^2(X, \mu)$ in Definition 18.31, with (X, μ) a semifinite measure space, the spectrum $\sigma(T_F) \subset \mathbb{C}$ is the so-called **essential range** of $F : X \to \mathbb{C}$, meaning the set

$$\left\{\lambda \in \mathbb{C} \;\middle|\; \mu(F^{-1}(B_\epsilon(\lambda))) > 0 \text{ for all } \epsilon > 0\right\}.$$

For a multiplication operator T_F on $L^2(X, \mu)$, we have

$$\langle u, T_F v\rangle_{L^2} = \langle u, Fv\rangle_{L^2} = \langle \bar{F}u, v\rangle_{L^2} = \langle T_{\bar{F}}u, v\rangle_{L^2},$$

thus

$$T_F^* = T_{\bar{F}}.$$

This implies:

PROPOSITION 18.35. *The multiplication operator $T_F : L^2(X, \mu) \to L^2(X, \mu)$ defined via a bounded measurable function $F : X \to \mathbb{C}$ is self-adjoint if and only if F is real valued, and unitary if and only if F takes values in the unit circle in $\mathbb{C}$.* □

We deduce from Proposition 18.35 that multiplication operators T_F have a special property, namely, they commute with their own adjoints:

$$T_F T_F^* = T_F T_{\bar{F}} = T_{F\bar{F}} = T_{\bar{F}F} = T_{\bar{F}} T_F = T_F^* T_F.$$

DEFINITION 18.36. A bounded linear operator $A \in \mathscr{L}(\mathcal{H})$ on a Hilbert space is called **normal** if $AA^* = A^*A$.

EXERCISE 18.37. For a complex Hilbert space $\mathcal{H}$, prove:

(a) The following conditions on $T \in \mathscr{L}(\mathcal{H})$ are equivalent:
  (i) T is normal;
  (ii) $T = A + iB$ for two self-adjoint operators $A, B \in \mathscr{L}(\mathcal{H})$ that commute with each other;
  (iii) $\|Tx\| = \|T^*x\|$ for every $x \in \mathcal{H}$.
  Hint: Consider $\|T(x + y)\|^2$ and $\|T(x + iy)\|^2$ for arbitrary $x, y \in \mathcal{H}$.

(b) If $T \in \mathscr{L}(\mathcal{H})$ is normal, then:
  (i) $\|T^2\| = \|T^*T\| = \|T\|^2$
  (ii) The spectral radius of T is $\|T\|$.
  (iii) Every eigenvector of T with eigenvalue $\lambda \in \mathbb{C}$ is also an eigenvector of T^* with eigenvalue $\bar{\lambda}$. *Hint: Consider $\|(\lambda - T)v\|^2$.*

(iv) If $v, w \in \mathcal{H}$ are eigenvectors of T with distinct eigenvalues, then $\langle v, w\rangle = 0$.
(v) If T is also compact, then $\mathcal{H}$ admits an orthonormal basis consisting of eigenvectors of T.

(c) If $T \in \mathscr{L}(\mathcal{H})$ is *unitary* (meaning $T^*T = TT^* = \mathbb{1}$), then its spectrum is contained in the unit circle $\{|\lambda| = 1\} \subset \mathbb{C}$.
Hint: Show $\|T\| = \|T^{-1}\| = 1$, and use the fact that operators with distance less than 1 *from the identity map are invertible.*

Unitary isomorphisms $U : \mathcal{H}_0 \to \mathcal{H}_1$ are inverse to their adjoints, so for any $A \in \mathscr{L}(\mathcal{H})$, one has

$$(UAU^{-1})^* = UA^*U^{-1},$$

and A and A^* will then commute if and only if UAU^{-1} and UA^*U^{-1} commute, showing that A is self-adjoint or normal if and only if the same is true of UAU^{-1}. This observation reveals some bad news: not every bounded linear operator $A \in \mathscr{L}(\mathcal{H})$ on a Hilbert space can admit a spectral representation, because the existence of such a representation implies that A is normal.

Here is the good news:

Theorem 18.38 (Spectral theorem for bounded normal operators). *Every normal operator $A \in \mathscr{L}(\mathcal{H})$ on a complex Hilbert space $\mathcal{H}$ admits a spectral representation; moreover, if $\mathcal{H}$ is separable, then the measure space (X, μ) in the representation can be chosen such that $\mu(X) < \infty$.*

Remark 18.39. We will give a detailed proof of Theorem 18.38 only in the separable case; the nonseparable case requires Zorn's lemma, but is otherwise based on the same set of ideas (see Remark 18.54 for a sketch). In the absence of separability, one cannot generally arrange $\mu(X) < \infty$, nor even σ-finiteness.

The proof of Theorem 18.38 has a lot of ingredients, some of which are interesting results in themselves, and we will introduce them gradually in the next few subsections. For technical reasons, it is easier to carry out most steps of the proof specifically for *self-adjoint* operators—the case where the function $F : X \to \mathbb{C}$ takes real values—but we will generalize the main result to normal operators in §18.15, and having done so, the intermediate steps then become valid for normal operators as well. The strategy of the proof in the self-adjoint case can be summarized as follows:

(1) For $A \in \mathscr{L}(\mathcal{H})$ self-adjoint and any $x \in \mathcal{H}$, there is a positive linear functional

$$C^0(\sigma(A)) \xrightarrow{\Lambda} \mathbb{C} : f \mapsto \langle x, f(A)x\rangle,$$

where $f(A) \in \mathscr{L}(\mathcal{H})$ is an operator uniquely determined by the obvious definition of $P(A) \in \mathscr{L}(\mathcal{H})$ for polynomial functions $P : \mathbb{R} \to \mathbb{C}$. This functional is also *positive*, meaning it satisfies $\Lambda(f) \geqslant 0$ whenever $f \geqslant 0$.

(2) According to the Riesz-Markov theorem, positive linear functionals on spaces of bounded continuous functions correspond to measures, so there exists a measure μ_x on $\sigma(A) \subset \mathbb{R}$ such that

$$\Lambda(f) = \int_{\sigma(A)} f\, d\mu_x \qquad \text{for all } f \in C^0(\sigma(A)).$$

We will construct the measure space (X, μ) out of $(\sigma(A), \mu_{x_i})$ for some countable subset $\{x_1, x_2, x_3, \dots\} \subset \mathcal{H}$ whose existence is guaranteed by the separability of $\mathcal{H}$.

The Riesz-Markov theorem (which is also sometimes called the Riesz representation theorem) is a nice result characterizing the dual space of C^0 on suitable topological spaces with Borel measures. Its proof would require a larger digression into measure theory than seemed reasonable to include in this course, so it will mostly have to be taken as a black box, but the missing details can be found e.g. in [**Sal16**, §3.3].

18.7. The continuous functional calculus. The spectrum of any bounded linear operator $A : \mathcal{H} \to \mathcal{H}$ is a closed and bounded subset $\sigma(A) \subset \mathbb{C}$, and is therefore compact, so continuous functions on $\sigma(A)$ are always bounded. We will denote by $C^0(\sigma(A))$ the space of continuous functions $\sigma(A) \to \mathbb{C}$, regarded as a Banach space with the usual C^0-norm.

The following definition is a precursor to the more general notion of a *Borel* functional calculus, to be discussed in §18.11. The existence of the latter will eventually be deduced from the spectral theorem, but our proof of the spectral theorem itself will make use of the precursor.

DEFINITION 18.40. A **continuous functional calculus** for a bounded linear operator $A \in \mathscr{L}(\mathcal{H})$ on a complex Hilbert space is a bounded linear map

$$C^0(\sigma(A)) \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A)$$

satisfying the following conditions:

(i) $(fg)(A) = f(A)g(A)$ and $\bar{f}(A) = f(A)^*$ for all $f, g \in C^0(\sigma(A))$;
(ii) For $f \equiv 1$, $f(A) = \mathbb{1} \in \mathscr{L}(\mathcal{H})$;
(iii) For the function $f(z) := z$ on $\sigma(A)$, $f(A) = A$.
(iv) For every $f \in C^0(\sigma(A))$, $\sigma(f(A)) = f(\sigma(A)) \subset \mathbb{C}$, and if $Ax = \lambda x$ for some $x \in \mathcal{H}$ and $\lambda \in \mathbb{C}$, then $f(A)x = f(\lambda)x$;
(v) If $f \in C^0(\sigma(A))$ takes nonnegative real values, then $\langle x, f(A)x\rangle \geqslant 0$ for all $x \in \mathcal{H}$;
(vi) $\|f(A)\| = \|f\|_{C^0}$ for all $f \in C^0(\sigma(A))$.

A few easy observations: If A has a continuous functional calculus, then all of the operators $f(A)$ for $f \in C^0(\sigma(A))$ commute with each other, since $fg = gf$ for $f, g \in C^0(\sigma(A))$. It follows that all of them are normal, since $f(A)^* = \bar{f}(A)$, and since this applies in particular for the function $f(z) = z$, A itself must therefore be normal. Moreover, the operator $f(A)$ is self-adjoint if and only if $f : \sigma(A) \to \mathbb{C}$ is real valued.

REMARK 18.41. Conditions (i) and (ii) make $f \mapsto f(A)$ into a so-called *algebraic $*$-homomorphism.*

LEMMA 18.42. *If a continuous functional calculus for $A \in \mathscr{L}(\mathcal{H})$ exists, then it is uniquely determined by conditions (i)–(iii) and the assumption that $C^0(\sigma(A)) \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A)$ is bounded.*

PROOF. Since $\sigma(A) \subset \mathbb{C}$ is compact, the Weierstrass approximation theorem tells us that every $f \in C^0(\sigma(A))$ can be approximated uniformly well on $\sigma(A)$ by a polynomial function $\mathbb{C} \to \mathbb{C} : x + iy \mapsto P(x, y)$ of the two real variables $x, y \in \mathbb{R}$, thus the functional calculus is uniquely determined by its definition on polynomials of this form. For $z = x + iy \in \mathbb{C}$, x and y can both be expressed as linear combinations of z and $\bar{z}$, thus the polynomials in question can also be written as finite summations $P(z) = \sum_{j,k} a_{j,k} z^j \bar{z}^k$ with complex coefficients $a_{j,k}$, and conditions (i)–(iii) then uniquely determine $P(A)$ as

$$P(A) = \sum_{j,k} a_{j,k} A^j (A^*)^k \in \mathscr{L}(\mathcal{H}) \tag{18.5}$$

□

We will see later that every bounded normal operator admits a continuous functional calculus; this can be deduced from the spectral theorem. For the moment, we'll settle for dealing with self-adjoint operators.

LEMMA 18.43. *Every bounded self-adjoint operator $A \in \mathscr{L}(\mathcal{H})$ admits a unique continuous functional calculus.*

PROOF. We have uniqueness already from Lemma 18.43, but the assumption that A is self-adjoint simplifies the discussion in one imporant respect. Since $\sigma(A) \subset \mathbb{R}$, every $f \in C^0(\sigma(A))$ can be approximated uniformly well on $\sigma(A)$ by a polynomial function $P(x) = \sum_{k=0}^n a_k x^k$ of *one* real variable $x \in \mathbb{R}$, with coefficients $a_k \in \mathbb{C}$; indeed, for polynomials in this form, properties (i)–(iii) dictate that we define

$$P(A) := \sum_{k=0}^n a_k A^k \in \mathscr{L}(\mathcal{H}).$$

There is a unique extension of the map $f \mapsto f(A)$ to the rest of $C^0(\sigma(A))$ if the linear operator $f \mapsto f(A)$ defined in this way on the space of polynomials is bounded, i.e. it will follow from the identity

$$\|P(A)\| = \|P\|_{C^0} = \sup_{\lambda \in \sigma(A)} |P(\lambda)| \tag{18.6}$$

for all polynomials $P : \mathbb{R} \to \mathbb{C}$. To prove (18.6), notice first that since $P(A)$ is normal, Exercise 18.37 implies

$$\|P(A)\| = r(P(A)) = \sup_{\lambda \in \sigma(P(A))} |\lambda|,$$

and the result now follows from the claim that

$$\sigma(P(A)) = P(\sigma(A)).$$

To see this, suppose $\lambda \in \sigma(A)$, and note that since $x \mapsto P(x) - P(\lambda)$ is a polynomial vanishing at $x = \lambda$, we can write

$$P(x) - P(\lambda) = (x - \lambda)Q(x) = Q(x)(x - \lambda)$$

for some polynomial Q. This implies

$$P(A) - P(\lambda) = (A - \lambda)Q(A) = Q(A)(A - \lambda),$$

and since $A - \lambda$ is not both injective and surjective, the same is therefore true of $P(A) - P(\lambda)$, proving $P(\lambda) \in \sigma(P(A))$. Conversely, suppose $\mu \in \sigma(P(A))$, and write

$$P(x) - \mu = a(x - \lambda_1) \dots (x - \lambda_r)$$

for some $a, \lambda_1, \dots, \lambda_r \in \mathbb{C}$. Then $P(A) - \mu = a(A - \lambda_1) \dots (A - \lambda_r)$, so the non-invertibility of $P(A) - \mu$ implies that at least one of the $A - \lambda_i$ must also fail to be invertible, meaning $\lambda_i \in \sigma(A)$, and since $P(\lambda_i) = \mu$, this proves $\mu \in P(\sigma(A))$.

With (18.6) established, density now implies that the map $f \mapsto f(A)$ determined by conditions (i)–(iii) extends uniquely from the space of polynomial functions to define a bounded linear map $C^0(\sigma(A)) \to \mathscr{L}(\mathcal{H})$, and it also automatically satisfies condition (vi) as a result of (18.6). We showed above that $\sigma(f(A)) = f(\sigma(A))$ holds whenever f is a polynomial function, and it is straightforward to verify for polynomials that $f(A)x = f(\lambda)x$ also holds whenever $Ax = \lambda x$; we leave it as an exercise to show by approximation that condition (iv) then holds also for all continuous functions on $\sigma(A)$. Finally, suppose $f : \sigma(A) \to [0, \infty)$, so $f = g^2$ for some real-valued continuous function $g : \sigma(A) \to \mathbb{R}$, thus $f(A) = g(A)^2$, where $g(A)$ is self-adjoint as a result of g being real-valued. This implies

$$\langle x, f(A)x \rangle = \langle x, g(A)g(A)x \rangle = \langle g(A)x, g(A)x \rangle \geqslant 0.$$

□

EXAMPLE 18.44. If $A \in \mathscr{L}(\mathcal{H})$ admits a continuous functional calculus and $\lambda \in \mathbb{C} \backslash \sigma(A)$, then the function

$$g(x) := \frac{1}{\lambda - x}$$

belongs to $C^0(\sigma(A))$, and $g \cdot (\lambda - x) = 1$ then implies $g(A)(\lambda - A) = \mathbb{1}$, thus $g(A) = R_\lambda(A)$. We can now extract from Lemma 18.43 the formula

$$\|R_\lambda(A)\| = \|g\|_{C^0} = \sup_{\mu \in \sigma(A)} \frac{1}{|\lambda - \mu|},$$

or to put it in more geometrically transparent terms:

PROPOSITION 18.45. *For any operator $A \in \mathscr{L}(\mathcal{H})$ admitting a continuous functional calculus and any $\lambda \in \rho(A)$,*

$$\|R_\lambda(A)\| = \frac{1}{\operatorname{dist}(\lambda, \sigma(A))}.$$

□

As it stands, Proposition 18.45 has been shown to hold for all bounded self-adjoint operators, but its validity will be extended to normal operators once we have derived the existence of a continuous functional calculus from the spectral theorem.

REMARK 18.46. Most steps in the proof of Lemma 18.43 go through with minor modifications if A is assumed normal but not self-adjoint, but the one that does not is the proof that $\sigma(P(A)) = P(\sigma(A))$, in which our argument required P to be expressible as a polynomial in only one variable. On the other hand, it is not difficult to show that for all normal $A \in \mathscr{L}(\mathcal{H})$ and all polynomials of the form $P(z) = \sum_{j,k} a_{j,k} z^k \bar{z}^k$, $Ax = \lambda x$ implies $P(A)x = P(\lambda)x$; the case $P(z) := \bar{z}$ is covered by Exercise 18.37(b)(iii), and one can then argue the rest by induction on the degree of the polynomial.

EXERCISE 18.47. For a normal operator $A \in \mathscr{L}(\mathcal{H})$, prove the following without assuming the existence of a continuous functional calculus or spectral representation:

(a) A has no residual spectrum, and $\lambda \in \sigma(A)$ if and only if there exists a sequence $x_n \in \mathcal{H}$ with $\|x_n\| = 1$ for all n and $(\lambda - A)x_n \to 0$.
(b) If $P(A)$ is defined by (18.5) for any polynomial function P of z and $\bar{z}$, then $P(\sigma(A)) \subset \sigma(P(A))$.

18.8. Spectral measures. Assume in the following that $A \in \mathscr{L}(\mathcal{H})$ admits a continuous functional calculus. For the next step in proving the spectral theorem, we consider the linear functional

$$C^0(\sigma(A)) \xrightarrow{\Lambda} \mathbb{C} : f \mapsto \langle x, f(A)x \rangle,$$

which is bounded since

$$|\langle x, f(A)x \rangle \leqslant \|x\|^2 \cdot \|f(A)\| = \|x\|^2 \cdot \|f\|_{C^0}.$$

Thanks to property (v) of the continuous functional calculus, this functional is also *positive*, meaning that its values on nonnegative real-valued functions are always nonnegative.

Here is the tool we need from measure theory:

THEOREM 18.48 (Riesz-Markov theorem). *For X a compact Hausdorff topological space and $\Lambda : C(X) \to \mathbb{C}$ a positive bounded linear functional on the space $C(X)$ of continuous functions with the* sup*-norm, there exists a unique regular finite measure μ defined on the Borel sets of X such that*

$$\Lambda(f) = \int_X f \, d\mu \qquad \text{for all} \qquad f \in C(X).$$

SKETCH OF THE PROOF. The word "regular" means that μ will be uniquely determined by its value on compact sets $K \subset X$. One then needs to set

$$\mu(K) := \inf \left\{ \Lambda(f) \;\middle|\; f : X \to [0, \infty) \text{ continuous with } f \geqslant \chi_K \right\},$$

and check that this defines a measure on X. We refer to [Sal16, §3.3] for the details. □

COROLLARY 18.49 (spectral measures). *For each $x \in \mathcal{H}$ and each operator $A \in \mathscr{L}(\mathcal{H})$ admitting a continuous functional calculus, there exists a unique finite regular Borel measure μ_x on the spectrum $\sigma(A) \subset \mathbb{C}$ such that for every $f \in C^0(\sigma(A))$,*

$$\langle x, f(A)x \rangle = \int_{\sigma(A)} f \, d\mu_x.$$

□

18.9. Diagonalization. We now complete the proof of Theorem 18.38 under the assumption that $A \in \mathscr{L}(\mathcal{H})$ admits a continuous functional calculus, so it will apply in particular whenever A is self-adjoint. For any normal operator A, let us call an element $x \in \mathcal{H}$ **cyclic** for A if the subspace of $\mathcal{H}$ spanned by all elements of the form $A^m (A^*)^n x$ for integers $m, n \geqslant 0$ is dense.

LEMMA 18.50. *If $A \in \mathscr{L}(\mathcal{H})$ preserves a linear subspace $W \subset \mathcal{H}$, then A^* preserves $W^\perp \subset \mathcal{H}$. In particular, if A and A^* both preserve W, then they also both preserve $W^\perp$.*

PROOF. If $v \in W^\perp$, then for all $w \in W$,

$$\langle A^* v, w \rangle = \langle v, Aw \rangle = 0$$

since $Aw \in W$, which proves $A^* v \in W^\perp$. □

LEMMA 18.51. *If $A \in \mathscr{L}(\mathcal{H})$ admits a continuous functional calculus and $x \in \mathcal{H}$ is cyclic for A, then A admits a spectral representation on the measure space*

$$(X, \mu) := (\sigma(A), \mu_x)$$

and is identified with the multiplication operator $T_F : L^2(X, \mu) \to L^2(X, \mu)$ for the function $F(\lambda) := \lambda$.

PROOF. The image of the bounded linear operator

$$C^0(\sigma(A)) \xrightarrow{T} \mathcal{H} : f \mapsto f(A)x$$

contains $A^m (A^*)^n x$ for every $m, n \geqslant 0$ and is thus dense by hypothesis, and we have

$$\|Tf\|^2 = \langle f(A)x, f(A)x \rangle = \langle x, \bar{f}(A) f(A) x \rangle = \langle x, |f|^2(A)x \rangle = \int_{\sigma(A)} |f|^2 \, d\mu_x = \|f\|_{L^2}^2.$$

Since $C^0(\sigma(A))$ is dense in $L^2(\sigma(A), \mu_x)$, it follows that T extends uniquely to an isometry $L^2(\sigma(A), \mu_x) \to \mathcal{H}$, which necessarily has closed image and is therefore also surjective; by Exercise 18.32, it is then a unitary isomorphism. Defining

$$U := T^{-1} : \mathcal{H} \to L^2(\sigma(A), \mu_x),$$

it remains to show that for the function $F(\lambda) := \lambda$ on $\sigma(A)$, $UAU^{-1} f = Ff$ for every $f \in L^2(\sigma(A), \mu_x)$. By density, it will suffice to prove that this holds for every $f \in C^0(\sigma(A))$, and indeed, when f is continuous, the properties of the continuous functional calculus imply

$$T(Ff) = (Ff)(A)x = F(A)f(A)x = Af(A)x = ATf,$$

which is equivalent to $UAU^{-1} f = Ff$. □

LEMMA 18.52. *If the Hilbert space $\mathcal{H}$ is separable and $A \in \mathscr{L}(\mathcal{H})$ is normal, then*

$$\mathcal{H} = \bigoplus_{n=1}^{N} \mathcal{H}_n$$

for some $N \in \mathbb{N} \cup \{\infty\}$ and closed subspaces $\mathcal{H}_n \subset \mathcal{H}$ satisfying:

(i) *$\mathcal{H}_m \perp \mathcal{H}_n$ for all $m \neq n$;*

(ii) *$A(\mathcal{H}_n) \subset \mathcal{H}_n$ and $A^*(\mathcal{H}_n) \subset \mathcal{H}_n$ for all n;*
(iii) *For every n, $\mathcal{H}_n$ contains an element x_n that is cyclic for $A|_{\mathcal{H}_n} \in \mathscr{L}(\mathcal{H}_n)$.*

PROOF. Suppose $\{y_1, y_2, \dots\} \subset \mathcal{H}$ is a countable dense subset. Set $x_1 := y_1$ and let $\mathcal{H}_1$ be the span of $\{A^m(A^*)^n x_1\}_{m,n \geqslant 0}$, which is preserved by both A and A^* since A is normal. If $\mathcal{H}_1 \neq \mathcal{H}$, set $x_2' := y_j$ for the smallest value of $j \in \mathbb{N}$ such that $y_j \notin \mathcal{H}_1$, and note that $x_2' \in x_2 + \mathcal{H}_1$ for a unique $x_2 \in \mathcal{H}_1^\perp$. We then define $\mathcal{H}_2$ as the span of $\{A^m(A^*)^n x_2\}_{m,n \geqslant 0}$, which is also preserved by both A and A^*, and note that by Lemma 18.50, A and A^* also preserve $\mathcal{H}_1^\perp$, implying $\mathcal{H}_2 \perp \mathcal{H}_1$. Notice also that $y_j = x_2' \in \mathcal{H}_1 \oplus \mathcal{H}_2$ by construction.

