Notes on Advanced Topics in Analysis

November 16, 2022

by Scott M. LaLonde

Math 5341 – Fall 2016, Fall 2018, Fall 2022
Math 5342 – Spring 2019
Math 5343 – Spring 2017

Contents

Preface

These notes were originally written to accompany the introductory graduate-level analysis sequence at the University of Texas at Tyler. A first draft emerged when I taught the sequence for the first time in the 2016–2017 academic year, and I subsequently revised and rewrote large parts of the notes while teaching those courses again in 2018–2019.

While I had taught undergraduate analysis courses before, these notes represent my first attempts at teaching graduate-level analysis. There was no single textbook that I felt completely comfortable following, so we ended up assimilating ideas from several different sources throughout the two incarnations of these courses. The suggested texts included the following:

- *Real Analysis* by Norman B. Haaser and Joseph A. Sullivan
- *Principles of Mathematical Analysis* by Walter Rudin
- *Introductory Real Analysis* by A. N. Kolmogorov and S. V. Fomin
- *Measure and Integral: An Introduction to Real Analysis* by Richard L. Wheeden and Antoni Zygmund
- *Real Analysis* by H. L. Royden and P. M. Fitzpatrick
- *Real Analysis: Modern Techniques and their Applications* by Gerald B. Folland

My vacillating approach to choosing a textbook grew from my desire to bridge the gap between a standard undergraduate course in real analysis and introductory courses in measure theory and functional analysis at the Ph.D. level. Consequently, the only prerequisite is a good foundation in undergraduate real analysis. (Some of the material—particularly in the later chapters—does require a solid understanding of some ideas from linear algebra as well.) These notes were originally intended for a masters-level course, but the early chapters could potentially be used for a second-semester undergraduate course as well.

The ultimate goal in the first half of these notes (which corresponds to the first semester of our sequence) is to study Lebesgue's theory of measure and integration on the real line. However, in order to fully appreciate the necessity of the Lebesgue integral, one must have a firm grasp on sequences and series of functions. This topic is not generally addressed in a first-semester undergraduate real analysis course, so I felt that it deserved mentioning in this course. Many of the important results occur within a much broader discussion about analysis on abstract metric spaces, which is yet another topic to which students should have some exposure. We also use this initial study of metric spaces as an opportunity to investigate some classical theorems from analysis, including the Baire Category Theorem, the Arzelà-Ascoli Theorem, and the Stone-Weierstrass Theorem.

The later chapters lack a single objective, but instead aim to introduce students to more advanced concepts that naturally spring from the study of the Lebesgue integral. One chapter focuses on abstract measure and integration, and we encounter the celebrated Fubini-Tonelli and Radon-Nikodym theorems along the way. Afterward, we transition to some topics in functional analysis, namely those pertaining to Banach and Hilbert spaces.

Finally, I would like to include a note to the students. For young mathematicians, analysis (together with abstract algebra) is one of the most important classes you will take. However, thanks to its abstract and technical nature, it is also one of the most difficult courses you will take. You should not get discouraged, though. Analysis is something that you must have a firm grasp on if you wish to study pure mathematics (or even applied math as well) in the future.

Acknowledgements

There are several people who have contributed to these notes in both direct and indirect ways, and the quality of the manuscript would have suffered without them. The organization of the first four chapters (as well as my decision to cover advanced topics in real analysis prior to the discussion of Lebesgue measure and integration on the real line) was largely inspired by a Michael Felland's graduate-level real analysis course at Clarkson University. Likewise, the backbone for the chapters on Banach and Hilbert spaces came from notes I took from Dana Williams (my eventual Ph.D. advisor) during an introductory functional analysis course at Dartmouth College.

I must also thank all of the students who have read these notes and taken my courses. They have pointed out numerous errors, and they have inspired me to include new examples and exposition on countless occasions through our in-class discussions. In no particular order, I would like to thank Humberto Bautista Serrano, Amanda Berry, Ali Chick, Kayla Cook, Rebecca Darby, Paulson Elekuru, Rebecca Darby, Gin Germany, Fletcher Larkin, Fariha Mahfuz, Yansy Perez, Betty Tran, Vincent Villalobos, and Jeremy Williamson.

Chapter 1

Introduction

Let us begin with a brief overview of this course. Our ultimate goal is to study Lebesgue measure and construct the Lebesgue integral on the real line (and in $\mathbf{R}^n$). We will now give a survey of where we are headed, as well as some of the ideas that we will need to deal with first.

1.1 Toward a Better Integration Theory

Though you may not have seen exactly why, the Riemann integral that you first encounter in real analysis (or really in calculus) has some serious drawbacks. There is one example that you have likely seen that gives some indication of the sorts of things that can go wrong.

Example 1.1.1. Dirichlet's function $\chi_{\mathbf{Q}} : [0,1] \to \mathbf{R}$ is defined by

$$\chi_{\mathbf{Q}}(x) = \begin{cases} 1 & \text{if } x \in \mathbf{Q} \\ 0 & \text{if } x \notin \mathbf{Q}. \end{cases}$$

Recall that $\chi_{\mathbf{Q}}$ is discontinuous everywhere. It also fails to be Riemann integrable. To see this, let $P = \{x_0, x_1, \ldots, x_n\}$ be a partition of $[0,1]$, and consider the associated upper and lower Darboux sums. Recall that the upper sum is

$$U(\chi_{\mathbf{Q}}, P) = \sum_{i=1}^{n} M_i(x_i - x_{i-1}),$$

where M_i denotes the supremum of $\chi_{\mathbf{Q}}$ on the interval $[x_{i-1}, x_i]$. It is easy to see that $M_i = 1$ for all i, so we have

$$U(\chi_{\mathbf{Q}}, P) = 1.$$

Similarly,

$$L(\chi_{\mathbf{Q}}, P) = \sum_{i=1}^{n} m_i(x_i - x_{i-1}),$$

where m_i is the infimum of $\chi_{\mathbf{Q}}$ on $[x_{i-1}, x_i]$. Clearly $m_i = 0$ for all i, so

$$L(\chi_{\mathbf{Q}}, P) = 0.$$

Thus for any partition P of $[0,1]$, we have

$$L(\chi_{\mathbf{Q}}, P) = 0 \neq 1 = U(\chi_{\mathbf{Q}}, P).$$

However, recall that a function $f : [0,1] \to \mathbf{R}$ is Riemann integrable if and only if given any $\varepsilon > 0$, there is a partition P of $[0,1]$ such that

$$U(f,P) - L(f,P) < \varepsilon.$$

Thus $\chi_{\mathbf{Q}}$ is not Riemann integrable.

It is not hard to see what goes wrong with Dirichlet's function. Since $\mathbf{Q}$ is dense in $\mathbf{R}$, we can't possibly choose a partition that makes the upper and lower Darboux sums sufficiently close to one another. In other words, the Riemann sums are simply impossible to control. Lebesgue found a clever way of getting around this issue—he realized that one should partition the *range* of a function instead of the domain, as we do with the Riemann integral. Given a function $f : [a,b] \to \mathbf{R}$, we can approximate the integral as follows: given an interval $[y_{i-1}, y_i]$ in the codomain, one needs to measure the size of the set

$$A_i = \big\{x \in [a,b] : y_{i-1} \leq f(x) \leq y_i\big\}.$$

We will denote the size of this set by $\mu(A_i)$. We then form the sum

$$\sum_{i=1}^{n} y_i^* \mu(A_i),$$

where y_i^* is a sample point from the interval $[y_{i-1}, y_i]$. Roughly speaking, we are approximating the area under the graph of f with a collection of horizontal "rectangles", with $\mu(A_i)$ giving us the "width" of the i^{th} rectangle. After taking some sort of limit of sums like this, we should obtain the true area under the curve, which we would denote by $\int f$.

Of course we will need to make this construction much more precise later on. We will see that the "size" $\mu(E)$ of a set E is called the **Lebesgue measure** of E, and it generalizes the notion of "length" from intervals to more complicated sets. We will also show that $\mu(\mathbf{Q}) = 0$, so the integral of Dirichlet's function over the unit interval is

$$\int \chi_{\mathbf{Q}} = 1 \cdot \mu(\mathbf{Q} \cap [0,1]) + 0 \cdot \mu(\mathbf{I} \cap [0,1]) = 1 \cdot 0 + 0 \cdot 1 = 0.$$

A natural question to ask is the following: who cares? Why would anyone ever care to integrate Dirichlet's function? No one is losing sleep over the fact that $\chi_{\mathbf{Q}}$ is not Riemann integrable. Lebesgue's theory is necessary for other reasons—the Riemann integral has other fatal flaws that are are a little more subtle.

Example 1.1.2. Here is the real reason that the Riemann integral is insufficient. Recall that $\mathbf{Q}$ is a countable set, so we can enumerate the rationals that lie in $[0,1]$, say

$$\mathbf{Q} \cap [0,1] = \{q_1, q_2, q_3, \ldots\}.$$

Now define a sequence of functions $f_n : [0,1] \to \mathbf{R}$ by

$$f_n(x) = \begin{cases} 1 & \text{if } x = q_k \text{ for some } k \leq n \\ 0 & \text{otherwise} \end{cases}$$

Notice that each f_n is Riemann integrable, since it has only finitely many discontinuities. Additionally, $f_n \to \chi_{\mathbf{Q}}$ *pointwise*, meaning that for each $x \in [0,1]$, we have $f_n(x) \to \chi_{\mathbf{Q}}(x)$. However,

$$\int f_n \not\to \int \chi_{\mathbf{Q}},$$

since the limit function $\chi_{\mathbf{Q}}$ is not even Riemann integrable.

The previous example illustrates the real issue with the Riemann integral: it does not play well with limits and convergence. The Lebesgue integral does not have this issue—there are several nice convergence theorems that tell us when we can interchange limits and integrals in the Lebesgue sense.

Before we can take up our study of measure theory, we need to investigate some concepts that are not always standard fare in an undergraduate analysis sequence. Foremost among these concepts is the notion of a sequence or series of functions. We simply cannot appreciate the convergence theorems that come with the Lebesgue integral otherwise. We will therefore begin with a study of abstract metric spaces in the next chapter, after which we will look at sequences and series of real-valued functions on metric spaces. After that, we can begin to develop Lebesgue's theory of measure and integration.

Chapter 2

Metric Spaces

The goal of this chapter is to revisit some ideas from your undergraduate courses in analysis, and approach those concepts from a more sophisticated perspective. In particular, we will study *metric spaces*, which are sets equipped with the minimum structure necessary to discuss limits and continuity. Consequently, they provide a general framework in which to study these ideas.

2.1 Basic Definitions and Examples

When you first encounter analysis, much of what you do only relies on the ability to measure the distance between two real numbers (or between two complex numbers, or two vectors in $\mathbf{R}^n$, depending on the setting). In particular, concepts like limits and continuity are defined in terms of "closeness." It seems then that much of what you learn in an undergraduate real analysis course should carry over to other settings. The most abstract such setting, where we are only able to measure distances, is that of *metric spaces*.

Definition 2.1.1. A **metric space** is a pair (X, d), where X is a nonempty set and $d : X \times X \to \mathbf{R}$ is a function (called a **metric**) satisfying the following properties:

1. (Positive definiteness) For all $x, y \in X$, $d(x, y) \geq 0$, and $d(x, y) = 0$ if and only if $x = y$.
2. (Symmetry) For all $x, y \in X$, $d(x, y) = d(y, x)$.
3. (Triangle inequality) For all $x, y, z \in X$, $d(x, y) \leq d(x, z) + d(z, y)$.

As we discussed in the introductory paragraph, you already know some simple examples of metric spaces from real and complex analysis.

Example 2.1.2. Define $d : \mathbf{R} \times \mathbf{R} \to \mathbf{R}$ by

$$d(x, y) = |x - y|.$$

Then $(\mathbf{R}, d)$ is a metric space. All the properties of a metric (especially the triangle inequality) are verified in an undergraduate analysis class. We will refer to this metric as the **standard metric** on $\mathbf{R}$.

Example 2.1.3. Define $d : \mathbf{C} \times \mathbf{C} \to \mathbf{R}$ by $d(z, w) = |z - w|$, where $|\cdot|$ denotes the modulus of a complex number:

$$|a + bi| = \sqrt{a^2 + b^2}.$$

Then d defines a metric on $\mathbf{C}$. This is generally proven in a basic complex analysis class. It will also follow from one of our later examples.

Now we will look at some examples of metric spaces that are perhaps less familiar. The first one is somewhat odd, so naturally it will become one of our recurring examples.

Example 2.1.4. Let X be any set, and define $d : X \times X \to \mathbf{R}$ by

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ 1 & \text{if } x \neq y. \end{cases}$$

Then d is a metric on X, called the **discrete metric**. Clearly d is positive definite by definition, and it is surely symmetric. We just need to verify the triangle inequality. Let $x, y, z \in X$. If $x = y$, then $d(x, y) = 0$ and the triangle inequality is trivial. Suppose then that $x \neq y$. Then we must also have either $x \neq z$ or $y \neq z$, meaning that either $d(x, z) = 1$ or $d(y, z) = 1$ (or perhaps both). Thus

$$d(x, y) = 1 \leq d(x, z) + d(z, y),$$

so the triangle inequality holds.

Example 2.1.5. We can turn $\mathbf{R}^n$ into a metric space in a natural way using something you have undoubtedly seen before. Recall that we already have a way of measuring distances in $\mathbf{R}^n$ via the Euclidean distance formula: given $x = (x_1, x_2, \ldots, x_n)$ and $y = (y_1, y_2, \ldots, y_n)$ in $\mathbf{R}^n$, we can define

$$d(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_n - y_n)^2}.$$

Of course we expect this to define a distance function, but we still need to prove that it satisfies the axioms for a metric. We certainly have $d(x, y) \geq 0$ for all $x, y \in \mathbf{R}^n$, and if $d(x, y) = 0$, then it must be the case that

$$(x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_n - y_n)^2 = 0.$$

But this can only happen if $x_i = y_i$ for $1 \le i \le n$, so $x = y$. Thus d is positive definite. It should also be clear that d is symmetric. The triangle inequality, however, is somewhat more complicated. We really need to prove that for all $x, y, z \in \mathbf{R}^n$,

$$\left(\sum_{i=1}^{n}(x_i - y_i)^2\right)^{1/2} \le \left(\sum_{i=1}^{n}(x_i - z_i)^2\right)^{1/2} + \left(\sum_{i=1}^{n}(z_i - y_i)^2\right)^{1/2}. \tag{2.1}$$

To make the computations simpler, we will define $a = x - z$ and $b = z - y$, so $x - y = a + b$. Then (2.1) becomes

$$\left(\sum_{i=1}^{n}(a_i + b_i)^2\right)^{1/2} \le \left(\sum_{i=1}^{n} a_i^2\right)^{1/2} + \left(\sum_{i=1}^{n} b_i^2\right)^{1/2}. \tag{2.2}$$

This inequality still should not be obvious at all. To prove it, we will need something called the **Cauchy-Schwarz inequality**. This inequality can be proven in myriad ways, often using complex numbers. We will give a clever vector-based proof from [Wad03]. (We will see another clever, though much more sophisticated proof of a general version of the Cauchy-Schwarz inequality in Chapter 6.)

Proposition 2.1.6 (Cauchy-Schwarz inequality)**.** *For all $a, b \in \mathbf{R}^n$,*

$$\sum_{i=1}^{n} a_i b_i \le \left(\sum_{i=1}^{n} a_i^2\right)^{1/2} \left(\sum_{i=1}^{n} b_i^2\right)^{1/2}.$$

Proof. Notice that the left side is precisely the dot product

$$a \cdot b = \sum_{i=1}^{n} a_i b_i,$$

while

$$\|a\| = \left(\sum_{i=1}^{n} a_i^2\right)^{1/2}, \quad \|b\| = \left(\sum_{i=1}^{n} b_i^2\right)^{1/2}$$

where $\|\cdot\|$ denotes the usual norm (or magnitude) of a vector in $\mathbf{R}^n$. Therefore, the Cauchy-Schwarz inequality can be stated more succinctly as

$$a \cdot b \le \|a\|\|b\|.$$

We will focus on proving this inequality. Note that it holds trivially when $b = 0$. Therefore, assume $b \ne 0$ and let

$$t = \frac{a \cdot b}{\|b\|^2}.$$

Then we have

$$(a - tb) \cdot b = a \cdot b - t\|b\|^2 = a \cdot b - a \cdot b = 0,$$

so

$$\begin{aligned} \|a - tb\|^2 &= (a - tb) \cdot (a - tb) \\ &= (a - tb) \cdot a - t(a - tb) \cdot b \\ &= a \cdot a - t(a \cdot b) - 0 \\ &= \|a\|^2 - \frac{(a \cdot b)^2}{\|b\|^2}. \end{aligned}$$

Since $\|a - tb\| \geq 0$, we have

$$\|a\|^2 \geq \frac{(a \cdot b)^2}{\|b\|^2},$$

or $(a \cdot b)^2 \leq \|a\|^2 \|b\|^2$. Taking square roots then yields the result. □

Now we can go back to our proof of the triangle inequality. If we square the right hand side of (2.2) we get

$$\begin{aligned} \left(\left(\sum_{i=1}^n a_i^2 \right)^{1/2} + \left(\sum_{i=1}^n b_i^2 \right)^{1/2} \right)^2 &= \sum_{i=1}^n a_i^2 + \sum_{i=1}^n b_i^2 + 2 \left(\sum_{i=1}^n a_i^2 \right)^{1/2} \left(\sum_{i=1}^n b_i^2 \right)^{1/2} \\ &\geq \sum_{i=1}^n a_i^2 + \sum_{i=1}^n b_i^2 + \sum_{i=1}^n 2a_i b_i \\ &= \sum_{i=1}^n (a_i + b_i)^2 \end{aligned}$$

by the Cauchy-Schwarz inequality. Now (2.1) follows by taking square roots.

During the verification of the triangle inequality in the last example, we used something special about the Euclidean metric on $\mathbf{R}^n$. Indeed, when proving the Cauchy-Schwarz inequality we observed that our metric was related to the *norm* or *magnitude* of a vector in $\mathbf{R}^n$. In fact, whenever we have a notion of magnitude for vectors, we automatically obtain a metric from it.

Let us first recall some facts about *vector spaces*. If $\mathbf{F}$ is a field (which we will always take to be $\mathbf{C}$ or $\mathbf{R}$), recall that a **vector space** over $\mathbf{F}$ is an abelian group $(V, +)$ together with a map $(\alpha, x) \mapsto \alpha x$ from $\mathbf{F} \times V \to V$ satisfying the following axioms:

1. $\alpha(v + w) = \alpha x + \alpha y$ for all $\alpha \in \mathbf{F}$ and all $x, y \in V$;

2. $(\alpha_1 + \alpha_2)x = \alpha_1 x + \alpha_2 x$ for all $\alpha_1, \alpha_2 \in \mathbf{F}$ and all $x \in V$;

3. $(\alpha_1\alpha_2)x = \alpha_1(\alpha_2 x)$ for all $\alpha_1, \alpha_2 \in \mathbf{F}$ and all $x \in V$;

4. $1 \cdot x = x$ for all $x \in V$, where $1 \in \mathbf{F}$ denotes the multiplicative identity.

The prototypical example of a vector space over $\mathbf{R}$ is the Euclidean space $\mathbf{R}^n$, though we will encounter some more exotic examples shortly.

Definition 2.1.7. Let V be a vector space over $\mathbf{R}$ or $\mathbf{C}$. A **norm** on V is a function $\|\cdot\| : V \to \mathbf{R}$ satisfying the following conditions:

1. (Positive definiteness) $\|x\| \geq 0$ for all $x \in V$, and $\|x\| = 0$ if and only if $x = 0$.

2. (Triangle inequality) $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in V$.

3. (Homogeneity) $\|\alpha x\| = |\alpha|\|x\|$ for all $\alpha \in \mathbf{R}$ and $x \in V$.

Example 2.1.8. In the previous example, we were using the **Euclidean norm** (also called the **2-norm**) on $\mathbf{R}^n$:

$$\|x\|_2 = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}.$$

As in our example, any normed vector space becomes a metric space in a natural way—to compute the distance between two vectors, we simply find the length of the vector between them. The proof of the following proposition is left as an exercise to the student.

Proposition 2.1.9. *If V is an $\mathbf{R}$-vector space and $\|\cdot\| : V \to \mathbf{R}$ is a norm on V, then*

$$d(x, y) = \|x - y\|$$

defines a metric on V.

There are actually many other ways to define norms on $\mathbf{R}^n$. The next two examples illustrate the most commonly studied ones besides the Euclidean norm.

Example 2.1.10. Define the **1-norm** (sometimes called the **taxicab** or **Manhattan norm**[1]) on $\mathbf{R}^n$ by

$$\|x\|_1 = |x_1| + |x_2| + \cdots + |x_n|.$$

[1]These alternative names come from the fact that in two dimensions, the 1-norm records the distance one would travel if they were required to move only in the x and y directions. It is exactly how one would measure distance when driving along a city grid, as opposed to the "as the crow flies" distance afforded by the Euclidean norm.

The proof that $\|\cdot\|_1$ defines a norm is fairly straightforward. Certainly we have $\|x\|_1 \geq 0$ for all $x \in \mathbf{R}^n$, and $\|x\|_1 = 0$ implies $x_i = 0$ for $1 \leq i \leq n$, so $x = 0$. The triangle inequality follows simply by applying the usual triangle inequality in $\mathbf{R}$ to each term in the sum: given $x, y \in \mathbf{R}^n$,

$$\begin{aligned}\|x+y\|_1 &= |x_1+y_1| + |x_2+y_2| + \cdots + |x_n+y_n| \\ &\leq |x_1| + |y_1| + |x_2| + |y_2| + \cdots + |x_n| + |y_n| \\ &= \|x\|_1 + \|y\|_1.\end{aligned}$$

Homogeneity is also easy: given $\alpha \in \mathbf{R}$ and $x \in \mathbf{R}^n$, we have

$$\begin{aligned}\|\alpha x\|_1 &= |\alpha x_1| + |\alpha x_2| + \cdots + |\alpha x_n| \\ &= |\alpha|\big(|x_1| + |x_2| + \cdots + |x_n|\big) \\ &= |\alpha|\|x\|_1.\end{aligned}$$

Thus $\|\cdot\|_1$ is a norm on $\mathbf{R}^n$, and consequently $d(x,y) = \|x-y\|_1$ defines a metric on $\mathbf{R}^n$.

Example 2.1.11. One can also define the ∞-**norm** on $\mathbf{R}^n$ by

$$\|x\|_\infty = \max_i |x_i|.$$

As with the 1-norm, it is easy to check that $\|\cdot\|_\infty$ is positive definite and homogeneous. The triangle inequality is also not hard to verify: if $x, y \in \mathbf{R}^n$, then

$$|x_j + y_j| \leq |x_j| + |y_j| \leq \max_i |x_i| + \max_i |y_i|$$

for all $1 \leq j \leq n$, so

$$\|x+y\|_\infty = \max_i |x_i + y_i| \leq \max_i |x_i| + \max_i |y_i| = \|x\|_\infty + \|y\|_\infty.$$

Thus the triangle inequality holds and $\|\cdot\|_\infty$ is a norm.

It will often behoove us to consider more exotic metric spaces beyond those that come from $\mathbf{R}^n$. In particular, certain things that work in $\mathbf{R}$ (or $\mathbf{R}^n$) may fail in general metric spaces. In order to properly investigate these issues when the time comes, we will need some more interesting examples at our disposal. We have one such example already, namely the discrete metric space associated to a set X. We can obtain three more examples by looking at infinite-dimensional versions of our three metric spaces associated to $\mathbf{R}^n$.

Example 2.1.12. Instead of considering $\mathbf{R}^n$, which consists of n-tuples of real numbers, we could study collections of infinite *sequences* of real numbers. However, we can't simply study the set of *all* sequences. We need to include extra conditions in order to guarantee that we obtain norms like the ones we have defined on $\mathbf{R}^n$.

The first space that we will define is an "infinite-dimensional" analogue of $\mathbf{R}^n$ equipped with the 1-norm:

$$\ell^1 = \left\{ (x_i)_{i=1}^{\infty} : x_i \in \mathbf{R} \text{ and } \sum_{i=1}^{\infty} |x_i| < \infty \right\}.$$

In other words, ℓ^1 is the set of all absolutely summable sequences of real numbers. It is straightforward to check that this set forms a vector space over $\mathbf{R}$, where the addition and scalar multiplication are defined entrywise: given $x = (x_i)_{i=1}^{\infty}$ and $y = (y_i)_{i=1}^{\infty}$ in ℓ^1, define

$$x + y = (x_i + y_i)_{i=1}^{\infty},$$

and if $\alpha \in \mathbf{R}$ we have $\alpha x = (\alpha x_i)_{i=1}^{\infty}$. Notice that the summability criterion is exactly what we need to extend the 1-norm to ℓ^1: for all $x \in \ell^1$, we can define

$$\|x\|_1 = \sum_{i=1}^{\infty} |x_i|.$$

Note that this sum is guaranteed to converge thanks to our definition of ℓ^1, so we have a well-defined function $\|\cdot\|_1 : \ell^1 \to \mathbf{R}$. To prove it is a norm, we can apply facts about the 1-norm on $\mathbf{R}^n$ to partial sums and then take limits. Notice first that for all $x \in \ell^1$ and each $n \geq 1$, we have

$$\sum_{i=1}^{n} |x_i| \geq 0,$$

so

$$\|x\|_1 = \lim_{n\to\infty} \sum_{i=1}^{n} |x_i| \geq 0.$$

Moreover, if $\|x\|_1 = 0$, then for all $n \geq 1$ we have

$$|x_n| \leq \sum_{i=1}^{n} |x_i| \leq \|x\|_1 = 0,$$

since the partial sums are monotonically increasing. Thus $|x_n| = 0$ for all $n \geq 1$, meaning that $x = 0$. Thus $\|\cdot\|_1$ is positive definite. Now let $x, y \in \ell^1$. Then for all $n \geq 1$ we have

$$\sum_{i=1}^{n} |x_i + y_i| \leq \sum_{i=1}^{n} |x_i| + \sum_{i=1}^{n} |y_i|$$

by the triangle inequality for the 1-norm on $\mathbf{R}^n$. Taking limits, we obtain

$$\|x + y\|_1 = \lim_{n\to\infty} \sum_{i=1}^{n} |x_i + y_i| \leq \lim_{n\to\infty} \sum_{i=1}^{n} |x_i| + \lim_{n\to\infty} \sum_{i=1}^{n} |y_i| = \|x\|_1 + \|y\|_1.$$

Therefore, the triangle inequality holds for $\|\cdot\|_1$. Finally, let $\alpha \in \mathbf{R}$ and $x \in \ell^1$. Then for all $n \geq 1$,

$$\sum_{i=1}^{n} |\alpha x_i| = |\alpha| \sum_{i=1}^{n} |x_i|,$$

so

$$\|\alpha x\|_1 = \lim_{n\to\infty} \sum_{i=1}^{n} |\alpha x_i| = |\alpha| \lim_{n\to\infty} \sum_{i=1}^{n} |x_i| = |\alpha| \|x\|_1.$$

Thus $\|\cdot\|_1$ is a norm on ℓ^1, so $d(x, y) = \|x - y\|_1$ defines a metric on ℓ^1.

Example 2.1.13. In a similar fashion to how we built ℓ^1, we could construct an analogue of $\mathbf{R}^n$ with the Euclidean norm by considering square summable sequences of real numbers. We denote this set by ℓ^2:

$$\ell^2 = \left\{ (x_i)_{i=1}^{\infty} : x_i \in \mathbf{R} \text{ and } \sum_{i=1}^{\infty} x_i^2 < \infty \right\}.$$

The definition guarantees that

$$\|x\|_2 = \left(\sum_{i=1}^{\infty} x_i^2 \right)^{1/2}$$

gives a well-defined function $\|\cdot\|_2 : \ell^2 \to \mathbf{R}$. It is left as an exercise (Exercise 2.1.4) to show that $\|\cdot\|_2$ defines a norm on ℓ^2.

Example 2.1.14. Finally, we can produce an infinite-dimensional version of the space $(\mathbf{R}^n, \|\cdot\|_\infty)$ by replacing the maximum with a supremum. Doing so forces us to consider *bounded* sequences of real numbers:

$$\ell^\infty = \left\{ (x_i)_{i=1}^{\infty} : x_i \in \mathbf{R} \text{ and } \sup_i |x_i| < \infty \right\}.$$

We can then equip ℓ^∞ with a norm by simply taking the supremum of each sequence:

$$\|x\|_\infty = \sup_i |x_i|.$$

The proof that $\|\cdot\|_\infty$ defines a norm is similar to the finite-dimensional case, except taking suprema is much more delicate than taking maxima.

We begin by observing that given any $x \in \ell^\infty$, $|x_i| \geq 0$ for all i, so $\sup |x_i| \geq 0$. Moreover, if $\|x\|_\infty = 0$, then we have

$$|x_i| \leq \sup |x_j| = 0$$

for all i, so $|x_i| = 0$ for all i. Thus $x = 0$ and $\|x\|_\infty$ is positive definite.

Next we will verify the triangle inequality. Let $x, y \in \ell^\infty$, and notice that for all i we have

$$|x_i + y_i| \le |x_i| + |y_i|.$$

Therefore, taking the supremum of the right hand side yields

$$|x_i + y_i| \le \sup |x_j| + \sup |y_j|$$

for all i. Now the right hand side is an upper bound for the left, so it must dominate the least upper bound of the left hand side. In other words,

$$\sup |x_j + y_j| \le \sup |x_j| + \sup |y_j|,$$

which says precisely that $\|x + y\|_\infty \le \|x\|_\infty + \|y\|_\infty$. Thus the triangle inequality holds.

Finally, suppose $\alpha \in \mathbf{R}$ and $x \in \ell^\infty$. Then by the properties of the supremum,

$$\|\alpha x\|_\infty = \sup |\alpha x_i| = \sup |\alpha||x_i| = |\alpha| \sup |x_i| = |\alpha| \|x\|_\infty,$$

so $\|\cdot\|_\infty$ is a norm on ℓ^∞.

There is one more example that deserves mentioning right now. This metric space is one that we will revisit and study in detail later on.

Example 2.1.15. Let $X = C[0,1]$, which denotes the set of all continuous functions $f : [0,1] \to \mathbf{R}$. For each $f \in X$, define

$$\|f\|_\infty = \sup_{x \in [0,1]} |f(x)|.$$

Notice that this supremum exists, since a continuous function on a closed interval is automatically bounded.[2] We claim that $\|\cdot\|_\infty$ defines a norm on $C[0,1]$.

It is obvious that $\|f\|_\infty \ge 0$ for all $f \in C[0,1]$. Moreover, if $\|f\|_\infty = 0$, then for all $x \in [0,1]$ we have

$$0 \le |f(x)| \le \sup_y |f(y)| = \|f\|_\infty = 0,$$

so $f(x) = 0$ for all $x \in [0,1]$. Thus $f = 0$, so $\|\cdot\|_\infty$ is positive definite. It is also straightforward to check that $\|\cdot\|_\infty$ is homogeneous: given $\alpha \in \mathbf{R}$ and $f \in C[0,1]$, we have

$$\|\alpha x\|_\infty = \sup_x |\alpha f(x)| = |\alpha| \sup_x |f(x)| = |\alpha| \|x\|_\infty.$$

All that remains is the triangle inequality. Let $f, g \in C[0,1]$. Then for all $x \in [0,1]$ we have

$$|f(x) + g(x)| \le |f(x)| + |g(x)| \le \sup_y |f(y)| + \sup_y |g(y)| = \|f\|_\infty + \|g\|_\infty$$

[2] In fact, a continuous function on $[0,1]$ must assume a maximum, we could actually replace "sup" with "max" here.

by the usual triangle inequality for $\mathbf{R}$. Thus $\|f\|_\infty + \|g\|_\infty$ is an upper bound for $|f(x) + g(x)|$ as x ranges over $[0, 1]$, meaning that

$$\|f + g\|_\infty = \sup_x |f(x) + g(x)| \leq \|f\|_\infty + \|g\|_\infty.$$

Thus the triangle inequality holds, so $\|\cdot\|_\infty$ is a norm.

Before closing out this section, let's introduce a couple more definitions that will be useful from time to time. First observe that if (X, d) is a metric space and $E \subseteq X$, then the restriction of d to E yields a metric on E. (More precisely, we can restrict the function d to $E \times E$.) Thus (E, d) is a metric space in its own right.

> **Definition 2.1.16.** Let (X, d) be a metric space, and let $E \subseteq X$. The metric space (E, d) obtained by restricting d to $E \times E$ is called a **subspace** of (X, d).

Example 2.1.17. We can think of $\mathbf{Q}$ as a subspace of $\mathbf{R}$ by equipping $\mathbf{Q}$ with the restriction of the standard metric:

$$d(r, s) = |r - s| \text{ for } r, s \in \mathbf{Q}.$$

In the last example, we could think of the inclusion $\iota : \mathbf{Q} \to \mathbf{R}$ as a map between metric spaces. By construction, it preserves the distance between two rational numbers when $\mathbf{Q}$ and $\mathbf{R}$ are equipped with their standard metrics. More generally, one can consider distance-preserving maps between arbitrary metric spaces.

> **Definition 2.1.18.** Suppose (X_1, d_1) and (X_2, d_2) are metric spaces. A function $f : X_1 \to X_2$ is called an **isometry** if
>
> $$d_2(f(x), f(y)) = d_1(x, y)$$
>
> for all $x, y \in X_1$.

It is an exercise to show that any isometry is necessarily injective. In the event that there is a bijective isometry $f : X_1 \to X_2$, we say that the metric spaces (X_1, d_1) and (X_2, d_2) are **isometric**.

Exercises for Section 2.1

Exercise 2.1.1 ([Rud76], Exercise 2.11)**.** Determine whether each of the functions $d : \mathbf{R} \times \mathbf{R} \to \mathbf{R}$ specifies a metric on $\mathbf{R}$.

(a) $d(x, y) = (x - y)^2$

(b) $d(x,y) = \sqrt{|x-y|}$

(c) $d(x,y) = \left|x^2 - y^2\right|$

(d) $d(x,y) = |x - 2y|$

(e) $d(x,y) = \dfrac{|x-y|}{1+|x-y|}$

Exercise 2.1.2. Let V be a $\mathbf{R}$-vector space. If $\|\cdot\| : V \to \mathbf{R}$ is a norm on V, prove that the map $d : V \times V \to \mathbf{R}$ defined by

$$d(x,y) = \|x - y\|$$

is a metric on V.

Exercise 2.1.3 ([HS91], Exercise 4.2.2)**.** Let (X,d) be a metric space. Prove the **quadrilateral inequality**:

$$|d(w,x) - d(x,y)| \leq d(w,z) + d(z,y)$$

for all $w, x, y, z \in X$.

Exercise 2.1.4. Earlier we defined the space of square-summable sequences of real numbers:

$$\ell^2 = \left\{ (x_i)_{i=1}^{\infty} : \sum x_i^2 < \infty \right\}.$$

Prove that

$$\|x\|_2 = \left(\sum_{i=1}^{\infty} x_i^2 \right)^{1/2}$$

defines a norm on ℓ^2. Conclude that ℓ^2 is a normed vector space, hence a metric space.

Hint: As with ℓ^1, first consider the n^{th} partial sum of the given series. Apply what we already proved about the Euclidean norm on $\mathbf{R}^n$, and then take limits.

Exercise 2.1.5. Let (X_1, d_1) and (X_2, d_2) be metric spaces, and suppose $f : X_1 \to X_2$ is an isometry. Prove that f is injective.

2.2 Topology of Metric Spaces

In real analysis, the standard metric on $\mathbf{R}$ opens up all kinds of tools for studying functions on the real line. In particular, one only needs a notion of distance to determine when a sequence converges, or when a function is continuous, among other things. Many of these concepts can also be phrased in terms of open and closed sets, both of which make sense in a general metric space. The initial building blocks are *open balls*, which are defined purely in terms of the metric. Throughout this section, let (X,d) denote an arbitrary metric space.

Definition 2.2.1. Given $x \in X$ and $r > 0$, the **open ball** of radius r centered at x is the set

$$B_r(x) = \{y \in X : d(x, y) < r\}.$$

Before going any further, let's think about what the open balls look like in some common examples of metric spaces. In some cases we will obtain a geometric picture to go along with our metric.

Example 2.2.2. If we equip $\mathbf{R}$ with its standard metric, the open balls are precisely the open intervals. That is, given $x \in \mathbf{R}$ and $r > 0$, the corresponding open ball is

$$B_r(x) = (x - r, x + r).$$

Example 2.2.3. Now consider $\mathbf{R}^2$ endowed with the Euclidean norm $\|\cdot\|_2$. For simplicity, we will try to describe the ball $B_1(0)$, which is called the **open unit ball**. This ball is precisely the set

$$B_1(0) = \{x \in \mathbf{R}^2 : \|x\|_2 < 1\},$$

and using the fact that $\|x\|_2 < 1$ if and only if $x_1^2 + x_2^2 < 1$, we see that $B_1(0)$ is just the open disk of radius 1 centered at the origin.

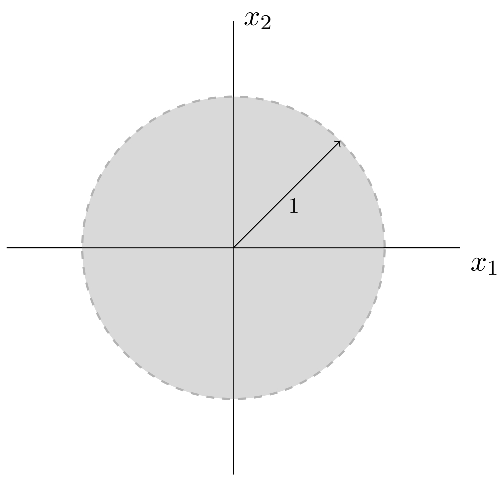

What if instead we consider $\mathbf{R}^2$ with the "taxicab" norm $\|\cdot\|_1$? Now the open unit ball looks like

$$B_1(0) = \{x \in \mathbf{R}^2 : |x_1| + |x_2| < 1\}.$$

In the first quadrant, $\|x\|_1 < 1$ simply means $x_1 + x_2 < 1$, so the ball is bounded by the line $x_2 = 1 - x_1$. In the second quadrant, we have $x_1 < 0$, so $\|x\|_1 < 1$ is equivalent to $-x_1 + x_2 < 1$, or $x_2 < 1 + x_1$. Continuing in this fashion, we see that the unit ball looks like a diamond centered at the origin.

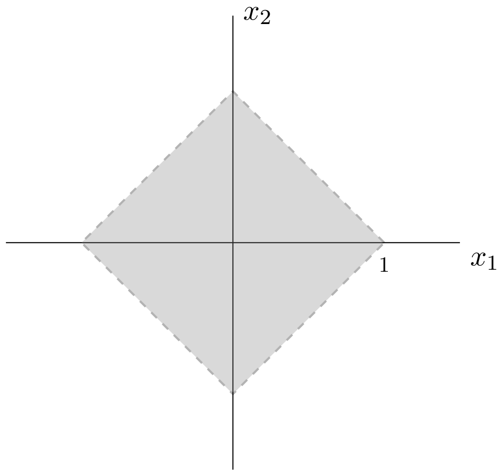

Finally, let's think about what the open balls in $\mathbf{R}^2$ look like when we use the norm $\|\cdot\|_\infty$. Since $\|x\|_\infty < 1$ if and only if $|x_1| < 1$ and $|x_2| < 1$, we have

$$B_1(0) = \{x \in \mathbf{R}^2 : |x_1|, |x_2| < 1\}.$$

Thus the ball is bounded by the lines $x_1 = \pm 1$ and $x_2 = \pm 1$, and we obtain a square centered at the origin.

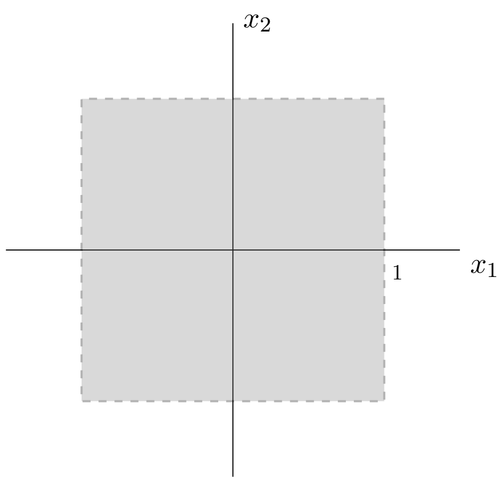

Note that in any of these three cases, an arbitrary open ball $B_r(x)$ can be obtained from the unit ball $B_1(0)$ via a translation and a dilation.

Example 2.2.4. Let X be a set equipped with the discrete metric. What do the open balls look like? The answer will depend on the radius r. If $x \in X$ and $r \leq 1$, the only point $y \in X$ satisfying $d(x, y) < r$ is $y = x$. That is,

$$B_r(x) = \{y \in X : d(x, y) < r\} = \{x\}.$$

If $r > 1$, then for all $y \in X$ we have $d(x, y) \leq 1 < r$, so $B_r(x) = X$.

The notion of "closeness" afforded by open balls in a metric space allows us to define what it means to be "inside" a set, in a certain sense.

Definition 2.2.5. Let $E \subseteq X$. A point $x \in X$ is an **interior point** of E if there is an $r > 0$ such that $B_r(x) \subseteq E$.

The set of all interior points of E is called the **interior** of E, which is denoted by E° or $\operatorname{int}(E)$.

Definition 2.2.6. A set $U \subseteq X$ is **open** if every point $x \in U$ is an interior point of U. Equivalently, U is open if $U = U^\circ$.

Example 2.2.7. In any metric space X, the empty set and X are both open.

One would expect there to be many more open sets than just $\emptyset$ and X.[3] Given the name, we might guess that any open ball is an example of an open set.

Proposition 2.2.8. *Let $x \in X$ and $r > 0$. Then $B_r(x)$ is an open set.*

Proof. Let $y \in B_r(x)$ and put $s = d(x, y)$. Then we claim that $B_{r-s}(y) \subseteq B_r(x)$. Given any point $z \in B_{r-s}(y)$, we have

$$d(z, x) \leq d(z, y) + d(y, x) < r - s + s = r$$

by the triangle inequality. Therefore, $z \in B_r(x)$. It follows that $B_{r-s}(y) \subseteq B_r(x)$, so $B_r(x)$ is open. $\square$

Likewise, since the interior of a set E consists of precisely the interior points of E, one would hope that E° is always an open set. This result is far from being tautological, however. It requires us to show that any interior point of E is necessarily an interior point of E°.

Proposition 2.2.9. *Let $E \subseteq X$. The interior E° is open.*

Proof. Let $x \in E^\circ$. Then x is an interior point of E, so there exists $r > 0$ such that $B_r(x) \subseteq E$. We need to show that $B_r(x) \subseteq E^\circ$. Let $y \in B_r(x)$. Since $B_r(x)$ is open by Proposition 2.2.8, there exists $s > 0$ such that

$$B_s(y) \subseteq B_r(x) \subseteq E,$$

so $y \in E^\circ$. Since $y \in B_r(x)$ was arbitrary, $B_r(x) \subseteq E^\circ$, so E° is open. $\square$

[3] Indeed, these are the only open sets precisely when X carries the indiscrete topology, which is metrizable if and only if X consists of a single point.

It is also useful to note that we can often build new open sets out of old ones via unions and intersections.[4]

Proposition 2.2.10. *Let $\{U_i\}_{i\in I}$ be a family of sets $U_i \subseteq X$ such that each U_i is open. Then the union $\bigcup_{i\in I} U_i$ is open.*

Proof. Let $x \in \bigcup_{i\in I} U_i$. Then $x \in U_{i_0}$ for some $i_0 \in I$, so there exists $r > 0$ such that

$$B_r(x) \subseteq U_{i_0} \subseteq \bigcup_{i\in I} U_i.$$

Thus x is an interior point of $\bigcup_{i\in I} U_i$, and it follows that the union is open. □

Note that in Proposition 2.2.10, there is no restriction on the size of the index set. The union of arbitrarily many open sets is always open. On the other hand, the intersection of a family of open sets is only guaranteed to be open if the family is finite.

Proposition 2.2.11. *Let $\{U_i\}_{i=1}^n$ be a finite collection of sets $U_i \subseteq X$ with each U_i open. Then the intersection $\bigcap_{i=1}^n U_i$ is open.*

Proof. Let $x \in \bigcap_{i=1}^n U_i$. Then $x \in U_i$ for $1 \leq i \leq n$, so for each i there is an $r_i > 0$ such that $B_{r_i}(x) \subseteq U_i$. Set

$$r = \min_{1\leq i\leq n} r_i.$$

Then $B_r(x) \subseteq B_{r_i}(x) \subseteq U_i$ for all i, and it follows that $B_r(x) \subseteq \bigcap_{i=1}^n U_i$. Hence the intersection is open. □

Example 2.2.12. To see why the finiteness condition in Proposition 2.2.11 is necessary, we need only look at $\mathbf{R}$. For each $n \in \mathbf{N}$, set

$$U_n = \left(-\tfrac{1}{n}, \tfrac{1}{n}\right).$$

Then each U_n is open, but

$$\bigcap_{n=1}^{\infty} U_n = \{0\},$$

which is not open.

[4] Propositions 2.2.10 and 2.2.11 together with Example 2.2.7 say that the open sets in a metric space (X, d) form a **topology** on X.

Example 2.2.13. Let (X, d) be a discrete metric space. We claim that *every* subset of X is open. Given $E \subseteq X$, write

$$E = \bigcup_{x \in E} \{x\}.$$

Recall that $B_1(x) = \{x\}$, so $\{x\}$ is open for each $x \in X$. Therefore, E is a union of open sets, hence open by Proposition 2.2.10.

With these basic results for open sets out of the way, we turn to the dual notion—that of a *closed* set.

Definition 2.2.14. A set $F \subseteq X$ is said to be **closed** if its complement F^c is open.

Example 2.2.15. Some simple examples of closed sets are the **closed balls**: given $x \in X$ and $r > 0$, define the closed ball of radius r centered at x to be

$$B_r[x] = \{y \in X : d(x, y) \leq r\}.$$

It is left as an exercise to show that any closed ball is a closed set.

We immediately have some results about unions and intersections of closed sets, which follow easily from the corresponding facts about open sets.

Proposition 2.2.16. *Let $\{F_i\}_{i \in I}$ be a family of closed sets in X. Then the intersection $\bigcap_{i \in I} F_i$ is closed.*

Likewise, if $\{F_i\}_{i=1}^n$ is a finite collection of closed sets in X, then $\bigcup_{i=1}^n F_i$ is closed.

Proof. The first assertion follows from Proposition 2.2.10 and De Morgan's law. Observe that if $\{F_i\}_{i \in I}$ is any collection of closed sets,

$$\left(\bigcap_{i \in I} F_i\right)^c = \bigcup_{i \in I} F_i^c,$$

is open by Proposition 2.2.10 since each F_i^c is open. Thus $\bigcap_{i \in I} F_i$ is closed.

Similarly, given a finite collection $\{F_i\}_{i=1}^\infty$ of closed sets, we have

$$\left(\bigcup_{i=1}^\infty F_i\right)^c = \bigcap_{i=1}^\infty F_i^c.$$

This set is open by Proposition 2.2.11 since each F_i^c is open. $\square$

Example 2.2.17. Given what we have seen for open sets, it should not be surprising that an infinite union of closed sets need not be closed. Consider the closed subsets of $\mathbf{R}$ defined by

$$F_n = \left[\tfrac{1}{n}, 2\right]$$

for $n \in \mathbf{N}$. Then each F_n is closed, but

$$\bigcup_{i=1}^{\infty} F_i = (0, 2]$$

is not closed.

Example 2.2.18. Let (X, d) be a discrete metric space, and let $E \subseteq X$. Then E^c is open, since all subsets of X are open. Thus E is closed, and it follows that all subsets of a discrete metric space are simultaneously open and closed.

Definition 2.2.19. A subset of a metric space is said to be **clopen** if it is both open and closed.

Remark 2.2.20. We have just seen that in a discrete space, all sets are clopen. In contrast, the only clopen subsets of $\mathbf{R}$ are the trivial ones – $\emptyset$ and $\mathbf{R}$. We will see later that the existence of nontrivial clopen sets in an arbitrary metric space X is closely tied to *connectedness.*

There is an alternative definition of closed set in terms of *limit points.* This definition will often be more useful than the one we have given, since it allows one to use sequences to characterize when a set is closed.

Definition 2.2.21. Let $E \subseteq X$. We say a point $x \in X$ is a **limit point** of E if for any $r > 0$, the open ball $B_r(x)$ contains a point $y \in E$ with $y \neq x$.

It is worth noting that a limit point of a set E may or may not belong to E. The limit points of E are precisely the points x that can be approximated arbitrarily closely by points of E other than x. This may very well be true for some (if not all) points of E, but also for points outside of E, as the following figure illustrates.

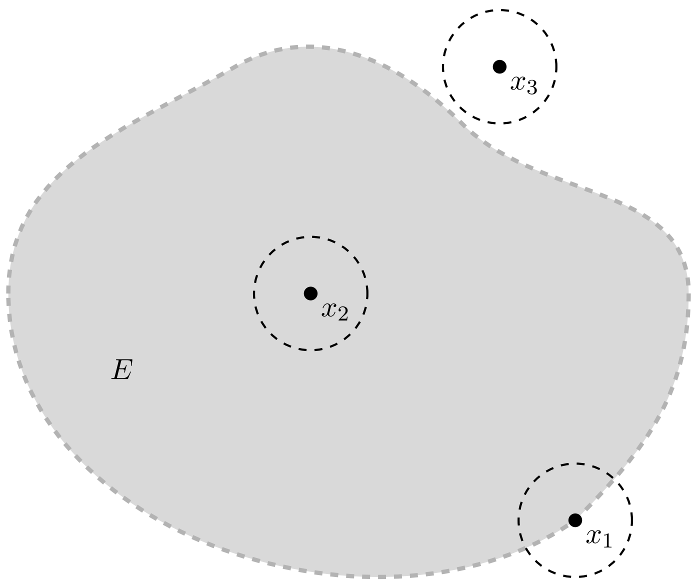

Notice that the point x_1 does not lie in E, but nevertheless any open ball centered at x_1 contains points of E. (Notice that any ball must also contain points from E^c.) Thus x_1 is a limit point of E. On the other hand, the point x_2 belongs to E, and it is also a limit point of E. Finally, we can find an open ball centered at x_3 that contains no points of E, so it cannot be a limit point of E.

The discussion so far is a little misleading in one regard. We saw above that the point x_2 both belongs to E and is a limit point of E. Indeed, it is clear from the way the picture is drawn that x_2 is an interior point of E. It need not be the case in general that every point of E is a limit point of E.

Example 2.2.22. Define $E \subseteq \mathbf{R}$ by

$$E = \left\{\tfrac{1}{n} : n \in \mathbf{N}\right\} \cup \{0\}.$$

It should be clear that 0 is a limit point of E, since given any $r > 0$, $1/n \in B_r(0)$ for any $n > 1/r$. Moreover, it is the only point of E that is also a limit point of E. Given $n \in \mathbf{N}$, let $x = \frac{1}{n}$ and set $r = x - \frac{1}{n+1}$. Then $B_r(x) \cap E = \{x\}$, so x is not a limit point of E.

We have now seen how a point $x \in E$ can fail to be a limit point of E—there must be an open ball centered at x that contains no other points of E.

Definition 2.2.23. Let $E \subseteq X$. A point $x \in E$ is an **isolated point** of E if there exists $r > 0$ such that $B_r(x) \cap E = \{x\}$.

Based on our earlier examples, it appears that there is a dichotomy between points of a set $E \subseteq X$—they are either limit points of E or isolated points of E. The proof of this fact is a straightforward exercise.

Proposition 2.2.24. *Let $E \subseteq X$. A point $x \in E$ is an isolated point of E if and only if it is not a limit point of E.*

The argument we gave in Example 2.2.22 to show that every point other than 0 was isolated can be strengthened to obtain the following result.

Proposition 2.2.25. *Let $E \subseteq X$ and suppose $x \in X$ is a limit point of E. Then every open ball centered at x contains infinitely many points of E.*

Proof. Suppose, to the contrary, that there exists $r > 0$ such that $B_r(x)$ contains only finitely many points of E distinct from x. That is,

$$B_r(x) \cap E = \{x_1, x_2, \ldots, x_n\}$$

with $x_i \neq x$ for $1 \leq i \leq n$. Put

$$r_0 = \min_{1 \leq i \leq n} d(x, x_i).$$

Then $0 < r_0 < r$ since $d(x, x_i) > 0$ for all i, and $x_i \notin B_{r_0}(x)$ for all i. Since $B_{r_0}(x) \subseteq B_r(x)$, it follows that $B_{r_0}(x)$ contains no point of E except possibly x itself. But this contradicts the assumption that x is a limit point of E. $\square$

Corollary 2.2.26. *A finite subset of a metric space has no limit points.*

As promised earlier, we now come to an alternative characterization of closed sets, which is phrased in terms of limit points.

Proposition 2.2.27. *A set $F \subseteq X$ is closed if and only if it contains all of its limit points.*

Proof. Suppose first that F is closed, and let $x \in F^c$. Since F^c is open, there exists $r > 0$ such that $B_r(x) \subseteq F^c$. But then $B_r(x)$ contains no points of F, so x cannot be a limit point of F. Thus any limit points of F must belong to F.

Conversely, suppose F contains all of its limit points, and let $x \in F^c$. Then x is not a limit point of F, so there exists $r > 0$ such that $B_r(x)$ contains no points of F. That is, $B_r(x) \subseteq F^c$, so x is an interior point of F^c. Since $x \in F^c$ was arbitrary, F^c is open, hence F is closed. $\square$

Corollary 2.2.28. *If $E \subseteq X$ is finite, then it is closed. In particular, singletons in a metric space are closed.*

Proof. If E is finite, then it has no limit points by Corollary 2.2.26. Therefore, E vacuously contains all of its limits points, so it is closed. $\square$

Now we come to one last idea regarding closed sets. First, let us introduce a bit of notation: given $E \subseteq X$, we let E' denote the set of all limit points of E.

Definition 2.2.29. Let $E \subseteq X$. The **closure** of E is the set $\overline{E} = E \cup E'$.

Since the closure of a set E is formed by adjoining the limit points of E, one might expect that $\overline{E}$ is the minimal way of building a closed set from E.

Proposition 2.2.30. *Let $E \subseteq X$.*

(a) *$\overline{E}$ is a closed set.*

(b) *$E = \overline{E}$ if and only if E is closed.*

(c) *$\overline{E}$ is the smallest closed set containing E. That is, if $E \subseteq F$ and F is closed, then $\overline{E} \subseteq F$.*

Proof. To prove (a), let $x \in (\overline{E})^c$. Then $x \notin E$ and x is not a limit point of E. Thus there exists $r > 0$ such that $B_r(x) \subseteq E^c$. We claim that in fact $B_r(x) \subseteq (\overline{E})^c$. To that end, suppose there is a limit point y of E in $B_r(x)$. Since $B_r(x)$ is open, there is an $s > 0$ such that $B_s(y) \subseteq B_r(x)$. Also, since y is a limit point of E, $B_s(y)$ must contain some point $z \in E$. But then $z \in B_r(x)$, which contradicts the fact that $B_r(x) \subseteq E^c$. Therefore, $B_r(x) \subseteq (\overline{E})^c$, so $\overline{E}$ is closed.

It follows immediately from (a) that if $E = \overline{E}$, then E is closed. Suppose then that E is closed. Then $E' \subseteq E$, so $\overline{E} = E \cup E' = E$. Thus (b) holds.

Finally, suppose F is a closed set with $E \subseteq F$. Then any limit point of E is necessarily a limit point of F, so $E' \subseteq F$ since F is closed. Thus $\overline{E} \subseteq F$. $\square$

Though we won't use the following concept much, it is worth mentioning in order to provide some closure to our earlier discussion about limit points. We saw that a limit point of a set E need not belong to E, provided it lies on the "edge" of E.

Definition 2.2.31. Let $E \subseteq X$. The **boundary** of E is the set $\partial E = \overline{E} \backslash E^\circ$.

Remark 2.2.32. There are some equivalent characterizations one could take for the boundary of a set. One could also define $\partial E = \overline{E} \cap \overline{E^c}$, or we could take ∂E to be the set of points x with the property that any open ball centered at x contains points of both E and E^c.

It is not hard to see that any limit point of E that does not belong to E is necessarily a boundary point of E. However, there are a couple of less-than-satisfying things that can happen. In particular, a boundary point of E must be a limit point of either E or E^c, but it need not be a limit point of both.

Example 2.2.33. Let $X = [0,1]$ equipped with the metric inherited from the standard metric on $\mathbf{R}$. Let $E = (0,1]$. Then $\overline{E} = [0,1]$ and $E^\circ = (0,1)$, so $\partial E = \{0,1\}$. Notice that both 0 and 1 are limit points of E, but neither is a limit point of E^c. In particular, $E^c = \{0\}$, which has no limit points in X.

Example 2.2.34. Let $X = \mathbf{R}$ with its standard metric, and let $x \in \mathbf{R}$. If we let $E = \{x\}$, then E is closed and $E^\circ = \emptyset$, so $\partial E = \{x\}$. However, x is not a limit point of E. It is a limit point of E^c, however.

We will finish our discussion of basic topology with one last topic that generalizes a familiar result from $\mathbf{R}$.

> **Definition 2.2.35.** Let $E \subseteq X$. We say E is **dense** in X if $\overline{E} = X$. Equivalently, E is dense in X if every point $x \in X$ is a point of E or a limit point of E.

Example 2.2.36. It is well-known that $\mathbf{Q}$ is dense in $\mathbf{R}$, since every real number is a limit point of the rationals.

The last example is quite interesting for other reasons. Recall that $\mathbf{R}$ is an uncountable set, while $\mathbf{Q}$ is countable. Thus $\mathbf{R}$ is much larger than $\mathbf{Q}$ from a set-theoretic standpoint. However, $\mathbf{Q}$ is dense in $\mathbf{R}$, so it manages to "fill up" the reals in a topological sense. That is, we have the ability to approximate any real number arbitrarily well with elements of $\mathbf{Q}$, which is a much more manageable set.

> **Definition 2.2.37.** A metric space X is **separable** if it contains a countable dense subset.

Example 2.2.38. As described above, $\mathbf{R}$ is separable. More generally, $\mathbf{R}^n$ is separable since $\mathbf{Q}^n$ is countable and dense. To see this, let $x \in \mathbf{R}^n$, and let $\varepsilon > 0$ be given. For each $1 \leq i \leq n$, choose a rational number q_i with $|x_i - q_i| < \varepsilon/\sqrt{n}$. Put $q = (q_1, q_2, \ldots, q_n)$. Then

$$\|x - q\|_2^2 = \sum_{i=1}^{n}(x_i - q_i)^2 < \sum_{i=1}^{n}\frac{\varepsilon^2}{n} = \varepsilon^2,$$

so $\|x - q\|_2 < \varepsilon$. It follows that x is a limit point of $\mathbf{Q}^n$.

Example 2.2.39. One can extend the arguments of the previous example to show that ℓ^2 is separable. The countable dense subset is

$$E = \left\{ (q_i)_{i=1}^{\infty} : q_i \in \mathbf{Q} \text{ for all } i \text{ and } q_i = 0 \text{ for all but finitely many } i \right\}.$$

It is an exercise to verify that E is countable and dense in ℓ^2.

Example 2.2.40. Let (X, d) be a discrete metric space. Since every subset of X is closed, the only dense subset is X itself. Thus X is separable if and only if X is countable.

Separability for metric spaces is closely related to another kind of countability condition. This condition involves identifying a family of "basic" open sets that generate all the open sets in the metric space via unions. One such family consists precisely of the open balls.

Proposition 2.2.41. *Any open set in a metric space X can be written as a union of open balls.*

Proof. Suppose $U \subseteq X$ is open. Then for each $x \in U$, there is an $r_x > 0$ such that $B_{r_x}(x) \subseteq U$. Thus

$$\bigcup_{x \in U} B_{r_x}(x) \subseteq U.$$

But each $x \in U$ belongs to the ball $B_{r_x}(x)$, so we clearly have the reverse containment as well. Thus $U = \bigcup_{x \in U} B_{r_x}(x)$. $\square$

Definition 2.2.42. A family $\mathcal{B} = \{B_i\}_{i \in I}$ of open sets in a metric space X is called a **base** if every open set in X can be written as a union of sets from $\mathcal{B}$.

The existence of a *countable base* is quite desirable, since it allows one to describe any open set using only countably many basic open sets.[5] It turns out that any separable metric space automatically has this property. The converse is also true, but it is left as an exercise.

Proposition 2.2.43. *If X is a separable metric space, then X has a countable base.*

[5] If a general topological space has a countable base, it is said to be **second countable**.

Proof. Since we already know the open balls form a base for X, we will try to refine this base and produce a countable family of open balls that do the job. In particular, suppose D is a countable dense subset of X, and consider the family

$$\mathcal{B} = \bigcup_{y \in D} \{B_q(y) : q \in \mathbf{Q}\}.$$

Then $\mathcal{B}$ is a union of countably many countable sets, so it is countable. It just remains to see that $\mathcal{B}$ is a base for X.

Let $U \subseteq X$ be an open set. Then for each $x \in U$ there is a $r_x > 0$ such that $B_{r_x}(x) \subseteq U$. Choose a rational number $0 < q_x < r_x$ for each x, and observe that $B_{q_x}(x) \subseteq B_{r_x}(x) \subseteq U$. Since D is dense in X, for each $x \in X$ we can find $y_x \in D$ such that $d(x, y_x) < q_x/2$. Then an application of the triangle inequality shows that

$$B_{q_x/2}(y_x) \subseteq B_{q_x}(x) \subseteq U$$

for each $x \in U$. Furthermore, $x \in B_{q_x/2}(y_x)$ for all $x \in U$. An argument like the one from the proof of Proposition 2.2.41 shows that

$$U = \bigcup_{x \in U} B_{q_x/2}(y_x).$$

Thus U is a union of elements from $\mathcal{B}$, so $\mathcal{B}$ is a base for X. □

Exercises for Section 2.2

Exercise 2.2.1. Let $E \subseteq X$. Prove E° is the largest open set contained in E. That is, if U is an open subset of X with $U \subseteq E$, then $U \subseteq E^\circ$.

Exercise 2.2.2 ([HS91], Exercise 4.3.3)**.** Let (X, d) be a metric space. Given $x \in X$ and $r > 0$, define the **closed ball** of radius r centered at x to be

$$B_r[x] = \{y \in X : d(x, y) \leq r\}.$$

(a) For all $x \in X$ and all $r > 0$, prove that $B_r[x]$ is a closed set.

(b) Show by example that $B_r[x] \neq \overline{B_r(x)}$ in general. That is, the closure of an open ball may not coincide with the corresponding closed ball.

Exercise 2.2.3 ([HS91], Exercise 4.3.7)**.** Let (X, d) be a metric space, and suppose $A, B \subseteq X$. If A is open and $A \cap \overline{B} \neq \emptyset$, show that $A \cap B \neq \emptyset$.

Exercise 2.2.4. Prove Proposition 2.2.24: Given $E \subseteq X$, a point $x \in E$ is an isolated point of E if and only if it is not a limit point of E.

Exercise 2.2.5 ([Rud76], Exercise 2.9 (e) & (f))**.** Let E be a subset of a metric space X.

(a) Do E and $\overline{E}$ always have the same interiors? Prove or give a counterexample.

(b) Do E and E° always have the same closures? Prove or give a counterexample.

Exercise 2.2.6 ([HS91], Exercise 4.3.5 modified)**.** Let A and B be subsets of a metric space X.

(a) Show that $A^\circ \cup B^\circ \subseteq (A \cup B)^\circ$ and $A^\circ \cap B^\circ = (A \cap B)^\circ$.

(b) Show that $\overline{A \cup B} = \overline{A} \cup \overline{B}$ and $\overline{A \cap B} \subseteq \overline{A} \cap \overline{B}$.

(c) Give examples to show that the inclusions in both (a) and (b) can be strict.

Exercise 2.2.7. Let $E \subseteq X$. Prove that the boundary ∂E is closed.

Exercise 2.2.8. Let $X = \ell^2$, and define $E \subseteq X$ by

$$E = \big\{(x_i)_{i=1}^\infty : x_i = 0 \text{ for all but finitely many indices } i\big\}.$$

(a) Prove that E is dense in ℓ^2.

(b) Now consider the subset of E consisting of sequences with rational terms:

$$E_{\mathbf{Q}} = \big\{(r_i)_{i=1}^\infty \in E : r_i \in \mathbf{Q} \text{ for all } i\big\}$$

Show that $E_{\mathbf{Q}}$ is countable.

(c) Show that $E_{\mathbf{Q}}$ is dense in ℓ^2. Conclude that ℓ^2 is a separable metric space.

Exercise 2.2.9 ([HS91], Exercise 4.3.11)**.** Show that if X is a metric space with a countable base, then X is separable.

2.3 Sequences

Let us first recall from real analysis that a sequence of real numbers is simply a function $f : \mathbf{N} \to \mathbf{R}$. Of course we usually write our sequences as lists, say $(x_n)_{n=1}^\infty$, where $x_n = f(n)$. One can then analyze convergence for sequences of real numbers only using the usual notion of distance on the real line. Certain phenomenon, such as monotonicity, will only make sense in $\mathbf{R}$, but convergence relies only on the presence of a metric.

Throughout this section, let (X, d) denote an arbitrary metric space. As in the reals, we will think of sequences as both functions $f : \mathbf{N} \to X$ and lists of points in X indexed by $\mathbf{N}$:

$$(x_n)_{n=1}^\infty, \text{ where } x_n = f(n).$$

The definition of convergence then looks like the familiar one from $\mathbf{R}$, albeit with the metric d substituted for the standard metric on $\mathbf{R}$.

Definition 2.3.1. A sequence $(x_n)_{n=1}^{\infty}$ in a metric space X is said to **converge** to a point $x \in X$ if given any $\varepsilon > 0$, there exists $N \in \mathbf{N}$ such that $n \geq N$ implies

$$d(x_n, x) < \varepsilon.$$

We will write $\lim_{n\to\infty} x_n = x$ or $x_n \to x$ to signify that $(x_n)_{n=1}^{\infty}$ converges to x.

We will begin with some elementary facts about sequences, all of which generalize familiar results from $\mathbf{R}$. Indeed, the proofs are almost identical to the ones from the real case.

Proposition 2.3.2. *Any sequence in a metric space X has at most one limit.*

Proof. Let $(x_n)_{n=1}^{\infty}$ be a sequence in X, and suppose $x_n \to x$ and $x_n \to y$ for some $x, y \in X$. Let $\varepsilon > 0$ be given. Then there exist $N_1, N_2 \in \mathbf{N}$ such that $n \geq N_1$ implies

$$d(x_n, x) < \tfrac{\varepsilon}{2}$$

and $n \geq N_2$ implies

$$d(x_n, y) < \tfrac{\varepsilon}{2}.$$

Now take $N = \max\{N_1, N_2\}$. Then for all $n \geq N$ we have

$$d(x, y) \leq d(x, x_n) + d(x_n, y) < \tfrac{\varepsilon}{2} + \tfrac{\varepsilon}{2} = \varepsilon.$$

Since this holds for all $\varepsilon > 0$, we must have $d(x, y) = 0$, so $x = y$. $\square$

We need a quick definition before we prove our next result. We say a sequence $(x_n)_{n=1}^{\infty}$ in a metric space is **bounded** if there exists a point $x \in X$ and a number $M > 0$ such that

$$d(x_n, x) \leq M$$

for all n. In other words, all the terms in the sequence lie inside some sufficiently large open ball.

Proposition 2.3.3. *If a sequence $(x_n)_{n=1}^{\infty}$ converges, then it is bounded.*

Proof. Assume $x_n \to x$ for some $x \in X$, and set $\varepsilon = 1$. Then we know there exists $N \in \mathbf{N}$ such that $d(x_n, x) < 1$ for all $n \geq N$. In other words, the open ball $B_1(x)$ contains all points that occur as terms in the sequence, except possibly $x_1, x_2, \ldots, x_{N-1}$. For $1 \leq i \leq N-1$, set $r_i = d(x_i, x)$, and put

$$M = \max\{r_1, r_2, \ldots, r_{N-1}, 1\}.$$

Then for all $n \in \mathbf{N}$ we have $d(x_n, x) \leq M$. Hence $(x_n)_{n=1}^{\infty}$ is bounded. $\square$

As one might guess, we can characterize the limit points of a set $E \subseteq X$ using sequences from E.

Proposition 2.3.4. *Let $E \subseteq X$. A point $x \in X$ is a limit point of E if and only if there is a sequence $(x_n)_{n=1}^{\infty}$ in E such that $x_n \to x$ and $x_n \neq x$ for all n.*

Proof. Suppose first that x is a limit point of E. Then for each $n \in \mathbf{N}$, there exists $x_n \in E$ such that $x_n \neq x$ and $d(x_n, x) < \frac{1}{n}$. It is then easy to see that the sequence $(x_n)_{n=1}^{\infty}$ converges to x by construction.

Conversely, suppose there is a sequence $(x_n)_{n=1}^{\infty}$ in E with the prescribed properties. Since $x_n \to x$, given any $\varepsilon > 0$, there exists $n \in \mathbf{N}$ such that $d(x_n, x) < \varepsilon$ and $x_n \neq x$. This shows that x is a limit point of E. $\square$

Even though an arbitrary sequence $(x_n)_{n=1}^{\infty}$ in a metric space may not converge, the set of points that occur as terms in the sequence,

$$\{x \in X : x = x_n \text{ for some } n \geq 1\}$$

may very well have plenty of limit points. These limit points arise precisely from the *subsequences* of $(x_n)_{n=1}^{\infty}$.

Definition 2.3.5. Let $(x_n)_{n=1}^{\infty}$ be a sequence, and suppose

$$n_1 < n_2 < n_3 < \cdots$$

is an increasing sequence of natural numbers. The sequence $(x_{n_k})_{k=1}^{\infty}$ is called a **subsequence** of $(x_n)_{n=1}^{\infty}$.

One should think of a subsequence as a way of selecting some of the terms from the original sequence, provided we adhere to the following two rules:

- We choose the terms in the same order in which they appeared in the original sequence.
- Each term of $(x_n)_{n=1}^{\infty}$ is selected at most once.

In this vein, we could equally well start with an increasing function $f : \mathbf{N} \to \mathbf{N}$ (which corresponds to selecting the indices $n_1, n_2, \ldots$) and define a subsequence by

$$(x_{f(k)})_{k=1}^{\infty}, \text{ where } f(k) = n_k.$$

The following proposition is a straightforward result about subsequences. Its proof is left as an exercise.

Proposition 2.3.6. *If a sequence $(x_n)_{n=1}^\infty$ converges to x, then every subsequence of $(x_n)_{n=1}^\infty$ also converges to x.*

As mentioned above, we can also characterize the limit points of a sequence in terms of subsequential limits.

Proposition 2.3.7. *Let $(x_n)_{n=1}^\infty$ be a sequence, and consider the set*

$$E = \{x \in X : x = x_n \text{ for some } n \geq 1\}.$$

A point $x \in X$ is a limit point of E if and only if there is a subsequence $(x_{n_k})_{k=1}^\infty$ converging to x with $x_{n_k} \neq x$ for all k.

Proof. Suppose first that x is a limit point of E. Then there exists $n_1 \in \mathbf{N}$ such that $d(x_{n_1}, x) < 1$ and $x_{n_1} \neq x$. Similarly, we can find $n_2 > n_1$ such that $d(x_{n_2}, x) < \frac{1}{n}$ and $x_{n_2} \neq x$. (Notice that we can force $n_2 > n_1$ because any open ball centered at x must contain infinitely many points of E.) Continuing inductively, we can construct a sequence $n_1 < n_2 < \cdots$ of natural numbers such that $d(x_{n_k}, x) < \frac{1}{k}$ and $x_{n_k} \neq x$ for all k. Thus $(x_{n_k})_{k=1}^\infty$ is the desired subsequence.

The converse is much less delicate. A subsequence $(x_{n_k})_{k=1}^\infty$ of $(x_n)_{n=1}^\infty$ with the specified properties is a sequence in E by definition, so $x = \lim_{k\to\infty} x_k$ is necessarily a limit point of E. $\square$

We will revisit subsequences a little later on. Many of the interesting phenomena for subsequences arise when studying compactness, so we will suspend this discussion until then.

Now we turn to an idea that also generalizes a well-known concept from sequences of real numbers. Recall that some sequences "desperately want" to converge, even though there may not be anything to which they can converge.

Definition 2.3.8. A sequence $(x_n)_{n=1}^\infty$ is said to be **Cauchy** if given any $\varepsilon > 0$, there exists $N \in \mathbf{N}$ such that

$$d(x_n, x_m) < \varepsilon$$

for all $n, m \geq N$.

In other words, a sequence is Cauchy if its terms can be made arbitrarily close to one another by going far enough out in the sequence. It is not hard to see that this condition will always be satisfied by convergent sequences.

Proposition 2.3.9. *If $(x_n)_{n=1}^\infty$ is convergent, then it is Cauchy.*

Proof. Let $\varepsilon > 0$, and choose $N \in \mathbf{N}$ such that $d(x_n, x) < \frac{\varepsilon}{2}$ for all $n \geq N$. Then for any $n, m \geq N$, we have

$$d(x_n, x_m) \leq d(x_n, x) + d(x, x_m) < \tfrac{\varepsilon}{2} + \tfrac{\varepsilon}{2} = \varepsilon$$

by the triangle inequality. Thus $(x_n)_{n=1}^\infty$ is a Cauchy sequence. $\square$

The converse to the last proposition is not true in a general metric space. For example, we know that there are Cauchy sequences in $\mathbf{Q}$ (equipped with its standard metric) that do not converge to rational numbers. (Take a sequence of successively better decimal approximations to $\sqrt{2}$, for example.) Thus metric spaces in which all Cauchy sequences converge are special.

Definition 2.3.10. A metric space (X, d) is said to be **complete** if every Cauchy sequence in X converges in X.

Example 2.3.11. It is well-known that $\mathbf{R}$ is complete with respect to its standard metric. On the other hand, $\mathbf{Q}$ is not complete (as we mentioned above).

Example 2.3.12. The Euclidean space $(\mathbf{R}^n, \|\cdot\|_2)$ is complete. This follows from an exercise in this section, together with the completeness of $\mathbf{R}$.

Suppose $(x_k)_{k=1}^\infty$ is a Cauchy sequence in $\mathbf{R}^n$. We will write $x_{k,i}$ to denote the ith coordinate of the vector x_k; that is,

$$x_k = (x_{k,1}, x_{k,2}, \ldots, x_{k,n}).$$

Let $\varepsilon > 0$ be given. Since $(x_k)_{k=1}^\infty$ is Cauchy, there exists $N \in \mathbf{N}$ such that for all $j, k \geq N$,

$$\|x_j - x_k\|_2 < \varepsilon.$$

But for $1 \leq i \leq n$ it is easy to check that

$$|x_{j,i} - x_{k,i}| \leq \|x_j - x_k\|_2 < \varepsilon.$$

Thus $(x_{k,i})_{k=1}^\infty$ is a Cauchy sequence in $\mathbf{R}$ for each i. Since $\mathbf{R}$ is complete, $x_{k,i} \to x_{0,i}$ for some $x_{0,i} \in \mathbf{R}$. Define

$$x_0 = (x_{0,1}, x_{0,2}, \ldots, x_{0,n}).$$

Then $x_k \to x_0$ by Exercise 2.3.3 since the sequence converges coordinatewise.

Example 2.3.13. It is possible to adapt the arguments of the previous example to show that ℓ^2 is complete. The details are left as an exercise. Note, however, that convergence of the individual coordinates in a sequence is not sufficient to guarantee convergence in ℓ^2.

Example 2.3.14. Any discrete metric space is complete. (Exercise 2.3.5)

There is one more example of a complete metric space that is worth mentioning right now. The proof incorporates some classical arguments from analysis, and we will revisit these ideas later.

Theorem 2.3.15. *For any closed interval* $[a, b]$*, the space* $C([a, b], \mathbf{R})$ *is complete with respect to the metric induced by the norm* $\|\cdot\|_\infty$.

Proof. Suppose $(f_n)_{n=1}^\infty$ is a Cauchy sequence in $C([a, b], \mathbf{R})$. Let $\varepsilon > 0$ be given, and choose $N \in \mathbf{N}$ such that $\|f_n - f_m\| < \varepsilon$ for all $n, m \geq N$. Notice that for each $x \in [a, b]$, we have

$$|f_n(x) - f_m(x)| \leq \sup_y |f_n(y) - f_m(y)| = \|f_n - f_m\|_\infty < \varepsilon$$

for all $n, m \geq N$. Thus $(f_n(x))_{n=1}^\infty$ is a Cauchy sequence in $\mathbf{R}$, so it converges. We can then define $f : [a, b] \to \mathbf{R}$ by

$$f(x) = \lim_{n\to\infty} f_n(x).$$

We claim that given $\varepsilon > 0$, there exists $N \in \mathbf{N}$ such that $n \geq N$ implies

$$|f_n(x) - f(x)| < \varepsilon$$

for all $x \in [a, b]$. That is, there is an N that works for all x simultaneously. Well, we know that $(f_n)_{n=1}^\infty$ is Cauchy, so we can find $N \in \mathbf{N}$ such that $n, m \geq N$ implies

$$|f_n(x) - f_m(x)| < \tfrac{\varepsilon}{2}$$

for all $x \in [a, b]$. We intend to hold n fixed and let $m \to \infty$. To be more precise, observe that for each $x \in [a, b]$, there is a natural number $N_x \geq N$ such that $|f_m(x) - f(x)| < \frac{\varepsilon}{2}$ for all $m \geq N_x$. Given $x \in [a, b]$, for all $n \geq N$ and any $m \geq N_x$ we have

$$|f_n(x) - f(x)| \leq |f_n(x) - f_m(x)| + |f_m(x) - f(x)| < \tfrac{\varepsilon}{2} + \tfrac{\varepsilon}{2} = \varepsilon.$$

Notice then that the N we have found is independent of $x \in [a, b]$, and our claim holds.

Now we aim to show that f is continuous on $[a,b]$. Let $x_0 \in [a,b]$ and $\varepsilon > 0$ be given. By the previous claim, we can find $N \in \mathbf{N}$ such that

$$|f_N(x) - f(x)| < \tfrac{\varepsilon}{3}$$

for all $x \in [a,b]$. Furthermore, since f_N is continuous there is a $\delta > 0$ such that $|x - x_0| < \delta$ implies

$$|f_N(x) - f_N(x_0)| < \tfrac{\varepsilon}{3}.$$

It then follows that whenever $|x - x_0| < \delta$,

$$\begin{aligned} |f(x) - f(x_0)| &\leq |f(x) - f_N(x)| + |f_N(x) - f_N(x_0)| + |f_N(x_0) - f(x_0)| \\ &< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \\ &= \varepsilon. \end{aligned}$$

Therefore, f is continuous at x_0. Since x_0 was arbitrary, f is continuous on $[a,b]$.