Now if $\mathcal{H}_1 \oplus \mathcal{H}_2 \neq \mathcal{H}$, we set $x_3' := y_j$ for the smallest j such that $y_j \notin \mathcal{H}_1 \oplus \mathcal{H}_2$, and then repeat the construction above to define $\mathcal{H}_3 \subset (\mathcal{H}_1 \oplus \mathcal{H}_2)^\perp$, and so forth. The process may or may not ever terminate; if it does not, then it produces an infinite sequence of closed subspaces $\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3, \dots$ such that each y_j is contained in a direct sum of finitely many of them, and since the points y_j are dense, this proves $\mathcal{H} = \bigoplus_{n=1}^\infty \mathcal{H}_n$. □

LEMMA 18.53. *Theorem 18.38 holds whenever $\mathcal{H}$ is separable and $A \in \mathscr{L}(\mathcal{H})$ admits a continuous functional calculus; in particular, it holds if A is self-adjoint.*

PROOF. Assuming $\mathcal{H}$ is separable, we can write $\mathcal{H} = \bigoplus_{n=1}^N \mathcal{H}_n$ and $x_n \in \mathcal{H}_n$ as in Lemma 18.52, where $N \in \mathbb{N} \cup \{\infty\}$, and rescale the elements x_n so that without loss of generality, $\|x_n\| = \frac{1}{2^n}$. Then

$$\mu_{x_n}(\sigma(A)) = \int_{\sigma(A)} 1\, d\mu_{x_n} = \langle x_n, 1(A)x_n \rangle = \langle x_n, x_n \rangle = \frac{1}{2^{2n}},$$

implying $\sum_{n=1}^N \mu(x_n) < \infty$. Using Lemma 18.51 to identify $\mathcal{H}_n$ with $L^2(\sigma(A), \mu_{x_n})$ for each n, we then obtain

$$\mathcal{H} = \bigoplus_{n=1}^N \mathcal{H}_n \cong \bigoplus_{n=1}^N L^2(\sigma(A), \mu_{x_n}) \cong L^2\left(\coprod_{n=1}^N (\sigma(A), \mu_{x_n})\right),$$

with $F : \coprod_{n=1}^N \sigma(A) \to \mathbb{R}$ defined as $\lambda \mapsto \lambda$ on each copy of $\sigma(A)$ in the disjoint union.[42] □

REMARK 18.54. In order to remove separability from the hypotheses of Lemma 18.53, one can define a partially ordered set $(S, \prec)$ whose elements are collections of pairwise orthogonal closed subspaces $\{\mathcal{H}_\alpha \subset \mathcal{H}\}_{\alpha \in I}$ satisfying the conditions

(i) $\mathcal{H}_\alpha \perp \mathcal{H}_\beta$ for all $\alpha \neq \beta$;
(ii) $A(\mathcal{H}_\alpha) \subset \mathcal{H}_\alpha$ and $A^*(\mathcal{H}_\alpha) \subset \mathcal{H}_\alpha$ for all α;
(iii) For every α, $\mathcal{H}_\alpha$ contains an element x_α that is cyclic for $A|_{\mathcal{H}_\alpha} \in \mathscr{L}(\mathcal{H}_\alpha)$,

ordered by inclusion. A minor variation on the argument in Lemma 18.52 shows that elements $\{\mathcal{H}_\alpha\}_{\alpha \in I} \in S$ with $\bigoplus_{\alpha \in I} \mathcal{H}_\alpha \neq \mathcal{H}$ are not maximal, and Zorn's lemma thus implies the existence of an orthogonal decomposition into subspaces on which A admits cyclic elements,

$$\mathcal{H} = \bigoplus_{\alpha \in I} \mathcal{H}_\alpha.$$

Lemma 18.51 now identifies A with a direct sum of multiplication operators on $\bigoplus_{\alpha \in I} L^2(\sigma(A), \mu_{x_\alpha})$, which can be identified with the space of L^2 functions on a single measure space (X, μ) defined as the disjoint union of all the finite measure spaces $(\sigma(A), \mu_{x_\alpha})$ for $\alpha \in I$. Since I may be an uncountable set in this situation, there is generally no way to arrange that μ is either finite or σ-finite, though it is semifinite (cf. Remark 18.30) since the measure on each individual component of the disjoint union is finite. For a fuller treatment of the nonseparable case, see e.g. [**BS18**, Theorem 5.82].

[42] The symbol $\coprod$ represents the *disjoint union* of measure spaces; in the situation at hand, the space being defined consists of multiple copies of $\sigma(A)$, one for each value of $n = 1, \dots, N$, with a different measure on each.

18.10. Commutation with normal operators. You might consider this section optional, as its contents will not be needed for extending the spectral theorem to normal operators, nor for anything that we'd like to prove about self-adjoint operators. The theorem below will be necessary, however, in order to establish the most general version of the Borel functional calculus for normal operators in §18.11, as well as for extending the theorem in §18.14 on simultaneous diagonalization to pairs of commuting normal operators that are not self-adjoint.

The following is a special case of a result known as Fuglede's theorem (see [**Fug50**]).

THEOREM 18.55. *For any bounded linear operators $A, B \in \mathscr{L}(\mathcal{H})$ on a Hilbert space $\mathcal{H}$ with A normal, if $AB = BA$, then $A^*B = BA^*$.*

The following proof is due to Rosenblum [**Ros58**]. It uses the exponential of a bounded linear operator $A \in \mathscr{L}(\mathcal{H})$, defined via the absolutely convergent power series

$$e^A := \mathbb{1} + A + \frac{1}{2!}A^2 + \frac{1}{3!}A^3 + \dots .$$

Note that we are not using any functional calculus in this definition, though whenever A does admit a continuous functional calculus, it is straightforward to show that $e^A = f(A)$ for the function $f(z) = e^z$. The same argument as in the finite-dimensional case shows that

$$e^{A+B} = e^A e^B = e^B e^A \qquad \text{whenever} \qquad AB = BA,$$

so in particular,

$$e^{-A} = (e^A)^{-1} \qquad \text{for all} \qquad A \in \mathscr{L}(\mathcal{H}).$$

Moreover, it is readily verified from the power series definition that

$$e^{(A^*)} = (e^A)^* \qquad \text{for all} \qquad A \in \mathscr{L}(\mathcal{H}),$$

and another important observation is that for any $A \in \mathscr{L}(\mathcal{H})$, the function

$$\mathbb{C} \to \mathscr{L}(\mathcal{H}) : z \mapsto e^{zA}$$

is complex analytic.

PROOF OF THEOREM 18.55. Given $AB = BA$, one shows first by induction on $n \geqslant 0$ that $A^n B = BA^n$, and applying this to the partial sums in $e^{\bar{z}A}B$ and $Be^{\bar{z}A}$, it follows that $e^{\bar{z}A}B = Be^{\bar{z}A}$ and hence

$$B = e^{-\bar{z}A} B e^{\bar{z}A} \qquad \text{for all} \qquad z \in \mathbb{C}. \tag{18.7}$$

Now consider the complex-analytic function $F : \mathbb{C} \to \mathscr{L}(\mathcal{H})$ defined by

$$F(z) := e^{zA^*} B e^{-zA^*}.$$

Using (18.7), we can rewrite this as

$$F(z) = e^{zA^*} e^{-\bar{z}A} B e^{\bar{z}A} e^{-zA^*} = e^{zA^* - \bar{z}A} B e^{\bar{z}A - zA^*},$$

where the second equality uses the assumption that A is normal, so that e.g. zA^* and $-\bar{z}A$ commute. We observe next that both of the exponentials on the right hand side are unitary operators, e.g.

$$(e^{zA^* - \bar{z}A})^* = e^{(zA^* - \bar{z}A)^*} = e^{\bar{z}A - zA^*} = (e^{zA^* - \bar{z}A})^{-1}.$$

Since unitary operators always have norm 1, it follows that

$$\|F(z)\| \leqslant \left\| e^{zA^* - \bar{z}A} \right\| \cdot \|B\| \cdot \left\| e^{\bar{z}A - zA^*} \right\| \leqslant \|B\|,$$

thus the function F is globally bounded. We deduce that for every $x, y \in \mathcal{H}$, the complex-valued analytic function $\mathbb{C} \to \mathbb{C} : z \mapsto \langle x, F(z)y \rangle$ is likewise bounded, and therefore constant, by Liouville's

theorem. This implies that F itself is constant, so in particular, we have $F(t) = F(0) = B$ for all $t \in \mathbb{R}$, meaning
$$e^{tA^*}B = Be^{tA^*}.$$
Both sides of this equation are differentiable functions of $t \in \mathbb{R}$ taking values in $\mathscr{L}(\mathcal{H})$, and differentiating them at $t = 0$ gives $A^*B = BA^*$. $\square$

18.11. The Borel functional calculus. The spectral theorem contains a hint that the functional calculus considered in §18.7 can be extended beyond continuous functions. Indeed, if we have a spectral representation identifying $A \in \mathscr{L}(\mathcal{H})$ with a multiplication operator $UAU^{-1} = T_F$ on $L^2(X, \mu)$ for some function $F : X \to \mathbb{C}$, then the formula

$$f(A) = U^{-1}T_{f\circ F}U \tag{18.8}$$

is easily shown to hold for any polynomial function $f : \mathbb{C} \to \mathbb{C}$, and by density, it remains valid when f is any continuous function defined on a compact set containing $F(X) \subset \mathbb{C}$. But the right hand side also makes sense for a much larger class of functions f, because multiplication operators

$$L^2(X, \mu) \xrightarrow{T_g} L^2(X, \mu) : u \mapsto gu$$

can be defined without assuming $g : X \to \mathbb{C}$ is continuous; any bounded measurable function will do. Assuming $F : X \to \mathbb{C}$ is measurable, the composition $f \circ F$ will be a well-defined and bounded measurable function whenever f is a bounded Borel-measurable function (i.e. the preimage of every open set should be a Borel set) defined on a domain in $\mathbb{C}$ that contains the image of F. The following lemma means we are free to take $\sigma(A)$ as the domain of f:

LEMMA 18.56. *Given a measure space (X, μ), a bounded measurable function $F : X \to \mathbb{C}$ and the resulting multiplication operator $T_F : L^2(X, \mu) \to L^2(X, \mu) : u \mapsto Fu$, $F(x)$ belongs to $\sigma(T_F)$ for almost every $x \in X$.*

PROOF. Recall from Exercise 18.34 that $\sigma(T_F)$ is the essential range of F, so if $\lambda \notin \sigma(T_F)$, it means there exists a neighborhood $\mathcal{U}_\lambda \subset \mathbb{C}$ of λ such that $\mu(F^{-1}(\mathcal{U})) = 0$. Since $\mathbb{C}\backslash\sigma(A)$ can be covered by countably many such neighborhoods, $\mu(F^{-1}(\mathbb{C}\backslash\sigma(A))) = 0$. $\square$

For a spectral representation $UAU^{-1} = T_F$ of $A \in \mathscr{L}(\mathcal{H})$, Lemma 18.56 allows us to change $F : X \to \mathbb{C}$ on a set of measure zero so that $F(X) \subset \sigma(A)$ without loss of generality. Under this assumption, (18.8) gives a reasonable definition of $f(A) \in \mathscr{L}(\mathcal{H})$ for any bounded Borel-measurable function $f : \sigma(A) \to \mathbb{C}$. For any Borel subset K in $\mathbb{R}^n$ or $\mathbb{C}^n$, let us denote

$$\mathscr{B}(K) := \left\{ f : K \to \mathbb{C} \;\middle|\; f \text{ is Borel measurable and bounded} \right\}.$$

Note that for functions $K \to \mathbb{C}$, "Borel measurable" is a slightly stricter condition than "Lebesgue measurable," requiring the preimage of every open subset of $\mathbb{C}$ to be an element of the Borel σ-algebra, rather than the larger σ-algebra of Lebesgue-measurable sets. A word of caution: Even if we replace the usual Lebesgue measure on K with its restriction to the Borel σ-algebra, $\mathscr{B}(K)$ is not the same thing as $L^\infty(K)$, because elements of the latter are only *equivalence classes* of functions defined almost everywhere, whereas the elements of $\mathscr{B}(K)$ are actual functions. The next lemma tells us, however, that $\mathscr{B}(K)$ is a reasonable "completion" of $C^0(K)$ with respect to a certain notion of convergence.

LEMMA 18.57. *For any compact subset $K \subset \mathbb{C}$, $\mathscr{B}(K)$ is the smallest subset $\mathcal{F} \subset \mathscr{B}(K)$ with the following properties:*

(i) $\mathcal{F}$ contains every constant function, and it also contains the restriction to K of the functions $z \mapsto z$ and $z \mapsto \bar{z}$;
(ii) For every $f, g \in \mathcal{F}$, $f + g$ and fg both lie in $\mathcal{F}$;

(iii) If $f_n \in \mathcal{F}$ is a sequence satisfying a uniform bound $|f_n| \leqslant C$ for all n and converging pointwise on K to a function $f : K \to \mathbb{C}$, then $f \in \mathcal{F}$.

PROOF SKETCH. Note first that by the first two properties, $\mathcal{F}$ must contain the restriction to K of every polynomial function $\mathbb{C} \to \mathbb{C}$, and the convergence property together with the Weierstrass approximation theorem then imply that it also contains $C^0(K)$. One then needs to verify the following claims:

(1) The set of all Borel sets $\Omega \subset K$ with $\chi_\Omega \in \mathcal{F}$ is a σ-algebra. The proof of this is elementary.
(2) For every open subset $\mathcal{U} \subset K$, $\chi_{\mathcal{U}} \in \mathcal{F}$. One sees this by using cutoff functions to construct a uniformly bounded sequence of continuous functions converging pointwise to $\chi_{\mathcal{U}}$.

The two claims imply that $\mathcal{F}$ contains the characteristic function of every Borel set in K, and it therefore also contains all Borel-measurable simple functions on K. The result then follows from the standard measure-theoretic result that every measurable function is the pointwise limit of a sequence of measurable simple functions, which can also be assumed uniformly bounded if the limit function is bounded. For details (in a slightly more general setting), see [**BS18**, Lemma 5.80]. $\square$

We now enhance Definition 18.40:

DEFINITION 18.58. A **Borel functional calculus** for a bounded linear operator $A \in \mathscr{L}(\mathcal{H})$ on a complex Hilbert space is a linear map

$$\mathscr{B}(\sigma(A)) \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A)$$

satisfying the following conditions:

(i) $(fg)(A) = f(A)g(A)$ and $\bar{f}(A) = f(A)^*$ for all $f, g \in \mathscr{B}(\sigma(A))$;
(ii) For $f \equiv 1$, $f(A) = \mathbb{1} \in \mathscr{L}(\mathcal{H})$;
(iii) For the function $f(z) := z$ on $\sigma(A)$, $f(A) = A$.
(iv) For every $f \in C^0(\sigma(A))$, $\sigma(f(A)) = f(\sigma(A)) \subset \mathbb{C}$, and for $f \in \mathscr{B}(\sigma(A))$, $x \in \mathcal{H}$ and $\lambda \in \mathbb{C}$, $Ax = \lambda x$ implies $f(A)x = f(\lambda)x$;
(v) If $f \in \mathscr{B}(\sigma(A))$ takes nonnegative real values, then $\langle x, f(A)x \rangle \geqslant 0$ for all $x \in \mathcal{H}$;
(vi) $\|f(A)\| = \|f\|_{C^0}$ for all $f \in C^0(\sigma(A))$, and for any sequence $f_n \in \mathscr{B}(\sigma(A))$ that is pointwise convergent $f_n \to f$ and satisfies a uniform bound $|f_n| \leqslant C$ for all n,
$$f_n(A)x \to f(A)x \qquad \text{for every } x \in \mathcal{H}.$$
(vii) If $AB = BA$ for some $B \in \mathscr{L}(\mathcal{H})$, then $f(A)B = Bf(A)$ for all $f \in \mathscr{B}(\sigma(A))$.

REMARK 18.59. We have formulated the conditions in Definition 18.58 so that a Borel functional calculus $\mathscr{B}(\sigma(A)) \to \mathscr{L}(\mathcal{H})$ restricts to $C^0(\sigma(A)) \subset \mathscr{B}(\sigma(A))$ as a continuous functional calculus, and for the same reasons as in the continuous case, the operators $f(A) \in \mathscr{L}(\mathcal{H})$ will always be normal, including A itself. Note that some of the properties are strictly stronger when $f \in \mathscr{B}(\sigma(A))$ is continuous, e.g. in (iv), one cannot generally expect $\sigma(f(A)) = f(\sigma(A))$ when $f : \sigma(A) \to \mathbb{C}$ is discontinuous, since its image may fail to be a compact set. Similarly, the continuity of the map $\mathscr{B}(\sigma(A)) \to \mathscr{L}(\mathcal{H})$ expressed in (vi) uses a much weaker notion of convergence than the C^0-norm, but the map still restricts to $C^0(\sigma(A))$ as an isometry.

We showed in Lemma 18.53 that any operator $A \in \mathscr{L}(\mathcal{H})$ admitting a continuous functional calculus admits a spectral representation. Here is a somewhat stronger converse for that result:

THEOREM 18.60. *If $A \in \mathscr{L}(\mathcal{H})$ admits a spectral representation, then it admits a unique Borel functional calculus.*

PROOF. Uniqueness holds for essentially the same reasons as in the continuous case: conditions (i)–(iii) determine $f(A)$ when f is a polynomial, and thanks to Lemma 18.57, condition (vi) uniquely determines the extension to all $f \in \mathscr{B}(\sigma(A))$. If a spectral representation is given, then we can establish existence by using (18.8) to define $f(A)$ for any $f \in \mathscr{B}(\sigma(A))$, i.e. we identify A unitarily with the multiplication operator $T_F : L^2(X,\mu) \to L^2(X,\mu)$ for a bounded measurable function $F : X \to \mathbb{C}$ with $F(X) \subset \sigma(A)$, and then define $f(A)$ by replacing F with $f \circ F$. Most of the stated properties follow from this definition as easy exercises—only the following are slightly less obvious:

Property (iv): Recall from Exercise 18.34 that $\sigma(A)$ is the essential range of the function $F : X \to \mathbb{C}$, meaning the set of all $\lambda \in \mathbb{C}$ such that open neighborhoods $\mathcal{U} \subset \mathbb{C}$ of λ always satisfy $\mu(F^{-1}(\mathcal{U})) > 0$. Similarly, $\sigma(f(A))$ is the essential range of $f \circ F$, and we have $(f \circ F)^{-1}(\mathcal{U}) = F^{-1}\big(f^{-1}(\mathcal{U})\big)$ for any $\mathcal{U} \subset \mathbb{C}$. Assume f is continuous, $\lambda \in \sigma(A)$ and $\lambda' = f(\lambda)$. Then for any open neighborhood $\mathcal{U} \subset \mathbb{C}$ of λ', continuity implies that $f^{-1}(\mathcal{U}) \subset \mathbb{C}$ is an open neighborhood of λ, implying $\mu(F^{-1}(f^{-1}(\mathcal{U}))) > 0$ and thus $\lambda' \in \sigma(f(A))$. Conversely, the condition $\lambda' \in \sigma(f(A))$ implies that every open neighborhood $\mathcal{V} \subset \mathbb{C}$ of λ' satisfies $\mu(F^{-1}(f^{-1}(\mathcal{V}))) > 0$, which is clearly false if λ' does not lie in the compact set $f(\sigma(A))$. The statement about eigenvalues is easier, and does not require $f \in \mathscr{B}(\sigma(A))$ to be continuous: It suffices to observe that for $u \in L^2(X,\mu)$, $T_F u = \lambda u$ means that u vanishes almost everywhere outside the set $F^{-1}(\lambda) \subset X$, and it follows that $T_{f \circ F} u = f(\lambda) u$.

Property (vi): Now that we know $\sigma(f(A)) = f(\sigma(A))$ for $f \in C^0(\sigma(A))$, the relation $\|f(A)\| = \|f\|_{C^0}$ follows by the same argument as in Lemma 18.43.[43] For the statement about pointwise convergence, it suffices to work in $L^2(X,\mu)$ and check that whenever $f_n \to f$ pointwise and $|f_n| \leqslant C$ for all n, we have

$$(f_n \circ F)u \xrightarrow{L^2} (f \circ F)u$$

for all $u \in L^2(X,\mu)$, that is

$$\int_X |f_n(F(x)) - f(F(x))|^2 \cdot |u(t)|^2 \, d\mu(x) \to 0$$

as $n \to \infty$. This follows from the dominated convergence theorem.