All that remains is to see that $f_n \to f$ in $C([a,b], \mathbf{R})$. We have really done all the required work already—by the earlier claim, given $\varepsilon > 0$ there exist $N \in \mathbf{N}$ such that $|f_n(x) - f(x)| < \varepsilon$ for all $x \in [a,b]$. But then we have

$$\|f_n - f\|_\infty = \sup_x |f_n(x) - f(x)| \leq \varepsilon$$

for all $n \geq N$, so $f_n \to f$. □

In many of the examples we have considered, the metric was induced by a norm on a vector space. A normed vector space that happens to be complete with respect to its natural metric has a special name. We will study these spaces in more detail later on.

Definition 2.3.16. A complete normed vector space is called a **Banach space**.

2.3.1 The Completion of a Metric Space

We will now investigate a phenomenon that we have already seen in the case of $\mathbf{R}$. Recall that $\mathbf{Q}$ is an incomplete metric space, but it can be "completed" to obtain the real numbers. That is, one can construct a larger number system in which every Cauchy sequence of rationals is convergent. Such a construction was carried out by Dedekind, who tried to remedy the failure of the least upper bound property for $\mathbf{Q}$ instead of working directly with Cauchy sequences. He defined real numbers to be *Dedekind cuts* of rational numbers, which represent the least upper bounds of subsets of $\mathbf{Q}$.

Dedekind's construction will not work in an arbitrary metric space, since it is unlikely that the least upper bound property will make sense there. We can still "complete" a general metric space, though we will need to use the Cauchy sequence approach in order to do so.

Definition 2.3.17. Let (X, d) be a metric space. We say a metric space (Y, d') is a **completion** of (X, d) if Y is complete and there is an isometry $\iota : X \to Y$ such that $\iota(X)$ is dense in Y.

The stipulation that X maps onto a *dense* subset Y is not just meant to mimic the fact that $\mathbf{Q}$ is dense in $\mathbf{R}$. It encodes the idea that Y should be the minimal complete metric space containing X as a subspace.

As we have already promised, every metric space has a completion. Furthermore, that completion is unique up to isometry.

Theorem 2.3.18. *Any metric space (X, d) has a unique completion $(\widetilde{X}, \tilde{d})$.*

The proof is quite technical, so we proceed via a series of lemmas. The general strategy is based on the proof of Theorem 7.4 in [KF70]. We begin by letting $\mathcal{X}$ denote the set of all Cauchy sequences in X. We then define a relation on $\mathcal{X}$ as follows: given two Cauchy sequences $(x_n), (y_n) \in \mathcal{X}$, we say $(x_n) \sim (y_n)$ if

$$\lim_{n\to\infty} d(x_n, y_n) = 0.$$

This relation is meant to identify Cauchy sequences if they are "trying to converge" to the same limit, even though that limit may not exist.

Lemma 2.3.19. *The relation $\sim$ defines an equivalence relation on $\mathcal{X}$.*

Proof. It should be obvious that $\sim$ is reflexive and symmetric. Transitivity will follow from a straightforward application of the triangle inequality. Suppose (x_n), (y_n), $(z_n) \in \mathcal{X}$ with $(x_n) \sim (y_n)$ and $(y_n) \sim (z_n)$. Then given $\varepsilon > 0$, there exists $N_1 \in \mathbf{N}$ such that $n \geq N_1$ implies

$$d(x_n, y_n) < \tfrac{\varepsilon}{2}.$$

Similarly, there is an $N_2 \in \mathbf{N}$ such that

$$d(y_n, z_n) < \tfrac{\varepsilon}{2}$$

for all $n \geq N_2$. If we let $N = \max\{N_1, N_2\}$, then $n \geq N$ implies

$$d(x_n, z_n) \leq d(x_n, y_n) + d(y_n, z_n) < \tfrac{\varepsilon}{2} + \tfrac{\varepsilon}{2} = \varepsilon,$$

so $(x_n) \sim (z_n)$. Thus $\sim$ is transitive. $\square$

Now we let $\widetilde{X} = \mathcal{X}/\sim$ denote the set of equivalence classes under $\sim$. This will provide the underlying set for the completion of X, but we need to equip it with a metric. To that end, define $\tilde{d} : \widetilde{X} \times \widetilde{X} \to \mathbf{R}$ by

$$\tilde{d}\big([(x_n)], [(y_n)]\big) = \lim_{n\to\infty} d(x_n, y_n).$$

Lemma 2.3.20. *The function $\tilde{d}$ defines a metric on $\widetilde{X}$.*

Proof. Before we can verify that $\tilde{d}$ is a metric, we need to check two things. First, we need to know that the limit on the right hand side even exists. That is, given two Cauchy sequences (x_n) and (y_n) in X, we claim that $\lim d(x_n, y_n)$ exists. This actually follows from the quadrilateral inequality (or really a modified version of it). Let $\varepsilon > 0$ be given, and choose $N \in \mathbf{N}$ such that $n, m \geq N$ implies

$$d(x_n, x_m) < \tfrac{\varepsilon}{2} \quad \text{and} \quad d(y_n, y_m) < \tfrac{\varepsilon}{2}.$$

Then the quadrilateral inequality implies

$$|d(x_n, y_n) - d(x_m, y_m)| \leq d(x_n, x_m) + d(y_n, y_m) < \varepsilon$$

for all $n, m \geq N$. Thus $(d(x_n, y_n))_{n=1}^{\infty}$ is a Cauchy sequence in $\mathbf{R}$, so it converges.

Now we need to check that $\tilde{d}$ is well-defined. That is, if we have $(x_n) \sim (x_n')$ and $(y_n) \sim (y_n')$ in $\mathcal{X}$, then we need

$$\tilde{d}\big([(x_n)], [(y_n)]\big) = \tilde{d}\big([(x_n')], [(y_n')]\big).$$

Well, we know that $d(x_n, x_n') \to 0$ and $d(y_n, y_n') \to 0$ as $n \to \infty$. Given $\varepsilon > 0$, there exists $N \in \mathbf{N}$ such that $d(x_n, x_n') < \frac{\varepsilon}{2}$ and $d(y_n, y_n') < \frac{\varepsilon}{2}$ for all $n \geq N$. Therefore,

$$\begin{aligned} d(x_n', y_n') &\leq d(x_n', x_n) + d(x_n, y_n) + d(y_n, y_n') \\ &< \tfrac{\varepsilon}{2} + d(x_n, y_n) + \tfrac{\varepsilon}{2} \\ &= d(x_n, y_n) + \varepsilon. \end{aligned}$$

A similar argument shows that

$$d(x_n, y_n) < d(x_n', y_n') + \varepsilon,$$

so

$$\big|d(x_n, y_n) - d(x_n', y_n')\big| < \varepsilon$$

for all $n \geq N$. Therefore $d(x_n, y_n)$ and $d(x_n', y_n')$ must converge to the same limit, so it follows that $\tilde{d}$ is well-defined.

The proof that $\tilde{d}$ satisfies the axioms for a metric is now fairly straightforward. Clearly $\tilde{d}(\tilde{x}, \tilde{y}) \geq 0$ for all $\tilde{x}, \tilde{y} \in \widetilde{X}$. If $\tilde{d}\big([(x_n)], [(y_n)]\big) = 0$, then

$$\lim_{n \to \infty} d(x_n, y_n) = 0,$$

so $(x_n) \sim (y_n)$ in $\mathcal{X}$. Thus $[(x_n)] = [(y_n)]$, and $\tilde{d}$ is positive definite. It is easy to see that $\tilde{d}$ is symmetric since d is. For the triangle inequality, observe that if $(x_n), (y_n), (z_n) \in \mathcal{X}$, then

$$\tilde{d}\big([(x_n)], [(y_n)]\big) = \lim_{n \to \infty} d(x_n, y_n)$$

$$\begin{aligned}
&\le \lim_{n\to\infty} \big[d(x_n, z_n) + d(z_n, y_n)\big] \\
&= \lim_{n\to\infty} d(x_n, z_n) + \lim_{n\to\infty} d(z_n, y_n) \\
&= \tilde{d}\big([(x_n)], [(z_n)]\big) + \tilde{d}\big([(z_n)], [(y_n)]\big).
\end{aligned}$$

Therefore, $\tilde{d}$ is a metric. □

Next, we will show that $\tilde{X}$ contains an isometric copy of X. Given $x \in X$, we let $(x) \in \mathfrak{X}$ denote the associated constant sequence. We can then define a map $\iota : X \to \widetilde{X}$ by

$$\iota(x) = [(x)].$$

Lemma 2.3.21. *The map $\iota : X \to \widetilde{X}$ is isometric. Moreover, $\iota(X)$ is dense in $\widetilde{X}$.*

Proof. Let $x, y \in X$. Observe that

$$\tilde{d}\big(\iota(x), \iota(y)\big) = \tilde{d}\big([(x)], [(y)]\big) = \lim_{n\to\infty} d(x, y) = d(x, y),$$

so ι is an isometry. Note that this implies ι is automatically injective.

Now we show $\iota(X)$ is dense in $\widetilde{X}$. Let $[(x_n)] \in \widetilde{X}$ and let $\varepsilon > 0$. Since (x_n) is Cauchy, choose $N \in \mathbf{N}$ such that $d(x_n, x_m) < \frac{\varepsilon}{2}$ for all $n, m \ge N$. Set $y = x_N$ and consider $\iota(y) = [(y)]$ in $\widetilde{X}$. Then we have

$$\tilde{d}\big([(y)], [(x_n)]\big) = \lim_{x\to\infty} d(y, x_n) \le \tfrac{\varepsilon}{2} < \varepsilon.$$

It follows that $B_\varepsilon([(x_n))$ contains a point of $\iota(X)$ for all $[(x_n)] \in \widetilde{X}$ and all $\varepsilon > 0$, so $\iota(X)$ is dense in $\widetilde{X}$. □

Lemma 2.3.22. *The metric space $(\widetilde{X}, \tilde{d})$ is complete.*

Proof. Let $(\tilde{x}_n)$ be a Cauchy sequence in $\widetilde{X}$. Since $\iota(X)$ is dense in $\widetilde{X}$, for each $n \in \mathbf{N}$ we can find $y_n \in X$ such that

$$\tilde{d}(\tilde{x}_n, \iota(y_n)) < \tfrac{1}{n}.$$

We claim that (y_n) is a Cauchy sequence in X. Given $\varepsilon > 0$, find an $N \in \mathbf{N}$ such that $N > 3/\varepsilon$ and

$$\tilde{d}(\tilde{x}_n, \tilde{x}_m) < \tfrac{\varepsilon}{3}$$

for all $n, m \ge N$. Then whenever $n, m \ge N$ we have

$$\begin{aligned}
d(y_n, y_m) &= \tilde{d}\big(\iota(y_n), \iota(y_m)\big) \\
&\le \tilde{d}\big(\iota(y_n), \tilde{x}_n\big) + \tilde{d}\big(\tilde{x}_n, \tilde{x}_m\big) + \tilde{d}\big(\tilde{x}_m, \iota(y_m)\big) \\
&< \tfrac{1}{n} + \tfrac{\varepsilon}{3} + \tfrac{1}{n} \\
&< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3}
\end{aligned}$$

$$= \varepsilon.$$

Now define $\tilde{x} = [(y_n)]$. Then for all n we have

$$\begin{aligned}\tilde{d}(\tilde{x}_n, \tilde{x}) &\leq \tilde{d}\big(\tilde{x}, \iota(y_n)\big) + \tilde{d}\big(\iota(y_n), \tilde{x}_n\big) \\ &< \lim_{k\to\infty} d(y_k, y_n) + \tfrac{1}{n} \\ &= \tfrac{1}{n},\end{aligned}$$

so $\tilde{x}_n \to \tilde{x}$ as $n \to \infty$. Therefore, $(\widetilde{X}, \tilde{d})$ is complete. □

To make the notation a little easier to handle, we will identify X and its image $\iota(X)$ in $\widetilde{X}$. That is, we will suppress the ι and think of X as a subset of $\widetilde{X}$.

All that remains is to see that the completion is unique. When we say it is unique, we mean any other completion of (X, d) is "isomorphic" to $(\widetilde{X}, \tilde{d})$ in the appropriate sense for metric spaces.

Lemma 2.3.23. *Suppose (Y, d') is another completion of (X, d). Then there is a surjective isometry $\varphi : Y \to \widetilde{X}$ that restricts to the identity map on X.*

Proof. Let $y \in Y$. Since Y is a completion of X, we know X is dense in Y, so there is a sequence (x_n) in X that converges to y. This sequence is necessarily Cauchy in X (and therefore in Y as well). Viewing X as a subset of $\widetilde{X}$, we must have that (x_n) is Cauchy in $\widetilde{X}$. Since $\widetilde{X}$ is complete, $x_n \to x$ for some $x \in \widetilde{X}$. Define $\varphi : Y \to \widetilde{X}$ by

$$\varphi(y) = x = \lim_{n\to\infty} x_n,$$

where the limit is taken in $\widetilde{X}$. We claim first that this definition does not depend on the approximating sequence (x_n). To see this, suppose (y_n) is another sequence in X such that $y_n \to y$ in Y. Then we must have $d(x_n, y_n) = d'(x_n, y_n) < \frac{\varepsilon}{2}$ and $\tilde{d}(x_n, x) < \frac{\varepsilon}{2}$ for sufficiently large n, so

$$\tilde{d}(y_n, x) \leq \tilde{d}(y_n, x_n) + \tilde{d}(x_n, x) < \tfrac{\varepsilon}{2} + \tfrac{\varepsilon}{2} = \varepsilon$$

eventually. Thus $y_n \to x = \varphi(y)$ in $\tilde{X}$, so φ is well-defined.

Now we claim that $\varphi : Y \to \tilde{X}$ is isometric. The previous computations show that this is certainly the case when we restrict φ to X, but we need to see that φ respects the metrics d' and $\tilde{d}$. Let $x, y \in Y$, and suppose (x_n) and (y_n) are sequences in Y with $x_n \to x$ and $y_n \to y$. Then by construction $x_n \to \varphi(x)$ and $y_n \to \varphi(y)$ in $\tilde{X}$, so

$$\tilde{d}\big(\varphi(x), \varphi(y)\big) = \lim_{n\to\infty} d(x_n, y_n) = d'(x, y).$$

Thus φ is an isometry, hence it is also injective. Surjectivity follows from the density of X in $\tilde{X}$: given $x \in \tilde{X}$, find a sequence (x_n) in X converging to x in $\widetilde{X}$. Then (x_n), viewed as a sequence in Y, converges to a point $y \in Y$. We then have $\varphi(y) = x$, so φ is surjective. Thus (Y, d') and $(\widetilde{X}, \tilde{d})$ are isometric. Moreover, φ clearly restricts to the identity on X. □

The uniqueness of the completion of a metric space lets us answer one last question: what if (X, d) was complete to begin with?

Proposition 2.3.24. *Let (X, d) be a complete metric space. Then any completion of (X, d) is isometric to (X, d).*

Proof. Let (Y, d') be a completion of (X, d). Since (X, d) is complete and X is certainly dense in itself, there is a bijective isometry $\varphi : Y \to X$. That is, (Y, d') is isometric to (X, d). $\square$

In other words, if we construct the completion of a complete metric space (X, d) via Cauchy sequences, we do not obtain anything new.

Exercises for Section 2.3

Exercise 2.3.1 ([HS91], Exercise 4.6.3)**.** Let $(x_n)_{n=1}^{\infty}$ be a sequence in a metric space (X, d). Show that if $x_n \to x$ for some $x \in X$ and $(x_{n_k})_{k=1}^{\infty}$ is a subsequence of $(x_n)_{n=1}^{\infty}$, then $x_{n_k} \to x$ as well.

Exercise 2.3.2. Let $(x_n)_{n=1}^{\infty}$ be a Cauchy sequence in a metric space (X, d), and suppose $(x_{n_k})_{k=1}^{\infty}$ is a subsequence that converges to a point $x \in X$. Prove that $x_n \to x$.

Exercise 2.3.3. Let $(x_k)_{k=1}^{\infty}$ be a sequence of points in $\mathbf{R}^n$, and write

$$x_k = (x_{k,1}, x_{k,2}, \ldots, x_{k,n})$$

for each k. (That is, $x_{k,i}$ denotes the ith entry of the vector x_k.) Show that $(x_k)_{k=1}^{\infty}$ converges with respect to the metric coming from the Euclidean norm $\|\cdot\|_2$ on $\mathbf{R}^n$ if and only if the sequence $(x_{k,i})_{k=1}^{\infty}$ converges in $\mathbf{R}$ for $1 \leq i \leq n$.

Exercise 2.3.4. Let $(x_n)_{n=1}^{\infty}$ be a sequence of points in ℓ^2. As in Problem 2, we write $x_{n,i}$ for the ith entry of the vector x_n:

$$x_n = (x_{n,1}, x_{n,2}, \ldots).$$

Prove that if $(x_n)_{n=1}^{\infty}$ converges in ℓ^2, then the sequence $(x_{n,i})_{n=1}^{\infty}$ converges for all i. Show by example that the converse fails—that is, find a sequence $(x_n)_{n=1}^{\infty}$ such that $(x_{n,i})_{n=1}^{\infty}$ converges for all i, but $(x_n)_{n=1}^{\infty}$ does not converge in ℓ^2.

Exercise 2.3.5 ([HS91], Exercise 4.6.7)**.** Let (X, d) be a metric space, where d is the discrete metric on X. Prove that (X, d) is complete. [**Hint:** What do the Cauchy sequences look like?]

Exercise 2.3.6. Prove that ℓ^2 is complete.

The following problem is not for the faint of heart, since it pulls in some ideas from abstract algebra. It is also quite technical.

Exercise 2.3.7. Earlier we described how to construct the completion of an arbitrary metric space via equivalence classes of Cauchy sequences. We also discussed how this approach is not sufficient for constructing $\mathbf{R}$ from $\mathbf{Q}$. In particular, one would need to check that the Cauchy completion of $\mathbf{Q}$ has the structure of a totally ordered field, as $\mathbf{R}$ is supposed to have. This exercise is meant to illustrate a way of reconciling the field structure of $\mathbf{R}$ with the Cauchy sequence approach to building completions.

(a) Let R denote the set of all Cauchy sequences of rational numbers. Prove that R is a commutative ring with identity with respect to the "pointwise" operations
$$(x_n) + (y_n) = (x_n + y_n)$$
and
$$(x_n) \cdot (y_n) = (x_n \cdot y_n).$$

(b) Define $I \subseteq R$ by
$$I = \big\{(x_n) \in R : x_n \to 0\big\}.$$
Prove that I is an ideal in R.

(c) Show that if J is an ideal in R that properly contains I, then J contains a unit. Use this to show that I is a maximal ideal, and conclude that R/I is a field.

(d) Define a function $\varphi : \mathbf{Q} \to R/I$ by
$$\varphi(x) = (x) + I,$$
where (x) denotes the constant sequence whose terms are all equal to x. Show that φ is an injective ring homomorphism, so $\mathbf{Q}$ is a subfield of R/I.

So far we have constructed a field extension R/I of $\mathbf{Q}$, but we don't yet know that it is a complete ordered field. To establish completeness, we can cheat a little bit. Recall that we defined the completion of a metric space X by placing an equivalence relation on the set of all Cauchy sequences in X: we declared $(x_n) \sim (y_n)$ if and only if $d(x_n, y_n) \to 0$ as $n \to \infty$.

(e) Given two Cauchy sequences $x = (x_n)$ and $y = (y_n)$ in R, show that $x \sim y$ if and only if $x - y \in I$. Use this fact to conclude that $R/I = \widetilde{\mathbf{Q}}$ (as sets), so R/I is a complete metric space containing $\mathbf{Q}$ as a dense subset.

(f) All that is left is to show that R/I is an ordered field. Given $(x_n), (y_n) \in R$, define $(x_n) + I \preceq (y_n) + I$ in R/I if either $(x_n - y_n) \in I$ or $x_n \leq y_n$ eventually. Prove that $\preceq$ defines a total order on R/I.

(g) Prove that for all $x, y, z \in R/I$, if $y \preceq z$ then $x + y \preceq x + z$. Also, show that if $x, y \succ 0$, then $xy \succ 0$. Conclude that R/I is an ordered field.

By the results of this exercise, R/I is a complete, totally ordered field containing $\mathbf{Q}$ as a subfield, so it is isomorphic to $\mathbf{R}$ by Theorem 1.19 of [Rud76].

2.4 The Baire Category Theorem

We are now going to investigate a classical theorem from analysis, which was proven by René-Louis Baire in the late 19th century. It pertains to Baire's description of the relative "size" of a subset of a metric space.

Of course we can describe the size of a set via its cardinality. However, Baire discovered a way of characterizing certain subsets of metric spaces as "small" or "large" in a way that incorporates the topology of the underlying space. (We will see a similar idea when we investigate Lebesgue measure later on.) Baire's theorem then says that complete metric spaces are necessarily large. If we think just in terms of cardinality, the familiar case of $\mathbf{R}$ is an informative example—when we complete $\mathbf{Q}$ to build $\mathbf{R}$, we obtain a much larger (i.e., uncountable) set. We will see shortly that there is a topological distinction between the "sizes" of $\mathbf{Q}$ and $\mathbf{R}$ as well.

We begin with Baire's notion of what it means for a set to be small. Since we are thinking in terms of topology, we may want to incorporate open sets or open balls into our definitions. Indeed, one should think of open sets as being somewhat large, or at least not small. Loosely speaking, we will deem a set E to be small if it has empty interior. More precisely, E should not be dense in any open set.

Definition 2.4.1. Let (X, d) be a metric space. A set $E \subseteq X$ is said to be **nowhere dense** if $(\overline{E})^\circ = \emptyset$.

Nowhere dense sets can be though of as being quite sparse in their parent metric space, as the next few examples show.

Example 2.4.2. In $\mathbf{R}$, any singleton is nowhere dense. Since singletons are closed, for any $x \in \mathbf{R}$ we have $\overline{\{x\}} = \{x\}$, but $\{x\}^\circ = \emptyset$.

Example 2.4.3. On the other hand, if (X, d) is a discrete metric space, then singletons are not nowhere dense. We still have $\overline{\{x\}} = \{x\}$ for all $x \in X$. However, recall that every subset of X is open, so $\{x\}^\circ = \{x\}$.

We can reconcile the last two examples with the following observation. The proof is an exercise.

Proposition 2.4.4. *Let (X, d) be a metric space, and let $x \in X$. The singleton $\{x\}$ is nowhere dense if and only if x is not an isolated point of X.*

Example 2.4.5. In $\mathbf{R}^2$, the set

$$E = \{(x, 0) : x \in \mathbf{R}\}$$

is nowhere dense. Notice that E is a line in $\mathbf{R}^2$ (namely the x-axis), so it is closed. (It is straightforward to show the complement is open.) Thus $\overline{E} = E$, so it suffices to show that $E^\circ = \emptyset$. Fix $(x, 0) \in E$, and let $r > 0$. Then the vector $(x, \frac{r}{2})$ does not belong to E, but it does lie in the open ball $B_r((x, 0))$. Thus $(x, 0)$ is not an interior point of E, and it follows that E is nowhere dense.

The argument from the previous example can be adapted to show that any "one-dimensional subset" of $\mathbf{R}^2$ (say, a smooth curve in $\mathbf{R}^2$) is nowhere dense in $\mathbf{R}^2$. It can also be modified to prove that a proper subspace of $\mathbf{R}^n$ is nowhere dense in $\mathbf{R}^n$.

Example 2.4.6. The set $\mathbf{Q}$ is not nowhere dense in $\mathbf{R}$. Note that $\overline{\mathbf{Q}} = \mathbf{R}$, and $\mathbf{R}^\circ = \mathbf{R}$. Nevertheless, we will deem $\mathbf{Q}$ to be a "small" set (in Baire's sense) soon.

It is perhaps not hard to see that if $E \subseteq X$ is nowhere dense, then E^c is dense in X. (We will prove this shortly.) However, the converse is not true, as the last example shows. If we consider $\overline{E}$ instead, we can formulate an equivalent condition to E being nowhere dense.

Proposition 2.4.7. *A set $E \subseteq X$ is nowhere dense in X if and only if $(\overline{E})^c$ is dense in X.*

Proof. Suppose first that E is nowhere dense, and let $x \in X$. It suffices to assume $x \in \overline{E}$. Then for any $r > 0$, the ball $B_r(x)$ is not contained in $\overline{E}$, so it must contain a point of $(\overline{E})^c$. Thus x is a limit point of $(\overline{E})^c$, so $(\overline{E})^c$ is dense in X.

Now assume conversely that $(\overline{E})^c$ is dense in X. Then given any $x \in X$ and $r > 0$, the open ball $B_r(x)$ must contain a point of $(\overline{E})^c$. This holds in particular if $x \in \overline{E}$. Thus for any $x \in \overline{E}$ and any $r > 0$, $B_r(x) \not\subseteq \overline{E}$. It follows that $(\overline{E})^\circ = \emptyset$, so E is nowhere dense. □

We can now state and prove Baire's theorem. There are two variants of the theorem, which are related via the last proposition. We will prove the first version, and then show the second version follows from the first. The proof is adapted from that of [Fol99, Theorem 5.9].

Theorem 2.4.8 (Baire Category Theorem)**.** *Let (X, d) be a complete metric space.*

1. *If $\{U_n\}_{n=1}^{\infty}$ is a countable family of open sets $U_n \subseteq X$ with each U_n dense in X, then the intersection $\bigcap_{n=1}^{\infty} U_n$ is dense in X.*
2. *We cannot express X as a countable union $\bigcup_{n=1}^{\infty} E_n$ where each E_n is nowhere dense in X.*

Proof. We will begin by proving the first statement. Suppose $\{U_n\}_{n=1}^{\infty}$ is a countable family of open, dense sets in X. Let $x_0 \in X$ and $r_0 > 0$, and set $W = B_{r_0}(x_0)$. It will suffice to show that W contains a point of $\bigcap_{n=1}^{\infty} U_n$.

Since U_1 is open and dense, $U_1 \cap W$ is nonempty and open. Thus there exists a point $x_1 \in U_1 \cap W$ and a number $0 < r_1 < 1$ such that

$$B_{r_1}(x_1) \subseteq U_1 \cap W.$$

By shrinking the open ball if necessary, we can actually arrange things so that[6]

$$B_{r_1}(x_1) \subseteq \overline{B_{r_1}(x_1)} \subseteq U_1 \cap W.$$

Similarly, since U_2 is open and dense in X, we know $U_2 \cap B_{r_1}(x_1)$ is nonempty and open. Thus there exists $x_2 \in U_2 \cap B_{r_1}(x_1)$ and $0 < r_2 < \frac{1}{2}$ such that

$$B_{r_2}(x_2) \subseteq \overline{B_{r_2}(x_2)} \subseteq U_2 \cap B_{r_1}(x_1).$$

We now continue inductively: for each successive $n \in \mathbf{N}$, we find an $x_n \in X$ and $0 < r_n < \frac{1}{2^{n-1}}$ such that

$$B_{r_n}(x_n) \subseteq \overline{B_{r_n}(x_n)} \subseteq U_n \cap B_{r_{n-1}}(x_{n-1}).$$

Now let $\varepsilon > 0$ be given, and choose $N \in \mathbf{N}$ such that $\frac{1}{2^{N-1}} < \frac{\varepsilon}{2}$. Then for all $n, m \geq N$, we have

$$x_n, x_m \in B_{r_N}(x_N),$$

so

$$d(x_n, x_m) < 2r_N < \varepsilon.$$

This shows that the sequence $(x_n)_{n=1}^{\infty}$ is Cauchy. Since X is complete, $x_n \to x$ for some $x \in X$. We claim that $x \in \bigcap_{n=1}^{\infty} U_n$ and x belongs to the original open ball W. Fix $n \in \mathbf{N}$, and observe that for all $k \geq n$ we have

$$x_k \in B_{r_k}(x_k) \subseteq B_{r_n}(x_n) \subseteq \overline{B_{r_n}(x_n)}.$$

[6] Observe that if $0 < r < r'$, then we have $B_r(x) \subseteq B_r[x] \subseteq B_{r'}(x)$. Since $B_r[x]$ is a closed set containing $B_r(x)$, we must also have $\overline{B_r(x)} \subseteq B_r[x] \subseteq B_{r'}(x)$.

It follows that $x \in \overline{B_{r_n}(x_n)}$ since this set is closed. Since this holds for any $n \in \mathbf{N}$, we have

$$x \in \overline{B_{r_n}(x_n)} \subseteq U_n \cap B_{r_{n-1}}(x_{n-1}) \subseteq U_n \cap W$$

for all $n \in \mathbf{N}$. It follows that

$$x \in \bigcap_{n=1}^{\infty} (U_n \cap W) = \left(\bigcap_{n=1}^{\infty} U_n \right) \cap W,$$

so W contains a point of the intersection $\bigcap_{n=1}^{\infty} U_n$. Thus $\bigcap_{n=1}^{\infty} U_n$ is dense in X.

Now we show that the second formulation of the Baire Category Theorem follows from the one we have just proven. Suppose $\{E_n\}_{n=1}^{\infty}$ is a countable family of nowhere dense sets in X. Then each set $(\overline{E}_n)^c$ is open and dense by Proposition 2.4.7, so

$$\bigcap_{n=1}^{\infty} (\overline{E}_n)^c$$

is dense. In particular, the intersection is nonempty. Then

$$\bigcup_{n=1}^{\infty} E_n \subseteq \bigcup_{n=1}^{\infty} \overline{E}_n = \left(\bigcap_{n=1}^{\infty} \overline{E}_n^c \right)^c \neq X.$$

Thus X is not a countable union of nowhere dense sets. □

Now we will investigate some applications of the Baire Category Theorem. The first one is straightforward, and it generalizes a result about perfect[7] subsets of $\mathbf{R}$.

Proposition 2.4.9. *If (X, d) is a complete metric space with no isolated points, then X is uncountable.*

Proof. Suppose to the contrary that X is countable, so we can write $X = \{x_n\}_{n=1}^{\infty}$. Since X has no isolated points, each singleton $\{x_n\}$ is nowhere dense by Proposition 2.4.4. Thus

$$X = \bigcup_{n=1}^{\infty} \{x_n\}$$

is a countable union of nowhere dense sets. This is impossible if X is complete, since it would violate the Baire Category Theorem. Thus X must be uncountable. □

The next application requires us to introduce some terminology first. One might wonder why the word *category* appears in the name of Baire's theorem. This has to do with Baire's classification of sets by their sizes. Baire defined a set $E \subseteq X$ to be

[7] Recall that a set $E \subseteq \mathbf{R}$ is **perfect** if every point of E is a limit point of E.

- of **first category** in X (or **meagre**) if E can be expressed as a countable union of nowhere dense sets, or
- of **second category** in X if it is not a set of first category.

Observe that the Baire Category Theorem says that a complete metric space is always a second category subset of itself.

Remark 2.4.10. It is worth noting that Baire's notion of category depends on the ambient metric space. For example, $\mathbf{Z}$ is first category in $\mathbf{R}$, since it is a countable set and singletons are nowhere dense in $\mathbf{R}$. However, if we consider $\mathbf{Z}$ as a metric space on its own (with the metric inherited from $\mathbf{R}$), then every point of $\mathbf{Z}$ is isolated. It follows that the only nowhere dense subset of $\mathbf{Z}$ is $\emptyset$, so $\mathbf{Z}$ is a second category set in itself.

Example 2.4.11. By the reasoning applied to the integers above, $\mathbf{Q}$ is a first category set in $\mathbf{R}$. Thus it is possible for a first category set (i.e., a "small" set in Baire's sense) to be dense in its parent metric space.

On the other hand, we claim that the irrationals $\mathbf{I} = \mathbf{R}\backslash\mathbf{Q}$ are not first category. Of course $\mathbf{I}$ is uncountable, so we can't simply take the nowhere dense sets to be the singletons. But how do we know that we are unable to write

$$\mathbf{I} = \bigcup_{n=1}^{\infty} E_n$$

for some family $\{E_n\}_{n=1}^{\infty}$ of nowhere dense sets? Well, if we could do so, then we could write $\mathbf{R} = \mathbf{Q} \cup \mathbf{I}$ as a countable union of nowhere dense sets, in violation of the Baire Category Theorem.

We have another interesting implication for $\mathbf{Q}$ thanks to the Baire Category Theorem. Recall that in a metric space, a countable union of closed sets need not be closed. The rationals are a prime example of this phenomenon—we can write

$$\mathbf{Q} = \bigcup_{x\in\mathbf{Q}} \{x\},$$

where the singletons $\{x\}$ are all closed subsets of $\mathbf{R}$, but $\mathbf{Q}$ is not closed. However, countable unions of closed sets have a special name: they are called F_σ-sets. The dual notion is a G_δ-set, which is a countable intersection of open sets.[8] It is easy to check (using De Morgan's law) that E is an F_σ if and only if E^c is a G_δ. It is certainly possible for sets to be both F_σ and G_δ. Can this happen for the rationals?

[8] You may have noticed that we often use F to denote a closed set. This is because the French word for "closed" is *fermé*. In French, *somme* means "sum" or "union", leading to the notation F_σ for countable unions of closed sets. Similarly, the German word for an open neighborhood is *Gebeit*, while *Durchschnitt* means "intersection", hence the G_δ notation.

Proposition 2.4.12. *The rationals do not form a G_δ-set.*

Proof. Suppose $\mathbf{Q}$ is a G_δ. Then the irrationals form an F_σ, so we can write

$$\mathbf{I} = \bigcup_{n=1}^{\infty} F_n$$

for some family $\{F_n\}_{n=1}^{\infty}$ of closed sets in $\mathbf{R}$. Since $\mathbf{I}$ has empty interior, so must F_n for all n. That is, each F_n is nowhere dense. However, we have already observed that $\mathbf{I}$ cannot be written as a countable union of nowhere dense sets, since this would violate the Baire Category Theorem. Therefore, $\mathbf{Q}$ is not a G_δ. $\square$

The last proposition might seem like a useless observation at first glance, but it does have an interesting consequence. Given a function $f : \mathbf{R} \to \mathbf{R}$, the set of points at which f is continuous is necessarily a G_δ set. Therefore, there is no function $f : \mathbf{R} \to \mathbf{R}$ that is continuous at every rational number and discontinuous at every irrational.

The Baire Category Theorem has many other interesting implications, most of which are far beyond the scope of this course. Some interesting ones are the following:

- In $C([a, b], \mathbf{R})$, the set of differentiable functions is first category. In fact, the set consisting of functions that have a derivative at some point of $[a, b]$ is meagre. Thus most functions are nowhere differentiable.
- Any basis for an infinite-dimensional Banach space is necessarily uncountable.
- There are several major results from functional analysis that rely on the Baire Category Theorem, namely the Open Mapping Theorem, the Closed Graph Theorem, and the Principle of Uniform Boundedness (also called the Banach-Steinhaus Theorem).

Exercises for Section 2.4

Exercise 2.4.1. Let (X, d) be a metric space. Given a point $x \in X$, show that the singleton $\{x\}$ is nowhere dense if and only if x is not an isolated point of X.

Exercise 2.4.2. Recall that Theorem 2.4.8 includes two different versions of the Baire Category Theorem—the first version is proved directly from the completeness of the metric space X, and the second version then follows from the first. However, you may have noticed that we did not address the converse. Indeed, the two versions that we stated are not equivalent—the second statement is strictly weaker than the first, and we could have actually used Theorem 2.4.8(1) to derive a much stronger result:

> *Let X be a complete metric space. If $\{E_n\}_{n=1}^{\infty}$ is a collection of sets that are each nowhere dense in X, then the union $\bigcup_{n=1}^{\infty} E_n$ has empty interior.*

Prove that this statement is equivalent to Theorem 2.4.8(1), and therefore can be considered an alternative formulation of the Baire Category Theorem.

Exercise 2.4.3. Recall that ℓ^∞ denotes the set of all bounded sequences of real numbers equipped with the norm

$$\|(x_i)_{i=1}^{\infty}\|_\infty = \sup_i |x_i|.$$

Though we have not proven it, ℓ^∞ is complete with respect to the metric induced from this norm.

(a) For each $n \in \mathbf{N}$, define a subset $E_n \subseteq \ell^\infty$ by

$$E_n = \bigl\{(x_i)_{i=1}^{\infty} \in \ell^\infty : x_i = 0 \text{ for all } i > n\bigr\}.$$

Prove that E_n is nowhere dense in ℓ^∞ for all $n \in \mathbf{N}$.

(b) For $n \geq 1$, define $e_n \in \ell^\infty$ by

$$e_n = (0, 0, \dots, 0, \underbrace{1}_{n^{\text{th}}}, 0, 0, \dots).$$

That is, the nth entry of e_n is 1, and all other entries are 0. Notice that $e_n \in E_n$ for each n. Prove that the set $\{e_1, e_2, \dots\}$ is linearly independent, and that

$$\operatorname{span}\{e_1, e_2, \dots\} \subseteq \bigcup_{n=1}^{\infty} E_n.$$

(c) Prove that $\bigcup_{n=1}^{\infty} E_n \neq \ell^\infty$. Conclude that the set $\{e_1, e_2, \dots\}$ does not form a basis for the vector space ℓ^∞.

2.5 Compactness

In introductory real analysis, one often studies the concept of a compact set in $\mathbf{R}$ (or $\mathbf{R}^n$). Such sets are "small", in the sense that they behave like finite sets in many ways. When working with compact sets, one can often use proof techniques that treat the set as if it were finite, in a certain sense. This phenomenon leads to all kinds of nice results in calculus and real analysis, with the Extreme Value Theorem perhaps being foremost among them.

In $\mathbf{R}^n$, there are several different characterizations of compactness, one of which is strikingly simple—a set is compact if and only if it is closed and bounded. Unfortunately, the equivalence of this condition with other, more complicated ones no longer holds in an arbitrary metric space. We are thus forced to take a more abstract description of compactness as our definition. As usual, we will let (X, d) denote an arbitrary metric space throughout this section.

Definition 2.5.1. Let $E \subseteq X$. An **open cover** for E is a family $\mathcal{U} = \{U_i\}_{i \in I}$ of open sets in X that covers E, in the sense that

$$E \subseteq \bigcup_{i \in I} U_i.$$

A **subcover** of an open cover $\mathcal{U}$ is a subset of $\mathcal{U}$ that still covers E.

Example 2.5.2. Let $X = \mathbf{R}$, and consider the open interval $E = (-1, 1)$. The collection of open sets $\mathcal{U} = \{U_n\}_{n=1}^{\infty}$ defined by

$$U_n = \left(-1 + \tfrac{1}{n}, 1 - \tfrac{1}{n}\right)$$

clearly covers E. Notice that any subcover of $\mathcal{U}$ is necessarily infinite—given any $n_0 \in \mathbf{N}$, choose x such that $1 - \frac{1}{n_0} < x < 1$. Then $x \notin U_n$ for any $n \leq n_0$, so the finite family $\{U_n\}_{n=1}^{n_0}$ cannot possibly cover $(-1, 1)$.

Using our "closed and bounded" notion of compactness in $\mathbf{R}$, we can easily see that the set E from the previous example is not compact. In general metric spaces, it is precisely the lack of a finite subcover that we want to avoid.

Definition 2.5.3. A set $K \subseteq X$ is **compact** if every open cover of K has a finite subcover.

The open cover definition of compactness probably seems quite odd to most people. However, the ability to "extract a finite subcover" from any open cover is quite handy when trying to prove results about compact sets. On the other hand, it is a criterion that is not so easy to verify when deciding whether a set is compact or not. As we have already described, we have other characterizations in $\mathbf{R}^n$ that are much easier to check. In particular, the following are equivalent for a set $K \subseteq \mathbf{R}^n$:

1. K is compact.
2. K is closed and bounded. (Heine-Borel)
3. Every sequence in K has a convergent subsequence. (Bolzano-Weierstrass)

Unfortunately, we no longer have the equivalence of conditions 1 and 2 in general metric spaces. One direction does still hold, however.

Definition 2.5.4. A set $E \subseteq X$ is **bounded** if there exist a point $x \in X$ and a number $M > 0$ such that $E \subseteq B_M[x]$.