Property (vii): Since A is normal, the condition $AB = BA$ implies via Theorem 18.55 that B also commutes with A^*, and it therefore commutes with all finite products of A and A^*, which means that it commutes with $f(A)$ for every polynomial function $f : \mathbb{C} \to \mathbb{C}$. This extends to all $f \in \mathscr{B}(\sigma(A))$ as a consequence of Lemma 18.57 and property (vi). □

Remark 18.61. Notice that the pointwise convergence of functions in Lemma 18.57 and Definition 18.58 is *everywhere*, not just "almost" everywhere. This is important in the proof of Theorem 18.60, because if $f_n \to f$ only *almost* everywhere, then it need not follow that $f_n \circ F \to f \circ F$ almost everywhere. For a similar reason, the operator $f(A) \in \mathscr{L}(\mathcal{H})$ may change if $f \in \mathscr{B}(\sigma(A))$ is modified on a set of measure zero.

18.12. Polar decomposition. Before continuing with the general development of spectral theory, we give a nice application of the functional calculus for self-adjoint operators.

Definition 18.62. An operator $A \in \mathscr{L}(\mathcal{H})$ on a Hilbert space is called **positive**, written

$$A \geqslant 0,$$

if $\langle x, Ax\rangle \geqslant 0$ for all $x \in \mathcal{H}$.

Exercise 18.63. For a complex Hilbert space $\mathcal{H}$ and operator $A \in \mathscr{L}(\mathcal{H})$, show:

[43] Lemma 18.43 was stated specifically for self-adjoint operators, but the only step in the proof that actually required self-adjointness was showing that $\sigma(f(A)) = f(\sigma(A))$ for f continuous. The use of a spectral representation enabled us to prove this by other means in the previous paragraph.

(a) If $A \geqslant 0$, then A is self-adjoint.
Hint: If $\langle x, Ax\rangle$ is real then $\langle x, Ax\rangle = \langle Ax, x\rangle$. Compute $\langle x + y, A(x + y)\rangle$ and $\langle x + iy, A(x + iy)\rangle$ for arbitrary $x, y \in \mathcal{H}$.
(b) If A is self-adjoint, then $A \geqslant 0$ if and only if $\sigma(A) \subset [0, \infty)$.
(c) If $\langle x, Ax\rangle > 0$ for all $x \neq 0 \in \mathcal{H}$, it need not follow that $0 \notin \sigma(A)$.

EXAMPLE 18.64. Any $A \in \mathscr{L}(\mathcal{H})$ gives rise to a positive operator $A^*A \geqslant 0$.

For a positive operator A, Exercise 18.63 implies that the square root function $\sqrt{\cdot}$ is well defined and continuous on $\sigma(A)$, since the latter is contained in $[0, \infty)$. We can therefore use the functional calculus to define $\sqrt{A} \in \mathscr{L}(\mathcal{H})$ so that

$$\sqrt{A}\sqrt{A} = A.$$

Equation (18.8) gives a formula for $\sqrt{A}$ in terms of any spectral representation $A = U^{-1}T_F U$, namely as

$$\sqrt{A} = U^{-1}T_{\sqrt{F}}U,$$

which makes sense because $F : X \to \mathbb{C}$ takes nonnegative real values almost everywhere. The function $\sqrt{F} : X \to \mathbb{R}$ is then also nonnegative by convention, and $\sqrt{A}$ is therefore also a positive operator,

$$\sqrt{A} \geqslant 0.$$

EXERCISE 18.65. Show that for any positive operator $A \in \mathscr{L}(\mathcal{H})$, $\ker \sqrt{A} = \ker A$.
Hint: This is almost obvious if A is a multiplication operator on $L^2(X, \mu)$.

One nice application of positive operators is to generalize the polar decomposition $z = \rho e^{i\theta}$ of complex numbers. The latter makes sense in part because for every $z \in \mathbb{C}$, $|z|^2 = \bar{z}z$ is a nonnegative real number, and therefore has a well-defined nonnegative square root $\rho := \sqrt{\bar{z}z}$. Example 18.64 generalizes this to operators on Hilbert spaces, and the role of $e^{i\theta}$ can be played by so-called "partial isometries":

DEFINITION 18.66. A bounded linear operator $U \in \mathscr{L}(\mathcal{H})$ on a Hilbert space is called a **partial isometry** if its restriction to the closed subspace $(\ker U)^\perp \subset \mathcal{H}$ defines an isometry

$$(\ker U)^\perp \xrightarrow{U} \operatorname{im} U \subset \mathcal{H}.$$

As usual, the word "isometry" in Definition 18.66 means the relation $\langle Ux, Uy\rangle = \langle x, y\rangle$ for all $x, y \in (\ker U)^\perp$. It follows easily from this condition that $\operatorname{im} U \subset \mathcal{H}$ must also be a closed subspace.

THEOREM 18.67 (polar decomposition). *Given a bounded linear operator $A \in \mathscr{L}(\mathcal{H})$ on a complex Hilbert space $\mathcal{H}$, let $P := \sqrt{A^*A}$. Then there is a unique operator $U \in \mathscr{L}(\mathcal{H})$ such that*

$$A = UP \qquad \text{and} \qquad \ker U = \ker A,$$

and moreover, U is then a partial isometry with $\operatorname{im} U = \overline{\operatorname{im} A} \subset \mathcal{H}$. In particular, U is a unitary isomorphism $\mathcal{H} \to \mathcal{H}$ if and only if A is injective and has dense image.

PROOF. We observe first that if $v \in \ker(A^*A) \subset \mathcal{H}$, then

$$0 = \langle v, A^*Av\rangle = \langle Av, Av\rangle = \|Av\|^2$$

and thus $v \in \ker A$, proving that $\ker A = \ker(A^*A)$. By Exercise 18.65, it follows that

$$\ker A = \ker P,$$

so P maps $\ker A$ to itself, and since P is self-adjoint, it therefore also maps $(\ker A)^\perp$ to itself, the latter injectively. We claim that $\operatorname{im} P$ is then a dense subspace of $(\ker A)^\perp$: indeed, if it is not, then there exists a nonzero element $v \in (\ker A)^\perp$ such that $\langle v, Pw\rangle = 0$ for all $w \in (\ker A)^\perp$, implying

$$0 = \langle v, Pw\rangle = \langle Pv, w\rangle$$

for all $w \in (\ker A)^\perp$, and since $Pv \in (\ker A)^\perp$, this implies $Pv = 0$, which can only happen if $v = 0$ since $P|_{(\ker A)^\perp}$ is injective.

Since we now know P to be a bijection from $(\ker A)^\perp$ to a dense subspace $\operatorname{im} P \subset (\ker A)^\perp$, the relation $A = UP$ will hold for some operator $U : \mathcal{H} \to \mathcal{H}$ if and only if U restricts to $\operatorname{im} P \subset \mathcal{H}$ as the composition

$$\begin{tikzcd} \operatorname{im} P \arrow[r, "P^{-1}"] \arrow[rr, bend right, "U"'] & (\ker A)^\perp \arrow[r, "A"] & \operatorname{im} A \subset \mathcal{H}. \end{tikzcd}$$

Defining $U : \operatorname{im} P \to \operatorname{im} A$ in this way, it is a bijection by construction, and we claim that it is also an isometry. Indeed, every $v \in \operatorname{im} P$ is of the form $v = Pw$ for a unique $w \in (\ker A)^\perp$, so the claim means that for every $w \in (\ker A)^\perp$, the elements $v := Pw$ and $Uv = Aw$ satisfy

$$\|Aw\| = \|Pw\|.$$

This is easily verified using the self-adjointness of $P = \sqrt{A^*A}$:

$$\|Pw\|^2 = \left\langle \sqrt{A^*A}w, \sqrt{A^*A}w \right\rangle = \left\langle w, \sqrt{A^*A}^2 w \right\rangle = \langle w, A^*Aw\rangle = \langle Aw, Aw\rangle = \|Aw\|^2.$$

With the claim established, U is now a bounded linear operator $\operatorname{im} P \to \mathcal{H}$ with image $\operatorname{im} A$, and thus has a unique continuous extension to the closure $\overline{\operatorname{im} P} = (\ker A)^\perp$, defining an isometry that maps the latter bijectively to the closure $\overline{\operatorname{im} A} \subset \mathcal{H}$. There is then exactly one extension of U to all of $\mathcal{H}$ such that $\ker U = \ker A$, and it is a partial isometry by construction. □

18.13. Spectral projections. Now that we have a definition of $f(A)$ for functions f that need not be continuous, we can apply it in particular to characteristic functions $\chi_\Omega \in \mathscr{B}(\mathbb{C})$ of Borel subsets $\Omega \subset \mathbb{C}$.

DEFINITION 18.68. Given an operator $A \in \mathscr{L}(\mathcal{H})$ that admits a spectral representation, the **spectral projection** associated to any Borel set $\Omega \subset \mathbb{C}$ is the operator

$$P_\Omega := \chi_\Omega(A) \in \mathscr{L}(\mathcal{H}).$$

In other words, if we use a spectral representation to identify $\mathcal{H}$ with $L^2(X, \mu)$ and A with a multiplication operator T_F for some function $F : X \to \mathbb{C}$, then P_Ω is identified with the multiplication operator

$$L^2(X,\mu) \to L^2(X,\mu) : u \mapsto (\chi_\Omega \circ F)u = \begin{cases} u & \text{on } F^{-1}(\Omega), \\ 0 & \text{elsewhere.} \end{cases}$$

Since χ_Ω is a nonnegative real-valued function, P_Ω is a positive self-adjoint operator, and it also satisfies

$$P_\Omega^2 = P_\Omega,$$

hence P_Ω is the orthogonal projection onto its image (cf. Exercise 3.31).

EXERCISE 18.69. Show that if $\Omega \cap \sigma(A) \subset \mathbb{C}$ consists of a finite set of eigenvalues of A, then $P_\Omega \in \mathscr{L}(\mathcal{H})$ is the orthogonal projection onto the direct sum of the corresponding eigenspaces, i.e.

$$\operatorname{im} P_\Omega = \bigoplus_{\lambda \in \Omega \cap \sigma(A)} \ker(\lambda - A).$$

The preceding exercise gives some useful intuition about the meaning of P_Ω in more general cases: Informally, we can think of $\operatorname{im} P_\Omega$ as the "approximate eigenspace" corresponding to the portion of the spectrum of A that lies in Ω.

EXERCISE 18.70. Prove the following properties of the spectral projections P_Ω for an operator $A \in \mathscr{L}(\mathcal{H})$:

(a) $P_\Omega = 0$ whenever $\Omega \cap \sigma(A) = \varnothing$.
(b) $P_\Omega = \mathbb{1}$ whenever $\sigma(A) \subset \Omega$.
(c) $\lambda \in \sigma(A)$ if and only if $P_\Omega \neq 0$ for all $\Omega \subset \mathbb{C}$ that contain open neighborhoods of λ.
(d) $P_{\Omega \cap \Omega'} = P_\Omega P_{\Omega'}$ for any two $\Omega, \Omega' \subset \mathbb{C}$.
(e) For any finite collection of pairwise disjoint sets $\Omega_1, \ldots, \Omega_m \subset \mathbb{C}$,

$$P_{\Omega_1} + \ldots + P_{\Omega_m} = P_{\Omega_1 \cup \ldots \cup \Omega_m}.$$

(f) For any sequence $\{\Omega_j \subset \mathbb{C}\}_{j=1}^\infty$ of pairwise disjoint sets and any $x \in \mathcal{H}$, the series $\sum_{j=1}^\infty P_{\Omega_j} x$ converges to

$$\sum_{j=1}^\infty P_{\Omega_j} x = P_\Omega x \qquad \text{where} \qquad \Omega := \bigcup_{j=1}^\infty \Omega_j.$$

REMARK 18.71. The last property in Exercise 18.70 is reminiscent of the countable additivity property of measures, and in part for this reason, the collection of orthogonal projections $\{P_\Omega \in \mathscr{L}(\mathcal{H})\}_{\Omega \subset \mathbb{C}}$ associated to $A \in \mathscr{L}(\mathcal{H})$ is sometimes called a **projection-valued measure**. It follows in fact from Exercise 18.70 that for any element $x \in \mathcal{H}$, the formula

$$\mu(\Omega) := \langle x, P_\Omega x \rangle$$

defines an actual finite measure on the Borel σ-algebra of $\mathbb{C}$, with every set that is disjoint from $\sigma(A)$ having measure zero. We will not go into the details here, but one can define integration of complex-valued functions with respect to a projection-valued measure, so that the integral itself takes values in $\mathscr{L}(\mathcal{H})$. For $f \in \mathscr{B}(\sigma(A))$, one then obtains the formula

$$f(A) = \int_{\sigma(A)} f(\lambda)\, dP_\lambda, \qquad \text{and in particular} \qquad A = \int_{\sigma(A)} \lambda\, dP_\lambda,$$

where the symbol dP_λ means integration with respect to the projection-valued measure $\{P_\Omega\}_{\Omega \subset \mathbb{C}}$. This generalizes the finite-dimensional formula

$$A = \sum_{\lambda \in \sigma(A)} \lambda \Pi_{E_\lambda}$$

for normal operators $A : \mathbb{C}^n \to \mathbb{C}^n$, where Π_{E_λ} denotes the orthogonal projection onto the eigenspace $E_\lambda := \ker(\lambda - A)$. For details, see [**RS80**, §VII.3].

18.14. Simultaneous diagonalization. The extension of Theorem 18.38 from self-adjoint to normal operators will be made possible by the answer to the following question: Under what conditions can a collection of multiple diagonalizable operators $\{A_i\}$ be diagonalized *simultaneously*? That is, do they admit a **common spectral representation**, meaning a single measure space (X, μ) and a single unitary isomorphism $U : \mathcal{H} \to L^2(X, \mu)$ such that $UA_iU^{-1} = T_{F_i}$ for a collection of bounded measurable functions $F_i : X \to \mathbb{C}$?

There is an obviously necessary condition for this: the multiplication operators T_{F_i} and T_{F_j} commute with each other for all $i \neq j$, so a common spectral representation for A_i and A_j clearly cannot exist unless A_i and A_j commute. This condition turns out to be sufficient:

THEOREM 18.72. *Suppose $A, B \in \mathscr{L}(\mathcal{H})$ are commuting operators that both admit spectral representations. Then there exists a semifinite measure space (X, μ) and a unitary isomorphism $U : \mathcal{H} \to L^2(X, \mu)$ such that $UAU^{-1} = T_F$ and $UBU^{-1} = T_G$ for two bounded measurable functions $F, G : X \to \mathbb{C}$. Moreover, (X, μ) can be assumed finite if $\mathcal{H}$ is separable.*

REMARK 18.73. The obvious generalization of Theorem 18.72 also holds for any finite collection of commuting operators, not just two. The proof is based on the same ideas as in the following, only the notation becomes more cumbersome.

The proof of Theorem 18.72 is based on the construction of a continuous functional calculus

$$C^0(\sigma(A) \times \sigma(B)) \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A, B).$$

The assumption that A and B each admit (separate) spectral representations implies that they are normal, so if they commute, then Theorem 18.55 implies that all *four* of the operators A, A^*, B, B^* commute with each other. There is then an obvious definition for $P(A, B)$ whenever P is a complex-valued polynomial function on $\mathbb{C} \times \mathbb{C}$, and this would extend uniquely to $C^0(\sigma(A) \times \sigma(B))$ if we could prove a bound of the form

$$\|P(A, B)\| \leqslant c \cdot \sup_{(\lambda,\mu) \in \sigma(A) \times \sigma(B)} |P(\lambda, \mu)|.$$

The corresponding step in the proof of Lemma 18.43 was easier, because in the setting of a single self-adjoint operator, we had the privilege of being able to work with polynomial functions of one real variable. Here, on the other hand, we have polynomial functions that depend on two complex variables and their complex conjugates, so describing the spectrum of $P(A, B)$ explicitly enough to estimate its spectral radius and therefore $\|P(A, B)\|$ is harder. We shall follow an alternative approach that uses spectral projections: instead of approximating continuous functions on $\sigma(A) \times \sigma(B)$ by polynomials, we can approximate them by functions that are constant on "rectangles" in $\mathbb{C}^2$.

LEMMA 18.74. *For two commuting operators $A, B \in \mathscr{L}(\mathcal{H})$ with spectral representations and any two Borel sets $\Omega, \Omega' \subset \mathbb{C}$, the spectral projections $\chi_\Omega(A)$ and $\chi_{\Omega'}(B)$ also commute.*

PROOF. Property (vii) of the Borel functional calculus (Definition 18.58) gives $\chi_\Omega(A)B = B\chi_\Omega(A)$, and applying the same result again with B in the former role of A then gives $\chi_{\Omega'}(B)\chi_\Omega(A) = \chi_\Omega(A)\chi_{\Omega'}(B)$. □

Let us refer to sets of the form

$$\left\{x + iy \;\middle|\; a \leqslant x < b \text{ and } c \leqslant y < d\right\} \subset \mathbb{C}, \qquad a < b,\ c < d$$

in the following as **half-open rectangles**, and define

$$\mathscr{R} \subset \mathscr{B}(\mathbb{C}^2)$$

as the vector space of all finite complex-linear combinations of characteristic functions $\chi_{\Omega \times \Omega'}$ for half-open rectangles $\Omega, \Omega' \subset \mathbb{C}$. The proof of the next lemma is an easy exercise:

LEMMA 18.75. *For any compact subsets $K, K' \subset \mathbb{C}$ and any continuous function $f : K \times K' \to \mathbb{C}$, there exists a sequence $f_n \in \mathscr{R}$ that converges uniformly to f on $K \times K'$.* □

By restriction, we can regard functions $f \in \mathscr{R}$ also as functions on $\sigma(A) \times \sigma(B) \subset \mathbb{C}^2$, and from the latter perspective, impose on them the sup-norm

$$\|f\| := \sup\left\{|f(z, w)| \;\middle|\; (z, w) \in \sigma(A) \times \sigma(B)\right\}.$$

Now define

$$\mathscr{R} \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A, B)$$

as the unique linear map such that for any two half-open rectangles $\Omega, \Omega' \subset \mathbb{C}$,

$$\chi_{\Omega \times \Omega'}(A, B) := \chi_{\Omega}(A)\chi_{\Omega'}(B).$$

This is well defined because e.g. whenever $\Omega = \Omega_1 \cup \Omega_2$ for two smaller disjoint half-open rectangles, the relation $\chi_{\Omega} = \chi_{\Omega_1} + \chi_{\Omega_2}$ translates into

$$\chi_{\Omega}(A)\chi_{\Omega'}(B) = \chi_{\Omega_1}(A)\chi_{\Omega'}(B) + \chi_{\Omega_2}(A)\chi_{\Omega'}(B)$$

due to Exercise 18.70(e). We also deduce from Exercise 18.70 the following properties, which are consistent with the notion of a continuous functional calculus:

(i) $(fg)(A, B) = f(A, B)g(A, B)$ and $\bar{f}(A, B) = f(A, B)^*$ for all $f, g \in \mathscr{R}$;
(ii) For $f \equiv 1$, $f(A, B) = \mathbb{1} \in \mathscr{L}(\mathcal{H})$.

The latter follows in particular from part (b) of Exercise 18.70 by taking $f := \chi_{\Omega \times \Omega'}$ for any half-open rectangles $\Omega, \Omega' \subset \mathbb{C}$ that are large enough to contain $\sigma(A)$ and $\sigma(B)$ respectively.

LEMMA 18.76. *The map $\mathscr{R} \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A, B)$ satisfies*

$$\|f(A, B)\| \leqslant \|f\|.$$

PROOF. By breaking up rectangles into smaller rectangles, one can write any $f \in \mathscr{R}$ as a finite linear combination of the form

$$f(x, y) = \sum_{i,j} c_{ij} \chi_{\Omega_i^A}(x)\chi_{\Omega_j^B}(y) \tag{18.9}$$

for some coefficients $c_{ij} \in \mathbb{C}$ and collections of *disjoint* half-open rectangles $\Omega_1^A, \dots, \Omega_m^A \subset \mathbb{C}$ and $\Omega_1^B, \dots, \Omega_n^B \subset \mathbb{C}$. In this representation,

$$f(A, B) = \sum_{i,j} c_{ij} \chi_{\Omega_i^A}(A)\chi_{\Omega_j^B}(B).$$

We are then free to assume $\Omega_i^A \cap \sigma(A)$ and $\Omega_j^B \cap \sigma(B)$ are both nonempty for every i, j, since Exercise 18.70(a) will otherwise imply that $\chi_{\Omega_i^A}(A)\chi_{\Omega_j^B}(B) = 0$. Under this assumption,

$$\|f\| = \max_{i,j} |c_{ij}|,$$

and we deduce

$$\|f(A, B)\| \leqslant \|f\| \cdot \left\| \sum_{i,j} \chi_{\Omega_i^A}(A)\chi_{\Omega_j^B}(B) \right\|.$$

Since the sets $\Omega_1^A, \dots, \Omega_m^A$ are all disjoint, we can use Exercise 18.70(e) to rewrite the summation on the right hand side as

$$\sum_{i,j} \chi_{\Omega_i^A}(A)\chi_{\Omega_j^B}(B) = \left(\sum_i \chi_{\Omega_i^A}(A) \right) \left(\sum_j \chi_{\Omega_j^B}(B) \right) = \chi_{\Omega_1^A \cup \dots \cup \Omega_m^A}(A)\chi_{\Omega_1^B \cup \dots \cup \Omega_n^B}(B).$$

As the composition of two orthogonal projections, the latter operator has norm at most 1, and we conclude $\|f(A, B)\| \leqslant \|f\|$. □

The estimate in Lemma 18.76 implies that $\mathscr{R} \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A, B)$ has a unique extension to a bounded linear operator $\bar{\mathscr{R}} \to \mathscr{L}(\mathcal{H})$, where $\bar{\mathscr{R}} \subset \mathscr{B}(\sigma(A) \times \sigma(B))$ is the closure of $\mathscr{R}$ in the sup-norm. By Lemma 18.75, $\bar{\mathscr{R}}$ contains $C^0(\sigma(A) \times \sigma(B))$, and we have therefore defined a bounded linear map

$$C^0(\sigma(A) \times \sigma(B)) \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A, B).$$

We claim that this map also satisfies the following property:

(iii) For the projection functions $\Pi_1(z,w) := z$ and $\Pi_2(z,w) := w$ on $\sigma(A) \times \sigma(B)$, we have $\Pi_1(A,B) = A$ and $\Pi_2(A,B) = B$.

To prove the statement about Π_1, choose $\Omega' \subset \mathbb{C}$ large enough to contain $\sigma(B)$, so that $\chi_{\Omega'}(B) = \mathbb{1}$ by Exercise 18.70(ii), and choose a sequence of functions $f_n : \mathbb{C} \to \mathbb{C}$, uniformly convergent on $\sigma(A)$ to $f(z) := z$, such that each f_n is a finite linear combination of characteristic functions of half-open rectangles. The functions $F_n(z,w) := f_n(z)\chi_{\Omega'}(w)$ then belong to $\mathscr{R}$ and are uniformly convergent on $\sigma(A) \times \sigma(B)$ to Π_1, thus $F_n(A,B) \to \Pi_1(A,B)$ in the operator norm. Using the Borel functional calculus of A, we can write $F_n(A,B) = f_n(A)\chi_{\Omega'}(B) = f_n(A)$, and since the functions f_n are uniformly bounded on the compact set $\sigma(A)$ and converge pointwise to f, properties (iii) and (vi) of the Borel functional calculus (Definition 18.58) imply $f_n(A)x \to f(A)x = Ax$ for every $x \in \mathcal{H}$. We conclude $\Pi_1(A,B)x = Ax$ for every $x \in \mathcal{H}$. The statement about Π_2 is proved in the same manner.

We now have the essential properties of a continuous functional calculus for functions of the commuting operators A and B, and we note that it automatically also has a positivity property: for any nonnegative real-valued continuous function $f \geqslant 0$ on $\sigma(A) \times \sigma(B)$, we can write $f = g^2$ with $g \geqslant 0$ to prove

$$\langle x, f(A,B)x\rangle = \langle x, g(A,B)g(A,B)x\rangle = \langle g(A,B)x, g(A,B)x\rangle \geqslant 0.$$

From here, the proof of Theorem 18.72 follows a familiar pattern. The Riesz-Markov theorem produces a finite Borel measure μ_x on $\sigma(A) \times \sigma(B)$ satisfying

$$\langle x, f(A,B)x\rangle = \int_{\sigma(A)\times\sigma(B)} f\, d\mu_x \qquad \text{for all} \qquad f \in C^0(\sigma(A) \times \sigma(B)),$$

and we then define $T : C^0(\sigma(A) \times \sigma(B)) \to \mathcal{H} : f \mapsto f(A,B)x$, which satisfies $\|Tf\| = \|f\|_{L^2}$ and therefore extends to an isometry

$$L^2(\sigma(A) \times \sigma(B), \mu_x) \xrightarrow{T} \mathcal{H}$$

whose image is the span of all elements of the form $A^m B^n (A^*)^p (B^*)^q x$ for integers $m,n,p,q \geqslant 0$. In this context, it is natural to say that $x \in \mathcal{H}$ is **cyclic** for A and B if the span of those elements is $\mathcal{H}$, and if that is true, then $U := T^{-1}$ identifies both A and B with multiplication operators on $L^2(\sigma(A) \times \sigma(B), \mu_x)$, namely via multiplication with the functions $(x,y) \mapsto x$ and $(x,y) \mapsto y$ respectively. If $\mathcal{H}$ is separable, one can now decompose $\mathcal{H}$ into countably many orthogonal closed subspaces admitting cyclic elements; in the nonseparable case, one instead obtains a possibly uncountable collection of such closed subspaces with the aid of Zorn's lemma, as in Remark 18.54. The result is a common spectral representation for A and B on a measure space (X,μ) that is a disjoint union of copies of $\sigma(A) \times \sigma(B)$ with various spectral measures μ_x. This completes the proof of Theorem 18.72.

18.15. Extension from self-adjoint to normal. In case you've been having trouble keeping track, let's recap the most important results that have been proved so far, and the logic by which we have proved them:

(1) By Lemma 18.43, every bounded self-adjoint operator $A \in \mathscr{L}(\mathcal{H})$ admits a continuous functional calculus.
(2) Lemma 18.53 then establishes that self-adjoint operators always have spectral representations.
(3) It follows via Theorem 18.60 that every self-adjoint operator also admits a unique Borel functional calculus, which can be used as in §18.13 to define spectral projections.
(4) By Theorem 18.72, any two commuting self-adjoint operators admit a common spectral representation.

The extension of all this to normal operators now proceeds as follows. By Exercise 18.37, any normal operator $A \in \mathscr{L}(\mathcal{H})$ can be written as

$$A = B + iC$$

for two commuting self-adjoint operators $B, C \in \mathscr{L}(\mathcal{H})$. Extracting from Theorem 18.72 a common spectral representation for B and C as multiplication operators $T_F, T_G : L^2(X, \mu) \to L^2(X, \mu)$ defined via functions $F, G : X \to \mathbb{R}$ respectively, the same unitary isomorphism $\mathcal{H} \to L^2(X, \mu)$ then identifies A with the multiplication operator T_{F+iG}. In other words:

(5) Theorem 18.72 for self-adjoint operators implies that bounded normal operators also admit spectral representations.
(6) It follows via Theorem 18.60 that every normal operator also admits a unique Borel functional calculus and spectral projections.
(7) Theorem 18.72 therefore also applies to any two commuting normal operators.

We summarize with a single statement:

THEOREM. *On a complex Hilbert space $\mathcal{H}$, a bounded linear operator admits a spectral representation if and only if it is normal, and two normal operators admit a common spectral representation if and only if they commute.* □

19. Unbounded operators

As a motivational example for what follows, consider the Laplace operator for complex-valued functions of n real variables,

$$-\Delta := -\sum_{j=1}^{n} \partial_j^2.$$

The reason for the minus sign will be clarified momentarily. We observe first that $-\Delta$ has some things in common with self-adjoint operators on $L^2(\mathbb{R}^n)$, e.g. for any two Schwartz functions $\varphi, \psi \in \mathscr{S}(\mathbb{R}^n)$, integration by parts gives

$$\langle \varphi, -\Delta\psi \rangle_{L^2} = -\sum_{j=1}^{n} \int_{\mathbb{R}^n} \bar{\varphi}\, \partial_j^2 \psi = \sum_{j=1}^{n} \int_{\mathbb{R}^n} \partial_j \bar{\varphi}\, \partial_j \psi = \langle \nabla\varphi, \nabla\psi \rangle_{L^2},$$

where we denote the $\mathbb{C}^n$-valued gradient of φ by $\nabla\varphi := (\partial_1\varphi, \ldots, \partial_n\varphi)$, and the boundary terms arising from integration by parts do not appear because φ and ψ decay to zero at infinity. Applying the same trick again to move all differentiation to φ instead of ψ, we obtain

$$\langle \varphi, -\Delta\psi \rangle_{L^2} = \langle \nabla\varphi, \nabla\psi \rangle_{L^2} = \langle -\Delta\varphi, \psi \rangle_{L^2}, \qquad \text{for all} \qquad \varphi, \psi \in \mathscr{S}(\mathbb{R}^n). \tag{19.1}$$

Now you can see why the minus sign is there: it also makes $-\Delta$ into a *positive* operator, in the sense that

$$\langle \varphi, -\Delta\varphi \rangle_{L^2} \geqslant 0 \qquad \text{for all} \qquad \varphi \in \mathscr{S}(\mathbb{R}^n).$$

If we want to discuss functions defined only on an open subset $\Omega \subset \mathbb{R}^n$, the same results hold for all compactly supported smooth functions $\varphi, \psi \in C_0^\infty(\Omega)$.

The usual argument from the linear algebra of self-adjoint operators then shows that if Δ has *eigenfunctions*, meaning solutions φ to the equation $\Delta\varphi = \lambda\varphi$ for $\lambda \in \mathbb{C}$, then the eigenvalues λ are always real and nonnegative, and eigenfunctions corresponding to two distinct eigenvalues are always L^2-orthogonal. We saw in Corollary 18.28 that for functions on a bounded open subset $\Omega \subset \mathbb{R}^n$, this really happens: in fact, $L^2(\Omega)$ admits a countable orthonormal basis of eigenfunctions of $-\Delta$, with a discrete set of positive eigenvalues of finite multiplicity that tend to infinity!

This result about eigenfunctions of $-\Delta$ was derived in a fairly indirect manner in §18.5: We showed that eigenfunctions of $-\Delta$ with eigenvalue $\lambda \in \mathbb{R}$ are equivalent to eigenfunctions of a certain *compact* self-adjoint operator $K : L^2(\Omega) \to L^2(\Omega)$ with eigenvalue $1/\lambda$, so that the existence of an orthonormal basis of eigenfunctions follows from the Hilbert-Schmidt theorem. In the present section, we aim to study the spectral properties of operators like $-\Delta$ more directly.

One problem immediately arises: The notions of eigenvectors and self-adjointness make sense for linear operators from an inner product space to *itself*, but $-\Delta$ is not a well-defined operator from $L^2(\mathbb{R}^n)$ or $L^2(\Omega)$ to itself: If you feed an arbitrary L^2-function into $-\Delta$, you will generally get a distribution, not a function. One *can* of course define $-\Delta$ on various dense subspaces such as $\mathscr{S}(\mathbb{R}^n) \subset L^2(\mathbb{R}^n)$ or $C_0^\infty(\Omega) \subset L^2(\Omega)$, but the resulting operator clearly does not admit a bounded extension to the rest of L^2; the unboundedness of the eigenvalues in the discussion above is a dead giveaway that this is impossible. The only reason the eigenfunction equation $-\Delta\varphi = \lambda\varphi$ makes sense at all is that its solutions are generally *better* than L^2-functions, e.g. in Corollary 18.28, they belong to the Sobolev space $H_0^1(\Omega)$, and one can apply methods from elliptic regularity theory as in Theorem 12.50 to show that they are actually smooth.

The inescapable conclusion is this: We should study eigenvalues and other spectral properties for linear operators of the form

$$\mathcal{H} \supset \mathcal{D} \xrightarrow{T} \mathcal{H}$$

on a Hilbert space $\mathcal{H}$, where the domain $\mathcal{D} \subset \mathcal{H}$ is a dense subspace and T *cannot* be extended continuously to the rest of $\mathcal{H}$. In this setting, the notion of a *closed* operator will serve as a reasonable substitute for boundedness, and we will see that the spectral theorem for bounded self-adjoint operators admits a very elegant extension.

19.1. Dense domains and closed operators. Our first task is to generalize the notion of a bounded linear operator between Banach spaces to operators that are only defined on dense subspaces.

DEFINITION 19.1. For two Banach spaces X, Y, an **unbounded operator** T from X to Y consists of a subspace $\mathcal{D} := \mathcal{D}(T) \subset X$ together with a linear map $T : \mathcal{D} \to Y$. We will generally abbreviate this with the notation

$$X \supset \mathcal{D} \xrightarrow{T} Y.$$

We say that the operator is **closed** if its graph

$$\Gamma_T := \left\{(x, Tx) \;\middle|\; x \in \mathcal{D}\right\} \subset X \oplus Y$$

is a closed subspace of the Banach space $X \oplus Y$. We say that it is **closable** if it admits a closed extension

$$X \supset \mathcal{D}' \xrightarrow{T'} Y,$$

i.e. a closed unbounded operator T' on a domain $\mathcal{D}' \subset X$ containing $\mathcal{D}$ such that $T'|_{\mathcal{D}} = T$. For a closable operator T, the smallest closed extension is called the **closure** of T and will be denoted by $\bar{T}$.

REMARK 19.2. The term "unbounded operator" as defined above is imperfect, because the definition does not exclude the possibility that $T : \mathcal{D} \to Y$ might be bounded with respect to the norm on X, in which case it admits a continuous extension to the closure of $\mathcal{D} \subset X$, which is often X. This does not happen in the cases that are typically of interest (cf. Remark 19.5 below), but explicitly excluding it would be inconvenient, so in practice, we'll be using the word "unbounded" in this context to mean that T is defined on a specified domain $\mathcal{D} \subset X$ that is allowed to be smaller than X, though T might also be bounded.

EXERCISE 19.3. Show that the following conditions on an unbounded operator $X \supset \mathcal{D} \xrightarrow{T} Y$ are equivalent:

(i) T is closable;

(ii) The closure $\bar{\Gamma}_T$ of $\Gamma_T \subset X \oplus Y$ is also the graph of an operator;
(iii) The conditions $(x, y) \in \bar{\Gamma}_T$ and $(x, y') \in \bar{\Gamma}_T$ imply $y = y'$;
(iv) For every $x \in \bar{\mathcal{D}} \subset X$, there exists at most one element $y \in Y$ arising as a limit of sequences Tx_n with $x_n \in \mathcal{D}$ converging to x.

Remark 19.4. If T is closed and its domain $\mathcal{D} \subset X$ is X, then the closed graph theorem (Theorem 16.11) implies that T is bounded. In this sense, closed unbounded operators are a natural generalization of bounded operators.

Remark 19.5. If the domain $\mathcal{D} \subset X$ of a closed operator T is dense but $\mathcal{D} \neq X$, then $T : \mathcal{D} \to Y$ must be a discontinuous map with respect to the norm on X. Indeed, continuity would imply that T admits an extension to a bounded linear operator $\bar{T} : X \to Y$, the graph of which would then be the closure of $\Gamma_T \subset X \oplus Y$, implying that T itself cannot be closed.

Exercise 19.6. On the domain $\mathcal{D}$ of an unbounded operator $X \supset \mathcal{D} \xrightarrow{T} Y$, one defines the **graph norm** by

$$\|x\|_T := \|x\| + \|Tx\|,$$

where the norms on the right hand side are the given norms on X and Y. This makes T a bounded linear operator $\mathcal{D} \to Y$ with respect to the graph norm on $\mathcal{D}$. Show that T is closed if and only if $(\mathcal{D}, \|\cdot\|_T)$ is a Banach space.

Remark 19.7. Suppose $X \supset \mathcal{D} \xrightarrow{T} Y$ is an unbounded operator whose domain is not dense in X, and such that its closure $\bar{\mathcal{D}} \subset X$ is complemented, so there exist a splitting

$$X = \bar{\mathcal{D}} \oplus X',$$

with $X' \subset X$ a closed subspace. There is then a unique extension of T to an unbounded operator

$$X \supset \mathcal{D} \oplus X' \xrightarrow{T'} Y$$

such that $T'|_{X'} = 0$. The enlarged domain $\mathcal{D} \oplus X' \subset X$ is now dense, and one can show that T is closed if and only if T' is closed.

Note that the assumption in Remark 19.7 that $\bar{\mathcal{D}} \subset X$ is complemented is *always* true in the setting of greatest interest to us, namely when X is a Hilbert space. We will almost always assume in the following that the domain $\mathcal{D}(T) \subset X$ of an unbounded operator is dense.

While closed unbounded operators $X \supset \mathcal{D} \xrightarrow{T} Y$ are not generally continuous, the condition of having a closed graph turns out to provide enough control for a reasonable theory, and it will be assumed in all the important results we prove about unbounded operators. The following example shows that whether an operator is closed or not can depend heavily on the choice of dense subspace that is considered to be its domain: not all choices of domain are equivalent, and the simplest choice is not always the right one.

Example 19.8. Let T_0 denote the Laplace operator $-\Delta$, regarded as an unbounded operator

$$L^2(\mathbb{R}^n) \supset \mathcal{D}_0 := \mathscr{S}(\mathbb{R}^n) \xrightarrow{T_0 := -\Delta} L^2(\mathbb{R}^n),$$

and let T_1 denote the same operator but with a different choice of domain, namely the Sobolev space $H^2(\mathbb{R}^n)$,

$$L^2(\mathbb{R}^n) \supset \mathcal{D}_1 := H^2(\mathbb{R}^n) \xrightarrow{T_1 := -\Delta} L^2(\mathbb{R}^n),$$

The domains in both cases are dense in $L^2(\mathbb{R}^n)$, and we have $\mathcal{D}_0 \subset \mathcal{D}_1$, but we claim

$$\bar{\Gamma}_{T_0} \supset \Gamma_{T_1},$$

implying that T_0 is not closed. To see this, consider an arbitrary element $(f, -\Delta f) \in \Gamma_{T_1}$, so $f \in H^2(\mathbb{R}^n)$, and notice that since $\mathscr{S}(\mathbb{R}^n)$ is dense in $H^2(\mathbb{R}^n)$ with respect to the H^2-norm, there exists a sequence $f_j \in \mathscr{S}(\mathbb{R}^n)$ with $f_j \xrightarrow{H^2} f$, implying $-\Delta f_j \xrightarrow{L^2} -\Delta f$ since $-\Delta$ defines a bounded linear operator $H^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$. The H^2-convergence of f_j also implies $f_j \xrightarrow{L^2} f$, thus

$$\Gamma_{T_0} \ni (f_j, -\Delta f_j) \to (f, -\Delta f) \quad \text{in } L^2(\mathbb{R}^n) \oplus L^2(\mathbb{R}^n),$$

proving that $(f, -\Delta f)$ lies in the closure of Γ_{T_0}.

We claim however that T_1 is closed, and thus $\bar{T}_0 = T_1$. Indeed, suppose $(f_j, -\Delta f_j) \in \Gamma_{T_1}$ is a sequence converging to $(f, g) \in L^2(\mathbb{R}^n) \oplus L^2(\mathbb{R}^n)$, meaning $f_j \xrightarrow{L^2} f$ and $-\Delta f_j \xrightarrow{L^2} g$. Taking Fourier transforms, it follows that $\widehat{f}_j \xrightarrow{L^2} \widehat{f}$ and

$$-\widehat{\Delta f_j} = 4\pi^2 |p|^2 \widehat{f}_j \xrightarrow{L^2} \widehat{g},$$

thus

$$(1 + |p|^2)\widehat{f}_j = \widehat{f}_j - \frac{1}{4\pi^2}\widehat{\Delta f_j} \xrightarrow{L^2} \widehat{f} + \frac{1}{4\pi^2}\widehat{g}.$$

This proves that $(1 + |p|^2)\widehat{f}_j$ is an L^2-Cauchy sequence, thus f_j is an H^2-Cauchy sequence, and its L^2-limit f is therefore also in $H^2(\mathbb{R}^n)$, with H^2-convergence $f_j \to f$. Since $-\Delta : H^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ is bounded, it follows that $-\Delta f_j \xrightarrow{L^2} -\Delta f$, thus $-\Delta f = g$ and $(f, g) \in \Gamma_{T_1}$.

For the purposes of spectral theory, the next result gives a good reason to confine our attention to closed operators:

Lemma 19.9. *For a closed operator $X \supset \mathcal{D} \xrightarrow{T} X$ and any scalar $\lambda \in \mathbb{K}$ such that $\lambda - T : \mathcal{D} \to X$ is bijective, its inverse*

$$\begin{tikzcd} X \arrow[r, "(\lambda - T)^{-1}"'] \arrow[rr, bend left, "R_\lambda(T)"] & \mathcal{D} \arrow[r, hook] & X \end{tikzcd}$$

determines a bounded linear operator $R_\lambda(T) \in \mathscr{L}(X)$.

Proof. We claim first that if T is closed then $\lambda - T$ is also a closed operator with the same domain $\mathcal{D} \subset X$ for every $\lambda \in \mathbb{K}$. Indeed, suppose $(x_n, \lambda x_n - T x_n)$ is a sequence in the graph of $\lambda - T$ converging to $(x, y) \in X \oplus X$, meaning $x_n \in \mathcal{D}$, $x_n \to x \in X$ and $\lambda x_n - T x_n \to y$, with the convergence of the latter two sequences in the norm of X. Then $T x_n = \lambda x_n - (\lambda x_n - T x_n) \to \lambda x - y$, so the assumption that T is closed implies $(x, \lambda x - y) \in \Gamma_T$, meaning $x \in \mathcal{D}$ and $Tx = \lambda x - y$. In other words, $(\lambda - T)x = y$, and (x, y) is therefore in the graph of $\lambda - T$.

Now if $\lambda - T : \mathcal{D} \to X$ is bijective, then the graph of $R_\lambda(T) : X \to X$ is

$$\begin{aligned} \Gamma_{R_\lambda(T)} &= \left\{ \left(x, (\lambda - T)^{-1} x\right) \in X \oplus X \;\middle|\; x \in X \right\} = \left\{ ((\lambda - T)y, y) \in X \oplus X \;\middle|\; y \in \mathcal{D} \right\} \\ &= \left\{ (x, y) \in X \oplus X \;\middle|\; (y, x) \in \Gamma_{\lambda - T} \right\}, \end{aligned}$$

which is a closed subspace of $X \oplus X$ since $\lambda - T$ is closed. It follows from the closed graph theorem (Theorem 16.11) that $R_\lambda(T)$ is bounded. ☐

19.2. Resolvent and spectrum. We assume from now on that the Banach space X is complex.

DEFINITION 19.10. For a complex-linear unbounded operator $X \supset \mathcal{D} \xrightarrow{T} X$ that is closed, we define its **resolvent set**

$$\rho(T) := \left\{ \lambda \in \mathbb{C} \;\middle|\; \lambda - T \text{ is a bijection } \mathcal{D} \to X \right\},$$

the complement of which is its **spectrum**

$$\sigma(T) := \mathbb{C} \backslash \rho(T).$$

For $\lambda \in \rho(T)$, the bounded linear operator $R_\lambda(T) \in \mathscr{L}(X)$ defined in Lemma 19.9 by inverting $\lambda - T$ is called the **resolvent** of T at λ.

We say that $\lambda \in \sigma(T)$ belongs to the **point spectrum** or is an **eigenvalue** of T if $\ker(\lambda - T)$ is a nontrivial subspace of $\mathcal{D}$, and it belongs to the **residual spectrum** if it is not an eigenvalue but the image of the linear map

$$\mathcal{D} \xrightarrow{\lambda - T} X$$

is not dense in X.