In other words, E is bounded if it can be placed inside a sufficiently large ball. There is another characterization of boundedness that is often useful.

Definition 2.5.5. Given a set $E \subseteq X$, we define the **diameter** of E to be

$$\operatorname{diam}(E) = \sup_{x,y \in E} d(x,y).$$

Note that if E is an open (or closed) ball, then $\operatorname{diam}(E)$ is just twice the radius. Therefore, it is not hard to check that a set E is bounded if and only if $\operatorname{diam}(E)$ is finite. If $E \subseteq B_M[x]$ for some $x \in X$ and $M > 0$, then it must be the case that $\operatorname{diam}(E) \leq 2M$. On the other hand, if $\operatorname{diam}(E) = M < \infty$, we have $E \subseteq B_M[x]$ for any $x \in E$.

Now we will show that one direction of the Heine-Borel theorem still holds—any compact set must be closed and bounded.

Proposition 2.5.6. *If $K \subseteq X$ is compact, then it is closed and bounded.*

Proof. To show that K is closed, we will show that K^c is open. Let $x \in K^c$. Then for each $y \in K$, we can find open balls U_y and V_y such that $y \in U_y$, $x \in V_y$, and $U_y \cap V_y = \emptyset$. Notice that the collection

$$\mathcal{U} = \{U_y\}_{y \in K}$$

is an open cover of K. Since K is compact, there is a finite subcover $\{U_{y_i}\}_{i=1}^n$. Furthermore, the set

$$V = \bigcap_{i=1}^n V_{y_i}$$

is an open neighborhood of x, and

$$V \subseteq \bigcap_{i=1}^n U_{y_i}^c = \left(\bigcup_{i=1}^n U_{y_i} \right)^c.$$

Since the right hand side contains K, we must have $V \cap K = \emptyset$. This shows that x is an interior point of K^c. It follows that K^c is open, so K is closed.

Now we show that K is bounded. Consider the open cover $\{B_1(x) : x \in K\}$. Since K is compact, there is a finite subcover $\{B_1(x_i)\}_{i=1}^n$. Now it is not hard to see that

$$\operatorname{diam}(K) \leq \operatorname{diam}(\{x_1, x_2, \ldots, x_n\}) + 2,$$

since each $x \in K$ is at most distance 1 from one of the x_i. In particular, if $x, y \in K$ with $x \in B_1(x_i)$ and $y \in B_1(x_j)$, then

$$d(x, y) \leq d(x, x_i) + d(x_i, x_j) + d(x_j, y) < d(x_i, x_j) + 2.$$

Therefore, $\operatorname{diam}(K) < \infty$, so K is bounded. $\square$

We will now investigate three examples that show how the converse to this proposition can fail in general.

Example 2.5.7. Take $X = \mathbf{Q}$, and consider the set

$$F = \{x \in \mathbf{Q} : x^2 < 2\}.$$

If we think of F as a subset of $\mathbf{R}$, we are just taking all the rational numbers x satisfying $-\sqrt{2} < x < \sqrt{2}$.

Notice that F is clearly bounded, since $\operatorname{diam}(F) = 2\sqrt{2}$. It is also closed *as a subset of* $\mathbf{Q}$. There are multiple ways to see this. For one, we can notice that $F = \mathbf{Q} \cap [-\sqrt{2}, \sqrt{2}]$, so F is relatively closed in $\mathbf{Q}$. If instead we try to show F^c is open, we can avoid appealing to the completion $\mathbf{R}$. Let $x \in F^c$, and assume without loss of generality that $x > 0$. Then we have $x^2 \geq 2$, and in fact we must have $x^2 > 2$, since there is no rational whose square is 2. Thus we can find a rational number $q > 0$ with $2 < q^2 < x^2$. If we set $r = x - q$, then $B_r(x) \subseteq F^c$. Thus F^c is open.

We have argued that F is closed and bounded, but it is not compact. To see this, consider the open cover $\mathcal{U} = \{U_n\}_{n=1}^\infty$ given by

$$U_n = \left(-\sqrt{2} + \tfrac{1}{n}, \sqrt{2} - \tfrac{1}{n}\right).$$

Then $\mathcal{U}$ has no finite subcover, for reasons similar to those of Example 2.5.2. (If we take a sequence of rationals "converging" to $\sqrt{2}$, the union of only finitely many of the U_n cannot possibly contain all the terms in the sequence.)

The previous example is somewhat contrived. Compactness failed largely because $\mathbf{Q}$ is not a complete metric space. We were able to construct a bounded set that is not closed in $\mathbf{R}$, but it did turn out to be closed in $\mathbf{Q}$. One might wonder if completeness is the crucial piece that allows one to prove the Heine-Borel theorem. Alas, this is not the case. We will now construct a closed, bounded subset of a *complete* metric space that fails to be compact.

Example 2.5.8. Though we have not proven it yet, it is not hard to show that ℓ^∞ is a Banach space with respect to its natural norm $\|\cdot\|_\infty$.[9] For each $n \in \mathbf{N}$, define $e_n \in \ell^\infty$ by

$$e_n = (0, 0, \ldots, 0, \underbrace{1}_{n\text{th}}, 0, \ldots).$$

Consider the set $\{e_n : n \in \mathbf{N}\}$. Notice that $\|e_n\|_\infty = 1$ for all n, so $E \subseteq B_1[0]$. Therefore, E is bounded. We could also have observed that $\|e_n - e_m\|_\infty = 1$ if $n \neq m$, so $\operatorname{diam}(E) = 1$.

This last observation has another important implication. Suppose $(x_k)_{k=1}^\infty$ is a sequence in E that converges to $x \in \ell^\infty$. Since $\|e_n - e_m\|_\infty = 1$ if $n \neq m$, any sequence consisting of points in E converges if and only if it is eventually constant. Moreover, we have $B_{1/2}(e_n) \cap E = \{e_n\}$ for all n, so every point of E is isolated. These two facts together show that E has no limit points, so it is vacuously closed.

Finally, we claim that E is not compact. Notice that the collection $\mathcal{U} = \{B_{1/2}(e_n)\}_{n=1}^\infty$ is an open cover of E. However, we already observed that $B_{1/2}(e_n) \cap E = \{e_n\}$ for all n, so no proper subset of $\mathcal{U}$ can cover E. In particular, there is no finite subcover. Thus E is not compact.

Remark 2.5.9. It is worth noting that the set E from the previous example can also be viewed as a subset of ℓ^1 or ℓ^2, and the same phenomena occur there as well. The only difference is that $\|e_n - e_m\|_1 = 2$ when $n \neq m$, while $\|e_n - e_m\|_2 = \sqrt{2}$.

Remark 2.5.10. In Example 2.5.8, we observed that

$$d(e_n, e_m) = \|e_n - e_m\|_\infty = \begin{cases} 1 & \text{if } n \neq m \\ 0 & \text{if } n = m. \end{cases}$$

In other words, E is really just a countable discrete metric space that happens to live inside ℓ^∞ as a subspace. It is easy to see that any infinite set equipped with the discrete metric would have the same issues that we observed in Example 2.5.8. In fact, any subset of a discrete metric space is compact if and only if it is finite. (See Exercise 2.5.1.)

The set in Example 2.5.8 failed miserably to be compact, since we chose a very nice open cover. Our cover consisted entirely of open balls, all with the same radius, and we still could not find a finite subcover. It is precisely this phenomenon that we need to avoid in order to guarantee compactness.

Definition 2.5.11. Let $E \subseteq X$. Given $\varepsilon > 0$, a collection of open balls $\{B_\varepsilon(x_i)\}_{i \in I}$, each of radius ε, that cover E is called an **ε-net** for E. We say E is **totally bounded** if for every $\varepsilon > 0$, there is a finite ε-net for E.

[9] Indeed, the proof is like the one that shows ℓ^2 is complete, but much easier.

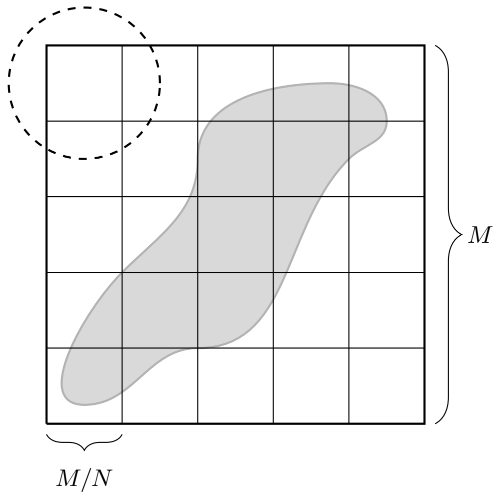

Figure 2.1: The set $E \subseteq \mathbf{R}^2$ (depicted in gray) can be placed inside a sufficiently large cube, since it is bounded. The cube is then subdivided into cells, each of which is small enough to fit inside an ε-ball. Thus there is a finite ε-net for the cube, hence there is one for E.

It is straightforward to prove that any compact set must be totally bounded. It will turn out that total boundedness is the correct substitute for boundedness in a general metric space. As a first step, we observe that compact sets are always totally bounded.

Proposition 2.5.12. *If $K \subseteq X$ is compact, then it is totally bounded.*

Proof. Let $\varepsilon > 0$ be given, and consider the ε-net $\{B_\varepsilon(x)\}_{x \in K}$. This ε-net forms an open cover for K, so there must be a finite subcover $\{B_\varepsilon(x_i)\}_{i=1}^n$ since K is compact. Thus K admits a finite ε-net, so it is totally bounded. □

Remark 2.5.13. In $\mathbf{R}^n$, a set is bounded if and only if it is totally bounded. To see this, suppose we have a bounded set $E \subseteq \mathbf{R}^n$. Then E is contained in an n-dimensional cube K of side length M for some $M > 0$. Given $\varepsilon > 0$, we can divide K up into smaller cubes, each of whose diagonals is smaller than 2ε. In particular, choose N such that

$$\frac{M}{N} < \frac{\varepsilon}{\sqrt{n}},$$

and then divide K into $(M/N)^n$ cubes of side length M/N. (See Figure 2.1.) Then each such cube has diameter $M\sqrt{n}/N \leq \varepsilon/2$, hence it can be placed inside an open ball of radius ε. It follows that there is a finite ε-net for K, so K is totally bounded.

We now know that total boundedness is boring in $\mathbf{R}^n$, since it is equivalent to the usual notion of boundedness. We also know that this equivalence can fail in infinite-dimensional spaces—we saw an example of a set in ℓ^∞ that is bounded, but not totally bounded. Now we will investigate a set in an infinite-dimensional space that does happen to be totally bounded.

Example 2.5.14. Define $H \subseteq \ell^2$ to be the set of all vectors $x = (x_1, x_2, \dots)$ satisfying

$$|x_i| \le \frac{1}{2^{i-1}}$$

for all i. This set is called the **Hilbert cube**. We claim that H is totally bounded. Let $\varepsilon > 0$ be given, and choose $n \in \mathbf{N}$ such that $\frac{1}{2^{n-1}} < \frac{\varepsilon}{2}$. For each $x \in H$, define a vector $x^* \in \ell^2$ by

$$x^* = (x_1, x_2, \dots, x_n, 0, 0, \dots).$$

Then observe that

$$\|x - x^*\|_2^2 = \sum_{i=n+1}^{\infty} x_i^2 \le \sum_{i=n+1}^{\infty} \frac{1}{4^{i-1}} = \sum_{i=n}^{\infty} \frac{1}{4^i}$$

for all $x \in H$. Since this is a geometric sum, we have

$$\|x - x^*\|_2^2 \le \frac{1}{4^n} \cdot \frac{4}{3} < \frac{1}{4^{n-1}} < \frac{\varepsilon^2}{4},$$

or $\|x - x^*\|_2 < \frac{\varepsilon}{2}$ for all $x \in H$. Now the set

$$E = \{x^* : x \in H\}$$

is bounded and belongs to an n-dimensional subspace of ℓ^2. In other words, we can view E as a subset of $\mathbf{R}^n$, and it follows from Remark 2.5.13 that E is totally bounded. Thus there is a finite $\frac{\varepsilon}{2}$-net $\{B_{\varepsilon/2}(y_i)\}_{i=1}^m$ for E. But now observe that if $x \in H$, we must have $x^* \in B_{\varepsilon/2}(y_i)$ for some i, and

$$\|x - y_i\|_2 \le \|x - x^*\|_2 + \|x^* - y_i\|_2 < \tfrac{\varepsilon}{2} + \tfrac{\varepsilon}{2} = \varepsilon.$$

Thus $x \in B_\varepsilon(y_i)$, and it follows that $\{B_\varepsilon(y_i)\}_{i=1}^m$ is a finite ε-net for H. Thus H is totally bounded.

As we have mentioned before, total boundedness is the correct substitute for boundedness when trying to characterize compactness in a general metric space. In order to prove this characterization, we will need a third notion of compactness. Recall our earlier example of $E = \{e_n : n \ge 1\}$ in ℓ^∞. This set was closed and bounded, but it failed to be compact precisely because it was not totally bounded. Notice that we could have also viewed E as a *sequence* in ℓ^∞, i.e. we could look at the sequence $(e_n)_{n=1}^\infty$. Recall that if $n \neq m$,

$$\|e_n - e_m\|_\infty = 1.$$

This implies that $(e_n)_{n=1}^{\infty}$ is not a Cauchy sequence. Moreover, no subsequence of $(e_n)_{n=1}^{\infty}$ could possibly be Cauchy! In other words, this sequence is bounded, but it has no convergent subsequences. This phenomenon cannot occur in $\mathbf{R}^n$, since the Bolzano-Weierstrass theorem guarantees that any bounded sequence must have a convergent subsequence. It appears then that the failure of the Heine-Borel theorem in general metric spaces is closely tied to the failure of the Bolzano-Weierstrass theorem. This leads us to define the following notion of compactness.

Definition 2.5.15. A set $K \subseteq X$ is **sequentially compact** if every sequence in K has a subsequence that converges in K.

We will see soon that a set is sequentially compact if and only if its compact. To prove it, we will need a couple of auxiliary results.

Proposition 2.5.16. *Any closed subset of a compact set is compact.*

Proof. Suppose $K \subseteq X$ is compact, and let $F \subseteq K$ be a closed set. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover of F. Since F is closed, F^c is open, and $\mathcal{U} \cup \{F^c\}$ is an open cover of K. Since K is compact, there is a finite subcover, say $\{U_i\}_{i=1}^n \cup \{F^c\}$. But then $\{U_i\}_{i=1}^n$ must cover F, and consequently F is compact. $\square$

Proposition 2.5.17. *If $K \subseteq X$ is compact, then any infinite subset of K has a limit point in K.*

Proof. Let $E \subseteq K$ be an infinite set. Since K is closed, any limit point of E will necessarily belong to K. Therefore, it suffices to show that E has a limit point.

Suppose to the contrary that E has no limit points. Then E is closed, hence compact. Furthermore, every point of E is isolated. Thus for each $x \in E$, there is an open ball U_x such that $U_x \cap E = \{x\}$. Then $\mathcal{U} = \{U_x\}_{x \in E}$ is an open cover of E, and no proper subset of $\mathcal{U}$ can cover E. In other words, $\mathcal{U}$ does not admit a finite subcover, since E is infinite. But this contradicts the fact that E is compact. Therefore, E must have a limit point. $\square$

We are now ready to prove our main theorem characterizing compactness in an arbitrary metric space. The proof is based on the one given in [KF70].

Theorem 2.5.18. *Let X be a metric space. The following are equivalent:*

1. *X is compact.*
2. *X is sequentially compact.*
3. *X is complete and totally bounded.*

Proof. ($1 \Rightarrow 2$) Assume first that X is compact, and let $(x_n)_{n=1}^{\infty}$ be a sequence in X. We will consider the set

$$E = \{x_n : n \geq 1\},$$

i.e., the set of points in X that occur as terms in the sequence. There are two possibilities—either E is finite, or it is infinite. If it is finite, then there is some $x \in E$ that occurs infinitely often the sequence. That is, there are infinitely many natural numbers $n_1 < n_2 < n_3 < \cdots$ such that $x_{n_k} = x$. In other words, $(x_n)_{n=1}^{\infty}$ contains the constant subsequence $(x_{n_k})_{k=1}^{\infty}$, which of course converges to x. Now suppose E is infinite. Then by Proposition 2.5.17, E has a limit point $x \in X$. It then follows from Proposition 2.3.7 that there is a subsequence $(x_{n_k})_{k=1}^{\infty}$ that converges to x. Thus X is sequentially compact.

($2 \Rightarrow 3$) Now assume that X is sequentially compact. We first claim that X must be complete. Any Cauchy sequence $(x_n)_{n=1}^{\infty}$ in X must have a convergent subsequence $(x_{n_k})_{k=1}^{\infty}$. If we let $x = \lim x_{n_k}$, then Exercise 2.3.2 guarantees that $x_n \to x$ as well. Therefore, X is complete.

Suppose X is not totally bounded. Then there exists $\varepsilon > 0$ such that there is no finite ε-net for X. Choose a point $x_1 \in X$. Since $B_\varepsilon(x_1)$ does not cover X, there is a point $x_2 \in X$ such that

$$d(x_1, x_2) \geq \varepsilon.$$

Again, $\{B_\varepsilon(x_1), B_\varepsilon(x_2)\}$ does not cover X, so there is a point $x_3 \in X$ with

$$d(x_1, x_3) \geq \varepsilon, \quad d(x_2, x_3) \geq \varepsilon.$$

Since there is no finite ε-net for X, we can continue inductively and produce a sequence $(x_n)_{n=1}^{\infty}$ such that $d(x_n, x_m) \geq \varepsilon$ for all $n, m \in \mathbf{N}$. It is then clear that $(x_n)_{n=1}^{\infty}$ has no convergent subsequences, contradicting the assumption that X is sequentially compact. Therefore, X is totally bounded.

($3 \Rightarrow 2$) Assume X is complete and totally bounded, and let $(x_n)_{n=1}^{\infty}$ be a sequence in X. Since X is totally bounded, there is a finite 1-net, i.e. a finite family of open balls of radius 1 that covers X. Then there must be some ball U_1 of radius 1 that contains infinitely many of the x_n. Choose one of these points and call it x_{n_1}. Similarly, there is a finite $\frac{1}{2}$-net for X, so we can find an open ball U_2 of radius $\frac{1}{2}$ such that $U_1 \cap U_2$ contains infinitely many of the x_n. Thus we can choose $x_{n_2} \in U_1 \cap U_2$ with $n_2 > n_1$. Now we continue inductively: for each $k \in \mathbf{N}$, find an open ball U_k of radius $\frac{1}{2^{k-1}}$ that contains infinitely many terms x_n, and choose $x_{n_k} \in U_1 \cap U_2 \cap \cdots \cap U_k$ such that $n_1 < n_2 < \cdots < n_k$.

We have thus far constructed a subsequence $(x_{n_k})_{k=1}^{\infty}$ of $(x_n)_{n=1}^{\infty}$, and we claim that this sequence is Cauchy. Let $\varepsilon > 0$ be given, and choose $N \in \mathbf{N}$ such that $\frac{1}{2^{N-1}} < \frac{\varepsilon}{2}$. Recall that for all $k \geq N$, we have $x_{n_k} \in U_N$ by construction. Thus for all $k, l \geq N$,

$$d(x_{n_k}, x_{n_l}) < \operatorname{diam}(U_N) = 2 \cdot \frac{1}{2^{N-1}} < 2 \cdot \frac{\varepsilon}{2} = \varepsilon.$$

Thus $(x_{n_k})_{k=1}^\infty$ is Cauchy. Since X is complete, this subsequence must converge. Hence X is sequentially compact.

($2 \Rightarrow 1$) Suppose X is sequentially compact, and let $\{U_i\}_{i\in I}$ be an open cover for X. We first claim that there exists $n \in \mathbf{N}$ such that every open ball of radius $\frac{1}{n}$ in X is contained in some U_i. Suppose this is not the case—then for each $n \in \mathbf{N}$, there is a point $x_n \in X$ such that $B_{1/n}(x_n) \nsubseteq U_i$ for all i. By assumption, the sequence $(x_n)_{n=1}^\infty$ has a convergent subsequence $(x_{n_k})_{k=1}^\infty$. If we let $x = \lim x_{n_k}$, then $x \in U_{i_0}$ for some $i_0 \in I$. Since U_{i_0} is open, there is an $\varepsilon > 0$ such that $B_\varepsilon(x) \subseteq U_{i_0}$. But $x_{n_k} \to x$, so there exists $N \in \mathbf{N}$ such that $d(x_{n_k}, x) < \frac{\varepsilon}{2}$ for all $k \geq N$. By possibly increasing N, we can arrange things so $\frac{1}{n_k} < \frac{\varepsilon}{2}$ when $k \geq N$. Thus we have

$$B_{1/n_k}(x_{n_k}) \subseteq B_\varepsilon(x) \subseteq U_{i_0}.$$

However, this contradicts the construction of our sequence $(x_n)_{n=1}^\infty$. Therefore, our initial claim must be correct.

Now choose $n \in \mathbf{N}$ as in the above claim. Since X is sequentially compact, it is totally bounded, so there exists a finite $\frac{1}{n}$-net $\{B_{1/n}(x_k)\}_{k=1}^m$ for X. But we have already established that each one of these balls belongs to some U_i, i.e., $B_{1/n}(x_k) \subseteq U_{i_k}$ for some $i_k \in I$. It follows that the family $\{U_{i_k}\}_{k=1}^m$ covers X. Thus we have produced a finite subcover, so X is compact. □

Though we did not phrase it this way, Theorem 2.5.19 applies to subsets of a metric space as well, provided the ambient metric space is complete. Recall that if $E \subseteq X$, then we can view E as a metric space in its own right by equipping it with the metric inherited from X. Thus E is compact in X precisely when it constitutes a compact metric space à la Theorem 2.5.19. Observe also that if X is complete, then any closed subset of X is complete as well. Thus the conditions "complete and totally bounded" can be replaced with "closed and totally bounded" for subspaces $E \subseteq X$. Thus we have the following version of Theorem 2.5.19 for subsets of a complete metric space.

Theorem 2.5.19. *Let X be a complete metric space, and suppose $K \subseteq X$. The following are equivalent:*

1. *K is compact.*
2. *K is sequentially compact.*
3. *K is closed and totally bounded.*

Things are complicated if X is not assumed to be complete. For example, take $X = \mathbf{Q}$ and $F = \{x \in \mathbf{Q} : x^2 < 2\}$, as in Example 2.5.7. Then F is closed *as a subset of* $\mathbf{Q}$ and it is totally bounded, but it is not compact (nor is it sequentially compact). The issue of course is that $\mathbf{Q}$ is not a complete metric space.

Exercises for Section 2.5

Exercise 2.5.1 ([HS91], Exercises 4.8.6 and 4.8.7)**.** Let (X, d) be a metric space.

(a) Prove that any finite subset of X is compact.

(b) If d is the discrete metric, prove that a set $K \subseteq X$ is compact if and only if K is finite.

Exercise 2.5.2 ([HS91], Exercise 4.10.2)**.** A metric space (X, d) is **locally compact** if given any $x \in X$, there is a compact set $K \subseteq X$ with $x \in K^\circ$. Prove that $\mathbf{R}^n$ is locally compact.

Exercise 2.5.3. Let (X, d) be a metric space.

(a) Suppose X is complete, and let $\{F_n\}_{n=1}^\infty$ be a nested sequence of closed sets in X, meaning that

$$F_1 \supseteq F_2 \supseteq F_3 \supseteq \cdots$$

Suppose further that

$$\lim_{n\to\infty} \operatorname{diam}(F_n) = 0.$$

Show that $\bigcap_{n=1}^\infty F_n \neq \emptyset$.

(b) Give an example of a complete metric space X and a nested sequence of closed sets $\{F_n\}_{n=1}^\infty$ such that $\bigcap_{n=1}^\infty F_n = \emptyset$. Reconcile this example with the result from part (a).

Exercise 2.5.4 ([HS91], Exercise 4.8.11)**.** A collection of sets $\mathcal{E} = \{E_i\}_{i\in I}$ has the **finite intersection property** if given any finite set $F \subseteq I$,

$$\bigcap_{i\in F} E_i \neq \emptyset.$$

Prove that a metric space X is compact if and only if any family $\{F_i\}_{i\in I}$ of closed sets with the finite intersection property has nonempty intersection, i.e.,

$$\bigcap_{i\in I} F_i \neq \emptyset.$$

Exercise 2.5.5 ([HS91], Exercise 4.9.2)**.** Let X be a metric space. Prove that if $E \subseteq X$ is totally bounded, then E is bounded.

2.6 Connectedness

We will now discuss one last fundamental property of metric spaces, called *connectedness*. You may recall that connectedness plays a crucial role in real analysis, as it underlies some major results, including the Intermediate Value Theorem. As usual, we will let (X, d) denote an arbitrary metric space.

Unfortunately, connectedness is most easily defined by first describing what it means to not be connected.

Definition 2.6.1. Two sets $A, B \subseteq X$ are said to be **separated** if $A \cap \overline{B} = \emptyset$ and $\overline{A} \cap B = \emptyset$.

Notice that it takes more than simply being disjoint for two sets to be separated. In addition to the requirement that $A \cap B = \emptyset$, neither set can contain limit points of the other. It is possible for A and B to have limit points in common, however.

Example 2.6.2. Consider the sets $A = (0, 1)$ and $B = (1, 2)$ in $\mathbf{R}$. Notice that $\overline{A} = [0, 1]$ and $\overline{B} = [1, 2]$, so A and B are separated. (Note that they both have 1 as a limit point, however.) On the other hand, the sets $A = (0, 1)$ and $B = [1, 2]$ are not separated, since $\overline{A} \cap B = \{1\}$.

Definition 2.6.3. A metric space X is **connected** if it cannot be written as the union of two nonempty separated sets. That is, $X \neq A \cup B$ where $A, B \neq \emptyset$ and A and B are separated.

We say a subset $E \subseteq X$ is **connected** if E is a connected metric space when viewed as a subspace of X.

Example 2.6.4. It is a well-known fact from real analysis that $\mathbf{R}$ is connected. More generally, a subset of $\mathbf{R}$ is connected if and only if it is an interval.

Example 2.6.5. In any metric space X, singletons are always connected. Indeed, suppose $x \in X$ and we write $\{x\} = A \cup B$ for some sets $A, B \subseteq X$ with $A \cap B = \emptyset$. Then $x \in A$ or $x \in B$; without loss of generality, assume $x \in A$. Then we must have $A = \{x\}$ and $B = \emptyset$, so $\{x\}$ is connected.

There is another characterization of connectedness that is often easier to verify. Suppose $X = A \cup B$, where A and B are nonempty and separated. Then we have $A \cap \overline{B} = \emptyset$, so $A \subseteq (\overline{B})^c$. In fact, since $X = A \cup B$, we have

$$A = B^c \supseteq (\overline{B})^c.$$

Therefore, $A = (\overline{B})^c$, and it follows that A is open. Similarly, $B = (\overline{A})^c$, so B is also open. Thus if X is not connected, we can express it as a union of two disjoint open sets. The converse holds as well—if $X = A \cup B$, where A and B are nonempty open sets with $A \cap B = \emptyset$, then A and B are necessarily separated. Indeed, notice that $A = B^c$, so A is closed. Thus

$$\overline{A} \cap B = A \cap B = \emptyset.$$

Similarly, $A \cap \overline{B} = \emptyset$. Therefore, we have proven the following result.[10]

Proposition 2.6.6. *A metric space X is connected if and only if it cannot be written as a union $X = A \cup B$, where A and B are nonempty open sets with $A \cap B = \emptyset$.*

Remark 2.6.7. The corresponding characterization of connectedness for subspaces looks a little bit different. If $E \subseteq X$ is not connected, we cannot write E as a union of two disjoint open sets unless E is open in X. However, we can say the following: $E \subseteq A \cup B$, where A and B are open, $A \cap E \neq \emptyset$ and $B \cap E \neq \emptyset$, and $A \cap B \cap E = \emptyset$. If we have such a setup, the sets $A \cap E$ and $B \cap E$ are relatively open, disjoint sets in E with

$$E = (A \cap E) \cup (B \cap E).$$

This is precisely what it means for the metric space E to be disconnected.

Now we will make another observation regarding connected (or disconnected) sets. Suppose we have a disconnected metric space X, so $X = A \cup B$ with A and B nonempty, open, and disjoint. We observed earlier that $A = B^c$ is also closed, as is $B = A^c$. In other words, A and B are both clopen sets. It turns out that the question of whether a metric space is connected is tied closely to the existence of nontrivial clopen sets.

Proposition 2.6.8. *A metric space X is connected if and only if its only clopen subsets are $\emptyset$ and X.*

Proof. Suppose first that X is not connected. Then there are nonempty, disjoint open sets A and B such that $X = A \cup B$. Since $A, B \neq \emptyset$, we also have $A, B \neq X$. Thus X contains nontrivial clopen sets.

On the other hand, suppose X contains a clopen set A with $A \neq \emptyset$ and $A \neq X$. Then $B = A^c$ is a nonempty clopen set, and $B \neq X$. Furthermore, $X = A \cup B$, so X is disconnected. □

[10] We could also easily prove the following characterization: X is connected if and only if it cannot be written as a union of two nonempty disjoint *closed* sets.

Example 2.6.9. Recall that if X is a discrete metric space, then every subset of X is clopen. In particular, singletons are clopen. Suppose $E \subseteq X$ and E consists of more than one point. Let $x \in E$, and put $A = \{x\}$. Then A is open, $B = E \backslash \{x\}$ is nonempty and open, and $E = A \cup B$. Thus E is disconnected. In other words, a nonempty subset of a discrete metric space is connected if and only if it is a singleton.

The phenomenon of Example 2.6.9 can occur in non-discrete metric spaces. Thus we have a special name for such spaces.

Definition 2.6.10. A metric space X is **totally disconnected** if the only connected subsets of X are the singletons.

Example 2.6.11. The rational numbers form a totally disconnected space with respect to their usual metric. Similarly, the irrationals are totally disconnected.

Example 2.6.12. The Cantor set (viewed as a subspace of $\mathbf{R}$) is totally disconnected.

Exercises for Section 2.6

Exercise 2.6.1. Show by example that the interior of a connected set need not be connected.

Exercise 2.6.2 ([HS91], Exercise 4.11.2)**.** Show that if E is connected and $E \subseteq F \subseteq \overline{E}$, then F is also connected. Conclude that, in particular, $\overline{E}$ is connected if E is connected.

2.7 Continuity

Now that we have discussed basic properties of metric spaces themselves, we will now turn our attention to functions on metric space. In particular, we will close this chapter with two sections containing results about continuous functions on metric spaces. The next chapter will consist of a thorough treatment of sequences and series of functions on metric spaces.

Recall from real analysis that a function $f : \mathbf{R} \to \mathbf{R}$ is **continuous** at a point $a \in \mathbf{R}$ if given any $\varepsilon > 0$, there exists a $\delta > 0$ such that $|x - a| < \delta$ implies

$$|f(x) - f(a)| < \varepsilon.$$

As we have already mentioned, this definition relies only on our ability to measure the distance between two real numbers. Therefore, the same definition works in any metric space, simply by substituting the appropriate metric in place of the standard metric on $\mathbf{R}$.

Definition 2.7.1. Let (X, d_1) and (Y, d_2) be metric spaces. A function $f : X \to Y$ is **continuous** at a point $x_0 \in X$ if given any $\varepsilon > 0$, there exists $\delta > 0$ such that $d_1(x, x_0) < \delta$ implies

$$d_2(f(x), f(x_0)) < \varepsilon.$$

We say f is **continuous on** X if it is continuous at every point of X.

Remark 2.7.2. We will usually say "f is continuous" to indicate that f is continuous on X.

There are other useful ways of characterizing continuity in a metric space. The first one is the *topological* definition of continuity, which says that the preimage of any open set is open.

Theorem 2.7.3. *A function $f : X \to Y$ is continuous if and only if $f^{-1}(U)$ is open in X for any open set $U \subseteq Y$.*

Proof. Suppose first that $f : X \to Y$ is continuous, and let $U \subseteq Y$ be an open set. Let $x_0 \in f^{-1}(U)$. Then $f(x_0) \in U$, and there is an $\varepsilon > 0$ such that $B_\varepsilon(f(x_0)) \subseteq U$ since U is open. Since f is continuous, there is a $\delta > 0$ such that $d_1(x, x_0) < \delta$ implies $d_2(f(x), f(x_0)) < \varepsilon$. In other words, if $x \in B_\delta(x_0)$, then

$$f(x) \in B_\varepsilon(f(x_0)) \subseteq U.$$

This shows that $x \in f^{-1}(U)$, and it follows that $B_\delta(x_0) \subseteq f^{-1}(U)$. Since $x_0 \in f^{-1}(U)$ was arbitrary, $f^{-1}(U)$ is open.

Conversely, suppose inverse images of open sets are open. Let $x_0 \in X$, and let $\varepsilon > 0$ be given. Put $U = B_\varepsilon(f(x_0))$. Then U is open, so $f^{-1}(U)$ is open in X. Consequently, there is a $\delta > 0$ such that $B_\delta(x_0) \subseteq f^{-1}(U)$. Equivalently, $d_1(x, x_0) < \delta$ implies $f(x) \in U$, which is the same as saying that $d_2(f(x), f(x_0)) < \varepsilon$. Thus f is continuous at x_0. It follows that f is continuous on X. □

By considering complements, we have the following immediate corollary regarding closed sets.

Corollary 2.7.4. *A function $f : X \to Y$ is continuous if and only if $f^{-1}(F)$ is closed in X for any closed set $F \subseteq Y$.*

Proof. Notice that if $F \subseteq Y$ is closed, then F^c is open. Thus

$$f^{-1}(F)^c = f^{-1}(F^c)$$

is open in X if f is continuous, so $f^{-1}(F)$ is closed. The converse is just as straightforward. □

Now we come to our second characterization of continuity, which is phrased in terms of convergent sequences.

Theorem 2.7.5. *A function $f : X \to Y$ is continuous at a point $x_0 \in X$ if and only if given any sequence $(x_n)_{n=1}^{\infty}$ in X converging to x_0, the sequence $(f(x_n))_{n=1}^{\infty}$ converges to $f(x_0)$ in Y.*

Proof. Assume first that f is continuous at x_0, and let $(x_n)_{n=1}^{\infty}$ be a sequence converging to x_0. Let $\varepsilon > 0$, and find $\delta > 0$ such that $d_1(x, x_0) < \delta$ implies $d_2(f(x), f(x_0)) < \varepsilon$. Now choose $N \in \mathbf{N}$ such that $n \geq N$ implies

$$d_1(x_n, x_0) < \delta.$$

Then for all $n \geq N$, we automatically have

$$d_2(f(x_n), f(x_0)) < \varepsilon,$$

so $f(x_n) \to f(x_0)$.

For the converse, assume that f is not continuous at x_0. Then there exists an $\varepsilon > 0$ such that for all $\delta > 0$ there is a point $x \in X$ with $d_1(x, x_0) < \delta$, but $d_2(f(x), f(x_0)) \geq \varepsilon$. In particular, for each $n \in \mathbf{N}$ we can find a point $x_n \in X$ such that $d_1(x_n, x_0) < \frac{1}{n}$ and $d_2(f(x_n), f(x_0)) \geq \varepsilon$. Then clearly $x_n \to x_0$, but $(f(x_n))_{n=1}^{\infty}$ does not converge to $f(x_0)$. $\square$

With these basic results on continuity out of the way, we will investigate the relationship between continuity and other topological concepts. We begin with compactness.

Theorem 2.7.6. *If $f : X \to Y$ is a continuous function and $K \subseteq X$ is compact, then $f(K)$ is compact.*

Proof. Let $\mathcal{U} = \{U_i\}_{i \in I}$ be an open cover for $f(K)$. Since f is continuous, $f^{-1}(U_i)$ is open for all $i \in I$, and we have

$$K \subseteq f^{-1}(f(K)) \subseteq f^{-1}\left(\bigcup_{i \in I} U_i\right) = \bigcup_{i \in I} f^{-1}(U_i).$$

Thus $\{f^{-1}(U_i)\}_{i \in I}$ is an open cover for K. Since K is compact, there is a finite subcover $\{f^{-1}(U_i)\}_{i=1}^{n}$. We claim that $\{U_i\}_{i=1}^{n}$ covers $f(K)$. Well, we have

$$f(K) \subseteq f\left(\bigcup_{i=1}^{n} f^{-1}(U_i)\right) = \bigcup_{i=1}^{n} f(f^{-1}(U_i)) = \bigcup_{i=1}^{n} U_i,$$

so $f(K)$ is compact. $\square$

As a byproduct, we can easily generalize a familiar result from real analysis (or from Calculus I, really).

Theorem 2.7.7 (Extreme Value Theorem). *If X is a compact metric space and $f : X \to \mathbf{R}$ is continuous, then f attains a maximum and a minimum value. That is, there exist points $x_1, x_2 \in X$ such that*

$$f(x_1) = \inf\{f(x) : x \in X\}$$

and

$$f(x_2) = \sup\{f(x) : x \in X\}.$$

Proof. By the previous theorem, $f(X)$ is a compact subset of $\mathbf{R}$, so it is closed and bounded. It is then a fact from real analysis that any such set must contain its infimum and supremum. □

As we will see momentarily, compactness is closely related to a stronger form of continuity, called *uniform continuity.*

Definition 2.7.8. A function $f : X \to Y$ is **uniformly continuous** if given any $\varepsilon > 0$, there exists $\delta > 0$ such that $d_1(x, y) < \delta$ implies $d_2(f(x), f(y)) < \varepsilon$ for all $x, y \in X$.

Notice that a function is uniformly continuous if given any ε, we can find a corresponding δ that works throughout X. It is easy to see that uniform continuity implies continuity, though the converse is not generally true. We will now look at several functions defined on subsets of $\mathbf{R}$ to explore the various things that can go wrong.

Example 2.7.9. The function $f : \mathbf{R} \to \mathbf{R}$ defined by $f(x) = x$ is easily seen to be uniformly continuous on $\mathbf{R}$. Given any $\varepsilon > 0$, take $\delta = \varepsilon$. Then $|x - y| < \delta$ clearly implies

$$|f(x) - f(y)| = |x - y| < \delta = \varepsilon.$$

Example 2.7.10. The function $f(x) = \sqrt{x}$ is uniformly continuous on $[0, \infty)$. Let $\varepsilon > 0$ be given, and let $\delta = \varepsilon^2$. If $x, y \in [0, \infty)$ satisfy $|x - y| < \delta$, then

$$\left|\sqrt{x} - \sqrt{y}\right|^2 \leq \left|\sqrt{x} + \sqrt{y}\right| \cdot \left|\sqrt{x} - \sqrt{y}\right| = |x - y| < \varepsilon^2.$$

Therefore,

$$\left|\sqrt{x} - \sqrt{y}\right| < \varepsilon$$

whenever $|x - y| < \delta$, so f is uniformly continuous on $[0, \infty)$.

Example 2.7.11. The function $f(x) = \frac{1}{x}$ is not uniformly continuous on the open interval $(0, 1)$, though it is continuous there. Intuitively, we can see that the outputs of the function are stretched further and further apart as the inputs are taken to be closer and closer to 0. The easiest way to rigorously show it involves appealing to the fact that uniformly continuous functions take Cauchy sequences to Cauchy sequences (Exercise 2.7.2). In particular, the let $(x_n)_{n=1}^{\infty}$ be any sequence in $(0, 1)$ that converges to 0 (such as $x_n = \frac{1}{n}$). Then $(x_n)_{n=1}^{\infty}$ is Cauchy, but

$$f(x_n) = \frac{1}{x_n}$$

goes off to infinity as $n \to \infty$. Thus the sequence $(f(x_n))_{n=1}^{\infty}$ is decidedly not Cauchy.