EXERCISE 19.11. For a closed operator $X \supset \mathcal{D} \xrightarrow{T} X$, prove:

(a) T is injective with closed image if and only if there exists a constant $c > 0$ such that $\|Tx\| \geqslant c\|x\|$.
Hint: The argument in the proof of Lemma 19.9 may be helpful.

(b) If $\lambda \in \sigma(T)$ is not in the residual spectrum of T, then there exists a sequence $x_n \in \mathcal{D}$ with $\|x_n\| = 1$ and $(\lambda - T)x_n \to 0$ in X.

REMARK 19.12. In the general formalism of unbounded operators, the domain $\mathcal{D} \subset X$ is not assumed to be equipped with any norm or topology other than what it inherits naturally from X, though in specific examples (such as $\mathcal{D}_1 = H^2(\mathbb{R}^n)$ in Example 19.8), it can happen that $\mathcal{D}$ is also a Banach or Hilbert space with a norm that is strictly stronger than the norm of X. In cases like that, the norm under discussion in general results and concepts such as Exercise 19.11 and Definition 19.1 is still the norm that is inherited from X.

CONVENTION 19.13. If T is not closed but is closable, we define $\sigma(T)$ to be the spectrum of its closure $\bar{T}$.

THEOREM 19.14. *For any closed complex-linear unbounded operator $X \supset \mathcal{D} \xrightarrow{T} X$, the resolvent set $\rho(T) \subset \mathbb{C}$ is open, and the function $\rho(T) \to \mathscr{L}(X) : \lambda \mapsto R_\lambda(T)$ is complex analytic.*

PROOF. The main idea is the same as in the bounded case, one just has to check that it really works. Given $\lambda_0 \in \rho(T)$ and $\mu \in \mathbb{C}$ sufficiently close to 0, let $\lambda := \lambda_0 + \mu$, and observe that $\lambda - T = (\lambda_0 - T) + \mu$ is a *bounded* small perturbation of $\lambda_0 - T$. The openness of $\rho(T)$ then follows from a general observation: If $A_0 : \mathcal{D} \to X$ is a (possibly unbounded) linear map with a bounded inverse $A_0^{-1} : X \to \mathcal{D} \hookrightarrow X$, and $B \in \mathscr{L}(X)$ is any bounded operator with $\|B\|$ sufficiently small, then we can regard $A_0^{-1}B$ as an element of $\mathscr{L}(X)$ with $\|A_0^{-1}B\| < 1$ and similarly assume $\|BA_0^{-1}\| < 1$, in which case the usual power series trick gives

$$(\mathbb{1} + A_0^{-1}B)^{-1} = \sum_{k=0}^{\infty} (-1)^k (A_0^{-1}B)^k \in \mathscr{L}(X),$$

with which we can define

$$(\mathbb{1} + A_0^{-1}B)^{-1}A_0^{-1} = A_0^{-1} - A_0^{-1}BA_0^{-1} + A_0^{-1}BA_0^{-1}BA_0^{-1} - \ldots = A_0^{-1}(\mathbb{1} + BA_0^{-1})^{-1} \in \mathscr{L}(X).$$

This operator has image in $\mathcal{D}$, and one easily checks that it is a two-sided inverse for $A_0 + B : \mathcal{D} \to X$. Applying this with $A_0 := \lambda_0 - T$ and $B := \mu$ proves that $\rho(T) \subset \mathbb{C}$ is open, and we also obtain the formula

$$R_\lambda(T) = \sum_{k=0}^{\infty} (-1)^k \mu^k R_{\lambda_0}^{k+1},$$

which shows that $\lambda \mapsto R_\lambda(T)$ is analytic. $\square$

As in the bounded case, Theorem 19.14 proves that the spectrum

$$\sigma(T) \subset \mathbb{C}$$

of a closed unbounded operator is always a closed set. In contrast to the setting of bounded operators, $\sigma(T)$ will *not* generally be bounded, as there is no way of guaranteeing that

$$\lambda - T = \lambda\left(\mathbb{1} - \frac{1}{\lambda}T\right)$$

is invertible for $|\lambda|$ large if we do not have a finite norm $\|T\|$. For the same reason, we cannot prove that $\|R_\lambda(T)\|$ becomes small as $|\lambda| \to \infty$. You may recall from Corollary 18.13 that our proof that $\sigma(T) \neq \varnothing$ in the bounded case depended on that fact; in the unbounded case, it can also happen that the spectrum is empty. The following example illustrates both pathologies.

Exercise 19.15. Denote by T_0 and T_1 the differential operator $i\frac{d}{dt}$ defined on the following domains in $L^2([0,1])$:

$$\mathcal{D}_0 := \mathcal{D}(T_0) := \left\{ f : [0,1] \to \mathbb{C} \;\middle|\; f \text{ is absolutely continuous and } f' \in L^2([0,1]) \right\},$$

and

$$\mathcal{D}_1 := \mathcal{D}(T_1) := \left\{ f \in \mathcal{D}_0 \;\middle|\; f(0) = 0 \right\}.$$

(a) Show that both domains $\mathcal{D}_0$ and $\mathcal{D}_1$ are dense in $L^2([0,1])$.
(b) Show that both operators T_0 and T_1 are closed.
Hint: You can use the fundamental theorem of calculus.
(c) For all $\lambda \in \mathbb{C}$, find an eigenfunction of T_0 with eigenvalue λ, proving

$$\sigma(T_0) = \mathbb{C}.$$

Do the eigenfunctions belong to the domain $\mathcal{D}_1$?
(d) For $\lambda \in \mathbb{C}$, the equation $(\lambda - T_1)f = g$ translates into the initial value problem

$$f'(t) = i\left[-\lambda f(t) + g(t)\right],$$
$$f(0) = 0,$$

with the understanding that for a given $g \in L^2([0,1])$, we are seeking absolutely continuous solutions $f : [0,1] \to \mathbb{C}$ that satisfy the differential equation almost everywhere. Adapt the Picard-Lindelöf theorem for this situation and prove that for every $g \in L^2([0,1])$, this problem really does have a unique absolutely continuous solution, thus $(\lambda - T_1) : \mathcal{D}_1 \to L^2([0,1])$ is bijective, and

$$\sigma(T_1) = \varnothing.$$

Remark 19.16. The domain $\mathcal{D}_0$ in Exercise 19.15 is secretly equivalent to the Sobolev space $W^{1,2}((0,1))$. This follows mostly from Exercises 13.51 and 13.56: The former gives an inclusion $\mathcal{D}_0 \hookrightarrow W^{1,2}((0,1))$, and also shows that every equivalence class in $W^{1,2}((0,1))$ has a unique representative that is absolutely continuous on compact subsets, and whose derivatives almost everywhere match their weak derivatives. The latter implies in turn that these continuous functions are also in $C^{0,\frac{1}{2}}((0,1))$, thus they are uniformly continuous on $(0,1)$, and therefore admit continuous extensions over $[0,1]$. One can deduce from the fundamental theorem of calculus that the extensions are also absolutely continuous.

19.3. Adjoints. From now on, we work in a complex Hilbert space $\mathcal{H}$. The issue of domains makes the notion of *adjoints* for unbounded operators on Hilbert spaces a bit subtle.

DEFINITION 19.17. Assume $\mathcal{H} \supset \mathcal{D}(T) \xrightarrow{T} \mathcal{H}$ is a densely-defined unbounded operator on the complex Hilbert space $\mathcal{H}$, i.e. its domain $\mathcal{D}(T) \subset \mathcal{H}$ is a dense subspace. The **adjoint** of T is then the unique unbounded operator

$$\mathcal{H} \supset \mathcal{D}(T^*) \xrightarrow{T^*} \mathcal{H}$$

such that:

(i) $\langle T^* x, y \rangle = \langle x, Ty \rangle$ for all $x \in \mathcal{D}(T^*)$ and $y \in \mathcal{D}(T)$;
(ii) T^* cannot be extended to any larger domain than $\mathcal{D}(T^*)$ such that condition (i) still holds.

In other words, the domain of the adjoint of T is the set

$$\mathcal{D}(T^*) = \left\{ x \in \mathcal{H} \;\middle|\; \text{there exists } z \in \mathcal{H} \text{ such that } \langle z, y \rangle = \langle x, Ty \rangle \text{ for all } y \in \mathcal{D}(T) \right\}.$$

Indeed, for any $x \in \mathcal{H}$ satisfying this condition, the linear functional $\Lambda : \mathcal{D}(T) \to \mathbb{C}$ defined by $\Lambda(y) := \langle x, Ty \rangle$ satisfies $|\Lambda(y)| \leqslant \|z\| \cdot \|y\|$ and is thus bounded, and since $\mathcal{D}(T) \subset \mathcal{H}$ is dense, Λ then has a unique extension to $\mathcal{H}$, defining an element of the dual space $\Lambda \in \mathcal{H}^*$. This forces us to define

$$T^* x := z \qquad \text{where} \qquad \Lambda = \langle z, \cdot \rangle,$$

i.e. T^* is the composition of the complex-antilinear map $\mathcal{D}(T^*) \to \mathcal{H}^* : x \mapsto \langle x, T\cdot \rangle$ with the natural complex-antilinear isomorphism $\mathcal{H}^* \to \mathcal{H}$ arising from the Riesz representation theorem. In particular, assuming $\mathcal{D}(T) \subset \mathcal{H}$ to be dense ensures that T^* is uniquely determined by T; this uniqueness would not otherwise be guaranteed.

DEFINITION 19.18. The operator $\mathcal{H} \supset \mathcal{D}(T) \xrightarrow{T} \mathcal{H}$ is **symmetric** if it satisfies

$$\langle Tx, y \rangle = \langle x, Ty \rangle \qquad \text{for all} \qquad x, y \in \mathcal{D}(T),$$

and it is **self-adjoint** if $T^* = T$, meaning that both the operator T^* and its domain $\mathcal{D}(T^*)$ are the same as T and $\mathcal{D}(T)$ respectively.

For bounded operators, symmetry and self-adjointness are the same thing, but for unbounded operators, this is not so. A symmetric operator need not be identical to its adjoint, because the latter may have a strictly larger domain $\mathcal{D}(T^*) \supset \mathcal{D}(T)$; the most that can be said in general is that if T is symmetric, then T^* is an *extension* of T.

EXERCISE 19.19. For a further variation on Exercise 19.15, consider the operator $T_2 := i\frac{d}{dt}$ on $L^2([0,1])$ with domain

$$\mathcal{D}_2 := \left\{ f \in \mathcal{D}_0 \;\middle|\; f(0) = f(1) = 0 \right\} \subset L^2([0,1]),$$

where $\mathcal{D}_0$ is again the space of absolutely continuous functions whose derivatives are L^2-functions defined almost everywhere. Use integration by parts to show that this operator satisfies

$$\langle f, T_2 g \rangle_{L^2} = \langle T_0 f, g \rangle_{L^2}$$

for all $f \in \mathcal{D}_0$ and $g \in \mathcal{D}_2$. This shows that T_2 is symmetric, but the domain of T_2^* also contains $\mathcal{D}_0$ and is thus strictly larger than that of T_2, thus T_2 is not self-adjoint.

EXERCISE 19.20. Show that the operator T_2 in Exercise 19.19 has the entirety of $\mathbb{C}$ as residual spectrum.

EXAMPLE 19.21. We saw in Example 19.8 that the operator

$$L^2(\mathbb{R}^n) \supset H^2(\mathbb{R}^n) \xrightarrow{-\Delta} L^2(\mathbb{R}^n)$$

is closed; we claim that it is also self-adjoint. We've seen already that $-\Delta$ is symmetric when its domain is taken to be $\mathscr{S}(\mathbb{R}^n)$; since the latter is dense in $H^2(\mathbb{R}^n)$, it is also symmetric on $H^2(\mathbb{R}^n)$. We therefore need to show that $H^2(\mathbb{R}^n)$ is also the domain of the adjoint $-\Delta^*$, i.e. that whenever there exist $f, g \in L^2(\mathbb{R}^n)$ satisfying

$$\langle g, h\rangle_{L^2} = \langle f, -\Delta h\rangle_{L^2} \qquad \text{for all} \qquad h \in H^2(\mathbb{R}^n), \tag{19.2}$$

it follows that $f \in H^2(\mathbb{R}^n)$. If (19.2) holds, then it holds in particular for all $h \in C_0^\infty(\mathbb{R}^n)$ and can then be interpreted to mean that the equation $-\Delta f = g$ holds in the sense of weak derivatives, and therefore also as a relation between two tempered distributions. Both sides then have well-defined Fourier transforms in $\mathscr{S}'(\mathbb{R}^n)$, giving the relation

$$4\pi^2|p|^2\hat{f} = \hat{g} \in \mathscr{S}'(\mathbb{R}^n),$$

and since f, g and their Fourier transforms are also L^2-functions, we can interpret this as an equality between two functions on $\mathbb{R}^n$, understood to hold almost everywhere. It follows that

$$\|f\|_{H^2}^2 = \int_{\mathbb{R}^n} (1 + |p|^2)|\hat{f}(p)|^2\, dp = \|\hat{f}\|_{L^2}^2 + \frac{1}{4\pi^2}\|\hat{g}\|_{L^2}^2 < \infty,$$

so $f \in H^2(\mathbb{R}^n)$.

The next example should give you a big hint as to what the spectral theorem for unbounded self-adjoint operators might look like.

EXAMPLE 19.22 (unbounded multiplication operators). Assume (X, μ) is a semifinite measure space and $F : X \to \mathbb{C}$ is a measurable function that is not assumed bounded. In the Hilbert space $L^2(X, \mu)$ of complex-valued L^2-functions on (X, μ), define the subspace

$$\mathcal{D}_F := \left\{u \in L^2(X, \mu) \;\middle|\; Fu \in L^2(X, \mu)\right\} \subset L^2(X, \mu).$$

We leave it as an exercise to show that $\mathcal{D}_F$ is dense in $L^2(X, \mu)$, and Exercise 19.23 below shows that the unbounded "multiplication" operator

$$L^2(X, \mu) \supset \mathcal{D}_F \xrightarrow{T_F} L^2(X, \mu) : u \mapsto Fu$$

is always closed. We claim that its adjoint is $T_{\bar{F}}$ with the same domain $\mathcal{D}(T_F^*) = \mathcal{D}_F$, thus T_F is self-adjoint whenever F is real valued. The relation

$$\langle \bar{F}u, v\rangle_{L^2} = \langle u, Fv\rangle_{L^2} \qquad \text{for all} \qquad u, v \in \mathcal{D}_F$$

is clear, so T_F is certainly *symmetric* if F is real valued. What remains to be shown is that whenever $u, v \in L^2(X, \mu)$ satisfy

$$\langle u, Fw\rangle_{L^2} = \langle v, w\rangle_{L^2} \qquad \text{for all} \qquad w \in \mathcal{D}_F,$$

it follows that $u \in \mathcal{D}_F$. To see this, define for each $N \in \mathbb{N}$ the characteristic function

$$\chi_N := \chi_{F^{-1}([-N,N])} : X \to [0, \infty),$$

so that $\chi_N F$ is bounded, and for all $w \in \mathcal{D}_F$, we then have

$$\langle \chi_N \bar{F}u, w\rangle_{L^2} = \langle u, F\chi_N w\rangle_{L^2} = \langle v, \chi_N w\rangle_{L^2} = \langle \chi_N v, w\rangle_{L^2},$$

which implies since $\mathcal{D}_F \subset L^2(X, \mu)$ is dense that $\chi_N \bar{F}u = \chi_N v$ almost everywhere for all $N \in \mathbb{N}$. By the monotone convergence theorem, we have $\|\chi_N v\|_{L^2} \to \|v\|_{L^2}$ and $\|\chi_N \bar{F}u\|_{L^2} \to \|\bar{F}u\|_{L^2}$ as $N \to \infty$, thus $\|\bar{F}u\|_{L^2} = \|v\|_{L^2} < \infty$, and since $\|\bar{F}u\|_{L^2} = \big\||F|u\big\|_{L^2} = \|Fu\|_{L^2}$, it follows that $u \in \mathcal{D}_F$.

Exercise 19.23. For the unbounded multiplication operator T_F in Example 19.22, show:

(a) T_F is closed.

(b) $\sigma(T_F)$ is the essential range of $F : X \to \mathbb{C}$, just as in the bounded case (cf. Exercise 18.34).

19.4. Spectral representations. We will not discuss unbounded *normal* operators in this course; the concept exists, but the issue of domains makes it tricky to make sense of conditions like $A^*A = AA^*$. We nonetheless have some important things to say about unbounded *self-adjoint* operators, and after Example 19.22, you can perhaps guess the statement of the main theorem.

Definition 19.24. A **spectral representation** of an unbounded operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ on a complex Hilbert space $\mathcal{H}$ consists of a semifinite measure space (X, μ), a measurable function $F : X \to \mathbb{C}$, and a unitary isomorphism

$$\mathcal{H} \xrightarrow{U} L^2(X, \mu)$$

such that

$$U(\mathcal{D}) = \mathcal{D}_F := \left\{ u \in L^2(X, \mu) \;\middle|\; Fu \in L^2(X, \mu) \right\}$$

and $UAU^{-1} =: T_F$ is the "multiplication" operator

$$L^2(X, \mu) \supset \mathcal{D}_F \xrightarrow{T_F} L^2(X, \mu) : u \mapsto Fu.$$

Example 19.25. For the Laplace operator $L^2(\mathbb{R}^n) \supset H^2(\mathbb{R}^n) \xrightarrow{-\Delta} L^2(\mathbb{R}^n)$ considered in Examples 19.8 and 19.21, the Fourier transform $\mathscr{F} : L^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ defines a spectral representation,

$$\mathscr{F}(-\Delta)\mathscr{F}^* = T_F \qquad \text{for} \qquad F : \mathbb{R}^n \to \mathbb{R} : p \mapsto 4\pi^2 |p|^2,$$

where the multiplication operator has domain

$$\mathscr{F}(H^2(\mathbb{R}^n)) = \left\{ u \in L^2(\mathbb{R}^n) \;\middle|\; |p|^2 u \in L^2(\mathbb{R}^n) \right\}.$$

Theorem 19.26 (Spectral theorem for unbounded self-adjoint operators). *Every unbounded self-adjoint operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ on a complex Hilbert space $\mathcal{H}$ admits a spectral representation, identifying it unitarily with the multiplication operator $L^2(X, \mu) \supset \mathcal{D}_F \xrightarrow{T_F} L^2(X, \mu)$ determined by some real-valued measurable function $F : X \to \mathbb{R}$ on a semifinite measure space (X, μ). Moreover, we can arrange $\mu(X) < \infty$ if $\mathcal{H}$ is separable.*

Maybe you can see now why it is necessary to be so nitpicky about domains: Theorem 19.26 clearly *requires* the operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ to be self-adjoint—not just symmetric—since by the argument in Example 19.22, the multiplication operator with which we want to identify it is self-adjoint. We've seen that there are also symmetric operators that are not self-adjoint, but for these, the theorem cannot be true.

Let us take note of some things that follow from Theorem 19.26 in conjunction with Exercise 19.23 that were not obvious up to this point:

Corollary 19.27. *For any unbounded self-adjoint operator A, the spectrum $\sigma(A)$ is a nonempty subset of $\mathbb{R}$.* □

Corollary 19.28. *A self-adjoint operator is bounded if and only if its spectrum $\sigma(A) \subset \mathbb{R}$ is bounded.* □

There is a cheap trick for reducing the proof of Theorem 19.26 to the spectral theorem for bounded *normal* operators. This is the secret reason—aside from issues of elegance and naturality—why we went through the effort of proving the main results of §18 for normal and not just self-adjoint operators. We'll need four lemmas, the last of which establishes the theorem.

LEMMA 19.29. *For any unbounded operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ on a complex Hilbert space, the adjoint A^* is closed.*

PROOF. Suppose $x_n \in \mathcal{D}(A^*)$ is a sequence converging to $x \in \mathcal{H}$ with $A^* x_n \to y \in \mathcal{H}$; then for every $z \in \mathcal{D} = \mathcal{D}(A)$,

$$\langle A^* x_n, z \rangle = \langle x_n, Az \rangle \to \langle x, Az \rangle,$$

and also

$$\langle A^* x_n, z \rangle \to \langle y, z \rangle,$$

implying $\langle y, z \rangle = \langle x, Az \rangle$ and thus $x \in \mathcal{D}(A^*)$, with $A^* x = y$. In other words, (x, y) belongs to the graph of A^*, which is therefore closed. □

COROLLARY 19.30. *Every densely-defined symmetric unbounded operator is closable.* □

Indeed, the adjoint of a symmetric operator is always a closed extension of it. In particular:

COROLLARY 19.31. *Every self-adjoint operator is closed.* □

LEMMA 19.32. *For a symmetric unbounded operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$, if there exists a sequence $x_n \in \mathcal{D}$ with $\|x_n\| = 1$ for all n and $(\lambda - A)x_n \to 0$ for some $\lambda \in \mathbb{C}$, then λ is real.*

PROOF. Let $\lambda = \alpha + i\beta \in \mathbb{C}$ for $\alpha, \beta \in \mathbb{R}$ with $\beta \neq 0$. Then for $x \in \mathcal{D}$, the result follows from the estimate

$$\begin{aligned}
\|(A - \lambda)x\|^2 &= \langle (A - \alpha)x - i\beta x, (A - \alpha)x - i\beta x \rangle \\
&= \|(A - \alpha)x\|^2 + \beta^2 \|x\|^2 - i\beta \langle (A - \alpha)x, x \rangle + i\beta \langle x, (A - \alpha)x \rangle \\
&= \|(A - \alpha)x\|^2 + \beta^2 \|x\|^2 \geqslant \beta^2 \|x\|^2.
\end{aligned}$$

□

EXAMPLE 19.33. The operator $T_2 = i\frac{d}{dt}$ of Exercise 19.19, with domain $\mathcal{D}_2 \subset L^2([0,1])$ taken to be the absolutely continuous functions that vanish at the end points and have derivatives in L^2, is symmetric but not self-adjoint, and by Exercise 19.20, its spectrum is the entirety of $\mathbb{C}$, and is all residual. This stands in contrast to the next lemma.

LEMMA 19.34. *For a self-adjoint operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$, the residual spectrum is empty.*

PROOF. Suppose $\lambda \in \sigma(A)$ is not an eigenvalue but the image of the operator $A - \lambda : \mathcal{D} \to \mathcal{H}$ is not dense. Then we find $v \neq 0 \in \mathcal{H}$ with $\langle (A - \lambda)x, v \rangle = 0$ for all $x \in \mathcal{D}$, and we therefore have

$$\langle v, Ax \rangle = \langle v, \lambda x \rangle = \langle \bar{\lambda} v, x \rangle \qquad \text{for all} \qquad x \in \mathcal{D},$$

which implies $v \in \mathcal{D}(A^*)$ and $A^* v = \bar{\lambda} v$. Since $A^* = A$, we conclude $\bar{\lambda}$ is an eigenvalue of A, so that Lemma 19.32 implies that it must be real, meaning λ is also an eigenvalue of A, which is a contradiction. □

Combining this result with Lemma 19.32 and Exercise 19.11 gives the following consequence, this time without needing to assume the result of Theorem 19.26:

COROLLARY 19.35. *For A self-adjoint, $\sigma(A) \subset \mathbb{R}$.* □

Now we come to the main step:

LEMMA 19.36. *Suppose $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ is a closed and symmetric unbounded operator whose spectrum contains neither λ nor $\bar{\lambda}$ for some $\lambda \in \mathbb{C}$. Then A is self-adjoint and satisfies the conclusions of Theorem 19.26.*

Proof. Define the bounded linear operators

$$T_+ := -R_\lambda(A) = (A-\lambda)^{-1} \in \mathscr{L}(X), \qquad \text{and} \qquad T_- := -R_{\bar\lambda}(A) = (A-\bar\lambda)^{-1} \in \mathscr{L}(X).$$

We claim that $T_+^* = T_-$ and $T_+T_- = T_-T_+$, thus both are normal. Indeed, for $x, y \in \mathcal{H}$, we have

$$\langle x, T_+y\rangle = \langle (A-\bar\lambda)T_-x, T_+y\rangle = \langle T_-x, (A-\lambda)T_+y\rangle = \langle T_-x, y\rangle,$$

where we've used the assumption that A is symmetric and the fact that $\operatorname{im}(T_\pm) \subset \mathcal{D}$. This proves $T_+^* = T_-$.

To establish commutativity, fix $x \in \mathcal{H}$, let $y = T_+T_-x$ and $z = T_-T_+x$. Since $\operatorname{im}(T_\pm) \subset \mathcal{D}$, we have $(A-\lambda)y = T_-x \in \mathcal{D}$ and thus $Ay \in \mathcal{D}$, so that $A^2y \in \mathcal{H}$ is well defined. For the same reasons, $A^2z \in \mathcal{H}$ is well defined, and we now have

$$\begin{aligned} x &= (A-\bar\lambda)(A-\lambda)y = (A^2 + |\lambda|^2 - 2(\operatorname{Re}\lambda)A)y \\ &= (A-\lambda)(A-\bar\lambda)z = (A^2 + |\lambda|^2 - 2(\operatorname{Re}\lambda)A)z. \end{aligned}$$

Since $A-\lambda$ and $A-\bar\lambda$ are both injective by hypothesis, the operator $A^2 + |\lambda|^2 - 2(\operatorname{Re}\lambda)A$ must also be injective, and we conclude $y = z$, completing the proof of the claim.

We can now extract from the spectral theorem for bounded normal operators a semifinite measure space (X, μ), which is finite if $\mathcal{H}$ is separable, along with a bounded measurable function $G : X \to \mathbb{C}$ and a unitary isomorphism $U : \mathcal{H} \to L^2(X, \mu)$ such that

$$UT_+U^{-1} = T_G : u \mapsto Gu.$$

Since T_+ is a bijection $\mathcal{H} \to \mathcal{D}$, it does not have 0 as an eigenvalue, so the set $G^{-1}(0) \subset X$ has measure zero in (X, μ), and the function $1/G$ is therefore defined almost everywhere. Moreover, $\operatorname{im}(T_+) = \mathcal{D}$, so that

$$U(\mathcal{D}) = \left\{ Gu \in L^2(X,\mu) \;\middle|\; u \in L^2(X,\mu) \right\}.$$

For $u \in U(\mathcal{D}) \subset L^2(X, \mu)$, we then have

$$\frac{1}{G}u = U(A-\lambda)U^{-1}u = UAU^{-1}u - \lambda u,$$

implying

$$UAU^{-1}u = Fu \qquad \text{where} \qquad F := \frac{1}{G} + \lambda.$$

The function F is well defined almost everywhere, which makes the multiplication operator T_F well defined on $U(\mathcal{D})$, the latter being the set of all functions u of the form Gv for $v \in L^2(X, \mu)$, which can also be characterized via the condition $Fu \in L^2(X, \mu)$.

Finally, we observe that a multiplication operator T_F can only be symmetric if $F \equiv \bar F$ almost everywhere, so after modifying F on a set of measure zero, we are free to assume that it is real valued. This completes the construction of a spectral representation for A, and we conclude also from Example 19.22 that A is self-adjoint. □

Proof of Theorem 19.26. Assuming A is self-adjoint, Corollary 19.31 implies that it is closed, and Corollary 19.35 implies that neither i nor $-i$ is in its spectrum; the result then follows immediately from Lemma 19.36. □

19.5. Criteria for self-adjointness. Lemma 19.36 has some other useful applications, such as a technical result that gives more practical methods for recognizing when an operator is self-adjoint. The following result and the related Theorem 19.40 are each sometimes referred to as the "basic criterion for self-adjointness". Both are often applied specifically with $\lambda := \pm i$, so that the two operators under consideration are $A \pm i$.

COROLLARY 19.37. *A closed symmetric operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ is self-adjoint if and only if for some $\lambda \in \mathbb{C}\backslash\mathbb{R}$, both of the operators $A - \lambda$ and $A - \bar{\lambda}$ from $\mathcal{D}$ to $\mathcal{H}$ are surjective.* $\square$

Recall from Corollary 19.30 that every densely-defined symmetric operator is closable.

DEFINITION 19.38. A densely-defined symmetric unbounded operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ is called **essentially self-adjoint** if its closure is self-adjoint.

EXERCISE 19.39. Show that if A is symmetric and closable, then its closure is also symmetric.

We now derive from Corollary 19.37 a weaker criterion that is somewhat easier to use in practice.

THEOREM 19.40. *A densely-defined symmetric operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ is essentially self-adjoint if and only if for some $\lambda \in \mathbb{C}\backslash\mathbb{R}$, both of the operators $A - \lambda$ and $A - \bar{\lambda}$ from $\mathcal{D}$ to $\mathcal{H}$ have dense image.*

PROOF. Suppose first that A is closed, in addition to being symmetric and densely defined. Then for $\lambda := \alpha + i\beta$ with $\alpha, \beta \in \mathbb{R}$ and $\beta \neq 0$, the operator $A' := A - \alpha$ with the same domain is also closed and symmetric, and the calculation in the proof of Lemma 19.32 implies that for all $x \in \mathcal{D}$,

$$\left\| \left((A - (\alpha \pm i\beta) \right) x \right\|^2 = \|(A' \mp i\beta)x\|^2 = \|A'x\|^2 + \beta^2\|x\|^2,$$

which is bounded below by a constant times the square of the graph norm $\|x\|_{A'} := \|x\| + \|A'x\|$ on $\mathcal{D}$, giving the estimate

$$\|(A' \mp i\beta)x\| \geqslant c\|x\|_{A'} \tag{19.3}$$

for some constant $c > 0$. By Exercise 19.6, the graph norm on $\mathcal{D}$ is complete, and moreover, the operators $A' \pm i\beta : \mathcal{D} \to \mathcal{H}$ are bounded with respect to the graph norm, so the estimate (19.3) implies via Exercise 17.25 that both have closed image.

With that understood, assume A is densely-defined and symmetric, but not necessarily closed. By Exercise 19.39, its closure $\bar{A}$ is then also symmetric, so the previous paragraph implies that $\operatorname{im}(\bar{A} - \lambda)$ and $\operatorname{im}(\bar{A} - \bar{\lambda})$ are closed in $\mathcal{H}$, and Corollary 19.37 tells us that A is then essentially self-adjoint if and only if these images are also dense in $\mathcal{H}$, which is true if and only if $\operatorname{im}(A - \lambda) \subset \mathcal{H}$ and $\operatorname{im}(A - \bar{\lambda}) \subset \mathcal{H}$ are dense. $\square$

EXAMPLE 19.41. A straightforward integration by parts shows that the operator

$$L^2(\mathbb{R}^n) \supset \mathscr{S}(\mathbb{R}^n) \xrightarrow{i\partial_j} L^2(\mathbb{R}^n)$$

is symmetric for any $j \in \{1, \ldots, n\}$. We claim that the operator $\mathscr{S}(\mathbb{R}^n) \xrightarrow{i\partial_j - \lambda} L^2(\mathbb{R}^n)$ has dense image for every $\lambda \in \mathbb{C}\backslash\mathbb{R}$, so that Theorem 19.40 implies that $i\partial_j$ is essentially self-adjoint. Indeed, if not, then there exists $g \neq 0 \in L^2(\mathbb{R}^n)$ such that

$$\langle i\partial_j\varphi - \lambda\varphi, g\rangle_{L^2} = 0 \qquad \text{for all} \qquad \varphi \in \mathscr{S}(\mathbb{R}^n).$$

Taking Fourier transforms, this is equivalent to the condition

$$0 = \langle -2\pi p_j\hat{\varphi} - \lambda\hat{\varphi}, \hat{g}\rangle_{L^2} = \langle (-2\pi p_j - \lambda)\hat{\varphi}, \hat{g}\rangle_{L^2} \qquad \text{for all} \qquad \hat{\varphi} \in \mathscr{S}(\mathbb{R}^n).$$

Since $\lambda \notin \mathbb{R}$, the polynomial function $p \mapsto -2\pi p_j - \lambda$ is nowhere zero, and $(-2\pi p_j \pm i)\widehat{\varphi}$ can therefore be an arbitrary function in the dense subspace $\mathscr{S}(\mathbb{R}^n) \subset L^2(\mathbb{R}^n)$, so this implies $\widehat{g} = 0$ and thus $g = 0$ almost everywhere.

The next exercise is a brainteaser—I recommend thinking about it a bit before you continue reading.

EXERCISE 19.42. What is the domain of the self-adjoint extension of the operator in Example 19.41?

The self-adjoint extension of $L^2(\mathbb{R}^n) \supset \mathscr{S}(\mathbb{R}^n) \xrightarrow{i\partial_j} L^2(\mathbb{R}^n)$ is easy to characterize once you notice the following detail: the Fourier transform provides us with an explicit spectral representation for it! Indeed, $\mathscr{F} : L^2(\mathbb{R}^n) \to L^2(\mathbb{R}^n)$ identifies $\mathscr{S}(\mathbb{R}^n) \subset L^2(\mathbb{R}^n)$ with itself while turning $i\partial_j$ into the multiplication operator, $\mathscr{F}(i\partial_j)\mathscr{F}^* = T_{2\pi p_j}$,

$$L^2(\mathbb{R}^n) \supset \mathscr{S}(\mathbb{R}^n) \xrightarrow{T_{2\pi p_j}} L^2(\mathbb{R}^n).$$

The self-adjoint extension of the multiplication operator has domain

$$\left\{ f \in L^2(\mathbb{R}^n) \;\middle|\; p \mapsto 2\pi p_j f(p) \text{ is an } L^2\text{-function on } \mathbb{R}^n \right\},$$

and applying $\mathscr{F}^*$ to this gives the domain

$$\mathcal{D} = \left\{ f \in L^2(\mathbb{R}^n) \;\middle|\; f \text{ has a weak partial derivative } \partial_j f \text{ in } L^2(\mathbb{R}^n) \right\}$$

for the self-adjoint extension of $L^2(\mathbb{R}^n) \supset \mathscr{S}(\mathbb{R}^n) \xrightarrow{i\partial_j} L^2(\mathbb{R}^n)$. If $n = 1$, then this domain is precisely the Sobolev space $H^1(\mathbb{R})$, though for $n > 1$, it is a strictly larger dense subspace of $L^2(\mathbb{R}^n)$.

19.6. Operators with compact resolvent. Before stating another theorem, let's review the essential reasons why the spectral theorem for unbounded self-adjoint operators holds. Self-adjointness of $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ implies two important things: first, that A is symmetric, and second that its spectrum is contained in $\mathbb{R}$. This implies in particular that $\pm i \notin \sigma(A)$, so that the *bounded* linear operators $T_\pm := R_{\pm i}(A)$ are normal, and a spectral representation of T_+ then gives rise to a spectral representation of A.

It's worth asking: What happens if the normal operators $T_\pm$ are not only bounded but also *compact*?

DEFINITION 19.43. An unbounded operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ is said to have **compact resolvent** if there exists a number $\lambda \in \rho(A) \subset \mathbb{C}$ such that $R_\lambda(A) \in \mathscr{L}(\mathcal{H})$ is a compact operator.

EXAMPLE 19.44. Suppose A is a closed operator of the form

$$L^2(\mathbb{T}^n) \supset H^s(\mathbb{T}^n) \xrightarrow{A} L^2(\mathbb{T}^n)$$

with $\sigma(A) \neq \mathbb{C}$, whose dense domain is the Sobolev space $H^s(\mathbb{T}^n)$ for some $s > 0$, and such that the linear map $A : H^s(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$ is bounded with respect to the H^s-norm. (Note that this doesn't make it a continuous map with respect to the ambient L^2-norm on $H^s(\mathbb{T}^n) \subset L^2(\mathbb{T}^n)$, since the H^s-norm is strictly stronger.) Then for any $\lambda \notin \sigma(A)$, the bijection $\lambda - A : H^s(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$ is also a bounded operator with respect to the H^s-norm, so that by the inverse mapping theorem, its inverse $(\lambda - A)^{-1} : L^2(\mathbb{T}^n) \to H^s(\mathbb{T}^n)$ is bounded. The resolvent $R_\lambda(A) : L^2(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$ is then the composition of a bounded operator with the compact inclusion $H^s(\mathbb{T}^n) \hookrightarrow L^2(\mathbb{T}^n)$,

$$\begin{tikzcd} L^2(\mathbb{T}^n) \arrow[r, "(\lambda-A)^{-1}"'] \arrow[rr, bend left, "R_\lambda(A)"] & H^s(\mathbb{T}^n) \arrow[r, hook] & L^2(\mathbb{T}^n) \end{tikzcd},$$

and is thus compact.

We supplement Example 19.44 with the observation that if A is self-adjoint, then the condition $\sigma(A) \neq \mathbb{C}$ comes for free, since $\sigma(A) \subset \mathbb{R}$.

THEOREM 19.45. *If $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ is an unbounded self-adjoint operator with compact resolvent, then its spectrum consists of countably-many real eigenvalues $\{\lambda_n \in \mathbb{R}\}_{n\in\mathbb{N}}$ such that $|\lambda_n| \to \infty$ as $n \to \infty$, each with finite multiplicity, and $\mathcal{H}$ admits an orthonormal basis consisting of eigenvectors of A.*

PROOF. Suppose $\mu \in \mathbb{C}\backslash\sigma(A)$ such that

$$R_\mu := R_\mu(A) \in \mathscr{L}(\mathcal{H})$$

is compact. Since $\sigma(A) \subset \mathbb{R}$, the complex conjugate of μ is also not in the spectrum, and the proof of Lemma 19.36 then shows that R_μ commutes with $R_\mu^* = R_{\bar{\mu}}(A)$; in particular, R_μ is normal. The spectral theorem for compact normal operators (Exercise 18.37(b)(v)) now gives an orthonormal basis of $\mathcal{H}$ consisting of eigenvectors of R_μ; moreover, by Theorem 18.22, $\sigma(R_\mu)$ is a discrete set of eigenvalues $\{\alpha_n \in \mathbb{C}\}_{n\in\mathbb{N}}$, each with finite multiplicity, except possibly for $0 \in \sigma(R_\mu)$, which is the only accumulation point of $\sigma(R_\mu)$. Since $R_\mu : \mathcal{H} \to \mathcal{D}$ is a bijection by construction, it cannot have 0 as an eigenvalue, and we conclude that *all* of its eigenvalues have finite multiplicity, and they form a nonvanishing sequence that converges to 0.

Now observe:

$$R_\mu v = \alpha v \quad \Leftrightarrow \quad v = \alpha(\mu - A)v \quad \Leftrightarrow \quad \frac{1}{\alpha}v = \mu v - Av \quad \Leftrightarrow \quad Av = \left(\mu - \frac{1}{\alpha}\right)v,$$

thus the eigenspace $\ker(\alpha - R_\mu)$ is also an eigenspace of A, namely $\ker(\lambda - A)$ with $\lambda := \mu - \frac{1}{\alpha}$, and λ must be real by Lemma 19.32 since A is symmetric. Writing $\lambda_n := \mu - 1/\alpha_n$, the convergence $\alpha_n \to 0$ then implies $|\lambda_n| \to \infty$. $\square$

EXAMPLE 19.46. The discussion of the Laplacian $-\Delta$ on $L^2(\Omega)$ at the beginning of this section can now be placed into a wider context. Assuming $\Omega \subset \mathbb{R}^n$ is a bounded open subset, we recall from Theorem 17.18 that $-\Delta$ defines a Banach space isomorphism $H_0^1(\Omega) \to H^{-1}(\Omega)$, and since $L^2(\Omega)$ is dense in $H^{-1}(\Omega)$, the subspace

$$\mathcal{D} := \left\{ f \in H_0^1(\Omega) \;\middle|\; \Delta f \in L^2(\Omega) \right\} \subset H_0^1(\Omega) \subset L^2(\Omega)$$

is dense in $H_0^1(\Omega)$, and therefore also in $L^2(\Omega)$ since $H_0^1(\Omega)$ contains $C_0^\infty(\Omega)$ and the latter is dense in $L^2(\Omega)$. The unbounded operator

$$L^2(\Omega) \supset \mathcal{D} \xrightarrow{-\Delta} L^2(\Omega)$$

is then symmetric and bijective, and the composition of its inverse with the inclusion $\mathcal{D} \hookrightarrow L^2(\Omega)$ can be written as the composition of three bounded operators, one of which is compact,

$$\begin{tikzcd} L^2(\Omega) \arrow[r, hook] \arrow[rrr, bend right, "R_0(-\Delta)"'] & H^{-1}(\Omega) \arrow[r, "-\Delta^{-1}"] & H_0^1(\Omega) \arrow[r, hook] & L^2(\Omega) \end{tikzcd}.$$

This defines a compact resolvent $R_0(-\Delta) \in \mathscr{L}(L^2(\Omega))$, and it also shows that the inverse of $-\Delta : \mathcal{D} \to L^2(\Omega)$ is bounded, implying via the closed graph theorem that the unbounded operator $L^2(\Omega) \supset \mathcal{D} \xrightarrow{-\Delta} L^2(\Omega)$ is closed, and moreover, $0 \notin \sigma(-\Delta)$. This operator then satisfies the hypotheses of Lemma 19.36 and is therefore self-adjoint. Feeding the compact resolvent $R_0(-\Delta)$ into Theorem 19.45 thus produces eigenfunctions of $-\Delta$ that live in $\mathcal{D} \subset L^2(\Omega)$ and form an orthonormal basis of $L^2(\Omega)$.

Example 19.47. Here is a variation on Example 19.46 in which the orthonormal basis of eigenfunctions can be seen explicitly. The operator

$$L^2(\mathbb{T}^n) \supset H^2(\mathbb{T}^n) \xrightarrow{-\Delta} L^2(\mathbb{T}^n)$$

has smooth eigenfunctions $f_k : \mathbb{T}^n \to \mathbb{C}$ defined by

$$f_k(x) := e^{2\pi i k\cdot x}, \qquad k \in \mathbb{Z}^n,$$

with $-\Delta f_k = 4\pi^2|k|^2 e^{2\pi i k\cdot x}$. We cannot invert $-\Delta$ in this case to find a compact resolvent, since $0 \in \sigma(-\Delta)$, but for any $\lambda \notin \sigma(-\Delta)$, $\lambda + \Delta$ defines a bounded operator $H^2(\mathbb{T}^n) \to L^2(\mathbb{T}^n)$, so that $R_\lambda(-\Delta)$ is the composition of its bounded inverse with the compact inclusion $H^2(\mathbb{T}^n) \hookrightarrow L^2(\mathbb{T}^n)$, and is therefore compact.

Example 19.48. Way back in §1.4, we studied the boundary value problem

$$\begin{aligned} \ddot{x}(t) + P(t)x(t) &= f(t), \\ x(0) = x(1) &= 0 \end{aligned} \tag{19.4}$$

for functions $x : [0,1] \to \mathbb{R}$, where $P, f : [0,1] \to \mathbb{R}$ are given continuous functions. Using the fact that invertibility is an open condition on bounded linear maps, we showed that this problem has a unique solution of class C^2 for any given f and P if P is sufficiently C^0-small.