Example 2.7.12. In the same vein as the last example, the function $f(x) = x^2$ is continuous on $\mathbf{R}$, but it is not uniformly continuous. Again, the outputs of the function are stretched apart as x becomes large. Let $x_0 \in \mathbf{R}$ and choose $\varepsilon > 0$. Then to make $\left|x^2 - x_0^2\right| < \varepsilon$, we would need

$$|x - x_0||x + x_0| < \varepsilon.$$

In order to make this happen, we must weigh $|x - x_0|$ versus $|x + x_0|$. If we initially pick $\delta = 1$, then $|x - x_0| < \delta$ implies

$$|x + x_0| \leq |x| + |x_0| \leq (|x_0| + 1) + |x_0| = 2|x_0| + 1,$$

so

$$\left|x^2 - x_0^2\right| = |x - x_0||x + x_0| \leq (2|x_0| + 1)|x - x_0|.$$

In order to make this quantity less than ε, we probably need to shrink δ. Indeed, we will choose $\delta = \min\{1, \varepsilon/(2|x_0| + 1)\}$. If $|x - x_0| < \delta$, then

$$\left|x^2 - x_0^2\right| = |x - x_0||x + x_0| < \frac{\varepsilon}{2|x_0| + 1}(|x| + |x_0|) \leq \frac{\varepsilon}{2|x_0| + 1}(2|x_0| + 1) = \varepsilon.$$

Thus our δ necessarily depends on x_0, and we can see that it needs to get smaller as $x_0 \to \infty$. Thus $f(x) = x^2$ is not uniformly continuous on $\mathbf{R}$.

On the other hand, suppose we consider $f(x) = x^2$ only on the interval $[0, 1]$. Notice that for $x, y \in [0, 1]$, we have

$$\left|x^2 - y^2\right| = |x - y||x + y| \leq 2|x - y|,$$

since $0 \leq x, y \leq 1$. Therefore, given $\varepsilon > 0$, we could take $\delta = \varepsilon/2$. Then $|x - y| < \delta$ would imply that $\left|x^2 - y^2\right| < \varepsilon$. Thus $f(x) = x^2$ is uniformly continuous on $[0, 1]$.

The last example is actually quite instructive—continuity need not imply uniform continuity in general, but it does when we restrict a function's domain to a compact set.

Theorem 2.7.13. *If X is compact and $f : X \to Y$ is continuous, then f is uniformly continuous.*

Proof. Suppose f is not uniformly continuous. Then there exists $\varepsilon > 0$ such that for all $n \in \mathbf{N}$, there are points $x_n, y_n \in X$ with $d_1(x_n, y_n) < \frac{1}{n}$, but

$$d_2(f(x_n), f(y_n)) \geq \varepsilon.$$

Thus we have two sequences $(x_n)_{n=1}^{\infty}$ and $(y_n)_{n=1}^{\infty}$, and since X is compact, they both have convergent subsequences. Let $(x_{n_k})_{k=1}^{\infty}$ be a convergent subsequence of $(x_n)_{n=1}^{\infty}$ with $x_{n_k} \to x$. Then a straightforward application of the triangle inequality shows that $y_{n_k} \to x$ as well. Since f is continuous, $f(x_{n_k}) \to f(x)$ and $f(y_{n_k}) \to f(x)$. Thus there exists $N \in \mathbf{N}$ such that $k \geq N$ implies

$$d_2(f(x_{n_k}), f(x)) < \tfrac{\varepsilon}{2} \quad \text{and} \quad d_2(f(y_{n_k}), f(x)) < \tfrac{\varepsilon}{2}.$$

Thus

$$d_2(f(x_{n_k}), f(y_{n_k})) \leq d_2(f(x_{n_k}), f(x)) + d_2(f(y_{n_k}), f(x)) < \varepsilon,$$

which contradicts how we constructed the two sequences. Therefore, f must be uniformly continuous. □

Now we turn to the issue of connectedness. We will see that the following result implies another important result from real analysis.

Theorem 2.7.14. *If X is connected and $f : X \to Y$ is continuous, then $f(X)$ is connected.*

Proof. Suppose $f(X)$ is disconnected. Then $f(X) \subseteq A \cup B$, where A and B are nonempty open sets satisfying $A \cap f(X) \neq \emptyset$, $B \cap f(X) \neq \emptyset$, and $A \cap B \cap f(X) = \emptyset$. Since f is continuous, $f^{-1}(A)$ and $f^{-1}(B)$ are open in X, and

$$f^{-1}(A) \cap f^{-1}(B) = f^{-1}(A \cap B \cap f(X)) = \emptyset.$$

Also, $f^{-1}(A)$ and $f^{-1}(B)$ are both nonempty, and

$$X = f^{-1}(A \cup B) = f^{-1}(A) \cup f^{-1}(B).$$

Therefore, X is not connected. □

Theorem 2.7.15 (Intermediate Value Theorem)**.** *Suppose X is connected and $f : X \to \mathbf{R}$ is continuous. If $a, b \in X$ with $f(a) < f(b)$ and $t \in \mathbf{R}$ satisfies $f(a) < t < f(b)$, then there exists $x \in X$ with $f(x) = t$.*

Proof. By the previous theorem, $f(X)$ is a connected subset of $\mathbf{R}$, hence it is an interval. Thus $f(a) < t < f(b)$ implies $t \in f(X)$ for any $a, b \in X$. □

Exercises for Section 2.7

Exercise 2.7.1. Let X be a metric space.

(a) ([HS91], Exercise 4.4.11) Fix a point $x_0 \in X$ and define $f : X \to \mathbf{R}$ by

$$f(x) = d(x, x_0).$$

Prove that f is continuous.

(b) ([HS91], Exercise 4.11.11) Suppose X is connected and contains at least two points. Prove that X is uncountable.

Exercise 2.7.2. Let (X, d_1) and (Y, d_2) be metric spaces,and suppose $f : X \to Y$ is a uniformly continuous function. If $(x_n)_{n=1}^\infty$ is any Cauchy sequence in X, prove that the sequence $(f(x_n))_{n=1}^\infty$ is Cauchy.

Exercise 2.7.3. Let (X_1, d_1) and (X_2, d_2) be metric spaces.

(a) If d_1 is the discrete metric on X_1, show that every function $f : X_1 \to X_2$ is continuous.

(b) Suppose d_2 is the discrete metric on X_2 and $f : X_1 \to X_2$ is continuous. What does f look like?

Exercise 2.7.4. Let V be a normed vector space over $\mathbf{R}$. A set $E \subseteq V$ is said to be **convex** if for any two points $x, y \in V$, the set

$$\big\{(1-t)x + ty : 0 \le t \le 1\big\}$$

is contained in E. (Note that if $V = \mathbf{R}^2$, for example, this set is just the line segment connecting x and y.)

(a) Prove that the open ball $B_1(0)$ is convex. (Since every open ball in V can be obtained from $B_1(0)$ via a translation and a dilation, we can conclude that open balls are always convex.)

(b) Show that any convex set $E \subseteq V$ is connected. (**Hint:** Assume $E = A \cup B$ where A and B are separated, and consider the function $\gamma : [0,1] \to V$ defined by

$$\gamma(t) = (1-t)x + ty$$

for two points $x \in A$ and $y \in B$.)

Note: To make this problem easier, you may freely assume that the vector space operations (i.e., addition and scalar multiplication) on V are continuous.

2.8 The Banach Fixed Point Theorem

We will close out this chapter by discussing a powerful theorem for complete metric spaces. It goes by many names, including the *Contraction Mapping Theorem* (or *Contraction Mapping Principle*) and the *Banach Fixed Point Theorem.* As we will see via example, this theorem is most often used to prove that certain equations (in a very broad sense) have unique solutions. Moreover, the proof of the theorem is actually constructive—it not only tells us that a solution exists, but how to find it (or really, how to approximate it).

Definition 2.8.1. Let (X, d_1) and (Y, d_2) be metric spaces. A function $f : X \to Y$ is a **contraction** if there is a constant $0 < \lambda < 1$ such that

$$d_2(f(x), f(y)) \leq \lambda \cdot d_1(x, y)$$

for all $x, y \in X$.

Example 2.8.2. The function $f : [0, 1] \to [0, 1]$ defined by $f(x) = \cos(x)$ is a contraction. Observe that $f'(x) = -\sin(x)$, so $|f'(x)| \leq \sin(1)$ for all $x \in [0, 1]$. If we set $\lambda = \sin(1)$, then the Mean Value Theorem guarantees that

$$|\cos(x) - \cos(y)| = \left|f'(c)\right||x - y|$$

for some $c \in (x, y)$, so

$$|\cos(x) - \cos(y)| \leq \lambda|x - y|.$$

Since $\lambda = \sin(1) < 1$, this shows that $\cos(x)$ is a contraction on $[0, 1]$.

Remark 2.8.3. The situation from the previous theorem can be generalized to other differentiable functions. If $f : [a, b] \to \mathbf{R}$ is differentiable and there is a constant $0 < \lambda < 1$ such that $|f'(x)| \leq \lambda$ for all $x \in [a, b]$, then f is a contraction. Again, we can use the Mean Value Theorem: if $x, y \in [a, b]$, then there is a $c \in (x, y)$ such that

$$|f(x) - f(y)| = \left|f'(c)\right||x - y|,$$

so

$$|f(x) - f(y)| \leq \lambda|x - y|.$$

We can even relax things a little further. Recall that a function $f : [a, b] \to \mathbf{R}$ is **Lipschitz** if there is a constant $\alpha > 0$ such that

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

for all $x, y \in [a, b]$. It is immediate that if f is Lipschitz with constant $\alpha < 1$, then f is a contraction.

Before we prove the main theorem, we will make one simple observation about contractions. (Note that the analogous result—with the same proof—holds for Lipschitz functions.)

Proposition 2.8.4. *Let (X, d_1) and (Y, d_2) be metric spaces. If $f : X \to Y$ is a contraction, then it is uniformly continuous.*

Proof. Suppose $f : X \to Y$ is a contraction with constant λ. Let $\varepsilon > 0$ be given, and take $\delta = \varepsilon$. For all $x, y \in X$, if $d_1(x, y) < \delta$ we have

$$d_2(f(x), f(y)) \leq \lambda \cdot d_1(x, y) < \lambda\delta < \delta = \varepsilon$$

since $\lambda < 1$. Thus f is uniformly continuous on X. □

Now we arrive at the main result, which says that a contraction that maps a complete metric space X to itself must have a fixed point. Moreover, the fixed point is unique.

Theorem 2.8.5 (Banach Fixed Point Theorem)**.** *Let X be a complete metric space, and suppose $f : X \to X$ is a contraction. Then f has a unique fixed point, i.e., there is a unique point $x_0 \in X$ satisfying $f(x_0) = x_0$.*

Proof. We begin by choosing a point $x_1 \in X$. Then we define $x_2 = f(x_1)$, $x_3 = f(x_2)$, and in general,

$$x_n = f(x_{n-1})$$

to build a sequence $(x_n)_{n=1}^{\infty}$ in X. We claim that this sequence is Cauchy. As a first step toward proving this assertion, first observe that for all $n \in \mathbf{N}$ we have

$$d(x_n, x_{n+1}) = d(f(x_{n-1}), f(x_n)) \leq \lambda \cdot d(x_{n-1}, x_n).$$

By applying the contraction condition repeatedly, we get

$$\begin{aligned} d(x_n, x_{n+1}) &\leq \lambda \cdot d(x_{n-1}, x_n) \\ &\leq \lambda^2 \cdot d(x_{n-2}, x_{n-1}) \\ &\;\;\vdots \\ &\leq \lambda^{n-1} \cdot d(x_1, x_2). \end{aligned}$$

Now suppose $n, m \in \mathbf{N}$ with $n < m$. Then

$$\begin{aligned} d(x_n, x_m) &\leq d(x_n, x_{n+1}) + d(x_{n+1}, x_{n+2}) + \cdots + d(x_{m-1}, x_m) \\ &\leq \lambda^{n-1} d(x_1, x_2) + \lambda^n d(x_1, x_2) + \cdots + \lambda^{m-2} d(x_1, x_2), \end{aligned}$$

so

$$d(x_n, x_m) \leq d(x_1, x_2) \cdot \sum_{i=n-1}^{m-2} \lambda^i < d(x_1, x_2) \cdot \sum_{i=n-1}^{\infty} \lambda^i.$$

The geometric sum on the right hand side converges to

$$\sum_{i=n-1}^{\infty} \lambda^i = \frac{\lambda^{n-1}}{1-\lambda},$$

so we have

$$d(x_n, x_m) < d(x_1, x_2) \cdot \frac{\lambda^{n-1}}{1-\lambda}.$$

Let $\varepsilon > 0$ be given. Since $\lambda < 1$, we can choose N such that $n \geq N$ implies

$$\lambda^{n-1} < \frac{1-\lambda}{d(x_1, x_2)} \cdot \varepsilon.$$

Then for all $n, m \geq N$, we have $d(x_n, x_m) < \varepsilon$, so $(x_n)_{n=1}^{\infty}$ is a Cauchy sequence.

Since X is complete, the Cauchy sequence $(x_n)_{n=1}^{\infty}$ converges to a point $x_0 \in X$. We claim that x_0 is the desired fixed point. To see this, observe that since f is continuous,

$$f(x_0) = \lim_{n\to\infty} f(x_n) = \lim_{n\to\infty} x_{n+1} = x_0.$$

Thus x_0 is a fixed point for f. To verify uniqueness, suppose $y \in X$ is another fixed point of f. Then

$$d(x_0, y) = d(f(x_0), f(y)) \leq \lambda \cdot d(x_0, y).$$

Since $\lambda < 1$, the only way this inequality can hold is if $d(x_0, y) = 0$. Thus $x_0 = y$, and the fixed point is unique. □

The Contraction Mapping Theorem/Banach Fixed Point Theorem has many fruitful applications. We will begin with a simple one pertaining to equation solving.

Example 2.8.6. We saw earlier that the function $f(x) = \cos(x)$ is a contraction on the complete metric space $X = [0, 1]$. The Banach Fixed Point Theorem thus guarantees that the equation

$$\cos(x) = x$$

has a unique solution in $[0, 1]$. How can we find it? Well, the proof of the theorem is constructive—it tells us how to generate a sequence that converges to the solution. Start with $x_1 = \frac{1}{2}$, say. By repeatedly applying the cosine function, we obtain better and better approximations to the solution. The first few terms (to six decimal places) are:

$$\begin{aligned} x_2 &= 0.877583 & x_{12} &= 0.741827 \\ x_3 &= 0.639012 & x_{13} &= 0.737236 \\ x_4 &= 0.802685 & x_{14} &= 0.740330 \\ x_5 &= 0.694778 & x_{15} &= 0.738246 \\ x_6 &= 0.768196 & x_{16} &= 0.739650 \\ x_7 &= 0.719165 & x_{17} &= 0.738705 \\ x_8 &= 0.752356 & x_{18} &= 0.739341 \\ x_9 &= 0.730081 & x_{19} &= 0.738912 \\ x_{10} &= 0.745120 & x_{20} &= 0.739201 \\ x_{11} &= 0.735006 & x_{21} &= 0.739007 \end{aligned}$$

which are approaching a value of approximately $0.739085\ldots$. (This number—the unique fixed point of the cosine function—is sometimes called the *Dottie number*.)

Notice that the terms in the sequence agree in the first decimal place after $n = 6$, and in the second decimal place after $n = 15$. Of course this is no accident, and we can even use the proof of the Banach Fixed Point Theorem to estimate the error in our approximation at each iteration. Notice that

$$d(x_n, x_0) = d(f(x_{n-1}), f(x_0)) \leq \lambda \cdot d(x_{n-1}, x_0),$$

and if we continue inductively as in the proof, we get

$$d(x_n, x_0) \leq \lambda^{n-1} \cdot d(x_1, x_0).$$

Thus if we have a crude estimate for $d(x_1, x_0)$, we can determine the rate of convergence. It looks like our initial guess was off by no more than 0.25, so we have

$$|x_n - x_0| \leq (\sin(1))^{n-1}(0.25). \tag{2.3}$$

For $n = 6$, we get

$$|x_6 - x_0| \leq 0.0888,$$

so all terms after that point must agree to at least the first decimal place. If we want agreement to two decimal places, we need to go to $n = 19$:

$$|x_{19} - x_0| \leq 0.0094.$$

Note that the terms seem to agree to within 0.01 well before that, however. We could even determine how far we need to go to obtain a given degree of accuracy. Suppose we want $|x_n - x_0| < \varepsilon$ for some $\varepsilon > 0$. Then setting the right hand side of (2.3) to be less than or equal to ε and solving for n, we obtain

$$n \geq \frac{\log(4\varepsilon)}{\log(\sin(1))} + 1.$$

Some of the most impressive and far-reaching applications of the Banach Fixed Point Theorem deal with equations involving unknown functions. Such problems can be solved by setting up a contraction (whose fixed point is the desired solution) on an appropriate space of functions. We begin with a fairly simple example of solving a so-called *functional equation.*

Example 2.8.7. Let $\mathcal{X}$ denote the space of all functions (not necessarily continuous) from $[0,1]$ to $[0,1]$. We can equip $\mathcal{X}$ with the supremum norm,

$$\|f\|_\infty = \sup_{x\in[0,1]} |f(x)|,$$

and it is fairly routine to show that $\mathcal{X}$ is complete with respect to the resulting metric $d(f,g) = \|f-g\|_\infty$. Given a constant μ with $0 \le \mu \le 1$, define a function $M : \mathcal{X} \to \mathcal{X}$ by

$$M(f)(x) = \begin{cases} \frac{1}{2}f(2\mu x) & \text{if } 0 \le x \le \frac{1}{2} \\ 1 - \frac{1}{2}f(2\mu(1-x)) & \text{if } \frac{1}{2} < x \le 1. \end{cases}$$

It is not hard to check that M is a contraction: given $f, g \in \mathcal{X}$, we have

$$\begin{aligned} \sup_{x\in[0,\frac{1}{2}]} |M(f)(x) - M(g)(x)| &= \sup_{x\in[0,\frac{1}{2}]} \left|\tfrac{1}{2}f(2\mu x) - \tfrac{1}{2}g(2\mu x)\right| \\ &= \tfrac{1}{2} \sup_{x\in[0,\frac{1}{2}]} |f(2\mu x) - g(2\mu x)| \\ &= \tfrac{1}{2} \sup_{x\in[0,\mu]} |f(x) - g(x)| \\ &\le \tfrac{1}{2} \sup_{x\in[0,1]} |f(x) - g(x)|. \end{aligned}$$

A similar computation shows that

$$\sup_{x\in[\frac{1}{2},1]} |M(f)(x) - M(g)(x)| \le \tfrac{1}{2} \sup_{x\in[0,1]} |f(x) - g(x)|.$$

It follows that $\|M(f) - M(g)\|_\infty \le \frac{1}{2}\|f-g\|_\infty$ for all $f, g \in \mathcal{X}$, so M is a contraction with constant $\lambda = \frac{1}{2}$. It follows from the Banach Fixed Point Theorem that M has a unique fixed point f, which is the unique solution to the functional equation

$$f(x) = \begin{cases} \frac{1}{2}f(2\mu x) & \text{if } 0 \le x \le \frac{1}{2} \\ 1 - \frac{1}{2}f(2\mu(1-x)) & \text{if } \frac{1}{2} < x \le 1. \end{cases} \tag{2.4}$$

This application of the Banach Fixed Point Theorem appears in [AL17]. The authors needed a function satisfying the functional equation (2.4) in order to solve a combinatorial problem involving patterns realized by certain dynamical systems.

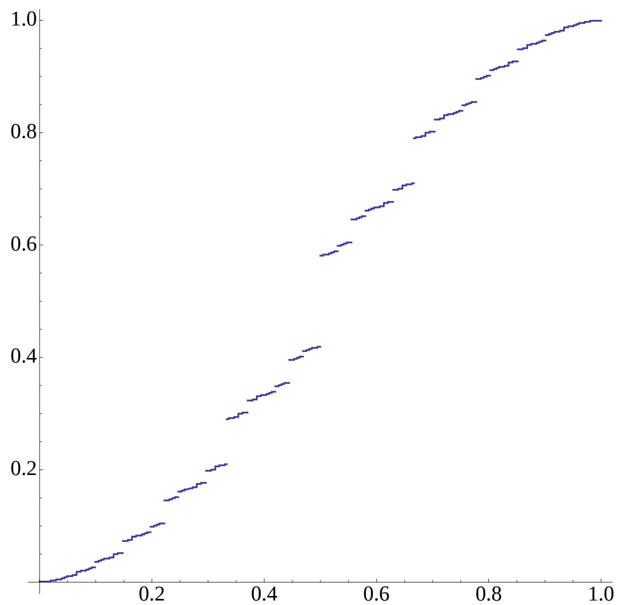

Figure 2.2: A graph of the fixed point f of the contraction M when $\mu = \frac{3}{4}$.

To gain some appreciation for the true power of the Banach Fixed Point Theorem, we will now use it to prove a familiar theorem on the existence and uniqueness of solutions to differential equations. In a first course on ordinary differential equations, one considers initial value problems of the form

$$y'(t) = f(t, y(t)), \quad y(t_0) = y_0, \tag{2.5}$$

where y is an unknown (differentiable) function defined on some interval, and f is defined on some neighborhood of the point (t_0, y_0) in $\mathbf{R}^2$ (and is usually assumed to be at least continuous there). You may have even seen a theorem that gives conditions guaranteeing the existence and uniqueness of solutions to an IVP like the one in (2.5). This theorem is often called the Picard-Lindelöf theorem, and it is usually stated more or less as follows in introductory differential equations textbooks: if both f and the partial derivative $\frac{\partial f}{\partial y}$ are continuous on an open rectangle of the form

$$\mathcal{R} = \big\{(t, y) \in \mathbf{R}^2 : -a < t - t_0 < a,\ -b < y - y_0 < b\big\},$$

then there is a unique solution to the IVP (2.5) on some interval containing t_0. We will actually prove a considerably stronger version of this theorem—instead of the continuity hypothesis on $\frac{\partial f}{\partial y}$, we need only require that f is "uniformly Lipschitz" in the variable y. We will also obtain an estimate for the size of the interval on which the solution is defined.

Theorem 2.8.8 (Picard-Lindelöf). *Suppose $U \subseteq \mathbf{R}^2$ is open and $f : U \to \mathbf{R}$ is continuous, and let $(t_0, y_0) \in U$. Assume there exist $a, b \in \mathbf{R}$ such that the rectangle*

$$\mathcal{R} = \{(t, y) \in \mathbf{R}^2 : |t - t_0| \le a, |y - y_0| \le b\}$$

is contained in U and f is uniformly Lipschitz on $\mathcal{R}$, meaning there is a constant $\alpha > 0$ such that

$$|f(t, x) - f(t, y)| \le \alpha |x - y|$$

for all $(t, x), (t, y) \in \mathcal{R}$. Suppose $M > 0$ satisfies $|f(t, y)| \le M$ for all $(t, y) \in \mathcal{R}$, and set

$$a_* = \min\{a, b/M\}.$$

If we let $I = [t_0 - a_*, t_0 + a_*]$, then there exists a unique function $y : I \to \mathbf{R}$ satisfying the initial value problem

$$y'(t) = f(t, y(t)), \quad y(t_0) = y_0 \tag{2.6}$$

whenever $|t - t_0| \le a_*$.

The key to the proof is the application of a method known as "Picard iteration" to zero in on a solution. We will shortly see that Picard iteration is simply an example of the process we used in the proof of the Banach Fixed Point Theorem—it involves iterating a carefully-chosen contraction on a certain function space. The first step in building this contraction is to convert the given differential equation to an integral equation instead.

Lemma 2.8.9. *A function $y(t)$ satisfies the initial value problem* (2.6) *if and only if it satisfies the Volterra integral equation*

$$y(t) = y_0 + \int_{t_0}^{t} f(s, y(s))\, ds. \tag{2.7}$$

Proof. Suppose first that y satisfies (2.7). Then differentiating via the Fundamental Theorem of Calculus yields

$$y'(t) = f(t, y(t)),$$

and we also have

$$y(t_0) = y_0 + \int_{t_0}^{t_0} f(s, y(s))\, ds = y_0 + 0 = y_0.$$

Thus y satisfies the IVP. On the other hand, suppose y satisfies (2.6). By integrating both sides of the differential equation we get

$$\int_{t_0}^{t} y'(s)\, ds = \int_{t_0}^{t} f(s, y(s))\, ds,$$

and the left hand side evaluates to

$$\int_{t_0}^{t} y'(s)\,ds = y(t) - y(t_0) = y(t) - y_0.$$

Putting these two equations together, we get

$$y(t) - y_0 = \int_{t_0}^{t} f(s, y(s))\,ds,$$

so (2.7) holds. □

The next step is to define a contraction in terms of an integral operator, which will have the solution to our Volterra equation as its fixed point. We first need a metric space upon which to define this contraction. To this end, we let $\mathcal{X}$ be the set of all continuous functions $y : I \to [y_0 - b, y_0 + b]$, and our proposed metric on $\mathcal{X}$ is defined by

$$d(x, y) = \sup_{t \in I} |x(t) - y(t)| e^{-2\alpha|t-t_0|}. \tag{2.8}$$

Of course we need to know that d is a metric, and that $\mathcal{X}$ is complete with respect to it. None of this is too hard to do, particularly because d is just a scaled-down version of the metric that we obtain from the usual supremum norm on $C(I, \mathbf{R})$.

Lemma 2.8.10. *The function d defined in* (2.8) *is a metric on $\mathcal{X}$, and $\mathcal{X}$ is complete with respect to d.*

Proof. Since $e^{-2\alpha|t-t_0|}$ is never zero, it is fairly clear that d is positive definite. It is also easy to see that d is symmetric. For the triangle inequality, observe that if $x, y, z \in \mathcal{X}$, then

$$|x(t) - y(t)| e^{-2\alpha|t-t_0|} \le |x(t) - z(t)| e^{-2\alpha|t-t_0|} + |z(t) - y(t)| e^{-2\alpha|t-t_0|}$$

for all $t \in I$. Taking the supremum of the right hand side yields

$$|x(t) - y(t)| e^{-2\alpha|t-t_0|} \le d(x, z) + d(z, y).$$

We can now safely take the supremum of the left hand side to obtain

$$d(x, y) \le d(x, z) + d(z, y).$$

Therefore, d is a metric.

The completeness takes a little more work, but it is still not hard. Suppose $(x_n)_{n=1}^{\infty}$ is a sequence in $\mathcal{X}$ which is Cauchy with respect to d. In fact, we claim this sequence is Cauchy with respect to the simpler metric

$$\rho(x, y) = \|x - y\|_{\infty}.$$

Let $\varepsilon > 0$ be given, and choose $N \in \mathbf{N}$ such that $n, m \geq N$ implies

$$d(x_n, x_m) < \frac{\varepsilon}{e^{2\alpha(a_* - t_0)}}$$

Then for all $t \in I$ we have

$$|x_n(t) - x_m(t)|e^{-2\alpha|t-t_0|} < \frac{\varepsilon}{e^{2\alpha(a_* - t_0)}},$$

or

$$|x_n(t) - x_m(t)| < \frac{\varepsilon}{e^{2\alpha(a_* - t_0)}} \cdot e^{2\alpha|t-t_0|} \leq \varepsilon$$

whenever $n, m \geq N$. Thus $(x_n)_{n=1}^{\infty}$ is Cauchy with respect to the supremum norm, so it converges to a continuous function x defined on I by Theorem 2.3.15. Notice that x necessarily takes values in $[y_0 - b, y_0 + b]$, so $x \in \mathcal{X}$.

All that remains is to see that $x_n \to x$ with respect to d. This is almost immediate, since

$$|x_n(t) - x(t)|e^{-2\alpha|t-t_0|} \leq |x_n(t) - x(t)|$$

for all $t \in I$, which implies that $d(x_n, x) \leq \|x_n - x\|_{\infty}$. It follows that $x_n \to x$ with respect to d. □

Now we define an operator $T : \mathcal{X} \to \mathcal{X}$ in what seems like the most natural way possible, given our problem:

$$T(y) = y_0 + \int_{t_0}^{t} f(t, y(s))\, ds$$

Note that a function y will be a solution to the Volterra equation precisely when it is a fixed point of T. Therefore, we need only verify that T is a contraction, and the Banach Fixed Point Theorem will do the rest of the work for us.

Lemma 2.8.11. *The function $T : \mathcal{X} \to \mathcal{X}$ defined above is a contraction on $(\mathcal{X}, d)$.*

Proof. It is not entirely obvious from the definition that T maps into $\mathcal{X}$, so we do need to check it. If $y \in \mathcal{X}$, the Fundamental Theorem of Calculus guarantees that $T(y)$ is continuous. Moreover, we have

$$|T(y) - y_0| = \left| \int_{t_0}^{t} f(t, y(s))\, ds \right| \leq |t - t_0| \cdot \sup_{t \in I} |f(t, y(t))| \leq a_* M \leq b,$$

so the function $T(y)$ takes values in the interval $[y_0 - b, y_0 + b]$.

Now we can get on with showing that T is a contraction. Let $y_1, y_2 \in \mathcal{X}$. Then for each $t \in I$ we have

$$|T(y_1)(t) - T(y_2)(t)| = \left| \int_{t_0}^{t} f(s, y_1(s))\, ds - \int_{t_0}^{t} f(s, y_2(s))\, ds \right|$$

$$
\begin{aligned}
&\leq \int_{t_0}^{t} |f(s, y_1(s)) - f(s, y_2(s))|\, ds \\
&\leq \alpha \int_{t_0}^{t} |y_1(s) - y_2(s)|\, ds \\
&= \alpha \int_{t_0}^{t} |y_1(s) - y_2(s)| e^{-2\alpha|s-t_0|} e^{2\alpha|s-t_0|}\, ds \\
&\leq \alpha \cdot d(y_1, y_2) \int_{t_0}^{t} e^{2\alpha|s-t_0|}\, ds \\
&= \tfrac{1}{2} d(y_1, y_2)(e^{2\alpha|t-t_0|} - 1) \\
&\leq \tfrac{1}{2} d(y_1, y_2) e^{2\alpha|t-t_0|}.
\end{aligned}
$$

Hence

$$|T(y_1)(t) - T(y_2)(t)| e^{-2\alpha|t-t_0|} \leq \tfrac{1}{2} d(y_1, y_2),$$

and taking suprema yields

$$d(T(y_1), T(y_2)) \leq \tfrac{1}{2} d(y_1, y_2).$$

Therefore, T is a contraction. □

Proof of the Picard-Lindelöf Theorem. Most of the work here is already done, and we need only put it all together. Since $T : \mathcal{X} \to \mathcal{X}$ is a contraction, it has a unique fixed point y by the Banach Fixed Point Theorem. It is clear that this function is a solution to the Volterra equation (2.7), hence to the initial value problem (2.6) by Lemma 2.8.9. Moreover, it is the unique solution to this IVP. □

Exercises for Section 2.8

Exercise 2.8.1 ([HS91], Exercise 5.2.6)**.** Let X be a complete metric space, and suppose $f : X \to X$ is a function such that

$$f^n = \underbrace{f \circ f \circ \cdots \circ f}_{n \text{ times}}$$

is a contraction for some positive integer n. Prove that f has a unique fixed point.

Exercise 2.8.2. Show that there is a unique function $f : [0,1] \to [0,1]$ satisfying the functional equation

$$f(x) = \begin{cases} \frac{1}{2} f(3x) & \text{if } 0 \leq x \leq \frac{1}{3} \\ \frac{1}{2} & \text{if } \frac{1}{3} < x < \frac{2}{3} \\ \frac{1}{2} + \frac{1}{2} f(3x - 2) & \text{if } \frac{2}{3} \leq x \leq 1. \end{cases}$$

This function is commonly known as the **Cantor function**.

Chapter 3

Sequences and Series of Functions

As part of our initial foray into the theory of metric spaces in the previous chapter, recall that we studied the vector space $C([a,b],\mathbf{R})$ consisting of all continuous real-valued functions defined on the interval $[a,b]$ equipped with the supremum norm:

$$\|f\|_\infty = \sup_{x\in[a,b]} |f(x)|.$$

Indeed, we showed that $C([a,b],\mathbf{R})$ is complete with respect to the metric induced by the norm $\|\cdot\|_\infty$. Roughly speaking, we will now replace $[a,b]$ with an arbitrary metric space X and consider the set

$$C(X,\mathbf{R}) = \big\{f : X \to \mathbf{R} \mid f \text{ is continuous}\big\}.$$

Though we will focus only on real-valued functions initially, we will eventually allow for complex-valued functions as well:

$$C(X) = C(X,\mathbf{C}) = \big\{f : X \to \mathbf{C} \mid f \text{ is continuous}\big\}.$$

There is one issue that prevents us from simply translating the facts we have already established for $C([a,b],\mathbf{R})$ *mutatis mutandi* to $C(X,\mathbf{R})$. Notice that if X is an arbitrary metric space, continuous functions on X need not be bounded. Thus the supremum norm does not make sense in this setting, and we cannot simply bring metric space theory to bear on $C(X,\mathbf{R})$. Nevertheless, we will see that many of our results for $C([a,b],\mathbf{R})$ do carry over to $C(X,\mathbf{R})$ for more general X once we have phrased them a little differently. Moreover, if we find ourselves in the case where X is compact then the supremum norm is well-defined, and $C(X,\mathbf{R})$ is a complete metric space. We will then be able to obtain some very nice analytic results in this compact setting.

3.1 Sequences of Functions

Throughout this section, let (X, d) denote a metric space. We will begin our investigation by studying sequences of real-valued functions on X, where the functions are not assumed to be continuous in general.

Let $(f_n)_{n=1}^{\infty}$ be a sequence of functions $f_n : X \to \mathbf{R}$. We have seen previously (namely in the proof of Theorem 2.3.15) that if the sequence $(f_n(x))_{n=1}^{\infty}$ of real numbers converges for each $x \in X$, then we can define a function $f : X \to \mathbf{R}$ by

$$f(x) = \lim_{n\to\infty} f_n(x).$$

The function f is called the **pointwise limit** of the sequence $(f_n)_{n=1}^{\infty}$, since it is defined by plugging individual points into the functions $(f_n)_{n=1}^{\infty}$. This observation leads us to one notion of convergence for a sequence of functions.

Definition 3.1.1. We say that a sequence $(f_n)_{n=1}^{\infty}$ of real-valued functions on X **converges pointwise** to a function $f : X \to \mathbf{R}$ if $f_n(x) \to f(x)$ for each $x \in X$.

Example 3.1.2. For each $n \in \mathbf{N}$, define $f_n : [0, 1] \to \mathbf{R}$ by $f(x) = x^n$. Notice that if $0 \le x < 1$, then $x^n \to 0$ as $n \to \infty$. If $x = 1$, then we have $x^n = 1$ for all n, so $x^n \to 1$. Therefore, the sequence $(f_n)_{n=1}^{\infty}$ converges pointwise to the function

$$f(x) = \begin{cases} 0 & \text{if } 0 \le x < 1 \\ 1 & \text{if } x = 1. \end{cases}$$

Notice that in the previous example, each f_n is continuous, but the pointwise limit f is not continuous. In other words, *the pointwise limit of a sequence of continuous functions need not be continuous!*

It makes sense to ask what else could go wrong as far as pointwise convergence is concerned. For example, must the pointwise limit of a sequence of differentiable functions be differentiable? Of course the answer is no, since the sequence $(f_n)_{n=1}^{\infty}$ from Example 3.1.2 converges pointwise to a non-differentiable function. Perhaps we can refine our question, however. Assume $f_n \to f$ pointwise, each f_n is differentiable, and f is differentiable. Must $f_n' \to f'$? Unfortunately, the answer is still no.

Example 3.1.3. Consider the sequence $(f_n)_{n=1}^{\infty}$, where $f_n : \mathbf{R} \to \mathbf{R}$ is defined by

$$f_n(x) = \frac{\sin nx}{\sqrt{n}}.$$

It is easy to show that $f_n \to 0$ pointwise. Let $\varepsilon > 0$ be given, and choose N such that $n \ge N$ implies $\frac{1}{\sqrt{n}} < \varepsilon$. Then

$$|f_n(x)| = \frac{|\sin nx|}{\sqrt{n}} \le \frac{1}{\sqrt{n}} < \varepsilon$$

for all $x \in \mathbf{R}$ when $n \geq N$. Furthermore, each f_n is differentiable on $\mathbf{R}$ with

$$f_n'(x) = \frac{n \cos nx}{\sqrt{n}} = \sqrt{n} \cos nx,$$

and the pointwise limit $f \equiv 0$ is certainly differentiable. However, notice that $f_n'(x) \to \infty$ as $n \to \infty$ for all $x \in \mathbf{R}$. Thus the sequence $(f_n')_{n=1}^{\infty}$ is not even convergent!

Perhaps we will have better luck with Riemann integrability. Unfortunately, we saw in the introductory chapter that a pointwise convergent sequence of Riemann integrable functions need not converge to a Riemann integrable function. Begin by enumerating the rational numbers in $[0,1]$:

$$\mathbf{Q} \cap [0,1] = \{r_1, r_2, r_3, \ldots\}.$$

Now define $f_n : [0,1] \to \mathbf{R}$ by

$$f_n(x) = \begin{cases} 1 & \text{if } x = r_i \text{ for some } i \leq n \\ 0 & \text{otherwise.} \end{cases}$$

If $x \in \mathbf{Q}$, then $f_n(x) \to 1$. Indeed, if $x \in \mathbf{Q}$ then $x = r_m$ for some m, so $f_n(x) = 1$ for all $n \geq m$. On the other hand, if $x \notin \mathbf{Q}$, then $f_n(x) = 0$ for all n. Therefore, $(f_n)_{n=1}^{\infty}$ converges pointwise to the function

$$\chi_{\mathbf{Q}}(x) = \begin{cases} 1 & \text{if } x \in \mathbf{Q} \\ 0 & \text{if } x \notin \mathbf{Q}, \end{cases}$$

which is not Riemann integrable. However, each f_n is Riemann integrable since it has only a finite number of discontinuities.

What if $(f_n)_{n=1}^{\infty}$ is a sequence of Riemann integrable functions that converges pointwise to a function f, and we assume f is Riemann integrable? Is it necessarily the case that $\int f_n \to \int f$? Unfortunately, the answer is still no.

Example 3.1.4. For each $n \in \mathbf{N}$, define $f_n : [0,1] \to \mathbf{R}$ by

$$f_n(x) = \begin{cases} 2^{2n} x & \text{if } 0 \leq x \leq \frac{1}{2^n} \\ 2^{2n}\left(\frac{1}{2^{n-1}} - x\right) & \text{if } \frac{1}{2^n} < x \leq \frac{1}{2^{n-1}} \\ 0 & \text{if } \frac{1}{2^{n-1}} < x \leq 1. \end{cases}$$

The first few terms in the sequence are plotted below.

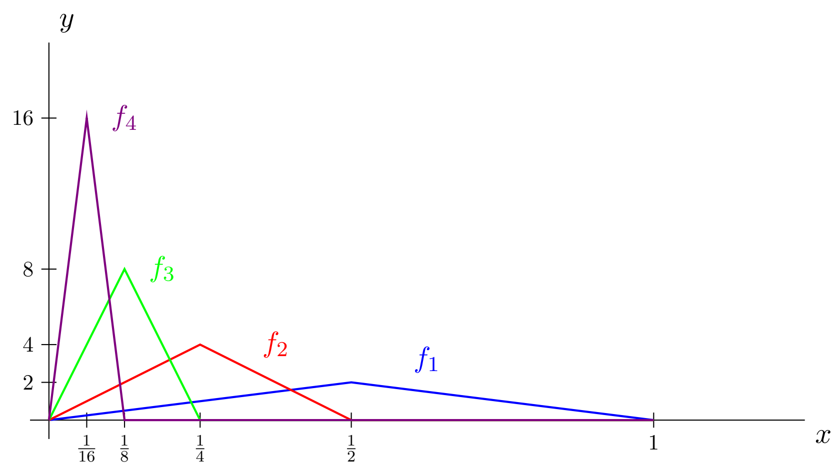

Observe that for each n, the region under the graph of f_n is a triangle with height 2^n and width $\frac{1}{2^{n-1}}$, so

$$\int_0^1 f_n(x)\,dx = \left(\tfrac{1}{2}\right)(2^n)\left(\tfrac{1}{2^{n-1}}\right) = 1$$

for all n. Also, notice that $f_n \to 0$ pointwise. To see this, let $x \in [0,1]$ and choose N such that $\frac{1}{2^{N-1}} < x$. Then $f_n(x) = 0$ for all $n \geq N$, so $f_n(x) \to 0$. Thus $\int \lim f_n = 0$, but $\int f_n \to 1$ as $n \to \infty$.

In each example, we have encountered issues when trying to interchange two limiting processes. For example, the question of whether a sequence of continuous functions converges to a continuous function amounts to asking if

$$\lim_{x \to x_0} \lim_{n \to \infty} f_n(x) = \lim_{n \to \infty} \lim_{x \to x_0} f_n(x)$$

for each $x_0 \in X$. Similarly, to show that f is differentiable and $f_n' \to f'$ requires something like

$$\lim_{n \to \infty} \lim_{x \to x_0} \frac{f_n(x) - f_n(x_0)}{x - x_0} = \lim_{x \to x_0} \lim_{n \to \infty} \frac{f_n(x) - f_n(x_0)}{x - x_0}$$

to hold. Lastly, for integrals we asked whether

$$\lim_{n \to \infty} \int f_n = \int \lim_{n \to \infty} f_n.$$

We arrived at a negative answer for each of these questions. Therefore, we will need to add some hypotheses in order to obtain positive results for the continuity, differentiability, and integrability of the limit of a sequence of functions. Believe it or not, pointwise convergence is important (and useful!). We just need to add some conditions to make things work the way we want.

When we begin our study of measure theory in the next chapter, we will obtain some very powerful results that allow us to interchange limits and integrals under fairly mild hypotheses. However, in order to say anything of import right now, we must consider a vast strengthening of pointwise convergence, called *uniform* convergence. This condition will rectify some of the issues that we have seen regarding continuity, differentiability, and Riemann integrability.

Definition 3.1.5. A sequence $(f_n)_{n=1}^{\infty}$ of real-valued functions on X is said to **converge uniformly** to a function $f : X \to \mathbf{R}$ if for any $\varepsilon > 0$, there is an $N \in \mathbf{N}$ such that $n \geq N$ implies

$$|f_n(x) - f(x)| < \varepsilon$$

for all $x \in X$.

Notice that uniform convergence certainly implies pointwise convergence, though the definition of uniform convergence requires an N that works for *all values of* x *simultaneously.* In the case that $X \subseteq \mathbf{R}$, the picture one should have in mind is that of an "ε-snake" around the graph of the limit function f. The sequence $(f_n)_{n=1}^{\infty}$ converges uniformly to f precisely when the graphs of the f_n eventually all lie inside the ε-snake.

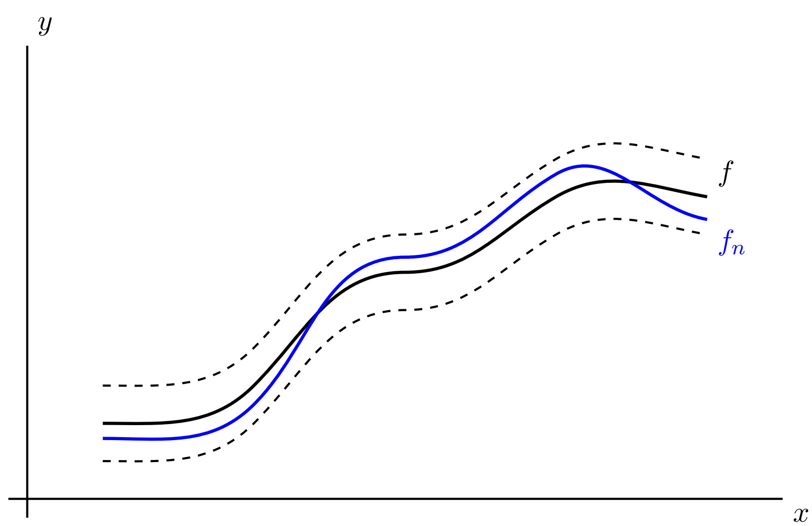

Remark 3.1.6. The definition of uniform convergence should be reminiscent of convergence with respect to the supremum norm. Indeed, if X is a compact metric space, then all continuous functions are bounded, and $C(X, \mathbf{R})$ is a normed vector space under

$$\|f\|_{\infty} = \sup_{x \in X} |f(x)|.$$

It is not hard to see that a sequence $(f_n)_{n=1}^{\infty}$ in $C(X, \mathbf{R})$ converges uniformly if and only if $\|f_n - f\|_{\infty} \to 0$ as $n \to \infty$. If X is not compact, then we cannot simply appeal to the supremum norm. We will instead need to develop appropriate substitutes for some of the tools that this norm would afford us.

We will now revisit some of our earlier examples and study them through the lens of uniform convergence.

Example 3.1.7. The sequence $(f_n)_{n=1}^{\infty}$ from Example 3.1.3, which was defined by

$$f_n(x) = \frac{\sin nx}{\sqrt{n}}$$

for $x \in \mathbf{R}$, converges uniformly to 0. Indeed, we observed in Example 3.1.3 that for each n, $|f_n(x)| \leq \frac{1}{\sqrt{n}}$ for all $x \in \mathbf{R}$. Since $\frac{1}{\sqrt{n}} \to 0$ as $n \to \infty$, it follows that $f_n \to 0$ uniformly.

Example 3.1.8. Let $(f_n)_{n=1}^{\infty}$ be the sequence defined on $[0,1]$ in Example 3.1.4. We observed earlier that $f_n \to 0$ pointwise on $[0,1]$, but we claim that $(f_n)_{n=1}^{\infty}$ does not converge uniformly. Indeed, observe that for each n,

$$\|f_n\|_{\infty} = \sup_{x \in [0,1]} |f_n(x)| = 2^n.$$

Thus $\|f_n\|_{\infty} \to \infty$ as $n \to \infty$, so the sequence does not converge uniformly to 0.

Example 3.1.9. The sequence defined on $[0,1]$ by $f_n(x) = x^n$ (as in Example 3.1.2) does not converge uniformly. Fix $n \in \mathbf{N}$ and observe that

$$\|f_n - f_{2n}\|_{\infty} = \sup_{x \in [0,1]} \left| x^n - x^{2n} \right| = \sup_{t \in [0,1]} \left| t - t^2 \right| = \frac{1}{4}.$$

If $m \geq 2n$, then $f_m \leq f_{2n} \leq f_n$, so

$$\|f_n - f_m\|_{\infty} \geq \frac{1}{4}.$$

This implies that $(f_n)_{n=1}^{\infty}$ is not a Cauchy sequence in $C([0,1], \mathbf{R})$, so it cannot converge uniformly.

In the last example, we used fact that a sequence in $C([0,1], \mathbf{R})$ converges uniformly if and only if it is Cauchy with respect to $\|\cdot\|_{\infty}$. Even if we are unable to appeal to the supremum norm (i.e., if X is not compact), we can still specify a Cauchy-like criterion that characterizes whether a sequence converges or not.

Theorem 3.1.10 (Cauchy criterion for uniform convergence)**.** *A sequence $(f_n)_{n=1}^{\infty}$ of real-valued functions on X converges uniformly if and only if for each $\varepsilon > 0$, there exists an N such that $n, m \geq N$ implies*

$$|f_n(x) - f_m(x)| < \varepsilon$$

for all $x \in X$.

Proof. Assume first that $f_n \to f$ uniformly, and let $\varepsilon > 0$ be given. Then there exists $N \in \mathbf{N}$ such that $n \geq N$ implies

$$|f_n(x) - f(x)| < \tfrac{\varepsilon}{2}$$

for all $x \in X$. If $n, m \geq N$, then we have

$$|f_n(x) - f_m(x)| \leq |f_n(x) - f(x)| + |f_m(x) - f(x)| < \varepsilon,$$

so the Cauchy criterion is satisfied.

Conversely, suppose $(f_n)_{n=1}^{\infty}$ satisfies the Cauchy criterion. Then for each $x \in X$, the sequence $(f_n(x))_{n=1}^{\infty}$ is Cauchy in $\mathbf{R}$, hence convergent. Thus we can define a function $f : X \to \mathbf{R}$ by

$$f(x) = \lim_{n \to \infty} f_n(x)$$

for all $x \in X$. By construction, $f_n \to f$ pointwise. We claim that the convergence is actually uniform. Let $\varepsilon > 0$ be given, and use the Cauchy criterion to find $N \in \mathbf{N}$ such that $n, m \geq N$ implies

$$|f_n(x) - f_m(x)| < \tfrac{\varepsilon}{2}.$$

Now hold n fixed, and let $m \to \infty$:

$$|f_n(x) - f(x)| \leq \tfrac{\varepsilon}{2} < \varepsilon.$$

Hence $f_n \to f$ uniformly. □

We will now see that uniform convergence rectifies some of the issues that appeared in our earlier examples. First, we address continuity. The proof is eerily similar to part of the proof of Theorem 2.3.15.

Theorem 3.1.11. *Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of continuous real-valued functions on X that converges uniformly to a function f. Then f is continuous.*

Proof. Fix $x_0 \in X$, and let $\varepsilon > 0$ be given. Since $f_n \to f$ uniformly, there exists N such that

$$|f_N(x) - f(x)| < \tfrac{\varepsilon}{3}$$

for all $x \in X$. Since f_N is continuous at x_0, there exists $\delta > 0$ such that $d(x, x_0) < \delta$ implies $|f_N(x) - f_N(x_0)| < \frac{\varepsilon}{3}$. Thus

$$\begin{aligned} |f(x) - f(x_0)| &\leq |f(x) - f_N(x)| + |f_N(x) - f_N(x_0)| + |f_N(x_0) - f(x_0)| \\ &< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \\ &= \varepsilon \end{aligned}$$

when $d(x, x_0) < \delta$. Therefore, f is continuous at x_0. Since $x_0 \in X$ was arbitrary, f is continuous on X. □

In the case that X is compact, we can combine the last two results to obtain a generalization of Theorem 2.3.15.

Corollary 3.1.12. *If X is a compact metric space, then $C(X, \mathbf{R})$ is complete with respect to the supremum norm.*

Remark 3.1.13. It is worth noting that exactly same sort of argument shows that $C(X) = C(X, \mathbf{C})$ is complete with respect to the supremum norm.

The converse of Theorem 3.1.11 is not true. That is, a sequence of continuous functions may converge pointwise to a continuous function without the convergence being uniform. In fact, the sequence from Example 3.1.4 converges pointwise to 0, though it does not converge uniformly. However, we can prove a partial converse to Theorem 3.1.11 when X is compact and the sequence is monotone. This result is known as Dini's theorem.

Theorem 3.1.14 (Dini)**.** *Let X be a compact metric space, and suppose $(f_n)_{n=1}^{\infty}$ is a sequence of continuous functions on X that converges pointwise to a continuous function f, and for each n we have*

$$f_n(x) \geq f_{n+1}(x)$$

for all $x \in X$. Then $f_n \to f$ uniformly.

Proof. To simplify matters, we first define a sequence $(g_n)_{n=1}^{\infty}$ by $g_n = f_n - f$ for each n. Then $g_n \to 0$ pointwise and $g_n \geq g_{n+1}$ for all n. To prove the theorem, it will suffice to show that $g_n \to 0$ uniformly.

Let $\varepsilon > 0$ be given, and for each $n \in \mathbf{N}$ put

$$K_n = \{x \in X : g_n(x) \geq \varepsilon\}.$$

Notice that $K_n = g_n^{-1}([\varepsilon, \infty))$. Since g_n is continuous for all n, each K_n is closed, hence compact. Moreover, if $x \in K_{n+1}$ then

$$g_n(x) \geq g_{n+1}(x) \geq \varepsilon,$$

so $x \in K_n$ as well. Thus $K_n \supseteq K_{n+1}$, so we have a nested sequence of compact sets:

$$K_1 \supseteq K_2 \supseteq K_3 \supseteq \cdots$$

For each $x \in X$, we know $g_n(x) \to 0$, so there exists n_x such that $x \notin K_{n_x}$. Thus

$$\bigcap_{n=1}^{\infty} K_n = \emptyset.$$

But a nested sequence of nonempty compact sets has nonempty intersection, so there must be an N such that $K_N = \emptyset$. This is equivalent to saying that $g_N(x) < \varepsilon$ for all $x \in X$. Furthermore, $K_n = \emptyset$ for all $n \geq N$, so we have $g_n(x) < \varepsilon$ for all $x \in X$ when $n \geq N$. It follows that $g_n \to 0$ uniformly. □

Remark 3.1.15. It is worth noting that we could have required the sequence in Dini's theorem to be increasing. In fact, by replacing f_n with $-f_n$, we can immediately obtain that result from the one that we have proven.

Example 3.1.16. The compactness assumption in Dini's theorem is essential. To see this, take $X = [0, \infty)$ and define $f_n : X \to \mathbf{R}$ for $n = 0, 1, 2, \ldots$ by

$$f_n(x) = \begin{cases} 0 & \text{if } x \leq n \\ \tanh(x - n) & \text{if } x > n. \end{cases}$$

Then each f_n is continuous, $f_n \to 0$ pointwise, and $f_n \geq f_{n+1}$ for each n. However, the convergence is not uniform. To see this, notice that

$$\sup_{x \geq 0} |f_n(x)| = \sup_{x \geq 0} \tanh(x - n) = 1.$$

Thus if we take $\varepsilon = \frac{1}{2}$, for any $n \in \mathbf{N}$ we can find $x \in [0, \infty)$ such that $f_n(x) \geq \varepsilon$.

Now we consider the role of uniform convergence in the study of differentiability and integrability. For this to really make sense, we will take $X = [a, b]$ to be a closed interval in the following two theorems.

Theorem 3.1.17. *Let $(f_n)_{n=1}^{\infty}$ be a sequence of differentiable functions on $[a, b]$. Suppose the sequence $(f_n')_{n=1}^{\infty}$ converges uniformly and that there exists a point $z \in [a, b]$ such that the sequence $(f_n(z))_{n=1}^{\infty}$ converges. Then $(f_n)_{n=1}^{\infty}$ converges uniformly to a differentiable function f, and*

$$f' = \lim_{n \to \infty} f_n'.$$

Proof. We first need to show that $(f_n)_{n=1}^{\infty}$ converges uniformly. To do so, we will use the Cauchy criterion for uniform convergence. Let $\varepsilon > 0$ be given. Since $(f_n(z))_{n=1}^{\infty}$ converges and $(f'_n)_{n=1}^{\infty}$ converges uniformly, we can find an N such that $n, m \geq N$ implies

$$|f_n(z) - f_m(z)| < \frac{\varepsilon}{2}$$

and

$$\left|f'_n(x) - f'_m(x)\right| < \frac{\varepsilon}{2(b-a)}$$

for all $x \in [a,b]$. Given $n, m \geq N$, define $g = f_n - f_m$. Then g is differentiable on $[a,b]$ with

$$\left|g'(x)\right| = \left|f'_n(x) - f'_m(x)\right| < \frac{\varepsilon}{2(b-a)}$$

for all $x \in [a,b]$. Thus the Mean Value Theorem implies

$$|g(x) - g(y)| \leq \frac{\varepsilon}{2(b-a)}|x-y| \leq \frac{\varepsilon}{2(b-a)}(b-a) = \frac{\varepsilon}{2} \tag{3.1}$$

for all $x, y \in [a,b]$. Therefore,

$$\begin{aligned} |f_n(x) - f_m(x)| &\leq |f_n(x) - f_m(x) + f_n(z) - f_m(z)| + |f_n(z) - f_m(z)| \\ &= |g(x) - g(z)| + |f_n(z) - f_m(z)| \\ &< \varepsilon \end{aligned}$$

for all $x \in [a,b]$, so the uniform Cauchy criterion holds. Hence $(f_n)_{n=1}^{\infty}$ converges uniformly to a function $f : [a,b] \to \mathbf{R}$.

We must now show that f is differentiable on $[a,b]$. Fix a point $x_0 \in [a,b]$ and define

$$\varphi_n(x) = \frac{f_n(x) - f_n(x_0)}{x - x_0} \quad n = 1, 2, \ldots$$

and

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

for $x \neq x_0$. Since each f_n is differentiable, we of course have

$$\lim_{x \to x_0} \varphi_n(x) = f'_n(x_0),$$

so we can set $\varphi_n(x_0) = f'_n(x_0)$. Thus φ_n is defined on all of $[a,b]$. Also, the computation in (3.1) shows that if we choose N so $|f'_n(x) - f'_m(x)| < \varepsilon$ for all $x \in [a,b]$ whenever $n, m \geq N$, then

$$\begin{aligned} |\varphi_n(x) - \varphi_m(x)| &= \frac{|f_n(x) - f_n(x_0) - f_m(x_0) + f_m(x_0)|}{|x - x_0|} \\ &\leq \frac{\varepsilon|x - x_0|}{|x - x_0|} \end{aligned}$$

$$= \varepsilon$$

for all $x \neq x_0$. However, we also have

$$|\varphi_n(x_0) - \varphi_m(x_0)| = \left|f_n'(x_0) - f_m'(x_0)\right| < \varepsilon,$$

so the sequence $(\varphi_n)_{n=1}^\infty$ converges uniformly. Indeed, it is easy to see that $\varphi_n \to \varphi$ pointwise when $x \neq x_0$, so $\varphi_n \to \varphi$ uniformly for $x \neq x_0$. It follows that[1]

$$\lim_{x\to x_0} \varphi(x) = \lim_{x\to x_0} \lim_{n\to\infty} \varphi_n(x) = \lim_{n\to\infty} \lim_{x\to x_0} \varphi_n(x) = \lim_{n\to\infty} f_n'(x_0).$$

In particular, the limit on the left hand side exists, which is exactly what it means for f to be differentiable at x_0. Furthermore,

$$f'(x_0) = \lim_{x\to x_0} \varphi(x) = \lim_{n\to\infty} f_n'(x_0).$$

Since $x_0 \in [a, b]$ was arbitrary, it follows that $f_n' \to f'$ uniformly. □

Thankfully, the proof of the corresponding result for integrals is far less delicate than the one for derivatives.

Theorem 3.1.18. *Suppose $(f_n)_{n=1}^\infty$ is a sequence of Riemann integrable functions on $[a, b]$ that converges uniformly to a function f. Then f is Riemann integrable, and*

$$\int_a^b f(x)\,dx = \lim_{n\to\infty} \int_a^b f_n(x)\,dx.$$

Proof. Let $\varepsilon > 0$ be given. Since $f_n \to f$ uniformly, we can find an N such that

$$|f_N(x) - f(x)| < \frac{\varepsilon}{3(b-a)}$$

for all $x \in [a, b]$. Since f_N is Riemann integrable, there is a partition P of $[a, b]$ such that

$$U(f_N, P) - L(f_N, P) < \frac{\varepsilon}{3}.$$

Since $\|f - f_N\|_\infty < \frac{\varepsilon}{3}$, we have

$$U(f - f_N, P) \leq (b-a) \cdot \sup_{x\in[a,b]} \big(f(x) - f_N(x)\big) < \frac{\varepsilon}{3(b-a)}(b-a) = \frac{\varepsilon}{3}.$$

Similarly,

$$L(f - f_N, P) \geq (b-a) \cdot \inf_{x\in[a,b]} \big(f(x) - f_N(x)\big) > -\frac{\varepsilon}{3(b-a)}(b-a) = -\frac{\varepsilon}{3}.$$

[1]Even though we didn't explicitly prove this fact, it lies at the heart of the proof that a uniform limit of continuous functions is continuous.

Therefore,

$$\begin{aligned} U(f,P) - L(f,P) &\le U(f - f_N, P) + U(f_N, P) - L(f_N, P) - L(f - f_N, P) \\ &< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \\ &= \varepsilon, \end{aligned}$$

so f is Riemann integrable on $[a,b]$.

Now let $\varepsilon > 0$ and choose N such that $n \ge N$ implies $\|f - f_n\|_\infty < \frac{\varepsilon}{2(b-a)}$. Then we have

$$\begin{aligned} \left| \int_a^b f_n(x)\,dx - \int_a^b f(x)\,dx \right| &\le \int_a^b |f_n(x) - f(x)|\,dx \\ &\le \frac{\varepsilon}{2(b-a)}(b-a) \\ &= \frac{\varepsilon}{2} \\ &< \varepsilon, \end{aligned}$$

so $\int_a^b f_n \to \int_a^b f$ as $n \to \infty$. □

Exercises for Section 3.1

Exercise 3.1.1 ([Rud76], Exercise 7.9)**.** Let $(f_n)_{n=1}^\infty$ be a sequence of continuous, real-valued functions on a metric space X that converges uniformly to a function f. Prove that

$$\lim_{n\to\infty} f_n(x_n) = f(x)$$

for every $x \in X$ and every sequence $(x_n)_{n=1}^\infty$ in X such that $x_n \to x$. Is the converse true?

Exercise 3.1.2 ([Rud76], Exercise 7.7)**.** For $n = 1, 2, 3, \ldots$, define $f_n : \mathbf{R} \to \mathbf{R}$ by

$$f_n(x) = \frac{x}{1 + nx^2}.$$

(a) Prove that the sequence $(f_n)_{n=1}^\infty$ converges uniformly to a function f.

(b) Show that the equation

$$f'(x) = \lim_{n\to\infty} f_n'(x)$$

holds for $x \ne 0$, but it fails if $x = 0$. Why doesn't this contradict our theorem about uniform convergence and differentiability?

3.2 Series of Functions

We now turn our attention to functions defined via infinite series. Recall that a series of real numbers converges precisely when the corresponding sequence of partial sums converges. Extending this notion to series of functions, we immediately see that we can bring all of our results about sequences of functions to bear on series.

Definition 3.2.1. Let $(f_n)_{n=1}^{\infty}$ be a sequence of functions on a metric space X. If the series

$$\sum_{n=1}^{\infty} f_n(x)$$

converges to a number $f(x)$ for each $x \in X$, we say the series **converges pointwise** to f. Equivalently, the series converges pointwise to f if its sequence of partial sums converges to f pointwise. If the sequence of partial sums converges uniformly to f, we say the series **converges uniformly** to f.

The Cauchy criterion for uniform convergence of sequences of functions yields a nice test for uniform convergence of series.

Theorem 3.2.2 (Weierstrass M-Test)**.** *Let $(f_n)_{n=1}^{\infty}$ be a sequence of functions, and suppose that for each n there is a constant M_n such that*

$$|f_n(x)| \leq M_n$$

for all $x \in X$. If the series $\sum_{n=1}^{\infty} M_n$ converges, then $\sum_{n=1}^{\infty} f_n$ converges uniformly.

Proof. Assume the series $\sum_{n=1}^{\infty} M_n$ converges, and let $\varepsilon > 0$ be given. By the Cauchy criterion for convergence of series of real numbers [Rud76, Theorem 3.22], there exists N such that $m \geq n \geq N$ implies

$$\sum_{k=n+1}^{m} M_n < \varepsilon.$$

Therefore, for all $x \in X$ we have

$$\left| \sum_{k=n+1}^{m} f_k(x) \right| \leq \sum_{k=n+1}^{m} |f_k(x)| \leq \sum_{k=n+1}^{m} M_n < \varepsilon \tag{3.2}$$

when $m \geq n \geq N$. But the left hand side is precisely the absolute value of the difference between the m^{th} and n^{th} partial sums of the series:

$$\sum_{k=1}^{m} f_k(x) - \sum_{k=1}^{n} f_k(x) = \sum_{k=n+1}^{m} f_k(x).$$

Thus (3.2) shows that the sequence of partial sums satisfies the Cauchy criterion for uniform convergence. Hence $\sum_{n=1}^{\infty} f_n$ converges uniformly by Theorem 3.1.10. □

As an application of our results on uniform convergence, we now prove some familiar facts from calculus regarding power series. Recall that a **power series** has the form

$$\sum_{n=0}^{\infty} c_n x^n,$$

where $c_n \in \mathbf{R}$ for all n. Such a series is guaranteed to converge for all x in some interval I (though it is possible this interval is degenerate, i.e. $I = \{0\}$). One can determine this interval using the Root Test: the series is guaranteed to converge absolutely for any x satisfying[2]

$$\limsup_{n\to\infty} \sqrt[n]{|c_n x^n|} < 1.$$

Observe that

$$\limsup_{n\to\infty} \sqrt[n]{|c_n x^n|} = \limsup_{n\to\infty} |x| \sqrt[n]{|c_n|} = |x| \cdot \limsup_{n\to\infty} \sqrt[n]{|c_n|}, \tag{3.3}$$

and if we put $L = \limsup \sqrt[n]{|c_N|}$, then the series converges precisely when

$$|x| < \frac{1}{L},$$

provided L is finite and nonzero. In this case, we put $R = \frac{1}{L}$ and call R the **radius of convergence** of the series. Thus $\sum_{n=1}^{\infty} c_n x^n$ converges for all x in the open interval $(-R, R)$. If $L = 0$, notice that (3.3) is always less than 1, so the series converges for all values of x. If $L = \infty$, the series converges only when $x = 0$.

Observe that our discussion so far only gives conditions for pointwise convergence of a power series. If we want to prove results about continuity or differentiability, we will need uniform convergence. Luckily, we can prove that any power series converges uniformly on compact subsets of its (open) interval of convergence.

Theorem 3.2.3. *Suppose there is an $R > 0$ such that the power series*

$$\sum_{n=0}^{\infty} c_n x^n$$

converges absolutely for all x satisfying $|x| < R$. Then for any $0 < r < R$, the series converges uniformly on the closed interval $[-r, r]$.

[2]See [Rud76, Theorem 3.33] for this grown-up version of the Root Test. Also recall that the *limit superior* of a sequence $(a_n)_{n=1}^{\infty}$ of real numbers is defined to be

$$\limsup_{n\to\infty} a_n = \lim_{n\to\infty} \sup_{m\geq n} a_m.$$

Equivalently, $\limsup a_n$ is the largest subsequential limit of $(a_n)_{n=1}^{\infty}$.

Proof. Fix $r < R$, and put $f_n(x) = c_n x^n$ and $M_n = |c_n| r^n$ for each n. For a given n, observe that

$$|f_n(x)| = |c_n x^n| \leq |c_n| r^n = M_n$$

when $|x| \leq r$. Moreover, the series

$$\sum_{n=0}^{\infty} M_n = \sum_{n=0}^{\infty} |c_n| r^n$$

converges, since $\sum_{n=0}^{\infty} c_n x^n$ converges absolutely when we take $x = r$. It then follows from the Weierstrass M-test that $\sum_{n=0}^{\infty} c_n x^n$ converges uniformly when $|x| \leq r$. $\square$

It turns out that uniform convergence on compacta is enough to prove results throughout the interval $(-R, R)$, even though the series may not converge uniformly on the entire interval. As an example, we now prove that any function defined via a convergent power series is necessarily continuous.

Corollary 3.2.4. *Suppose the power series $\sum_{n=0}^{\infty} c_n x^n$ converges when $|x| < R$, and define*

$$f(x) = \sum_{n=0}^{\infty} c_n x^n.$$

Then f is continuous on the interval $(-R, R)$.

Proof. Let $x_0 \in (-R, R)$, and choose f such that $|x| \leq r < R$. (In particular, we could take $r = |x_0|$.) Then $\sum_{n=0}^{\infty} c_n x^n$ converges uniformly on the closed interval $[-r, r]$, so f is continuous on $[-r, r]$. In particular, f is continuous at x_0. Since x_0 was arbitrary, f is continuous throughout $(-R, R)$. $\square$

Another fact that we use freely in Calculus II is the ability to differentiate a power series term-by-term.

Corollary 3.2.5. *If the series $\sum_{n=0}^{\infty} c_n x^n$ converges when $|x| < R$, then the function $f(x) = \sum_{n=0}^{\infty} c_n x^n$ is differentiable on $(-R, R)$, and its derivative is given by*

$$f'(x) = \sum_{n=1}^{\infty} n c_n x^{n-1}$$

for $|x| < R$.

Proof. For any $r < R$, the series $\sum_{n=0}^{\infty} c_n x^n$ converges uniformly on the closed interval $[-r, r]$. Moreover, if we let

$$f_N(x) = \sum_{n=0}^{N} c_n x^n$$

denote the N^{th} partial sum of the series, then f_N is differentiable with

$$f_N'(x) = \sum_{n=1}^{N} n c_n x^n.$$

We claim the sequence $(f_N')_{N=1}^{\infty}$ converges uniformly on $[-r, r]$. Observe that

$$\limsup_{n\to\infty} \sqrt[n]{n|c_n|} = \limsup_{n\to\infty} \sqrt[n]{|c_n|}$$

since $\sqrt[n]{n} \to 1$ as $n \to \infty$. This implies that the series

$$\sum_{n=1}^{\infty} n c_n x^{n-1} \tag{3.4}$$

has the same radius of convergence as the original series $\sum_{n=0}^{\infty} c_n x^n$. Thus the series (3.4) converges to a function g when $|x| < R$, hence it converges uniformly on $[-r, r]$. It follows from Theorem 3.1.17 that f is differentiable on $[-r, r]$ with $f' = g$. Since $r < R$ was arbitrary, we can conclude that f is differentiable throughout $(-R, R)$ with

$$f'(x) = \sum_{n=1}^{\infty} n c_n x^{n-1}$$

for all $x \in (-R, R)$. □

Remark 3.2.6. If we apply the previous corollary repeatedly, we can see that a function defined by a power series has derivatives of all orders. Furthermore, it is easy to check that $f^{(n)}(0) = n! c_n$, so the coefficients of the original power series expansion are determined entirely by the derivatives of f:

$$c_n = \frac{f^{(n)}(0)}{n!}.$$

Of course these are the familiar Taylor coefficients that one encounters in calculus.

Corollary 3.2.7. *If the series $\sum_{n=0}^{\infty} c_n x^n$ converges when $|x| < R$, then the function $f(x) = \sum_{n=0}^{\infty} c_n x^n$ is Riemann integrable on any closed interval $[a, b] \subseteq (-R, R)$, with is given by*

$$\int_a^b f(x)\,dx = \sum_{n=0}^{\infty} \left(c_n \int_a^b x^n\,dx \right)$$

Proof. Given a closed interval $[a, b] \subseteq (-R, R)$, we can find $r < R$ such that $[a, b] \subseteq [-r, r]$. Then the series converges uniformly on $[-r, r]$, hence on $[a, b]$. Moreover, each partial sum is Riemann integrable with

$$\int_a^b \sum_{n=0}^N c_n x^n \, dx = \sum_{n=0}^N c_n \left(\int_a^b x^n \, dx \right).$$

It follows from Theorem 3.1.18 that f is Riemann integrable and

$$\int_a^b f = \lim_{N\to\infty} \sum_{n=0}^N c_n \left(\int_a^b x^n \, dx \right) = \sum_{n=0}^\infty c_n \left(\int_a^b x^n \, dx \right),$$

as desired. □

Exercises for Section 3.2

Exercise 3.2.1. Prove that the series

$$\sum_{n=0}^\infty \frac{2^n \cos nx}{n!}$$

converges uniformly on $\mathbf{R}$.

Exercise 3.2.2 ([Rud76], Exercise 7.11)**.** Suppose $(f_n)_{n=1}^\infty$ and $(g_n)_{n=1}^\infty$ are sequences of functions on a metric space X such that:

- The sequence of partial sums of $\sum f_n$ is uniformly bounded.
- $g_n \to 0$ uniformly.
- $g_1(x) \geq g_2(x) \geq g_3(x) \geq \cdots$ for all $x \in X$.

Prove that $\sum f_n g_n$ converges uniformly.

3.3 The Arzelà-Ascoli Theorem

In the next two sections we will narrow our focus to sequences of functions on a *compact* metric space X, and we will obtain two deep theorems from classical analysis. Recall that if X is a compact metric space, then continuous functions on X are necessarily bounded, and the supremum norm is well-defined for $C(X, \mathbf{R})$. Moreover, this norm makes $C(X, \mathbf{R})$ into a complete metric space. The first theorem that we will study stems from a question that might seem, at first glance, to be quite odd: what conditions guarantee a subset of $C(X, \mathbf{R})$ is compact? Of course we already have such a condition—a set $\mathcal{F} \subseteq C(X, \mathbf{R})$ is compact if and only if it is closed and totally bounded. Perhaps we can come up with conditions that are more intrinsic to $C(X, \mathbf{R})$ that ensure compactness, however.

This question should appear strange as stated. We are asking when a *collection of functions* on a compact metric space is itself compact. However, the question becomes more natural if we think along the lines of sequential compactness: under what conditions must a sequence $(f_n)_{n=1}^{\infty}$ in $C(X, \mathbf{R})$ have a uniformly convergent subsequence? Such a result seems like it would be much more applicable in practice. (It indeed has applications to differential equations and calculus of variations, for example.) Consequently, we will take this approach as we head toward our main result, which is called the Arzelà-Ascoli theorem. Once we have proven this theorem, we will return to the original question regarding compactness.

As a first step toward the main result, observe that if $\mathcal{F} \subseteq C(X, \mathbf{R})$ is compact, then it must be closed and bounded. This leads us to define two notions of boundedness for families of functions in $C(X, \mathbf{R})$.

Definition 3.3.1. Let $\mathcal{F} \subseteq C(X, \mathbf{R})$ be a family of functions on X.

1. We say $\mathcal{F}$ is **pointwise bounded** if for each $x \in X$, there is a constant $M_x > 0$ such that
$$|f(x)| \leq M_x$$
for all $f \in \mathcal{F}$. Equivalently, the set $\{f(x) : f \in \mathcal{F}\} \subseteq \mathbf{R}$ is bounded for each $x \in X$.

2. We say $\mathcal{F}$ is **uniformly bounded** if there exists an $M > 0$ such that for each $x \in X$,
$$|f(x)| \leq M$$
for all $f \in \mathcal{F}$. Equivalently, $\|f\|_\infty \leq M$ for all $f \in \mathcal{F}$.

Notice that a set $\mathcal{F} \subseteq C(X, \mathbf{R})$ is uniformly bounded precisely when it is bounded with respect to the supremum norm. That is, uniform boundedness is precisely the usual notion of boundedness in the metric space $C(X, \mathbf{R})$. On the other hand, it seems as though it would be much easier to verify that a set is pointwise bounded than it would be to produce a uniform bound. Fortunately, we will see later on that pointwise boundedness implies uniform boundedness under the right hypotheses.

Example 3.3.2. Consider the sequence $(f_n)_{n=1}^{\infty}$ on $X = [0, 1]$ defined by $f_n(x) = x^n$. Notice that for all $x \in [0, 1]$,

$$|f_n(x)| = |x^n| \leq 1$$

for all n. Thus $(f_n)_{n=1}^{\infty}$ is uniformly bounded.

Example 3.3.3. Let $(f_n)_{n=1}^{\infty}$ be the sequence of functions $f_n : [0,1] \to \mathbf{R}$ defined by

$$f_n(x) = \begin{cases} 2^{2n}x & \text{if } 0 \le x \le \frac{1}{2^n} \\ 2^{2n}\left(\frac{1}{2^{n-1}} - x\right) & \text{if } \frac{1}{2^n} < x \le \frac{1}{2^{n-1}} \\ 0 & \text{if } \frac{1}{2^{n-1}} < x \le 1. \end{cases}$$

We saw in Example 3.1.4 that $f_n \to 0$ pointwise. In other words, the sequence $(f_n(x))_{n=1}^{\infty}$ is convergent, hence bounded, for each $x \in [0,1]$. Thus $(f_n)_{n=1}^{\infty}$ is pointwise bounded. However, it is not uniformly bounded. Notice that

$$\|f_n\|_{\infty} = 2^n$$

for each n, so the sequence $\{\|f_n\|_{\infty}\}_{n=1}^{\infty}$ is unbounded.

Proposition 3.3.4. *Let $(f_n)_{n=1}^{\infty}$ be a sequence of functions on a compact metric space X.*

1. *If $(f_n)_{n=1}^{\infty}$ converges pointwise, then it is pointwise bounded.*
2. *If $(f_n)_{n=1}^{\infty}$ converges uniformly, then it is uniformly bounded.*

Proof. Suppose $(f_n)_{n=1}^{\infty}$ converges pointwise. Then for each $x \in X$, $(f_n(x))_{n=1}^{\infty}$ is a convergent sequence of real numbers, so it is bounded. But this is precisely what it means to be pointwise bounded.

On the other hand, saying $(f_n)_{n=1}^{\infty}$ converges uniformly is the same as saying it converges with respect to the supremum norm on $C(X, \mathbf{R})$. But we know that a convergent sequence in a metric space is bounded, so $(f_n)_{n=1}^{\infty}$ is bounded in $C(X, \mathbf{R})$, which is identical to uniform boundedness. □

Of course the converse to the last proposition is false. The sequence from Example 3.3.2 is uniformly bounded, but it does not converge uniformly. (Recall that $(f_n)_{n=1}^{\infty}$ converges pointwise to a discontinuous function, which implies that the convergence cannot be uniform.) In fact, $(f_n)_{n=1}^{\infty}$ has no uniformly convergent subsequences either. The computations in Example 3.1.9 show that no subsequence of $(f_n)_{n=1}^{\infty}$ is uniformly Cauchy. Thus uniform boundedness of a sequence does not guarantee the existence of a uniformly convergent subsequence—in other words, the Bolzano-Weierstrass theorem fails to hold in $C(X, \mathbf{R})$. We clearly need to add a hypothesis besides uniform boundedness.

Definition 3.3.5. A family of functions $\mathcal{F} \subseteq C(X, \mathbf{R})$ is **equicontinuous** if for each $\varepsilon > 0$, there exists $\delta > 0$ such that whenever $x, y \in X$ with $d(x, y) < \delta$,

$$|f(x) - f(y)| < \varepsilon$$

for all $f \in \mathcal{F}$

Notice that if $\mathcal{F}$ is an equicontinuous family of functions, then every function $f \in \mathcal{F}$ is uniformly continuous. This is not a very deep observation at all—it is automatic that every $f \in C(X, \mathbf{R})$ is uniformly continuous, since X is compact. Equicontinuity is much stronger than the requirement that each function be uniformly continuous—it says that given any $\varepsilon > 0$, there is a single $\delta > 0$ that works for *every* $f \in \mathcal{F}$. This seems like a fairly stringent requirement, so one might wonder how common such families are. They are actually quite common, as the next result shows.

Proposition 3.3.6. *Let X be a compact metric space. If $(f_n)_{n=1}^{\infty}$ is a uniformly convergent sequence in $C(X, \mathbf{R})$, then $(f_n)_{n=1}^{\infty}$ is equicontinuous.*

Proof. Let $\varepsilon > 0$ be given. Since $(f_n)_{n=1}^{\infty}$ converges uniformly, it is uniformly Cauchy. Thus we can find a natural number N such that $n \geq N$ implies

$$\|f_n - f_N\|_\infty < \frac{\varepsilon}{3}.$$

Since f_N is continuous and X is compact, f_N is uniformly continuous. Thus we can find $\delta_0 > 0$ such that whenever $x, y \in X$ with $d(x, y) < \delta_0$,

$$|f_N(x) - f_N(y)| < \frac{\varepsilon}{3}.$$

Thus for all $n \geq N$ we have

$$\begin{aligned} |f_n(x) - f_n(y)| &\leq |f_n(x) - f_N(x)| + |f_N(x) - f_N(y)| + |f_N(y) - f_n(y)| \\ &< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \\ &= \varepsilon \end{aligned}$$

when $d(x, y) < \delta_0$.