Spectral theory allows us to say something about this problem without assuming P to be C^0-small. Let us allow solutions to be *complex*-valued functions $x : [0,1] \to \mathbb{C}$, so that $C^k([0,1])$ and $L^2([0,1])$ become complex vector spaces, and define the domain

$$\mathcal{D} := \left\{ x \in C^1([0,1]) \;\middle|\; \dot{x} \text{ is absolutely continuous and } \ddot{x} \in L^2([0,1]) \right\},$$

with the usual understanding that as the derivative of an absolutely continuous function, $\ddot{x}$ is defined almost everywhere. Assuming P is still real-valued, but allowing it to be any function of class L^2 on $[0,1]$, not necessarily continuous, the unbounded operator

$$L^2([0,1]) \supset \mathcal{D} \xrightarrow{T} L^2([0,1]) : x \mapsto \ddot{x} + Px$$

can be shown to be self-adjoint, and we claim that it has compact resolvent. To see this, pick any $\lambda \in \mathbb{C}\backslash\mathbb{R}$, so that self-adjointness implies $\lambda \notin \sigma(T)$, and $\lambda - T : \mathcal{D} \to L^2([0,1])$ is therefore bijective. Define a norm on $\mathcal{D} \subset L^2([0,1])$ by

$$\|x\|_{\mathcal{D}} := \|x\|_{C^1} + \|\ddot{x}\|_{L^2}.$$

This norm makes $\mathcal{D}$ into a Banach space so that $\lambda - T : \mathcal{D} \to L^2([0,1])$ becomes a bounded operator. Moreover, the inclusion $\mathcal{D} \hookrightarrow L^2([0,1])$ is compact with respect to the norm $\|\cdot\|_{\mathcal{D}}$, because any sequence $x_k \in \mathcal{D}$ satisfying a uniform bound $\|x_k\|_{\mathcal{D}} \leqslant C$ is also uniformly C^1-bounded, so that the Arzelà-Ascoli theorem gives a C^0-convergent subsequence, which then also converges in L^2.

Applying Theorem 19.45, we now obtain an orthonormal basis of $L^2([0,1])$ consisting of eigenfunctions $x_\lambda \in \mathcal{D}$ with eigenvalues $\lambda \in \mathbb{R}$ that are solutions to the problem

$$\begin{aligned} \ddot{x}_\lambda(t) + [P(t) - \lambda]\, x_\lambda(t) &= 0, \\ x_\lambda(0) = x_\lambda(1) &= 0. \end{aligned}$$

For any $\lambda \in \mathbb{R}$ outside the discrete set of eigenvalues, we can instead say that the problem

$$\ddot{x}(t) + [P(t) - \lambda]\, x(t) = f(t),$$
$$x(0) = x(1) = 0$$

has a unique solution $x \in \mathcal{D}$ for every $f \in L^2([0,1])$. Moreover, if P and f in this problem are both continuous, then one deduces from the differential equation that $x : [0,1] \to \mathbb{C}$ if of class C^2.

Exercise 19.49. Prove that the operator in Example 19.48 really is self-adjoint.
Hint: Interpret the condition defining the domain of its adjoint in terms of weak derivatives.

Remark 19.50. Similarly to Remark 19.16, the domain $\mathcal{D}$ in Example 19.48 is secretly the Sobolev space $W^{2,2}((0,1))$.

19.7. Functional calculus. Now that we have spectral representations for unbounded self-adjoint operators

$$\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H},$$

it will not surprise you to learn that one can also develop a functional calculus. Writing $A = U^{-1} T_F U$ and $\mathcal{D} = U^{-1}(\mathcal{D}_F)$ for a unitary operator $U : \mathcal{H} \to L^2(X,\mu)$ and a measurable function $F : X \to \sigma(A) \subset \mathbb{R}$ on a semifinite measure space (X,μ), we can define

$$f(A) := U^{-1} T_{f \circ F} U \in \mathscr{L}(\mathcal{H}) \qquad \text{for} \qquad f \in \mathscr{B}(\sigma(A)),$$

where $\mathscr{B}(\sigma(A))$ is again the space of bounded Borel-measurable functions $\sigma(A) \to \mathbb{C}$. Note: $f(A)$ is a *bounded* operator $\mathcal{H} \to \mathcal{H}$, since f is a bounded function. It is straightforward to verify from this definition that the map $f \mapsto f(A)$ has most of the properties that one expects from a functional calculus, but there is one important exception: If A is not bounded, then neither is its spectrum $\sigma(A) \subset \mathbb{R}$, which means that the function $f(\lambda) := \lambda$ does not belong to $\mathscr{B}(\sigma(A))$. This should not be surprising, since A also does not belong to $\mathscr{L}(\mathcal{H})$, so the condition "$f(A) = A$ for $f(\lambda) = \lambda$" has no place in a functional calculus that turns bounded functions into bounded operators. The following exercise demonstrates what can be proved instead:

Exercise 19.51. Consider the multiplication operator

$$L^2(X,\mu) \supset \mathcal{D}_F \xrightarrow{T_F} L^2(X,\mu) : u \mapsto Fu$$

introduced in Example 19.22, with $F : X \to \mathbb{C}$ assumed to be a measurable function taking values in $\sigma(T_F)$. Show that if we define $f(T_F) := T_{f \circ F} \in \mathscr{L}(\mathcal{H})$ for $f \in \mathscr{B}(\sigma(T_F))$, then for any sequence $f_n \in \mathscr{B}(\sigma(T_F))$ that converges pointwise to the function $f(\lambda) := \lambda$ and satisfies $|f_n| \leqslant |f|$ for all n, we have

$$f_n(T_F)u \xrightarrow{L^2} T_F u \qquad \text{for all} \qquad u \in \mathcal{D}_F.$$

The exercise produces item (iii) in the following theorem, whose proof is a straightforward adaptation of arguments we have already carried out in the bounded case.

Theorem 19.52. *For any unbounded self-adjoint operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$, there exists a linear map*

$$\mathscr{B}(\sigma(A)) \to \mathscr{L}(\mathcal{H}) : f \mapsto f(A)$$

satisfying the following conditions:

(i) $(fg)(A) = f(A)g(A)$ and $\bar{f}(A) = f(A)^*$ for all $f, g \in \mathscr{B}(\sigma(A))$;
(ii) For $f \equiv 1$, $f(A) = \mathbb{1} \in \mathscr{L}(\mathcal{H})$;
(iii) For any sequence $f_n \in \mathscr{B}(\sigma(A))$ converging pointwise to the function $f(\lambda) := \lambda$ and satisfying $|f_n| \leqslant |f|$ for all n,

$$f_n(A)x \to Ax \qquad \textit{for every } x \in \mathcal{D}.$$

(iv) *For every* $f \in C_b^0(\sigma(A))$, $\sigma(f(A)) = \overline{f(\sigma(A))} \subset \mathbb{C}$, *and for* $f \in \mathscr{B}(\sigma(A))$, $x \in \mathcal{D}$ *and* $\lambda \in \mathbb{R}$, $Ax = \lambda x$ *implies* $f(A)x = f(\lambda)x$;

(v) *If* $f \in \mathscr{B}(\sigma(A))$ *takes nonnegative real values, then* $\langle x, f(A)x\rangle \geqslant 0$ *for all* $x \in \mathcal{H}$;

(vi) $\|f(A)\| = \|f\|_{C^0}$ *for all* $f \in C_b^0(\sigma(A))$, *and for any sequence* $f_n \in \mathscr{B}(\sigma(A))$ *that is pointwise convergent* $f_n \to f$ *and satisfies a uniform bound* $|f_n| \leqslant C$ *for all* n,

$$f_n(A)x \to f(A)x \qquad \text{for every } x \in \mathcal{H}.$$

□

REMARK 19.53. If A is unbounded then $\sigma(A) \subset \mathbb{R}$ is not a bounded set, and continuous functions on $\sigma(A)$ need not be bounded. For this reason, we've written $C_b^0(\sigma(A))$ in Theorem 19.52 so that we are only considering bounded functions. A similar issue is that for $f \in C_b^0(\sigma(A))$, $f(\sigma(A)) \subset \mathbb{C}$ may fail to be compact, in which case it cannot be identical to $\sigma(f(A))$ in property (iv); however, it is easy to check that replacing $f(\sigma(A))$ by its closure repairs this problem.

The uniqueness of the functional calculus is less obvious than in the bounded case, because when $\sigma(A)$ is unbounded, $\mathscr{B}(\sigma(A))$ does not contain the polynomial functions. The online mathematical community seems to feel unanimously that the functional calculus *is* unique and no one wants to take the trouble of writing down a complete proof. I decided not to include the word "unique," because we don't really need it in the following.

EXERCISE 19.54. Use the functional calculus to show that the formula

$$\|R_\lambda(A)\| = \frac{1}{\operatorname{dist}(\lambda, \sigma(A))}$$

also holds when A is an unbounded self-adjoint operator and $\lambda \in \rho(A)$.

The following application of the functional calculus is made possible by the fact that self-adjoint operators A have spectrum contained in $\mathbb{R}$, so that the function $\lambda \mapsto e^{i\lambda}$ is always bounded on $\sigma(A)$.

COROLLARY 19.55. *Every unbounded self-adjoint operator* $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ *gives rise to a family of unitary operators*

$$\left\{U(t) := e^{itA} \in \mathscr{L}(\mathcal{H})\right\}_{t\in\mathbb{R}}$$

with the following properties:

(1) $U(0) = \mathbb{1}$;

(2) $U(s+t) = U(s)U(t)$ *for all* $s, t \in \mathbb{R}$;

(3) *For every* $x \in \mathcal{H}$, *the function* $\mathbb{R} \to \mathcal{H} : t \mapsto U(t)x$ *is continuous.*

□

REMARK 19.56. In the setting of Corollary 19.55, the domain $\mathcal{D} \subset \mathcal{H}$ is precisely the set of all $x \in \mathcal{H}$ such that

$$\frac{1}{t}(e^{itA}x - x)$$

has a well-defined limit in $\mathcal{H}$ as $t \to 0$, and this limit is then iAx. In other words, the equation

$$\left.\frac{d}{dt}e^{itA}x\right|_{t=0} = iAx$$

holds for $x \in \mathcal{D}$, and $\mathcal{D}$ is precisely the space of all $x \in \mathcal{H}$ for which the left hand side is well defined. To see this, one can work directly in a spectral representation, so that without loss of generality, $\mathcal{H} = L^2(X, \mu)$, $\mathcal{D} = \mathcal{D}_F$ and $A = T_F$ for some measure space (X, μ) and measurable function $F : X \to \sigma(A) \subset \mathbb{R}$. Then e^{itA} is the multiplication operator $T_{e^{itF}}$, and the limit in question becomes the L^2-limit as $t \to 0$ of functions of the form

$$\frac{1}{t}(e^{itF}u - u) \in L^2(X, \mu),$$

for $u \in L^2(X, \mu)$. These functions converge pointwise to iFu, so if there is also L^2-convergence, it implies that iFu is in $L^2(X, \mu)$ and thus $u \in \mathcal{D}_F$. Conversely, if we assume $u \in \mathcal{D}_F$, then since $|e^{is} - 1| \leqslant |s|$ for all $s \in \mathbb{R}$, we can feed the estimate

$$\left|\frac{1}{t}(e^{itF} - 1)\right| \leqslant \frac{|tF|}{|t|} = |F|$$

into the dominated convergence theorem, i.e. it gives us the estimate

$$\left|\frac{1}{t}(e^{itF}u - u) - iFu\right|^2 \leqslant (|F| + |iF|)^2 \, |u|^2 = 4|Fu|^2,$$

where the right hand side is an integrable function, so dominated convergence gives

$$\lim_{t \to 0} \int_X \left|\frac{1}{t}(e^{itF}u - u) - iFu\right|^2 \, d\mu = 0.$$

19.8. Self-adjoint generators of unitary transformations. To wrap up our survey on unbounded operators, I'd like to describe a beautiful converse of Corollary 19.55 that is of fundamental significance in quantum mechanics. From a purely mathematical perspective, the theorem also makes the following point: even if you are only interested in *bounded* (and in particular unitary) operators, you may sometimes be *forced* to deal with unbounded self-adjoint operators.

Let us first discuss the bounded analogue of Corollary 19.55. For $A \in \mathscr{L}(\mathcal{H})$, we can define $e^A \in \mathscr{L}(\mathcal{H})$ without the aid of any functional calculus, because the power series

$$e^A := \mathbb{1} + A + \frac{1}{2!}A^2 + \frac{1}{3!}A^3 + \dots$$

converges absolutely. We made use of this in §18.10 for proving that any operator that commutes with a normal operator also commutes with its adjoint. The relation

$$e^{A+B} = e^A e^B = e^B e^A$$

also holds whenever A and B commute, which makes e^{-A} the inverse of e^A, and as a consequence, e^{iA} is unitary whenever A is self-adjoint. Moreover, for any $A \in \mathscr{L}(\mathcal{H})$ the map

$$\mathbb{R} \to \mathscr{L}(\mathcal{H}) : t \mapsto e^{tA}$$

is continuous with respect to the operator norm on $\mathscr{L}(\mathcal{H})$: indeed, since $e^{(t_0+t)A} = e^{t_0 A} e^{tA}$, it suffices to observe that

$$\lim_{t \to 0} e^{tA} = \mathbb{1},$$

which holds because

$$\|e^{tA} - \mathbb{1}\| \leqslant \sum_{n=1}^{\infty} \frac{1}{n!}\|tA\|^n = \|tA\| \cdot \sum_{n=0}^{\infty} \frac{1}{n+1}\frac{\|tA\|^n}{n!} \leqslant \|tA\| \cdot \sum_{n=0}^{\infty} \frac{\|tA\|^n}{n!} = \|tA\| \cdot e^{\|tA\|} \tag{19.5}$$

The bounded analogue of Corollary 19.55 follows easily from these properties:

Theorem 19.57. *Every bounded self-adjoint operator $A \in \mathscr{L}(\mathcal{H})$ gives rise to a family of unitary operators*

$$\left\{U(t) := e^{itA} \in \mathscr{L}(\mathcal{H})\right\}_{t \in \mathbb{R}}$$

with the following properties:

(1) $U(0) = \mathbb{1}$;

(2) $U(s+t) = U(s)U(t)$ *for all* $s,t \in \mathbb{R}$*;*
(3) The map $\mathbb{R} \to \mathscr{L}(\mathcal{H}) : t \mapsto U(t)$ *is continuous.*

□

The following may seem at first like a nitpicky observation, but it is important: The estimate (19.5), showing that $t \mapsto U(t)$ is a continuous map $\mathbb{R} \to \mathscr{L}(\mathcal{H})$, is heavily dependent on the fact that A has a finite norm. It is not a coincidence that in Corollary 19.55, we did not claim that $t \mapsto U(t)$ is continuous with respect to the operator norm, but instead a weaker condition.

Definition 19.58. Given a Banach space X, the **strong operator topology**[44] on $\mathscr{L}(X)$ is the locally convex topology generated by the collection of seminorms

$$\{\|A\|_x := \|Ax\|\}_{x \in X}\,.$$

In other words, a sequence $A_n \in \mathscr{L}(X)$ converges to $A \in \mathscr{L}(X)$ in the strong operator topology if and only if $A_n x \to Ax$ for every $x \in X$; similarly, a map $\mathbb{R} \to \mathscr{L}(X) : t \mapsto A(t)$ is called **strongly continuous** if and only if the map

$$\mathbb{R} \to X : t \mapsto A(t)x$$

is continuous for every $x \in X$. This is the notion of continuity that appears in Corollary 19.55. It is true if $\mathbb{R} \to \mathscr{L}(X) : t \mapsto A(t)$ is also continuous in the operator norm, but the converse is false; in general, there are also strongly continuous maps $\mathbb{R} \to \mathscr{L}(X)$ that are *not* continuous in the operator norm. We will see in fact that in familiar Hilbert spaces such as $L^2(\mathbb{R}^n)$, there are simple and natural examples.

Before we get to that, let's discuss the converse of Theorem 19.57:

Theorem 19.59. *Suppose* $\mathbb{R} \to \mathscr{L}(\mathcal{H}) : t \mapsto U(t)$ *is a continuous (in the operator norm) family of unitary operators such that* $U(0) = \mathbb{1}$ *and* $U(s+t) = U(s)U(t)$ *for all* $s,t \in \mathbb{R}$*. Then there exists a unique bounded self-adjoint operator* $A \in \mathscr{L}(\mathcal{H})$ *such that* $U(t) = e^{itA}$ *for all* $t \in \mathbb{R}$*.*

In other words, bounded self-adjoint operators are the "infinitesimal generators" of continuous 1-parameter groups of unitary transformations. There is a simple proof of this theorem that uses some elementary ideas from the theory of smooth manifolds. We will sketch the finite-dimensional version, and then comment briefly on how it generalizes to infinite dimensions. The basic idea is a geometric observation: the space of unitary operators is a smooth submanifold of $\mathscr{L}(\mathcal{H})$, and its tangent space at $\mathbb{1} \in \mathscr{L}(\mathcal{H})$ is the space of anti-self-adjoint operators, which becomes the space of self-adjoint operators after multiplying by i.

Proof of Theorem 19.59 when $\dim \mathcal{H} < \infty$. We assume $\mathcal{H} = \mathbb{C}^n$ with the standard Hermitian inner product, so $\mathscr{L}(\mathcal{H})$ is the space of matrices $\mathbb{C}^{n\times n}$. Some of the objects appearing in the following are naturally complex vector spaces, but we shall regard them as *real* vector spaces for the purposes of discussing smooth maps and smooth submanifolds.

Let $\Sigma \subset \mathbb{C}^{n\times n}$ denote the real subspace consisting of self-adjoint (i.e. Hermitian) matrices; note that this is only a *real* (not a complex) subspace, and $i\Sigma$ is a complementary real subspace, namely the space of anti-self-adjoint matrices, satisfying

$$\mathbb{C}^{n\times n} = \Sigma \oplus i\Sigma.$$

The unitary group $\mathrm{U}(n) \subset \mathbb{C}^{n\times n}$ is the level set $F^{-1}(\mathbb{1})$ for the smooth map

$$\mathbb{C}^{n\times n} \xrightarrow{F} \Sigma : A \mapsto A^*A.$$

[44] As you might guess, there is also a *weak* operator topology, which is related to the weak topology on X. This is why the use of the word "strong" makes sense, even though it is actually a weaker topology than the usual Banach space topology on $\mathscr{L}(X)$.

The derivative of this map at $\mathbb{1}$ is the real-linear map

$$dF(\mathbb{1}) : \mathbb{C}^{n\times n} \to \Sigma : B \mapsto B + B^*,$$

which is surjective, so by the implicit function theorem, a neighborhood of $\mathbb{1}$ in $\mathrm{U}(n) \subset \mathbb{C}^{n\times n}$ is a smooth submanifold, and its tangent space at $\mathbb{1}$ is

$$T_{\mathbb{1}}\,\mathrm{U}(n) = \ker dF(\mathbb{1}) = i\Sigma \subset \mathbb{C}^{n\times n}.$$

We observe next that restricting the matrix exponential to $i\Sigma \subset \mathbb{C}^{n\times n}$ gives a smooth map

$$i\Sigma \xrightarrow{\Phi} \mathrm{GL}(n,\mathbb{C}) : A \mapsto e^A$$

whose image lies in $\mathrm{U}(n)$, and the derivative $d\Phi(0) : i\Sigma \to \mathbb{C}^{n\times n}$ of this map at $0 \in i\Sigma$ is just the inclusion $i\Sigma \hookrightarrow \mathbb{C}^{n\times n}$, which is an isomorphism onto $T_{\mathbb{1}}\,\mathrm{U}(n)$. Applying the inverse function theorem for smooth maps between manifolds, we conclude that there exist open neighborhoods $\mathcal{O} \subset i\Sigma$ of 0 and $\mathcal{U} \subset \mathrm{U}(n)$ of $\mathbb{1}$ such that the map

$$\begin{gathered} i\Sigma \supset \mathcal{O} \to \mathcal{U} \subset \mathrm{U}(n), \\ A \mapsto e^A \end{gathered}$$

is a smooth bijection with a smooth inverse, i.e. a diffeomorphism.

For the continuous family $U(t) \in \mathrm{U}(n)$, we can now fix $\epsilon > 0$ such that $U(t) \in \mathcal{U}$ for all $t \in (-\epsilon, \epsilon)$, and find a unique continuous path $A(t) \in i\mathcal{O} \subset \Sigma$ with $A(0) = 0$ such that

$$U(t) = e^{iA(t)}, \qquad \text{for} \qquad t \in (-\epsilon,\epsilon).$$

Fix $t_0 \in (0,\epsilon)$, and notice that for all $N \in \mathbb{N}$, we have

$$U(t_0/N) = e^{iA(t_0/N)} = e^{iA(t_0)/N},$$

since $U(t_0/N)^N = U(t_0)$ and $e^{NiB} = (e^{iB})^N$ for every $B \in \Sigma$. This proves $A(t_0/N) = A(t_0)/N$ for every $N \in \mathbb{N}$, and by taking $U(t_0/N)$ to any power $k \in \{-N,\dots,N\}$, one similarly finds

$$A(kt_0/N) = \frac{k}{N}A(t_0).$$

This shows that the functions $t \mapsto A(t_0 t)$ and $t \mapsto tA(t_0)$ match on a dense subset of the interval $[-1,1]$, and since both are continuous, it follows that they are identical. The conclusion is that the formula $U(t) = e^{itA}$ holds for some $A \in \Sigma$ with t in a sufficiently small neighborhood of 0.

In particular, we now see that the map $t \mapsto U(t)$ is smooth near $t = 0$. One then uses the relation $U(s+t) = U(s)U(t)$ to deduce that $t \mapsto U(t)$ is in fact smooth everywhere, and it satisfies the differential equation

$$\dot{U}(t) = \frac{d}{ds}U(t+s)\bigg|_{s=0} = \frac{d}{ds}U(s)U(t)\bigg|_{s=0} = \dot{U}(0)U(t) = iAU(t).$$

The latter is also satisfied by $t \mapsto e^{iAt}$, and it follows that for any $x \in \mathbb{C}^n$, the function $z(t) := U(t)x - e^{iAt}x \in \mathbb{C}^n$ satisfies $\dot{z}(t) = iAz(t)$ and thus

$$\frac{d}{dt}\|z(t)\|^2 = \langle iAz(t), z(t)\rangle + \langle z(t), iAz(t)\rangle = -i\langle Az(t), z(t)\rangle + i\langle z(t), Az(t)\rangle = 0.$$

Since $z(0) = 0$, we conclude $z(t) = 0$ for all t, and thus $U(t)x = e^{itA}x$. □

REMARK 19.60. The fundamental ingredients used in the proof of Theorem 19.59 were as follows:

- The product rule for vector-valued functions, used in the last step to differentiate $\|z(t)\|^2$;
- The concept of a smooth submanifold of a vector space;

- The implicit function theorem, namely for showing that a regular level set is a smooth submanifold;
- The inverse function theorem, for showing that a map between two smooth submanifolds is a local diffeomorphism.

The quick way to make Theorem 19.59 also work in infinite dimensions is to say the following: all four of those ingredients have well-behaved generalizations to the setting of infinite-dimensional Banach spaces. One starts by defining the notion of differentiability for maps $f : \mathcal{U} \to Y$ where X, Y are Banach spaces and $\mathcal{U} \subset X$ is an open subset; the derivative at a point $x \in \mathcal{U}$, when it exists, is then required to be a *bounded* linear operator $df(x) \in \mathscr{L}(X, Y)$. Since $\mathscr{L}(X, Y)$ is also a Banach space, one can also discuss the differentiability of a derivative $df : \mathcal{U} \to \mathscr{L}(X, Y)$, and in this way, one inductively obtains a notion of smoothness for maps $\mathcal{U} \to Y$. The notion of a *smooth Banach manifold* then makes sense: it is a space M endowed with local "coordinate charts" that identify small open subsets of M with open subsets of a Banach space, such that the resulting "coordinate transformations" are all smooth bijections between open subsets of Banach spaces.

The implicit and inverse function theorems in this setting are not so different from the finite-dimensional case, with one caveat: for showing that a level set $f^{-1}(y) \subset \mathcal{U}$ is a smooth submanifold, it is not enough to know that $df(x) : X \to Y$ is *surjective* at every point $x \in f^{-1}(y)$. One also needs to know that it admits a bounded right-inverse, or equivalently, that its kernel is *complemented.* If you look at standard ways of deducing the finite-dimensional implicit function theorem from the inverse function theorem, you'll notice that this assumption is always used, and it's not an issue in finite dimensions since *all* subspaces are complemented, but we've seen that in infinite-dimensional Banach spaces, closed complements do not come for free. In the case of interest for Theorem 19.59, the surjective operator in question is

$$\mathscr{L}(\mathcal{H}) \to \Sigma : B \mapsto B + B^*,$$

where $\mathcal{H}$ is a Hilbert space and $\Sigma \subset \mathscr{L}(\mathcal{H})$ is the real-linear subspace consisting of self-adjoint operators, which is a closed subspace. Its kernel is thus the space $i\Sigma \subset \mathscr{L}(\mathcal{H})$ of anti-self-adjoint operators, which is also a closed subspace, and has Σ as a closed complement, so the extra condition needed for the implicit function theorem is easily verified.

The inverse function theorem works in the Banach space setting for essentially the same reason as in finite dimensions: it is a consequence of the Banach fixed point theorem (also known as the *contraction mapping principle*). If you work through a standard proof in finite dimensions, but take care to avoid using any choices of finite-dimensional coordinates at every step, then you'll find that it generalizes to the infinite-dimensional case without much effort. The Picard-Lindelöf theorem similarly generalizes to apply to ODEs for functions taking values in a Banach space; we could have used this to simplify the last step by arguing that $U(t)$ and e^{iAt} must be identical because they satisfy the same initial value problem.

The standard reference for the details on differential calculus in Banach spaces is [**Lan93**, Chapters XIII-IV], as well as the first two chapters of [**Lan99**] for the basic notions of smooth Banach manifolds.

Theorem 19.59 is an elegant result, but now I can tell you the bad news, which will force unbounded operators back into the picture. Some of the most natural examples arising of 1-parameter groups of unitary transformations $U(t) \in \mathscr{L}(\mathcal{H})$ are strongly continuous, but *not* continuous in the operator norm.

EXAMPLE 19.61. Let $e_1, \ldots, e_n \in \mathbb{R}^n$ denote the standard basis. For each $j = 1, \ldots, n$, we define the family of *spatial translation* operators $U_j(t) := \tau_{te_j}$,

$$L^2(\mathbb{R}^n) \xrightarrow{U_j(t)} L^2(\mathbb{R}^n), \qquad \big(U_j(t)f\big)(x) := f(x + te_j).$$

We saw in Theorem 8.4 that the map $\mathbb{R}^n \to L^2(\mathbb{R}^n) : v \mapsto \tau_v f$ is continuous for every $f \in L^2(\mathbb{R}^n)$, thus $U_j(t)$ is a strongly continuous family. However, one can easily cook up examples to show that

$$\sup_{f \neq 0} \frac{\|U_j(t)f - f\|_{L^2}}{\|f\|_{L^2}}$$

need not become arbitrarily small as $t \to 0$, so $U(t)$ does not converge to $U(0) = \mathbb{1}$ in the operator norm.

To understand better what is happening in Example 19.61, we can use an explicit spectral representation of the operators $U_j(t)$, given by the Fourier transform: for $f \in \mathscr{S}(\mathbb{R}^n)$, a change of variables gives

$$\widehat{U_j(t)f}(p) = \int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} f(x + te_j)\, dx = e^{2\pi i p \cdot te_j} \int_{\mathbb{R}^n} e^{-2\pi i p \cdot x} f(x)\, dx = e^{2\pi i p_j t} \hat{f}(p),$$

thus

$$\mathscr{F} U_j(t) \mathscr{F}^* = T_{e^{2\pi i p_j t}},$$

and by density, this formula is valid on all of $L^2(\mathbb{R}^n)$. Now compare: We saw in Example 19.41 that the unbounded operator

$$L^2(\mathbb{R}^n) \supset \mathscr{S}(\mathbb{R}^n) \xrightarrow{i\partial_j} L^2(\mathbb{R}^n)$$

is essentially self-adjoint, and the Fourier transform also gives an explicit spectral representation for its self-adjoint extension, with

$$\mathscr{F}(i\partial_j)\mathscr{F}^* = T_{-2\pi p_j}.$$

Using the functional calculus for unbounded self-adjoint operators, we obtain from this the formula

$$U_j(t) = e^{it(-i\partial_j)} =: e^{t\partial_j},$$

so that we can interpret the unbounded self-adjoint operator $-i\partial_j$ as the infinitesimal generator of the translations in the e_j direction. In this special case, we have obtained a converse of Corollary 19.55, and it turns out that this can always be done. The difference from Theorem 19.59 is that if the family of unitary operators is strongly but not norm continuous, then the resulting self-adjoint generator will be unbounded:

THEOREM 19.62 (Stone's theorem). *Suppose $\mathbb{R} \to \mathscr{L}(\mathcal{H}) : t \mapsto U(t)$ is a strongly continuous family of unitary operators such that $U(0) = \mathbb{1}$ and $U(s + t) = U(s)U(t)$ for all $s, t \in \mathbb{R}$. Then there exists a unique densely-defined self-adjoint operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$ such that $U(t) = e^{itA}$ for all $t \in \mathbb{R}$.*

REMARK 19.63. In light of Theorems 19.57 and 19.59, the self-adjoint operator A appearing in Stone's theorem will be bounded if and *only* if the given family of unitary operators $U(t)$ defines a continuous map $\mathbb{R} \to \mathscr{L}(\mathcal{H})$ with respect to the operator norm.

Before proving Stone's theorem, let us briefly comment on its physical meaning. In quantum mechanics, the possible states of a physical system are elements in a Hilbert space, and *symmetries* of the system are therefore defined via transformations that preserve the structure of the Hilbert space, i.e. unitary transformations. For instance, the unitarity of the operators in Example 19.61 expresses the fact the laws of physics are invariant under spatial translations. By a basic principle known as *Noether's theorem*—which is valid in classical as well as quantum mechanics—there is a correspondence between symmetries and conserved quantities: every 1-parameter group of symmetries of a physical system gives rise to a measurable quantity that depends on the state of the system and remains constant as the state evolves in time. This is, however, a rather "classical" way to express things: in quantum mechanics, measurable quantities—also known as "observables"—do not have precise values, but every state in $\mathcal{H}$ determines a probability distribution for the possible values of each observable in that state. In particular, each observable is represented by a self-adjoint operator on $\mathcal{H}$, whose eigenvectors (if any exist) are interpreted as the states in which the observable has a precise value, namely the eigenvalue. Stone's theorem gives us a recipe to find the self-adjoint operator associated to each conserved quantity, e.g. in classical mechanics, the conserved quantity associated to spatial translation symmetry is *momentum*, and for this reason, Example 19.61 reveals $-i\partial_j$ to be (up to issues of normalization) the self-adjoint operator that represents momentum in quantum mechanics.

Proof of Theorem 19.62. The uniqueness of A follows from Remark 19.56, which shows how to recover both A and its domain from the family $U(t)$: the domain will have to be

$$\left\{ x \in \mathcal{H} \;\middle|\; \lim_{t \to 0} \frac{U(t)x - x}{t} \text{ exists in } \mathcal{H} \right\},$$

and for x in this subspace, A will be given by

$$Ax = -i \lim_{t \to 0} \frac{U(t)x - x}{t} \in \mathcal{H}.$$

We prove existence via a series of three claims.

Claim 1: There exists a dense subspace $\mathcal{D}_0 \subset \mathcal{H}$ and an unbounded symmetric operator $\mathcal{H} \supset \mathcal{D}_0 \xrightarrow{A} \mathcal{H}$ such that

$$Ax = -i \lim_{t \to 0} \frac{U(t)x - x}{t} \qquad \text{for all} \qquad x \in \mathcal{D}_0, \tag{19.6}$$

and moreover,

$$A(\mathcal{D}_0) \subset \mathcal{D}_0 \qquad \text{and} \qquad U(t)(\mathcal{D}_0) \subset \mathcal{D}_0 \text{ for all } t \in \mathbb{R}.$$

One can define $\mathcal{D}_0$ using an idea similar to *mollification*: given any $x \in \mathcal{H}$, we use the action of $U(t)$ to define a "smoothening" of x so that the limit on the right hand side of (19.6) is guaranteed to exist. Indeed, given $x \in \mathcal{H}$ and $\varphi \in C_0^\infty(\mathbb{R})$, let

$$x_\varphi := \int_{-\infty}^{\infty} \varphi(s) U(s) x \, ds \in \mathcal{H}.$$

Note that since $U(t)$ is a strongly continuous family, the integrand in this definition is a continuous function valued in $\mathcal{H}$, and Riemann integrals for such functions can be defined in the same way as for functions valued in finite-dimensional vector spaces.[45] Since $U(s)x$ is close to x near $s = 0$, an approximate identity $\rho_j : \mathbb{R} \to [0, \infty)$ with shrinking support gives rise to a convergent sequence

$$x_{\rho_j} \to x \text{ in } \mathcal{H} \qquad \text{as } j \to \infty,$$

thus the space $\mathcal{D}_0 \subset \mathcal{H}$ of all finite linear combinations of elements of the form x_φ for $x \in \mathcal{H}$ and $\varphi \in C_0^\infty(\mathbb{R})$ is dense in $\mathcal{H}$. Now compute:

$$\begin{aligned}
\frac{U(t)x_\varphi - x_\varphi}{t} &= \frac{1}{t}\left(\int_{-\infty}^{\infty} \varphi(s) U(s+t) x \, ds - \int_{-\infty}^{\infty} \varphi(s) U(s) x \, ds \right) \\
&= \frac{1}{t}\left(\int_{-\infty}^{\infty} \varphi(s-t) U(s) x \, ds - \int_{-\infty}^{\infty} \varphi(s) U(s) x \, ds \right) \\
&= \int_{-\infty}^{\infty} \frac{\varphi(s-t) - \varphi(s)}{t} U(s) x \, ds \to \int_{-\infty}^{\infty} -\varphi'(s) U(s) x \, ds = x_{-\varphi'} \qquad \text{as} \qquad t \to 0.
\end{aligned}$$

[45] A precise definition of the notion of Riemann integration used here can be found in [Lan93, §XIII.1]. Defining Lebesgue integration for functions valued in general Banach spaces is harder, and unnecessary for our purposes, but it can also be done, and *is* done in Chapter VI of [Lan93].

The convergence of the integrals in the last step follows because for $\varphi \in C_0^\infty(\mathbb{R})$, the convergence of the difference quotients to $-\varphi'$ is uniform and everything has compact support. This formula shows that if $A : \mathcal{D}_0 \to \mathcal{H}$ is defined via (19.6), then Ax_φ also belongs to $\mathcal{D}_0$ for every $x \in \mathcal{H}$ and $\varphi \in C_0^\infty(\mathbb{R})$; similarly, acting on x_φ by $U(t)$ just translates the parametrization of φ, thus A and $U(t)$ both preserve $\mathcal{D}_0$. One now uses the unitarity of $U(t)$ to verify that A is symmetric.

Claim 2: The operator $\mathcal{H} \supset \mathcal{D}_0 \xrightarrow{A} \mathcal{H}$ defined above is essentially self-adjoint.

By the basic criterion, Theorem 19.40, it will suffice to show that for every $\lambda \in \mathbb{C}\backslash\mathbb{R}$, the image of $A - \lambda : \mathcal{D}_0 \to \mathcal{H}$ is dense. Indeed, if it isn't, then there exists $v \neq 0 \in \mathcal{H}$ such that $\langle (A - \lambda)x, v\rangle = 0$ for all $x \in \mathcal{D}_0$. Since $x \in \mathcal{D}_0$, we can use the formula for A from Claim 1 and the fact that $U(t)$ preserves $\mathcal{D}_0$ to compute

$$\frac{d}{dt}U(t)x = \frac{d}{ds}U(t+s)x\bigg|_{s=0} = \frac{d}{ds}U(s)U(t)x\bigg|_{s=0} = iAU(t)x,$$

so that the complex-valued function $f(t) := \langle U(t)x, v\rangle$ satisfies

$$f'(t) = \langle iAU(t)x, v\rangle = \langle i\lambda U(t)x, v\rangle = -i\bar{\lambda} f(t).$$

It follows that $f(t) = f(0)e^{-i\bar{\lambda}t}$ for all $t \in \mathbb{R}$, and this has to be a globally bounded function since the unitarity of $U(t)$ implies

$$|f(t)| \leqslant \|U(t)\| \cdot \|x\| \cdot \|v\| \leqslant \|x\| \cdot \|v\|.$$

Since $\lambda \notin \mathbb{R}$, the only way for $f(t)$ to be bounded is if $f(0) = 0$, which means $\langle x, v\rangle = 0$ for all $x \in \mathcal{D}_0$, contradicting the fact that $\mathcal{D}_0$ is dense.

We conclude that the closure of $\mathcal{H} \supset \mathcal{D}_0 \xrightarrow{A} \mathcal{H}$ is a self-adjoint operator $\mathcal{H} \supset \mathcal{D} \xrightarrow{A} \mathcal{H}$, defined on a dense domain $\mathcal{D} \subset \mathcal{H}$ that contains $\mathcal{D}_0$. By Corollary 19.55, this gives rise to a strongly continuous family of unitary operators e^{iAt}.

Claim 3: $U(t) = e^{iAt}$ for all $t \in \mathbb{R}$.

Fix $x \in \mathcal{D}_0$, so we know $\frac{U(t)x - x}{t} \to iAx$ as $t \to 0$, and thus by the calculation in Claim 2,

$$\frac{d}{dt}U(t)x = iAU(t)x$$

for all $t \in \mathbb{R}$. Since $\mathcal{D}_0 \subset \mathcal{D}$, the same calculation works for e^{iAt} due to Remark 19.56, implying $e^{iAt}x \in \mathcal{D}$ for all $t \in \mathbb{R}$ and

$$\frac{d}{dt}e^{iAt}x = iAe^{iAt}x.$$

The path $z(t) := U(t)x - e^{iAt}x \in \mathcal{D} \subset \mathcal{H}$ thus satisfies the differential equation

$$\dot{z}(t) = iAz(t),$$

which implies

$$\frac{d}{dt}\|z(t)\|^2 = \langle iAz(t), z(t)\rangle + \langle z(t), iAz(t)\rangle = -i\langle Az(t), z(t)\rangle + i\langle z(t), Az(t)\rangle = 0.$$

Since $z(0) = 0$, we conclude $z(t) = 0$ for all t, and thus $U(t)x = e^{itA}x$. This is true for all x in a dense subspace of $\mathcal{H}$, so it implies $U(t) = e^{itA}$. $\square$

The end.

Bibliography

[AF03] R. A. Adams and J. J. F. Fournier, *Sobolev spaces*, 2nd ed., Pure and Applied Mathematics (Amsterdam), vol. 140, Elsevier/Academic Press, Amsterdam, 2003.

[AE01] H. Amann and J. Escher, *Analysis. III*, Grundstudium Mathematik. [Basic Study of Mathematics], Birkhäuser Verlag, Basel, 2001 (German, with German summary).

[Ban31] S. Banach, *Über die Baire'sche Kategorie gewisser Funktionenmengen*, Studia Mathematica **3** (1931), 174–179.

[BS18] T. Bühler and D. A. Salamon, *Functional analysis*, Graduate Studies in Mathematics, vol. 191, American Mathematical Society, Providence, RI, 2018.

[Cla36] J. A. Clarkson, *Uniformly convex spaces*, Trans. Amer. Math. Soc. **40** (1936), no. 3, 396–414.

[DM72] H. Dym and H. P. McKean, *Fourier series and integrals*, Academic Press, New York-London, 1972. Probability and Mathematical Statistics, No. 14.

[FH15] P. M. Fitzpatrick and B. R. Hunt, *Absolute continuity of a function and uniform integrability of its divided differences*, Amer. Math. Monthly **122** (2015), no. 4, 362–366.

[Fug50] B. Fuglede, *A commutativity theorem for normal operators*, Proc. Nat. Acad. Sci. U.S.A. **36** (1950), 35–40.

[Har16] G. H. Hardy, *Weierstrass's non-differentiable function*, Trans. Amer. Math. Soc. **17** (1916), no. 3, 301–325.

[Hei05] J. Heinonen, *Lectures on Lipschitz analysis*, Report. University of Jyväskylä Department of Mathematics and Statistics, vol. 100, University of Jyväskylä, Jyväskylä, 2005.

[Hir94] M. W. Hirsch, *Differential topology*, Springer-Verlag, New York, 1994.

[Hor66] J. Horváth, *Topological vector spaces and distributions. Vol. I*, Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1966.

[Ill03] S. Illman, *The very-strong C^∞ topology on $C^\infty(M,N)$ and K-equivariant maps*, Osaka J. Math. **40** (2003), no. 2, 409–428.

[Joh10] J. Johnsen, *Simple proofs of nowhere-differentiability for Weierstrass's function and cases of slow growth*, J. Fourier Anal. Appl. **16** (2010), no. 1, 17–33.

[Kat04] Y. Katznelson, *An introduction to harmonic analysis*, 3rd ed., Cambridge Mathematical Library, Cambridge University Press, Cambridge, 2004.

[Lan93] S. Lang, *Real and functional analysis*, 3rd ed., Springer-Verlag, New York, 1993.

[Lan99] ———, *Fundamentals of differential geometry*, Springer-Verlag, New York, 1999.

[LL01] E. H. Lieb and M. Loss, *Analysis*, 2nd ed., Graduate Studies in Mathematics, vol. 14, American Mathematical Society, Providence, RI, 2001.

[Maz31] S. Mazurkiewicz, *Sur les fonctions non dérivables*, Studia Mathematica **3** (1931), 92–94.

[RS80] M. Reed and B. Simon, *Methods of modern mathematical physics. I*, 2nd ed., Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York, 1980. Functional analysis.

[RSN90] F. Riesz and B. Sz.-Nagy, *Functional analysis*, Dover Books on Advanced Mathematics, Dover Publications, Inc., New York, 1990. Translated from the second French edition by Leo F. Boron; Reprint of the 1955 original.

[Ros58] M. Rosenblum, *On a theorem of Fuglede and Putnam*, J. London Math. Soc. **33** (1958), 376–377.

[Roy88] H. L. Royden, *Real analysis*, 3rd ed., Macmillan Publishing Company, New York, 1988.

[RF10] H. L. Royden and P. M. Fitzpatrick, *Real analysis*, 4th ed., Pearson, 2010.

[Rud76] W. Rudin, *Principles of mathematical analysis*, 3rd ed., McGraw-Hill Book Co., New York-Auckland-Düsseldorf, 1976. International Series in Pure and Applied Mathematics.

[Rud87] ———, *Real and complex analysis*, 3rd ed., McGraw-Hill Book Co., New York, 1987.

[Sal16] D. A. Salamon, *Measure and integration*, EMS Textbooks in Mathematics, European Mathematical Society (EMS), Zürich, 2016. MR3469972

[Sch99] E. Schechter, *Handbook of analysis and its foundations*, Academic Press, Inc., San Diego, CA, 1999.

[Shi18] N. Shioji, *Simple proofs of the uniform convexity of L^p and the Riesz representation theorem for L^p*, Amer. Math. Monthly **125** (2018), no. 8, 733–738.

[Tay96] M. E. Taylor, *Partial differential equations. I*, Applied Mathematical Sciences, vol. 115, Springer-Verlag, New York, 1996.

[Tes] G. Teschl, *Topics in Linear and Nonlinear Functional Analysis*, Graduate Studies in Mathematics, American Mathematical Society, Providence, RI. to appear; manuscript available at https://www.mat.univie.ac.at/~gerald/ftp/book-fa/.

[Wen23] C. Wendl, *Topology I–III* (2023). Notes from the course at HU Berlin, available at https://www.mathematik.hu-berlin.de/~wendl/Sommer2025/Topologie3/lecturenotes.pdf.

[Wen19] ———, *Parameterabhängige Lebesgue-Integrale* (2019). Skript zur Vorlesung Analysis III an der HU Berlin, verfügbar unter https://www.mathematik.hu-berlin.de/~wendl/Winter2019/Analysis3/Skript_parameter.pdf.

[Whi66] R. Whitley, *Mathematical Notes: Projecting m onto c_0*, Amer. Math. Monthly **73** (1966), no. 3, 285–286.