Now we consider the case where $n < N$. Each f_n is uniformly continuous on X, so for $n = 1, 2, \ldots, N-1$ there exists $\delta_n > 0$ such that $d(x, y) < \delta_n$ implies

$$|f_n(x) - f_n(y)| < \varepsilon.$$

Set

$$\delta = \min\{\delta_0, \delta_1, \delta_2, \ldots, \delta_{N-1}\}.$$

It is then clear that this δ works for each f_n, so $(f_n)_{n=1}^{\infty}$ is equicontinuous. □

We are almost ready to prove the Arzelà-Ascoli theorem. However, we first need to prove two lemmas that will allow us to exploit the compactness of X. The first lemma is a general statement about compact metric spaces.

Lemma 3.3.7. *Any compact metric space is separable.*

Proof. Let X be a compact metric space. Then X is totally bounded, so for each $n \in \mathbf{N}$ there is a finite $\frac{1}{n}$-net

$$\mathcal{B}_n = \left\{B\left(x_{n,i}, \tfrac{1}{n}\right)\right\}_{i=1}^{i_n}$$

that covers X. Let $D_n = \{x_{n,i}\}_{i=1}^{i_n}$ be the set consisting of the centers of the balls in $\mathcal{B}_n$ for each n. Then define

$$D = \bigcup_{n=1}^{\infty} D_n.$$

Notice that D is a countable union of finite sets, hence countable. Furthermore, we claim that D is dense in X. Given $x \in X$ and $\varepsilon > 0$, choose $n \in \mathbf{N}$ such that $\frac{1}{n} < \varepsilon$. Since $\mathcal{B}_n$ covers X, we have $x \in B_{1/n}(x_{n,i})$ for some i. Hence

$$d(x, x_{n,i}) < \tfrac{1}{n} < \varepsilon,$$

and it follows that D is dense in X. Thus X is separable. □

The second lemma is a selection theorem for a sequence of functions on a countable set. The statement and proof are largely lifted from [Rud76, Theorem 7.23].

Lemma 3.3.8. *Let E be a countable set, and let $(f_n)_{n=1}^{\infty}$ be a pointwise bounded sequence of functions on E. Then there is a subsequence $(f_{n_k})_{k=1}^{\infty}$ that converges pointwise on E.*

Proof. Since E is countable, we will write it as

$$E = \{x_1, x_2, x_3, \ldots\}.$$

By assumption, $(f_n)_{n=1}^{\infty}$ is pointwise bounded, so $(f_n(x_1))_{n=1}^{\infty}$ is a bounded sequence of real numbers. Therefore, the Bolzano-Weierstrass theorem guarantees that it has a convergent subsequence. In other words, there is a subsequence $(f_{1,k})_{k=1}^{\infty}$ of $(f_n)_{n=1}^{\infty}$ such that $(f_{1,k}(x_1))_{k=1}^{\infty}$ converges. Similarly, the $(f_{1,k}(x_2))_{k=1}^{\infty}$ is bounded, so we can find a subsequence $(f_{2,k})_{k=1}^{\infty}$ of $(f_{1,k})_{k=1}^{\infty}$ such that the sequence $(f_{2,k}(x_2))_{k=1}^{\infty}$ converges. We can now proceed inductively, and for each $m \in \mathbf{N}$ find a sequence $(f_{m,k})_{k=1}^{\infty}$ such that:

- $(f_{m,k})_{k=1}^{\infty}$ is a subsequence of $(f_{m-1,k})_{k=1}^{\infty}$ for each m;
- $(f_{m,k}(x_m))_{k=1}^{\infty}$ converges. (In fact, $(f_{m,k}(x))_{k=1}^{\infty}$ converges for $x = x_1, x_2, \ldots, x_m$.)

Now we form the sequence $(f_{m,m})_{n=1}^{\infty}$. If we arrange our "sequence of subsequences" in a grid, are simply selecting the elements on the diagonal:

$$\begin{array}{ccccc}
\boxed{f_{1,1}} & f_{1,2} & f_{1,3} & f_{1,4} & \cdots \\
f_{2,1} & \boxed{f_{2,2}} & f_{2,3} & f_{2,4} & \cdots \\
f_{3,1} & f_{3,2} & \boxed{f_{3,3}} & f_{3,4} & \cdots \\
f_{4,1} & f_{4,2} & f_{4,3} & \boxed{f_{4,4}} & \cdots \\
\vdots & \vdots & \vdots & \vdots & \ddots
\end{array}$$

Notice that $(f_{m,m})_{n=1}^{\infty}$ is a subsequence of $(f_n)_{n=1}^{\infty}$, and that $(f_{m,m}(x_i))_{n=1}^{\infty}$ converges for each $x_i \in E$. To verify the latter assertion, notice that when $m \geq i$, $f_{m,m}$ appears as a term in the sequence $(f_{i,k})_{k=1}^{\infty}$. In other words, $(f_{m,m})_{m=i}^{\infty}$ is a subsequence of $(f_{i,k})_{k=1}^{\infty}$. Since $(f_{i,k}(x_i))_{k=1}^{\infty}$ converges, it follows that the sequence $(f_{m,m}(x_i))_{n=1}^{\infty}$ converges. Since $x_i \in E$ was arbitrary, the subsequence $(f_{m,m})_{m=1}^{\infty}$ converges pointwise on E. □

Theorem 3.3.9 (Arzelà-Ascoli)**.** *Let X be a compact metric space, and suppose $(f_n)_{n=1}^{\infty}$ is a pointwise bounded, equicontinuous sequence in $C(X, \mathbf{R})$. Then:*

1. *$(f_n)_{n=1}^{\infty}$ is uniformly bounded.*
2. *$(f_n)_{n=1}^{\infty}$ has a uniformly convergent subsequence.*

Proof. We first show the sequence is uniformly bounded. Since $(f_n)_{n=1}^{\infty}$ is equicontinuous, we can find $\delta > 0$ such that $d(x, y) < \delta$ implies

$$|f_n(x) - f_n(y)| < 1$$

for all n. Since X is compact, hence totally bounded, we can cover X with a finite collection of δ-balls. That is, there exist finitely many points $x_1, x_2, \ldots, x_m$ such that the family $\{B_\delta(x_i)\}_{i=1}^m$ covers X. Since $(f_n)_{n=1}^{\infty}$ is pointwise bounded, for $i = 1, 2, \ldots, m$ there exists $M_i > 0$ such that

$$|f_n(x_i)| \leq M_i$$

for all n. Let $M = \max\{M_1, M_2, \ldots, M_m\}$. Then given $x \in X$, we have $x \in B_\delta(x_i)$ for some i, so

$$|f_n(x)| \leq |f_n(x) - f_n(x_i)| + |f_n(x_i)| < 1 + M$$

for all n. Thus $|f_n(x)| \leq M + 1$ for all n and all $x \in X$, which implies that $(f_n)_{n=1}^{\infty}$ is uniformly bounded.

Now recall that since X is compact, it is separable by our first lemma. Let $E \subseteq X$ be a countable dense subset. By the second lemma, there is a subsequence $(f_{n_k})_{k=1}^{\infty}$ that converges pointwise on E. To simplify notation, we will set $g_k = f_{n_k}$ for each k. Our goal is to show that $(g_k)_{k=1}^{\infty}$ converges uniformly on X.

Let $\varepsilon > 0$ be given. The sequence $(g_k)_{k=1}^{\infty}$ is equicontinuous, so we can find $\delta > 0$ such that $d(x, y) < \delta$ implies

$$|g_k(x) - g_k(y)| < \frac{\varepsilon}{3}$$

for all k. Since E is dense in X, the collection

$$\left\{B_\delta(x)\right\}_{x \in E}$$

is an open cover of X. Since X is compact, there is a finite subcover $\{B_\delta(x_i)\}_{i=1}^m$. We know that $(g_k)_{k=1}^\infty$ converges pointwise on E, so $(g_k(x_i))_{k=1}^\infty$ converges for $1 \le i \le m$. Thus we can find an N such that $k, l \ge N$ implies

$$|g_k(x_i) - g_l(x_i)| < \frac{\varepsilon}{3}$$

for $i = 1, \ldots, m$. (This works because we are only considering finitely many points.) If $x \in X$, we have $x \in B_\delta(x_i)$ for some i, so

$$|g_k(x) - g_k(x_i)| < \frac{\varepsilon}{3}.$$

Therefore, for all $k, l \ge N$,

$$\begin{aligned} |g_k(x) - g_l(x)| &\le |g_k(x) - g_k(x_i)| + |g_k(x_i) - g_l(x_i)| + |g_l(x_i) - g_l(x)| \\ &< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \\ &= \varepsilon \end{aligned}$$

for all $x \in X$. Thus $(g_k)_{k=1}^\infty$ is uniformly Cauchy, so it converges uniformly. □

Using the Arzelà-Ascoli theorem, we can now return to our original question regarding compactness in $C(X, \mathbf{R})$.

Theorem 3.3.10. *Let X be a compact metric space. A set $\mathcal{F} \subseteq C(X, \mathbf{R})$ is compact if and only if it is uniformly closed, pointwise bounded, and equicontinuous.*

Proof. Suppose first that $\mathcal{F} \subseteq C(X, \mathbf{R})$ is compact. Then it is closed and bounded as a subset of $C(X, \mathbf{R})$ (i.e., it is uniformly closed and uniformly bounded), which implies that $\mathcal{F}$ is closed and pointwise bounded. We just need to establish equicontinuity.

Let $\varepsilon > 0$ be given. Since $\mathcal{F}$ is compact, hence totally bounded, we can cover $\mathcal{F}$ with a finite number of $\frac{\varepsilon}{3}$-balls. That is, there exist functions $f_1, f_2, \ldots, f_m \in C(X, \mathbf{R})$ such that $\{B_{\varepsilon/3}(f_i)\}_{i=1}^m$ covers $\mathcal{F}$. Since X is compact, each f_i is uniformly continuous. Thus for each i we can find $\delta_i > 0$ such that $d(x, y) < \delta_i$ implies

$$|f_i(x) - f_i(y)| < \frac{\varepsilon}{3}$$

for $1 \le i \le m$. Set $\delta = \min\{\delta_1, \delta_2, \ldots, \delta_m\}$. Given $f \in \mathcal{F}$, we have $f \in B_{\varepsilon/3}(f_i)$ for some i, so

$$\begin{aligned} |f(x) - f(y)| &\le |f(x) - f_i(x)| + |f_i(x) - f_i(y)| + |f_i(y) - f(y)| \\ &< \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \\ &= \varepsilon \end{aligned}$$

when $d(x, y) < \delta$. Therefore, $\mathcal{F}$ is equicontinuous.

Now suppose $\mathcal{F}$ is closed, pointwise bounded, and equicontinuous. Then any sequence $(f_n)_{n=1}^{\infty}$ has a uniformly convergent subsequence by the Arzelà-Ascoli theorem. Since $\mathcal{F}$ is closed, the limit of this subsequence must belong to $\mathcal{F}$. Therefore, $\mathcal{F}$ is sequentially compact, hence compact. □

Both versions of the Arzelà-Ascoli theorem also hold for the space $C(X)$ of complex-valued functions on X, and the proof is nearly identical to the one we have already given.

Theorem 3.3.11 (Complex Arzelà-Ascoli)**.** *Let X be a compact metric space.*

1. *If a sequence $(f_n)_{n=1}^{\infty}$ in $C(X)$ is pointwise bounded and equicontinuous, then it has a uniformly convergent subsequence.*
2. *A set $\mathcal{F} \subseteq C(X)$ is compact if and only if it is uniformly closed, pointwise bounded, and equicontinuous.*

Exercises for Section 3.3

Exercise 3.3.1. Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of differentiable functions $f_n : [a, b] \to \mathbf{R}$ with uniformly bounded derivatives, meaning there exists $M > 0$ such that for all n,

$$\left|f_n'(x)\right| \leq M$$

for all $x \in [a, b]$. Prove that $(f_n)_{n=1}^{\infty}$ is equicontinuous.

Exercise 3.3.2. Let X be a (not necessarily compact) metric space, and let $(f_n)_{n=1}^{\infty}$ be a sequence of continuous functions $f_n : X \to \mathbf{R}$. If $f_n \to f$ uniformly for some function $f : X \to \mathbf{R}$ and the sequence is equicontinuous, prove that f is uniformly continuous on X.

Exercise 3.3.3 ([Rud76], Exercise 7.15)**.** Suppose $f : \mathbf{R} \to \mathbf{R}$ is continuous, and define $f_n : [0, 1] \to \mathbf{R}$ by

$$f_n(x) = f(nx), \quad n = 1, 2, 3, \ldots$$

If we assume the sequence $(f_n)_{n=1}^{\infty}$ is equicontinuous on $[0, 1]$, what can you say about f?

Exercise 3.3.4 ([Rud76], Exercise 7.16)**.** Suppose $(f_n)_{n=1}^{\infty}$ is an equicontinuous sequence of functions on a compact metric space X, and $(f_n)_{n=1}^{\infty}$ converges pointwise to a function f. Prove that $f_n \to f$ uniformly.

Exercise 3.3.5. Fix a function $g \in C([0,1],\mathbf{R})$, and define $T : C([0,1],\mathbf{R}) \to C([0,1],\mathbf{R})$ by

$$T(f)(x) = \int_0^x f(t)g(t)\,dt.$$

If $(f_n)_{n=1}^{\infty}$ is a uniformly bounded sequence in $C([0,1],\mathbf{R})$, prove that $(T(f_n))_{n=1}^{\infty}$ has a uniformly convergent subsequence.

3.4 The Stone-Weierstrass Theorem

We are about to close our discussion of continuous functions on metric spaces with another classical result from analysis, which is known as the Stone-Weierstrass theorem. The original form of this theorem (proven by Weierstrass) encodes an idea that we all use (at least subconsciously) throughout our mathematical lives—the idea of approximating continuous functions with polynomials.

More precisely, Weierstrass proved the following result: given a continuous function $f : [a,b] \to \mathbf{R}$ on some interval $[a,b] \subseteq \mathbf{R}$ and any $\varepsilon > 0$, there exists a polynomial p satisfying

$$|f(x) - p(x)| < \varepsilon$$

for all $x \in X$. If we let $P([a,b])$ denote the set of all polynomials (with real coefficients) on $[a,b]$, Weierstrass's theorem can be interpreted as saying that $P([a,b])$ is dense in $C([a,b],\mathbf{R})$ with respect to the supremum norm. Stone greatly generalized this theorem—he determined conditions that guarantee a collection of functions is dense in $C(X,\mathbf{R})$, where X is a compact metric space. In order to understand Stone's generalization and consequently prove the Stone-Weierstrass theorem, we need to determine the essential properties of $P([a,b])$ that make the original result of Weierstrass work.

Let X be a compact metric space. We have observed repeatedly that $C(X,\mathbf{R})$ is a complete normed vector space (i.e., a Banach space) with respect to the supremum norm $\|\cdot\|_\infty$. It is actually more than that—in addition to the vector space operations, we can multiply functions pointwise: if $f, g \in C(X,\mathbf{R})$, then so is fg, where

$$(fg)(x) = f(x)g(x), \quad x \in X.$$

It is easy to check that multiplication interacts nicely with the addition and scalar multiplication on $C(X,\mathbf{R})$. Moreover, the constant function $f \equiv 1$ is a multiplicative identity, so $C(X,\mathbf{R})$ is a commutative, unital ring. A set that is simultaneously a ring and a vector space (over a field $\mathbb{F}$) is called an **algebra** (over $\mathbb{F}$). We have just argued that $C(X,\mathbf{R})$ is an algebra over $\mathbf{R}$.

Example 3.4.1. The set $P([a,b])$ of polynomials on an interval $[a,b]$ is easily seen to be an algebra. The sum or product of two polynomials is again a polynomial, and a scalar multiple of a polynomial is still a polynomial. In other words, $P([a,b])$ is a **subalgebra** of $C([a,b],\mathbf{R})$.

Since polynomials do not make sense on an arbitrary metric space X, we will just consider general subalgebras of $C(X, \mathbf{R})$ instead. Our goal is to determine conditions that guarantee a subalgebra $\mathcal{A} \subseteq C(X, \mathbf{R})$ is dense in $C(X, \mathbf{R})$. That is, what are the crucial properties of $P([a,b])$ that allowed Stone to generalize the Weierstrass approximation theorem?

Definition 3.4.2. Let $\mathcal{A}$ be a subalgebra of $C(X, \mathbf{R})$.

1. We say that $\mathcal{A}$ **separates points** if for any pair of distinct points $x, y \in X$, there is a function $f \in \mathcal{A}$ such that $f(x) \neq f(y)$.
2. We say $\mathcal{A}$ is **nowhere vanishing** if for each $x \in X$, there is a function $f \in \mathcal{A}$ such that $f(x) \neq 0$.

Example 3.4.3. The algebra of polynomials $\mathcal{A} = P([a,b])$ separates points and vanishes nowhere. To see that $\mathcal{A}$ separates points, it suffices to consider the function $f(x) = x$. Notice that if $x_1, x_2 \in [a,b]$ with $x_1 \neq x_2$, then $f(x_1) \neq f(x_2)$. Similarly, we need only look at the constant function $g \equiv 1$ to see that $\mathcal{A}$ is nowhere vanishing—for any $x \in [a,b]$, we have $g(x) = 1 \neq 0$.

Remark 3.4.4. The previous example shows that there are sometimes simpler criterion that guarantee a subalgebra separates points and vanishes nowhere. Notice that if $\mathcal{A}$ contains an injective function, then it automatically separates points. Similarly, if $\mathcal{A}$ contains a nonzero constant function, then it vanishes nowhere. In fact, if $\mathcal{A}$ contains a constant function, then it contains *all* constant functions on X, since $\mathcal{A}$ is an algebra.

The properties described in Definition 3.4.2 are precisely the ones that guarantee a subalgebra $\mathcal{A} \subseteq C(X, \mathbf{R})$ is dense. Before we can prove this fact, we need a couple of preliminary results.

Lemma 3.4.5. *If $\mathcal{A}$ is a subalgebra of $C(X, \mathbf{R})$, then its closure $\overline{\mathcal{A}}$ is also a subalgebra of $C(X, \mathbf{R})$.*

Proof. Let $f, g \in \overline{\mathcal{A}}$. Then there are sequences $(f_n)_{n=1}^\infty$ and $(g_n)_{n=1}^\infty$ in $C(X, \mathbf{R})$ such that $f_n \to f$ and $g_n \to g$ uniformly. It is then straightforward to check that $f_n + g_n \to f + g$ uniformly: given $\varepsilon > 0$, if we choose N such that

$$\|f_n - f\|_\infty < \frac{\varepsilon}{2} \quad \text{and} \quad \|g_n - g\|_\infty < \frac{\varepsilon}{2},$$

for all $n \geq N$, then

$$\|(f_n + g_n) - (f + g)\|_\infty \leq \|f_n - f\|_\infty + \|g_n - g\|_\infty < \varepsilon$$

when $n \geq N$. Since $\mathcal{A}$ is an algebra, $f_n + g_n \in \mathcal{A}$ for all n, so it follows that $f+g \in \overline{\mathcal{A}}$. Similarly, one can check that $f_n g_n \to fg$ uniformly, and that $cf_n \to cf$ uniformly for any constant $c \in \mathbf{R}$. Thus $\overline{\mathcal{A}}$ is closed under multiplication and scalar multiplication as well, so $\overline{\mathcal{A}}$ is an algebra. □

Lemma 3.4.6. *Suppose $\mathcal{A} \subseteq C(X, \mathbf{R})$ is a subalgebra that separates points and vanishes nowhere. For any pair of distinct points $x_1, x_2 \in X$ and any constants $c_1, c_2 \in \mathbf{R}$, there is a function $f \in \mathcal{A}$ satisfying*

$$f(x_1) = c_1 \quad \textit{and} \quad f(x_2) = c_2.$$

Proof. Let $x_1, x_2 \in X$ with $x_1 \neq x_2$, and let $c_1, c_2 \in \mathbf{R}$. Since $\mathcal{A}$ separates points, there is a function $g \in \mathcal{A}$ such that $g(x_1) \neq g(x_2)$, and functions $h, k \in \mathcal{A}$ with $h(x_1) \neq 0$ and $k(x_2) \neq 0$. Define $u, v : X \to \mathbf{R}$ by

$$u(x) = g(x)k(x) - g(x_1)h(x)$$

and

$$v(x) = g(x)h(x) - g(x_2)h(x).$$

Then clearly $u, v \in \mathcal{A}$. Furthermore, $u(x_1) = 0$ and

$$u(x_2) = \big(g(x_2) - g(x_1)\big)k(x_2) \neq 0.$$

Similarly, $v(x_2) = 0$ and

$$v(x_1) = \big(g(x_1) - g(x_2)\big)h(x_1) \neq 0.$$

Now define $f \in \mathcal{A}$ by

$$f(x) = \frac{c_1 v(x)}{v(x_1)} + \frac{c_2 u(x)}{u(x_2)}.$$

Then it is easy to check that $f(x_1) = c_1$ and $f(x_2) = c_2$. □

We also need a very specific polynomial approximation in order to get us started. We will follow the lead of Folland (see [Fol99, Lemma 4.47]). Consider the function $f : (-\infty, 1] \to \mathbf{R}$ defined by

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

Though we eschew the proof, f is analytic on the interval $(-1, 1)$. Its Taylor series (centered at 0) can be computed using Newton's binomial theorem:

$$\sqrt{1-x} = \sum_{n=0}^{\infty} \binom{\frac{1}{2}}{n} (-x)^n = 1 + \sum_{n=1}^{\infty} \frac{\frac{1}{2}(-\frac{1}{2})(-\frac{3}{2}) \cdots (\frac{3}{2} - n)}{n!} (-1)^n x^n$$

for $-1 < x < 1$, which after some massaging becomes[3]

$$\sqrt{1-x} = 1 - \sum_{n=1}^{\infty} \frac{(2n-3)!!}{2^n n!} x^n = 1 - \sum_{n=1}^{\infty} \frac{(2n-2)!}{2^{2n-1} n!(n-1)!} x^n, \quad -1 < x < 1.$$

Recall that if $r < 1$, the series converges uniformly on $[-r, r]$ by Theorem 3.2.3.[4] In fact, one can argue that this series converges at $x = 1$ and $x = -1$, so it follows from Abel's Theorem (see [Wad03, Theorem 7.27]) that the series converges uniformly on $[0, 1]$. Thus the function $\sqrt{1-x}$ can be approximated uniformly by polynomials (namely the Taylor polynomials coming from this series) on $[-1, 1]$.

We now parlay our polynomial approximation to $\sqrt{1-x}$ into the one we really want. Observe that for each $x \in [-1, 1]$, we have

$$|x| = \sqrt{x^2} = \sqrt{1-(1-x^2)}$$

Substituting $1 - x^2$ into our Taylor series for $\sqrt{1-x}$, we get

$$|x| = 1 - \sum_{n=1}^{\infty} \frac{(2n-2)!}{2^{2n-1} n!(n-1)!} (1-x^2)^n.$$

This series still converges uniformly when $-1 \le x \le 1$, so its partial sums provide uniform polynomial approximations to the absolute value function. That is, given any $\varepsilon > 0$, there exists a polynomial $p : [-1, 1] \to \mathbf{R}$ such that

$$\big| |x| - p(x) \big| < \varepsilon$$

for all $-1 \le x \le 1$. By replacing $p(x)$ with $p(x) - p(0)$, we can guarantee that $p(0) = 0$ as well.

With these two approximations in hand, we are now ready to state and prove Stone's generalization of the Weierstrass approximation theorem.

---

[3]Recall that the **double factorial** is defined to be

$$n!! = n(n-2)(n-4)\cdots 6 \cdot 4 \cdot 2$$

if n is even and

$$n!! = n(n-2)(n-4)\cdots 5 \cdot 3 \cdot 1$$

if n is odd, with the convention that $0!! = 1$ and $(-1)!! = 1$. In general, one can also write

$$(2n-1)!! = \frac{(2n)!}{2^n n!}$$

to simplify the double factorial of an odd integer.

[4]Though the series converges uniformly, that theorem says nothing about whether the series converges uniformly *to* $\sqrt{1-x}$. Indeed, the Taylor series for a function, even if it exists, need not converge to the function. One usually needs to verify directly that the series converges to the function, using either Taylor's inequality or some other means.

Theorem 3.4.7 (Stone-Weierstrass). *Let X be a compact metric space, and suppose $\mathcal{A}$ is a subalgebra of $C(X, \mathbf{R})$. If $\mathcal{A}$ separates points and is nowhere vanishing, then $\mathcal{A}$ is uniformly dense in $C(X, \mathbf{R})$.*

Proof. The ultimate goal is to show that if $f \in C(X, \mathbf{R})$ and $\varepsilon > 0$, there is a function $g \in \mathcal{A}$ with $\|f - g\|_\infty < \varepsilon$. We will bootstrap our way up to this conclusion via a series of claims.

Claim 1: If $f \in \overline{\mathcal{A}}$, then $|f| \in \overline{\mathcal{A}}$.
Proof of Claim 1: Suppose $f \in \overline{\mathcal{A}}$. If $f(x) = 0$ for all x, then there is nothing to prove. Therefore, assume f is not the zero function, and set $g(x) = f(x)/\|f\|_\infty$ for each $x \in X$. Then $g \in \overline{\mathcal{A}}$, and $-1 \leq g(x) \leq 1$ for all x.

Let $\varepsilon > 0$ be given. By the earlier discussion, there is a polynomial $p : [-1, 1] \to \mathbf{R}$ satisfying $p(0) = 0$ and

$$\big||t| - p(t)\big| < \frac{\varepsilon}{\|f\|_\infty}$$

for all $t \in [-1, 1]$. Now define $h : X \to \mathbf{R}$ by $h(x) = p(g(x))$. Since $p(0) = 0$ and p is a polynomial, $h \in \overline{\mathcal{A}}$. Moreover, for all $x \in X$ we have

$$\big|h(x) - |g(x)|\big| < \varepsilon,$$

so

$$\big|\|f\|_\infty h(x) - \|f\|_\infty |g(x)|\big| = \big|\|f\|_\infty h(x) - |f(x)|\big| < \|f\|_\infty \cdot \frac{\varepsilon}{\|f\|_\infty} = \varepsilon.$$

In other words,

$$\big\|\|f\|_\infty h - |f|\big\|_\infty < \varepsilon.$$

Since the function $\|f\|_\infty h \in \overline{\mathcal{A}}$, this shows that $|f| \in \overline{\mathcal{A}}$ as well.

Claim 2: If $f, g \in \overline{\mathcal{A}}$, then $f \vee g, f \wedge g \in \overline{\mathcal{A}}$, where $f \vee g$ and $f \wedge g$ are the functions defined by

$$(f \vee g)(x) = \max\{f(x), g(x)\}, \quad (f \wedge g)(x) = \min\{f(x), g(x)\}$$

for all $x \in X$.
Proof of Claim 2: It should not be hard to convince oneself that the maximum of two real numbers a and b can be expressed as

$$\max\{a, b\} = \frac{a + b}{2} + \frac{|a - b|}{2}.$$

Using this observation, it is straightforward to check that

$$f \vee g = \frac{f + g}{2} + \frac{|f - g|}{2}.$$

Similar reasoning shows that

$$f \wedge g = \frac{f+g}{2} - \frac{|f-g|}{2}$$

as well. But then it follows from Claim 2 and the fact that $\overline{\mathcal{A}}$ is an algebra that $f \vee g, f \wedge g \in \overline{\mathcal{A}}$ for any $f, g \in \overline{\mathcal{A}}$.

Claim 3: For any function $f \in C(X, \mathbf{R})$, any point $x \in X$, and any $\varepsilon > 0$, there is a function $g_x \in \overline{\mathcal{A}}$ satisfying

$$g_x(x) = f(x) \quad \text{and} \quad g_x(y) > f(y) - \varepsilon$$

for all $y \in X$.
Proof of Claim 3: For any $t \in X$, there exists $h_t \in \overline{\mathcal{A}}$ such that

$$h_t(x) = f(x), \quad h_t(t) = f(t)$$

by Lemma 3.4.6. Since each h_t is continuous, there is an open set U_t containing t such that

$$h_t(y) > f(y) - \varepsilon$$

for all $y \in U_t$.[5] Since X is compact, there exists a finite collection of points $t_1, t_2, \ldots, t_n \in X$ such that the collection $\{U_{t_i}\}_{i=1}^n$ covers X. Now set

$$g_x = h_{t_1} \vee h_{t_2} \vee \cdots \vee h_{t_n}.$$

By inductively applying Claim 3, we see that $g_x \in \overline{\mathcal{A}}$. Furthermore, for all $y \in X$ we have

$$g_x(y) > f(y) - \varepsilon$$

by construction—for each $y \in X$, we have $h_{t_i}(y) > f(y) - \varepsilon$ for some y, since the U_{t_i} cover X. Finally, it is immediate from the definition of the h_{t_i} that

$$g_x(x) = \max\{h_{t_1}(x), h_{t_2}(x), \ldots, h_{t_n}(x)\} = f(x)$$

as well. Thus g_x has the desired properties.

Claim 4: Given $f \in C(X, \mathbf{R})$ and $\varepsilon > 0$, there exists a function $g \in \overline{\mathcal{A}}$ such that

$$\|f - g\|_\infty < \varepsilon.$$

Proof of Claim 4: For each $x \in X$, the function g_x afforded by Claim 4 is continuous, so there is an open set V_x such that

$$g_x(y) < f(y) + \varepsilon$$

[5] To make this clearer, define $H = h_t - f$. Then $H(t) = 0$, so continuity allows us to find an open neighborhood U_t of t on which H is bounded below by $-\varepsilon$, namely $U_t = H^{-1}((-\varepsilon, \infty))$.

for all $t \in V_x$. Since X is compact, there is a finite collection of points $x_1, x_2, \ldots, x_m \in X$ such that $\{V_{x_i}\}_{i=1}^m$ covers X. Now set

$$g = g_{x_1} \wedge g_{x_2} \wedge \cdots \wedge g_{x_m}.$$

Then $g \in \overline{\mathcal{A}}$ by Claim 3, and

$$g(y) < f(y) + \varepsilon$$

for all $y \in X$ by construction. Since $g_{x_i}(y) > f(y) - \varepsilon$ for all $y \in X$ for $i = 1, 2, \ldots, m$, it follows that

$$g(y) > f(y) - \varepsilon.$$

Therefore,

$$|g(y) - f(y)| < \varepsilon$$

for all $y \in X$, so the claim holds.

The main result now follows from Claim 4, which shows that $\overline{\mathcal{A}}$ is dense in $C(X, \mathbf{R})$. But $\overline{\mathcal{A}}$ is closed, so $\overline{\mathcal{A}} = C(X, \mathbf{R})$. That is, $\mathcal{A}$ is dense in $C(X, \mathbf{R})$. □

We have already observed that if $[a, b] \subseteq \mathbf{R}$ is a closed interval, then the algebra of polynomials on $[a, b]$ separates points and vanishes nowhere. Therefore, we immediately recover the Weierstrass approximation theorem.

Corollary 3.4.8 (Weierstrass)**.** *For any closed interval $[a, b] \subseteq \mathbf{R}$, the algebra of polynomials $P([a, b])$ is dense in $C([a, b], \mathbf{R})$.*

We now turn our attention to a related question that is often useful in practice: what if we consider continuous functions $f : X \to \mathbf{C}$ instead? That is, does the Stone-Weierstrass theorem hold for $C(X) = C(X, \mathbf{C})$ as well? Unfortunately, a subalgebra of $C(X)$ that separates points and vanishes nowhere need not be dense in $C(X)$. (An example is described in Exercise 3.4.3.) We need an extra condition to make the theorem work for complex-valued functions.

Definition 3.4.9. A subalgebra $\mathcal{A} \subseteq C(X)$ is **self-adjoint** if for all $f \in \mathcal{A}$, the function $\overline{f} \in \mathcal{A}$, where

$$\overline{f}(x) = \overline{f(x)}$$

for all $x \in X$.

Theorem 3.4.10 (Stone-Weierstrass, complex version)**.** *Let X be a compact metric space. If $\mathcal{A} \subseteq C(X)$ is a self-adjoint subalgebra that separates points and vanishes nowhere, then $\mathcal{A}$ is dense in $C(X)$.*

Proof. Let $\mathcal{A}_{\mathbf{R}} \subseteq \mathcal{A}$ denote the set of all real-valued functions in $\mathcal{A}$. Notice that $\mathcal{A}_{\mathbf{R}}$ is nonempty, since $0 \in \mathcal{A}_{\mathbf{R}}$. Now observe that $\mathcal{A}_{\mathbf{R}}$ separates points—if $x_1, x_2 \in X$ with $x_1 \neq x_2$, there exists $f \in \mathcal{A}$ with

$$f(x_1) \neq f(x_2)$$

since $\mathcal{A}$ separates points. If we write $f = u + iv$ with $u, v \in \mathcal{A}_{\mathbf{R}}$, then

$$u = \tfrac{1}{2}\big(f + \overline{f}\big) \quad \text{and} \quad v = \tfrac{1}{2}\big(f - \overline{f}\big)$$

both belong to $\mathcal{A}_{\mathbf{R}}$, since $\mathcal{A}$ is self-adjoint. Moreover, we must have either $u(x_1) \neq u(x_2)$ or $v(x_1) \neq v(x_2)$, so $\mathcal{A}_{\mathbf{R}}$ separates points. It is also easy to check that $\mathcal{A}_{\mathbf{R}}$ vanishes nowhere: given $x \in X$, there exists $f \in \mathcal{A}$ such that $f(x) \neq 0$, so $|f(x)| \neq 0$. But observe that

$$|f| = f \cdot \overline{f}$$

belongs to $\mathcal{A}_{\mathbf{R}}$. It now follows from the real version of the Stone-Weierstrass theorem that $\mathcal{A}_{\mathbf{R}}$ is dense in $C(X, \mathbf{R})$.

Now let $f \in C(X)$, and write $f = u + iv$ with $u, v \in \mathcal{A}_{\mathbf{R}}$. Let $\varepsilon > 0$ be given. Then there exist functions $u_0, v_0 \in \mathcal{A}_{\mathbf{R}}$ such that

$$\|u - u_0\|_\infty < \frac{\varepsilon}{\sqrt{2}}, \quad \|v - v_0\|_\infty < \frac{\varepsilon}{\sqrt{2}}.$$

Put $g = u_0 + iv_0$. Then

$$\begin{aligned} \|f - g\|_\infty &= \sup_{x \in X} |f(x) - g(x)| \\ &= \sup_{x \in X} \sqrt{[u(x) - u_0(x)]^2 + [v(x) - v_0(x)]^2} \\ &< \varepsilon, \end{aligned}$$

so $\mathcal{A}$ is dense in $C(X)$. □

Exercises for Section 3.4

Exercise 3.4.1. Let X be a metric space containing at least two distinct points. For each $x \in X$, define a function $f_x : X \to \mathbf{R}$ by

$$f_x(y) = d(x, y).$$

Let $\mathcal{F} = \{f_x : x \in X\}$. Prove that $\mathcal{F}$ separates points and vanishes nowhere.

Exercise 3.4.2. Let $X = [0, 2\pi]$, and let $\mathcal{A} \subseteq C(X, \mathbf{R})$ denote the set of all functions of the form

$$f(x) = \sum_{n=0}^{N} a_n \cos nx + b_n \sin nx, \quad a_n, b_n \in \mathbf{R}.$$

(a) Prove that $\mathcal{A}$ is a subalgebra of $C(X, \mathbf{R})$.

(b) Define the subalgebra $\mathcal{B} \subseteq C(X, \mathbf{R})$ of **periodic functions** by

$$\mathcal{B} = \{f \in C(X, \mathbf{R}) : f(0) = f(2\pi)\}.$$

Let $\mathbf{T} = \{z \in \mathbf{C} : |z| = 1\}$ denote the unit circle in the complex plane. By composing with the map

$$[0, 2\pi] \to \mathbf{T}, \quad t \mapsto e^{it},$$

we can identify $\mathcal{B}$ with $C(\mathbf{T}, \mathbf{R})$. Thus the Stone-Weierstrass theorem implies that any subalgebra $\mathcal{A} \subseteq \mathcal{B}$ that separates points (excluding 0 and 2π) and vanishes nowhere must be dense in $\mathcal{B}$. Use these facts to prove that $\mathcal{A}$ is dense in $\mathcal{B}$.

Exercise 3.4.3. Let $\mathbf{T} = \{z \in \mathbf{C} : |z| = 1\}$ denote the unit circle in the complex plane. Recall that every element of $\mathbf{T}$ can be written in the form e^{it}, where $t \in \mathbf{R}$. Let $\mathcal{A} \subseteq C(\mathbf{T})$ denote the algebra of functions of the form

$$f(e^{it}) = \sum_{n=0}^{N} a_n e^{int}, \quad a_n \in \mathbf{C}.$$

(a) Show that $\mathcal{A}$ separates points and vanishes nowhere.

(b) Show that for every function $f \in \mathcal{A}$,

$$\int_0^{2\pi} f(e^{it}) e^{it}\, dt = 0, \tag{3.5}$$

and that the same is true for all $f \in \overline{\mathcal{A}}$.

(c) Give an example of a function $g \in C(\mathbf{T})$ for which (3.5) does not hold. Conclude that $\mathcal{A}$ is not dense in $C(\mathbf{T})$.

Chapter 4

Lebesgue Measure and Integration

We are now prepared to address the problems that we laid out in the introductory chapter regarding integration on the real line. Recall that the Riemann integral has some fatal drawbacks, particularly with respect to convergence of sequences of functions. Lebesgue discovered a way to rectify these issues by overhauling Riemann's theory of integration. His idea was to partition the *range* of a function rather than its domain, as one usually does when building Riemann sums. As we will soon see, Lebesgue's idea amounts to approximating an arbitrary function with piecewise constant functions, the integrals of which are relatively easy to define.

Suppose we have a (nonnegative) constant function $f \equiv c$ defined on some set $E \subseteq \mathbf{R}$. If E happens to be an interval $[a, b]$, the integral of f over E is simply

$$\int_E f = c \cdot (b - a),$$

since the "area under the curve" is just a rectangle. If E is a more complicated set, then $\int_E f$ should still represent the area under the graph of f. However, this area could be hard to determine for an arbitrary set E. Morally, it ought to still be the product of c and the "length" of E. Therefore, Lebesgue's theory of integration requires us to generalize the notion of length from intervals to more interesting subsets of $\mathbf{R}$.

The need to "measure" subsets of $\mathbf{R}$ leads us to the concept of **Lebesgue measure**. We seek a function $\mu : \mathcal{P}(\mathbf{R}) \to [0, \infty]$ that has the following properties.

1. For any interval $[a, b]$, we have $\mu([a, b]) = b - a$. That is, the measure of an interval is simply its length.

2. If $E_1, E_2 \subseteq \mathbf{R}$ are disjoint, then it would make sense to have

$$\mu(E_1 \cup E_2) = \mu(E_1) + \mu(E_2).$$

In fact, we want this property to hold for *countable* unions of disjoint sets: if $\{E_i\}_{i=1}^{\infty}$ is a countable collection of pairwise disjoint sets, then

$$\mu\left(\bigcup_{i=1}^{\infty} E_i\right) = \sum_{i=1}^{\infty} \mu(E_i).$$

We will call this property **countable additivity**.

3. We would expect μ to be **translation-invariant**, meaning the measure of a set does not change if we shift the set by some fixed amount.

We will see that it is impossible to define a function on $\mathcal{P}(\mathbf{R})$ with all of these properties. Our solution will be to restrict the domain of μ; we will find a proper subset $\mathcal{L}$ of $\mathcal{P}(\mathbf{R})$ on which there is a function μ with the requisite properties. The sets that belong to $\mathcal{L}$ will be deemed **measurable sets**.

The content of this chapter is based on several sources. At times we will closely parallel the approach of [RF10], but we will also draw ideas from [WZ77] and [Fol99].

4.1 Outer Measure

We need to determine how to construct a function μ that satisfies the properties we described in the introduction. We will first construct a precursor to Lebesgue measure, which we call the **Lebesgue outer measure**, or simply **outer measure**. To see how it will be built, we begin with an instructive example.

Example 4.1.1. Recall that the Cantor set C is defined as follows. We begin by setting $C_0 = [0, 1]$, and then for each $n \in \mathbf{N}$ we construct a set C_n by removing the middle third of each interval appearing in C_{n-1}. In other words, we have

$$\begin{aligned}
C_0 &= [0, 1] \\
C_1 &= \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right] \\
C_2 &= \left[0, \tfrac{1}{9}\right] \cup \left[\tfrac{2}{9}, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, \tfrac{7}{9}\right] \cup \left[\tfrac{8}{9}, 1\right] \\
&\vdots
\end{aligned}$$

We then set

$$C = \bigcap_{n=1}^{\infty} C_n.$$

How can we determine the "length" of the Cantor set? Well, for each n we see that C_n is just a union of disjoint intervals. Hence the "length" of C_n should just be the sum of the lengths of these intervals. If we let $v([a, b]) = b - a$ denote the length of an interval, we have

$$v(C_0) = 1$$

$$\begin{aligned} v(C_1) &= \tfrac{1}{3} + \tfrac{1}{3} = \tfrac{2}{3} \\ v(C_2) &= \tfrac{1}{9} + \tfrac{1}{9} + \tfrac{1}{9} + \tfrac{1}{9} = \tfrac{4}{9} \\ &\vdots \\ v(C_n) &= \left(\tfrac{2}{3}\right)^n. \end{aligned}$$

Since $v(C_n) \to 0$ as $n \to \infty$, we would expect the measure of C to be

$$\mu(C) = \lim_{n\to\infty} v(C_n) = 0.$$

In the example above, we estimated the measure of the Cantor set by covering it with intervals. It is precisely this idea that we will use to define outer measure. First, we introduce some notation for the sake of convenience. For brevity, we will let

$$\mathcal{E} = \left\{ I \subseteq \mathbf{R} : I = [a,b] \text{ for some } a, b \in \mathbf{R} \right\}$$

denote the set of all closed, bounded intervals in $\mathbf{R}$.[1] Also, we denote the length of a closed interval $I = [a,b]$ by

$$v(I) = b - a,$$

as in the previous example.

Definition 4.1.2. Given a set $E \subseteq \mathbf{R}$, we define the **Lebesgue outer measure** of E to be

$$\mu^*(E) = \inf \left\{ \sum_{k=1}^{\infty} v(I_k) : I_k \in \mathcal{E} \text{ and } E \subseteq \bigcup_{k=1}^{\infty} I_k \right\}.$$

Notice that we determine the outer measure of a set E by considering all possible coverings of E by countable collections of closed intervals. This definition makes sense for any subset of $\mathbf{R}$, so we have a function $\mu^* : \mathcal{P}(\mathbf{R}) \to [0,\infty]$. However, it is not the finished product—we will need to make some modifications in order to obtain the true definition of Lebesgue measure.

Example 4.1.3. Any singleton has outer measure 0. Let $x \in \mathbf{R}$, and given $\varepsilon > 0$, define $I_\varepsilon = [x - \frac{\varepsilon}{2}, x + \frac{\varepsilon}{2}]$. Then we have $\{x\} \subseteq I_\varepsilon$ and $v(I_\varepsilon) = \varepsilon$. Thus $\mu^*(\{x\}) \leq \varepsilon$ for all $\varepsilon > 0$, meaning that

$$\mu^*(\{x\}) = 0.$$

Example 4.1.4. The work we did in Example 4.1.1 shows that the Cantor set has outer measure 0, since

$$0 \leq \mu^*(C) \leq \inf_n v(C_n) = 0.$$

[1] Don't read too much into this notation just yet. We have chosen the symbol $\mathcal{E}$ for "elementary set". We will soon consider more general notions of outer measure an arbitrary set X, and we will cover sets with elements of some collection $\mathcal{E} \subseteq \mathcal{P}(X)$ of elementary sets.

Since Lebesgue measure is supposed to generalize the notion of length for intervals, we would expect the outer measure of an interval to agree with its length.

Proposition 4.1.5. *If $I \subseteq \mathbf{R}$ is a closed, bounded interval, then $\mu^*(I) = v(I)$.*

Proof. Since I constitutes a covering of itself, we certainly have $\mu^*(I) \leq v(I)$. Suppose that $\{I_k\}_{k=1}^{\infty}$ is an arbitrary covering of I by closed, bounded intervals. Let $\varepsilon > 0$ be given, and for each k find $J_k \in \mathcal{E}$ such that

$$I_k \subseteq J_k^\circ$$

and

$$v(J_k) \leq (1+\varepsilon) v(I_k).$$

Since $\{J_k^\circ\}_{k=1}^{\infty}$ is an open cover of the compact interval I, there is a finite subcover $\{J_k^\circ\}_{k=1}^{N}$. Certainly we have

$$v(I) \leq \sum_{k=1}^{N} v(J_k) \leq (1+\varepsilon) \sum_{k=1}^{N} v(I_k) \leq (1+\varepsilon) \sum_{k=1}^{\infty} v(I_k),$$

which holds for all $\varepsilon > 0$. Thus

$$v(I) \leq \sum_{k=1}^{\infty} v(I_k).$$

Since $\{I_k\}_{k=1}^{\infty}$ is an arbitrary covering of I, it follows that $v(I) \leq \mu^*(I)$. Hence $\mu^*(I) = v(I)$. □

Another reasonable property that we would expect the outer measure to have is the following, which we term *monotonicity.*

Proposition 4.1.6. *If $E_1 \subseteq E_2$, then $\mu^*(E_1) \leq \mu^*(E_2)$.*

Proof. If $\{I_k\}_{k=1}^{\infty} \subseteq \mathcal{E}$ is a covering of E_2 by intervals, then it is also a covering of E_1. Thus

$$\mu^*(E_1) \leq \sum_{k=1}^{\infty} v(I_k).$$

Since this holds for all coverings of E_2, we must have $\mu^*(E_1) \leq \mu^*(E_2)$. □

We mentioned in the introduction that Lebesgue measure should be countably additive. As a first step toward this result, we show that outer measure is *subadditive.*

Proposition 4.1.7. *If $\{E_i\}_{i=1}^{\infty}$ is a countable collection of sets, then*

$$\mu^*\left(\bigcup_{i=1}^{\infty} E_i\right) \leq \sum_{i=1}^{\infty} \mu^*(E_i).$$

Proof. For simplicity, let $E = \bigcup_{i=1}^{\infty} E_i$. If $\mu^*(E_i) = \infty$ for some n, then the inequality clearly holds. Therefore, we can assume $\mu^*(E_i) < \infty$ for all i.

Let $\varepsilon > 0$ be given. For each i, let $\{I_{k_i}^i\}_{k_i=1}^{\infty} \subseteq \mathcal{E}$ be a covering of E_i satisfying

$$\sum_{k_i=1}^{\infty} v(I_{k_i}^i) \leq \mu^*(E_i) + \frac{\varepsilon}{2^i}.$$

Then $\bigcup_{i=1}^{\infty}\{I_{k_i}^i\}_{k_i=1}^{\infty}$ is a covering of E, and

$$\sum_{i=1}^{\infty}\sum_{k_i=1}^{\infty} v(I_{k_i}^i) \leq \sum_{i=1}^{\infty}\left(\mu^*(E_i) + \tfrac{\varepsilon}{2^i}\right) = \varepsilon + \sum_{i=1}^{\infty} \mu^*(E_i).$$

Thus

$$\mu^*(E) \leq \varepsilon + \sum_{i=1}^{\infty} \mu^*(E_i)$$

for all $\varepsilon > 0$, and the result follows. □

Corollary 4.1.8. *If $E \subseteq \mathbf{R}$ is countable, then $\mu^*(E) = 0$.*

Proof. Suppose E is countable, and write

$$E = \bigcup_{x \in E} \{x\}.$$

Then by the previous proposition, we have

$$\mu^*(E) \leq \sum_{x \in E} \mu^*(\{x\}) = 0$$

since singletons have outer measure 0. □

The last two properties that we will investigate relate outer measure to the standard topology on the real line.

Proposition 4.1.9. *Let $E \subseteq \mathbf{R}$. For each $\varepsilon > 0$, there is an open set $U \subseteq \mathbf{R}$ such that $E \subseteq U$ and*

$$\mu^*(U) \leq \mu^*(E) + \varepsilon.$$

Proof. Let $\varepsilon > 0$ be given, and suppose $\{I_k\}_{k=1}^\infty \subseteq \mathcal{E}$ is a covering of E satisfying

$$\sum_{k=1}^\infty v(I_k) \leq \mu^*(E) + \frac{\varepsilon}{2}.$$

For each k, choose $J_k \in \mathcal{E}$ such that $I_k \subseteq J_k^\circ$ and

$$v(J_k) \leq v(I_k) + \frac{\varepsilon}{2^{k+1}}.$$

Set $U = \bigcup_{k=1}^\infty J_k^\circ$. Then U is open and

$$\begin{aligned} \mu^*(U) &\leq \sum_{k=1}^\infty v(J_k) \\ &\leq \sum_{k=1}^\infty \left(v(I_k) + \frac{\varepsilon}{2^{k+1}} \right) \\ &= \frac{\varepsilon}{2} + \sum_{k=1}^\infty v(I_k) \\ &\leq \mu^*(E) + \varepsilon, \end{aligned}$$

as desired. □

In particular, note that Proposition 4.1.9 implies that

$$\mu^*(E) = \inf\big\{\mu^*(U) : U \text{ is open and } E \subseteq U\big\}.$$

In other words, μ^* is **outer regular**—the outer measure of a set E can be computed by approximating E from the outside with open sets. If we instead replace "open" with "G_δ", then we can actually hit the outer measure of E on the nose.

Proposition 4.1.10. *If $E \subseteq \mathbf{R}$, there is a G_δ-set A such that $E \subseteq A$ and $\mu^*(A) = \mu^*(E)$.*

Proof. For each $n \in \mathbf{N}$, there is an open set U_n such that $E \subseteq U_n$ and

$$\mu^*(U_n) \leq \mu^*(E) + \frac{1}{n}$$

by the Proposition 4.1.9. Put $A = \bigcap_{n=1}^{\infty} U_n$. Then A is a G_δ-set containing E, and

$$\mu^*(E) \leq \mu^*(A) \leq \mu^*(U_n) \leq \mu^*(E) + \frac{1}{n}$$

for all $n \in \mathbf{N}$. Therefore, $\mu^*(E) = \mu^*(A)$. □

Exercises for Section 4.1

Exercise 4.1.1 ([WZ77], Exercise 3.4 modified)**.** Define a Cantor-like set $K \subseteq [0,1]$ as follows: begin with $K_0 = [0,1]$, and define K_n inductively by removing the middle fifth of each interval in K_{n-1}. For example, the first few such sets are

$$\begin{aligned} K_0 &= [0,1] \\ K_1 &= [0, \tfrac{2}{5}] \cup [\tfrac{3}{5}, 1] \\ K_2 &= [0, \tfrac{4}{25}] \cup [\tfrac{6}{25}, \tfrac{2}{5}] \cup [\tfrac{3}{5}, \tfrac{19}{25}] \cup [\tfrac{21}{25}, 1] \\ &\vdots \end{aligned}$$

Now set $K = \bigcap_{n=1}^{\infty} K_n$. Compute the outer measure of K.

Exercise 4.1.2 ([RF10], Exercise 2.6)**.** Let E denote the set of irrational numbers in $[0,1]$. Prove that $\mu^*(E) = 1$.

Exercise 4.1.3 ([WZ77], Exercise 3.5 modified)**.** Define a Cantor-like set $K \subseteq [0,1]$ as follows: begin with $K_0 = [0,1]$, and define K_n inductively by removing an interval of length $\frac{1}{2 \cdot 3^k}$ from the middle of each interval in K_{n-1}. For example, the first few such sets are

$$\begin{aligned} K_0 &= [0,1] \\ K_1 &= [0, \tfrac{5}{12}] \cup [\tfrac{7}{12}, 1] \\ K_2 &= [0, \tfrac{13}{72}] \cup [\tfrac{17}{72}, \tfrac{5}{12}] \cup [\tfrac{7}{12}, \tfrac{55}{72}] \cup [\tfrac{59}{72}, 1] \\ &\vdots \end{aligned}$$

Now set $K = \bigcap_{n=1}^{\infty} K_n$. Compute the outer measure of K.

4.2 Lebesgue Measurability

We now begin to single out the collection of sets on which Lebesgue measure will actually be defined. We will use the definition from [RF10], which is originally due to Carathéodory.

Definition 4.2.1. A set $E \subseteq \mathbf{R}$ is **Lebesgue measurable** if for every set $A \subseteq \mathbf{R}$ we have

$$\mu^*(A) = \mu^*(A \cap E) + \mu^*(A \cap E^c).$$

Notice that since $A = (A \cap E) \cup (A \cap E^c)$ and outer measure is subadditive, we always have

$$\mu^*(A) \leq \mu^*(A \cap E) + \mu^*(A \cap E^c).$$

Therefore, to check that a set E is measurable, it suffices to prove that

$$\mu^*(A \cap E) + \mu^*(A \cap E^c) \leq \mu^*(A)$$

for all $A \subseteq \mathbf{R}$.

Remark 4.2.2. Another definition for measurability, which is perhaps more intuitive, is given in [WZ77]: a set E is measurable if for any $\varepsilon > 0$, there is an open set U containing E such that

$$\mu^*(U \backslash E) < \varepsilon.$$

In other words, a set is measurable if it is close to being an open set. Each definition has its advantages. For example, it is immediate from the second definition that any open set is measurable. On the other hand, it is much easier to obtain basic properties of Lebesgue measure directly from the Carathéodory condition. We will see later that these two definitions are equivalent.

Definition 4.2.3. If $E \subseteq \mathbf{R}$ is measurable, we define its **Lebesgue measure**, denoted $\mu(E)$, to be equal to the outer measure of E:

$$\mu(E) = \mu^*(E).$$

In other words, all we have done to construct Lebesgue measure is restrict the domain of μ^* to a particular subset of $\mathcal{P}(\mathbf{R})$. We will denote the set of all Lebesgue measurable sets by $\mathcal{L}$. We will see later that $\mathcal{L}$ is a *proper* subset of $\mathcal{P}(\mathbf{R})$.

As mentioned above, we can obtain some properties of Lebesgue measure somewhat quickly from the Carathédory condition for measurability.

Proposition 4.2.4. *A set $E \subseteq \mathbf{R}$ is measurable if and only if E^c is measurable.*

Proof. Clearly the Carathéodory condition is symmetric in E and E^c. That is, it holds for E if and only if it holds for E^c. $\square$

Proposition 4.2.5. *If $\mu^*(E) = 0$, then E is measurable.*

Proof. Let $A \subseteq \mathbf{R}$. Since $A \cap E \subseteq E$, we have

$$\mu^*(A \cap E) \leq \mu^*(E) = 0.$$

Hence $\mu^*(A \cap E) = 0$. Similarly, $A \cap E^c \subseteq A$, so

$$\mu^*(A \cap E^c) \leq \mu^*(A).$$

Therefore,

$$\begin{aligned} \mu^*(A) &\leq \mu^*(A \cap E) + \mu^*(A \cap E^c) \\ &= 0 + \mu^*(A \cap E^c) \\ &\leq \mu^*(A). \end{aligned}$$

Thus $\mu^*(A) = \mu^*(A \cap E) + \mu^*(A \cap E^c)$, so E is measurable. □

As an initial step toward the countable additivity result that we have promised, we now prove that the union of a finite collection of measurable sets is measurable.

Proposition 4.2.6. *Let $\{E_i\}_{i=1}^n$ be a finite collection of measurable sets. Then $\bigcup_{i=1}^n E_i$ is measurable.*

Proof. We proceed by induction. If $n = 1$, the result is trivial. For the inductive step, it suffices to first show that the union of two measurable sets is measurable.

Suppose $E_1, E_2 \subseteq \mathbf{R}$ are measurable. Since E_2 is measurable, for any $A \subseteq \mathbf{R}$ we have

$$\begin{aligned} \mu^*(A \cap E_1^c) &= \mu^*(A \cap E_1^c \cap E_2) + \mu^*(A \cap E_1^c \cap E_2^c) \\ &= \mu^*(A \cap E_1^c \cap E_2) + \mu^*(A \cap (E_1 \cup E_2)^c). \end{aligned}$$

Also, notice that

$$A \cap (E_1 \cup E_2) = (A \cap E_1) \cup (A \cap E_2) = (A \cap E_1) \cup (A \cap E_2 \cap E_1^c)$$

so

$$\mu^*(A \cap (E_1 \cup E_2)) \leq \mu^*(A \cap E_1) + \mu^*(A \cap E_2 \cap E_1^c).$$

Putting this all together, we get

$$\begin{aligned} \mu^*(A) &= \mu^*(A \cap E_1) + \mu^*(A \cap E_1^c) \\ &= \mu^*(A \cap E_1) + \mu^*(A \cap E_1^c \cap E_2) + \mu^*(A \cap (E_1 \cup E_2)^c) \end{aligned}$$

$$\geq \mu^*(A \cap (E_1 \cup E_2)) + \mu^*(A \cap (E_1 \cup E_2)^c)$$

since E_1 is measurable. It follows that $E_1 \cup E_2$ is measurable.

Now for the inductive hypothesis we assume $\bigcup_{i=1}^{n-1} E_i$ is measurable. Then we observe that

$$\bigcup_{i=1}^{n} E_i = \left(\bigcup_{i=1}^{n-1} E_i\right) \cup E_n$$

is a union of two measurable sets, so it is measurable by the discussion above. □

Corollary 4.2.7. *The intersection of a finite family of measurable sets is measurable.*

Proof. The result is immediate from Proposition 4.2.6, De Morgan's law, and the fact that a set is measurable if and only if its complement is measurable. □

Next, we would like to extend the conclusion of Proposition 4.2.6 to countable unions of measurable sets. Before we can do this, we need an additional result regarding finite unions.

Proposition 4.2.8. *Let* $\{E_i\}_{i=1}^n$ *be a finite collection of pairwise disjoint measurable sets, and let* $A \subseteq \mathbf{R}$ *be any set. Then*

$$\mu^*\left(A \cap \bigcup_{i=1}^{n} E_i\right) = \sum_{i=1}^{n} \mu^*(A \cap E_i).$$

Proof. Again, we proceed by induction on n. The result clearly holds when $n = 1$. Assume it holds for $n - 1$. Since E_n is measurable, we have

$$\mu^*\left(A \cap \bigcup_{i=1}^{n} E_i\right) = \mu^*\left(A \cap \bigcup_{i=1}^{n} E_i \cap E_n\right) + \mu^*\left(A \cap \bigcup_{i=1}^{n} E_i \cap E_n^c\right).$$

However, notice that

$$A \cap \left(\bigcup_{i=1}^{n} E_i\right) \cap E_n = A \cap E_n,$$

since the E_i are pairwise disjoint. Similarly,

$$A \cap \left(\bigcup_{i=1}^{n} E_i\right) \cap E_n^c = A \cap \left(\bigcup_{i=1}^{n-1} E_i\right).$$

Therefore,

$$\mu^*\left(A \cap \bigcup_{i=1}^{n} E_i\right) = \mu^*(A \cap E_n) + \mu^*\left(A \cap \bigcup_{i=1}^{n-1} E_i\right)$$
$$= \mu^*(A \cap E_n) + \sum_{i=1}^{n-1} \mu^*(A \cap E_i)$$

by the inductive hypothesis, and the result follows. □

By taking $A = \mathbf{R}$, we immediately obtain the following useful corollary.

Corollary 4.2.9. *Let $\{E_i\}_{i=1}^n$ be a finite collection of pairwise disjoint measurable sets. Then*

$$\mu\left(\bigcup_{i=1}^{n} E_i\right) = \sum_{i=1}^{n} \mu(E_i).$$

Using Proposition 4.2.8, we can now establish that countable unions of measurable sets are measurable. First we need to make an observation, which will lend itself to a proof technique that we will use repeatedly. Suppose $\{E_i\}_{i=1}^\infty$ is a countable collection of measurable sets. We can construct a collection $\{A_i\}_{i=1}^\infty$ of *disjoint* measurable sets as follows: set

$$A_1 = E_1$$
$$A_2 = E_2 \backslash E_1$$
$$A_3 = E_3 \backslash (E_1 \cup E_2),$$

and in general,

$$A_i = E_i \backslash (E_1 \cup E_2 \cup \cdots \cup E_{i-1}).$$

Notice that each A_i is measurable, since we know measurability is preserved under finite unions, intersections, and complements. If $i < j$, then we have $A_i \subseteq E_i$, while $E_i \subseteq A_j^c$, so $A_i \cap A_j = \emptyset$. Thus the A_i are pairwise disjoint. Furthermore, it is clear that

$$\bigcup_{i=1}^{\infty} A_i \subseteq \bigcup_{i=1}^{\infty} E_i$$

since $A_i \subseteq E_i$ for all i. On the other hand, if $x \in \bigcup_{i=1}^\infty E_i$, then $x \in E_i$ for some i. Let i_0 be the smallest such i. Then $x \notin E_j$ for $j < i_0$, so $x \in A_{i_0}$. It follows that $x \in \bigcup_{i=1}^\infty A_i$, so

$$\bigcup_{i=1}^{\infty} A_i = \bigcup_{i=1}^{\infty} E_i.$$

Theorem 4.2.10. *Let $\{E_i\}_{i=1}^{\infty}$ be a countable collection of measurable sets. Then $\bigcup_{i=1}^{\infty} E_i$ is measurable.*

Proof. For simplicity, we let $E = \bigcup_{i=1}^{\infty} E_i$. Replace the collection $\{E_i\}_{i=1}^{\infty}$ with a collection of disjoint sets $\{A_i\}_{i=1}^{\infty}$ whose union is E, as in the discussion above. For each $n \in \mathbf{N}$, put

$$F_n = \bigcup_{i=1}^{n} A_i.$$

Then F_n is measurable and $E^c \subseteq E_n^c \subseteq F_n^c$ for all n. Therefore, for any $A \subseteq \mathbf{R}$ we have

$$\begin{aligned}\mu^*(A) &= \mu^*(A \cap F_n) + \mu^*(A \cap F_n^c) \\ &\geq \mu^*(A \cap F_n) + \mu^*(A \cap E^c)\end{aligned}$$

for each n. By Proposition 4.2.8, we have

$$\mu^*(A \cap F_n) = \mu^*\left(\bigcup_{i=1}^{n} A \cap A_i\right) = \sum_{i=1}^{n} \mu^*(A \cap A_i)$$

so

$$\mu^*(A) \geq \sum_{i=1}^{n} \mu^*(A \cap A_i) + \mu^*(A \cap E^c).$$

The left hand side does not depend on n, so

$$\begin{aligned}\mu^*(A) &\geq \sum_{i=1}^{\infty} \mu^*(A \cap A_i) + \mu^*(A \cap E^c) \\ &\geq \mu^*(A \cap E) + \mu^*(A \cap E^c).\end{aligned}$$

Therefore, E is measurable. □

The proof of the following corollary is identical to that of Corollary 4.2.7.

Corollary 4.2.11. *If $\{E_i\}_{i=1}^{\infty}$ is a countable family of measurable sets, then $\bigcap_{i=1}^{\infty} E_i$ is measurable.*

We have thus far shown that the set $\mathcal{L}$ of Lebesgue measurable sets is closed under the operations of taking complements, countable unions, and countable intersections. A collection of sets with these properties has a special name.

Definition 4.2.12. Let X be a set. A collection of sets $\Sigma \subseteq \mathcal{P}(X)$ is called a **σ-algebra** if it is closed under complements and countable unions.

By combining Proposition 4.2.4 and Theorem 4.2.10, we have already proven the following theorem regarding Lebesgue measurable sets.

Theorem 4.2.13. *The set $\mathcal{L}$ of Lebesgue measurable sets is a σ-algebra.*

As we have already observed with measurable sets, notice that a σ-algebra is automatically closed under countable intersections thanks to De Morgan's law.

4.2.1 The Borel σ-algebra

We now consider another natural σ-algebra of subsets of $\mathbf{R}$. First, let X be a set, and suppose $\{\Sigma_i\}_{i\in I}$ is a collection of σ-algebras on X. It is straightforward to check (Exercise 4.2.2) that

$$\Sigma = \bigcap_{i\in I} \Sigma_i$$

is a σ-algebra. Therefore, if $E \subseteq \mathcal{P}(X)$, we can intersect all the σ-algebras containing E to obtain a σ-algebra $\Sigma(E)$, which is necessarily the smallest possible σ-algebra containing E. We call $\Sigma(E)$ the **σ-algebra generated by** E. The σ-algebra we plan to investigate is the one generated by the collection of all open sets in $\mathbf{R}$.

Definition 4.2.14. The smallest σ-algebra $\mathcal{B}$ containing all the open subsets of $\mathbf{R}$ is called the **Borel σ-algebra**. The elements of $\mathcal{B}$ are called **Borel sets**.

Notice that the Borel σ-algebra necessarily contains all closed sets, all G_δ sets, and all F_σ sets. Consequently, it contains countable unions of G_δ sets (called $G_{\delta\sigma}$ sets), countable intersections of F_σ sets (called $F_{\sigma\delta}$ sets), and so on. We will further investigate the hierarchy of Borel sets later on. First we will show that every open set in $\mathbf{R}$ is measurable, from which it will follow that $\mathcal{B} \subseteq \mathcal{L}$.

Proposition 4.2.15. *Every open interval (a, b) in $\mathbf{R}$ is measurable.*

Proof. We begin by showing that for all $a \in \mathbf{R}$, the interval (a, ∞) is measurable. Let $A \subseteq \mathbf{R}$. Notice that if $\mu^*(A) = \infty$, then the Carathéodory condition holds trivially. Therefore, we may assume that $\mu^*(A) < \infty$. For simplicity, we define

$$A_1 = (a, \infty) \cap A, \quad A_2 = (a, \infty)^c \cap A.$$

Let $\varepsilon > 0$ be given, and choose a covering $\{I_k\}_{k=1}^{\infty}$ of A by closed intervals such that

$$\sum_{k=1}^{\infty} v(I_k) \leq \mu^*(A) + \varepsilon.$$

For each k, put $I'_k = I_k \cap [a, \infty)$ and $I''_k = I_k \cap (-\infty, a]$. Then $v(I_k) = v(I'_k) + v(I''_k)$ and $\{I'_k\}_{k=1}^{\infty}$ and $\{I''_k\}_{k=1}^{\infty}$ cover A_1 and A_2, respectively. Thus

$$\mu^*(A_1) \leq \mu^*\left(\bigcup_{k=1}^{\infty} I'_k\right) \leq \sum_{k=1}^{\infty} v(I'_k)$$

and

$$\mu^*(A_2) \leq \mu^*\left(\bigcup_{k=1}^{\infty} I''_k\right) \leq \sum_{k=1}^{\infty} v(I''_k),$$

so

$$\mu^*(A_1) + \mu^*(A_2) \leq \sum_{k=1}^{\infty} \bigl[v(I'_k) + v(I''_k)\bigr] = \sum_{k=1}^{\infty} v(I_k) \leq \mu^*(A) + \varepsilon.$$

Since this holds for all $\varepsilon > 0$, it follows that $\mu^*(A_1) + \mu^*(A_2) \leq \mu^*(A)$, so (a, ∞) is measurable.

Now suppose $a, b \in \mathbf{R}$ with $a < b$. Then (a, ∞) is measurable, and we can write

$$(a, b) = (a, \infty) \cap (-\infty, b),$$

where $(-\infty, b) = [b, \infty)^c$ is measurable since

$$[b, \infty) = \{b\} \cup (b, \infty)$$

is measurable. Consequently, (a, b) is measurable. □

Corollary 4.2.16. *Every open set in* $\mathbf{R}$ *is measurable.*

Proof. Since $\mathbf{R}$ has a countable base consisting of open intervals, every open set can be written as a countable union of measurable sets. □

Since the Lebesgue σ-algebra $\mathcal{L}$ contains all open sets, it must contain the Borel σ-algebra. In other words:

Corollary 4.2.17. *Every Borel set in* $\mathbf{R}$ *is measurable.*

We have just seen that $\mathcal{B} \subseteq \mathcal{L}$. However, it turns out that there are Lebesgue measurable sets that are not Borel. There are a couple of ways to see this, and we will investigate one of them now. In the next two propositions, we show that $\mathcal{B}$ and $\mathcal{L}$ have different cardinalities. We write $\operatorname{card} A$ to represent the cardinality of a set A, and we use $\mathfrak{c}$ to denote the cardinality of $\mathbf{R}$.

Proposition 4.2.18. *The Borel σ-algebra $\mathcal{B}$ has cardinality $\mathfrak{c}$.*

Proof. The proof relies on an argument by transfinite induction. Recall that Ω denotes the set of all countable ordinals. Let $\mathcal{E}_0$ denote the collection of all open intervals in $\mathbf{R}$. Observe that $\mathcal{E}_0$ generates the Borel σ-algebra; that is, $\Sigma(\mathcal{E}_0) = \mathcal{B}$. Put $\mathcal{E}_1 = \mathcal{E}_0 \cup \{E : E^c \in \mathcal{E}_0\}$, and for each countable ordinal α, we do the following:

1. If α is a successor ordinal, define
$$\mathcal{E}_\alpha = \left\{ E \subseteq \mathbf{R} : E = \bigcup_{j=1}^{\infty} E_j \text{ where } \{E_j\}_{j=1}^{\infty} \subseteq \mathcal{E}_{\alpha-1} \text{ or } E^c \in \mathcal{E}_{\alpha-1} \right\}.$$
2. If α is a limit ordinal, define
$$\mathcal{E}_\alpha = \bigcup_{\beta<\alpha} \mathcal{E}_\beta.$$

We claim that $\mathcal{E}_\alpha \subseteq \Sigma(\mathcal{E}_0) = \mathcal{B}$ for every countable ordinal α. Clearly $\mathcal{E}_1 \subseteq \mathcal{B}$. Suppose α is a successor ordinal, and assume $\mathcal{E}_{\alpha-1} \subseteq \mathcal{B}$ and $E \in \mathcal{E}_\alpha$. Then either $E^c \in \mathcal{E}_{\alpha-1}$ or E is a countable union of elements of $\mathcal{E}_{\alpha-1}$, so $E \in \mathcal{B}$. Therefore, $\mathcal{E}_\alpha \subseteq \mathcal{B}$. Now suppose α is a limit ordinal and $\mathcal{E}_\beta \subseteq \mathcal{B}$ for all $\beta < \alpha$. If $E \in \mathcal{E}_\alpha$, then $E \in \mathcal{E}_\beta$ for some $\beta < \alpha$, so $E \in \mathcal{B}$ by assumption. Thus $\mathcal{E}_\alpha \subseteq \mathcal{B}$ for every $\alpha \in \Omega$ by transfinite induction. Hence $\bigcup_{\alpha\in\Omega} \mathcal{E}_\alpha \subseteq \mathcal{B}$.

Suppose that $\{E_j\}_{j=1}^{\infty} \subseteq \bigcup_{\alpha\in\Omega} \mathcal{E}_\alpha$. Then $E_j \subseteq \mathcal{E}_{\alpha_j}$ for some ordinal α_j. Every countable subset of Ω has a supremum in Ω, so we can put $\alpha = \sup \alpha_j$. Then $E_j \in \mathcal{E}_\alpha$ for all j, hence $\bigcup_{j=1}^{\infty} E_j \in \mathcal{E}_{\alpha+1}$. Thus $\bigcup_{\alpha\in\Omega} \mathcal{E}_\alpha$ is closed under countable unions, and it is clearly closed under complementation, so it is a σ-algebra. Since it contains $\mathcal{E}_0$, we have $\mathcal{B} \subseteq \bigcup_{\alpha\in\Omega} \mathcal{E}_\alpha$. But $\mathcal{B}$ is the smallest σ-algebra containing $\mathcal{E}_0$, so $\mathcal{B} = \bigcup_{\alpha\in\Omega} \mathcal{E}_\alpha$. Since $\operatorname{card} \mathcal{E}_\alpha \leq \mathfrak{c}$ for all $\alpha \in \Omega$ and $\operatorname{card} \Omega \leq \mathfrak{c}$, it follows that $\operatorname{card} \mathcal{B} \leq \mathfrak{c}$. But $\mathcal{E}_0 \subseteq \mathcal{B}$ and $\operatorname{card} \mathcal{E}_0 = \mathfrak{c}$, so $\operatorname{card} \mathcal{B} = \mathfrak{c}$. $\square$

Proposition 4.2.19. *The cardinality of the σ-algebra $\mathcal{L}$ is* $\operatorname{card} \mathcal{P}(\mathbf{R}) = 2^{\mathfrak{c}}$.

Proof. Let C denote the Cantor set. Recall that $\mu(C) = 0$, so every subset of C has outer measure zero and is thus measurable. But C is uncountable, so $\operatorname{card} \mathcal{P}(C) = \operatorname{card} \mathcal{P}(\mathbf{R})$. We have already argued that $\mathcal{P}(C) \subseteq \mathcal{L} \subseteq \mathcal{P}(\mathbf{R})$, so we must have $\operatorname{card} \mathcal{L} = \operatorname{card} \mathcal{P}(\mathbf{R})$. $\square$

Corollary 4.2.20. *There are Lebesgue measurable sets in* $\mathbf{R}$ *that are not Borel.*

Proof. Since $\operatorname{card} \mathcal{B} = \mathfrak{c} < 2^{\mathfrak{c}} = \operatorname{card} \mathcal{L}$, it follows that $\mathcal{B}$ is a proper subset of $\mathcal{L}$. $\square$

We will see another approach to demonstrating the existence of a non-Borel, Lebesgue measurable set later on.

Exercises for Section 4.2

Exercise 4.2.1 ([WZ77], Exercise 3.10)**.** Show that if $E_1, E_2 \subseteq \mathbf{R}$ are measurable, then

$$\mu(E_1) + \mu(E_2) = \mu(E_1 \cup E_2) + \mu(E_1 \cap E_2).$$

Exercise 4.2.2. Let X be a set, and let $\{\Sigma_i\}_{i \in I}$ be a collection of σ-algebras $\Sigma_i \subseteq P(X)$. Prove that

$$\Sigma = \bigcap_{i \in I} \Sigma_i$$

is a σ-algebra.

4.3 Further Properties of Lebesgue Measure

We now explore some additional useful properties of Lebesgue measure. First among them is the long-promised property of *countable additivity.*

Theorem 4.3.1. *Let* $\{E_i\}_{i=1}^{\infty}$ *be a countable collection of pairwise disjoint measurable sets. Then*

$$\mu\left(\bigcup_{i=1}^{\infty} E_i\right) = \sum_{i=1}^{\infty} \mu(E_i).$$

Proof. Notice first that we already have

$$\mu\left(\bigcup_{i=1}^{\infty} E_i\right) \le \sum_{i=1}^{\infty} \mu(E_i)$$

since outer measure is countably subadditive. On the other hand, for all $n \in \mathbf{N}$ we have

$$\bigcup_{i=1}^{n} E_i \subseteq \bigcup_{i=1}^{\infty} E_i,$$

so

$$\mu\left(\bigcup_{i=1}^{n} E_i\right) \le \mu\left(\bigcup_{i=1}^{\infty} E_i\right).$$

We know Lebesgue measure is additive on finite collections of disjoint sets, so the left hand side equals $\sum_{i=1}^{n} \mu(E_i)$. Thus

$$\sum_{i=1}^{n} \mu(E_i) \leq \mu\left(\bigcup_{i=1}^{\infty} E_i\right).$$

This inequality holds for all n, so letting $n \to \infty$ on the left hand side, we obtain

$$\sum_{i=1}^{\infty} \mu(E_i) \leq \mu\left(\bigcup_{i=1}^{\infty} E_i\right),$$

and the result follows. □

The next result allows one to easily compute the measure of a set that can be realized as a "limit" of a nested sequence of sets. It is often called *continuity of measure*.

Theorem 4.3.2. *Let $\{E_i\}_{i=1}^{\infty}$ be a countable collection of measurable sets.*

1. *If $E_1 \subseteq E_2 \subseteq E_3 \subseteq \cdots$, then*

$$\mu\left(\bigcup_{i=1}^{\infty} E_i\right) = \lim_{i \to \infty} \mu(E_i).$$

2. *If $E_1 \supseteq E_2 \supseteq E_3 \supseteq \cdots$ and $\mu(E_1) < \infty$, then*

$$\mu\left(\bigcap_{i=1}^{\infty} E_i\right) = \lim_{i \to \infty} \mu(E_i).$$

Proof. Suppose first that $E_1 \subseteq E_2 \subseteq E_3 \subseteq \cdots$ and let $E = \bigcup_{i=1}^{\infty} E_i$. If $\mu(E_i) = \infty$ for any i, then (1) holds trivially. Therefore, we may assume that $\mu(E_i) < \infty$ for all i. As in the discussion preceding Theorem 4.2.10, we replace $\{E_i\}_{i=1}^{\infty}$ with a disjoint family of measurable sets $\{A_i\}_{i=1}^{\infty}$ by setting $A_1 = E_1$ and

$$A_i = E_i \setminus \left(\bigcup_{j=1}^{i-1} E_j\right) = E_i \setminus E_{i-1}$$

for all $i \geq 2$. Since $\bigcup_{i=1}^{\infty} A_i = E$, we have

$$\mu(E) = \sum_{i=1}^{\infty} \mu(A_i)$$

$$= \mu(A_1) + \sum_{i=2}^{\infty} \big[\mu(E_i) - \mu(E_{i-1})\big],$$

where we have used the fact that

$$E_i = E_{i-1} \cup (E_i \backslash E_{i-1})$$

to write

$$\mu(E_i \backslash E_{i-1}) = \mu(E_i) - \mu(E_{i-1}).$$

The above sum is telescoping, and we have

$$\mu(E) = \mu(E_1) + \lim_{i\to\infty} \mu(E_i) - \mu(E_1) = \lim_{i\to\infty} \mu(E_i).$$

This establishes (1).

Now suppose $E_1 \supseteq E_2 \supseteq E_3 \supseteq \cdots$ and $\mu(E_1) < \infty$, and let $E = \bigcap_{i=1}^{\infty} E_i$. For each i, define $B_i = E_1 \backslash E_i$. Then $B_1 \subseteq B_2 \subseteq B_3 \subseteq \cdots$, so we have

$$\mu\left(\bigcup_{i=1}^{\infty} B_i\right) = \lim_{i\to\infty} \mu(B_i)$$

by (1). Now notice that

$$\begin{aligned}
\bigcup_{i=1}^{\infty} B_i &= \bigcup_{i=1}^{\infty} (E_1 \cap E_i^c) \\
&= E_1 \cap \left(\bigcup_{i=1}^{\infty} E_i^c\right) \\
&= E_1 \cap \left(\bigcap_{i=1}^{\infty} E_i\right)^c .
\end{aligned}$$

Thus

$$\mu\left(\bigcup_{i=1}^{\infty} B_i\right) = \mu(E_1) - \mu\left(\bigcap_{i=1}^{\infty} E_i\right),$$

while

$$\lim_{i\to\infty} \mu(B_i) = \lim_{i\to\infty} \big[\mu(E_1) - \mu(E_i)\big] = \mu(E_1) - \lim_{i\to\infty} \mu(E_i).$$

Hence

$$\mu(E_1) - \mu\left(\bigcap_{i=1}^{\infty} E_i\right) = \mu(E_1) - \lim_{i\to\infty} \mu(E_i),$$

and since $\mu(E_1)$ is finite, we have

$$\mu\left(\bigcap_{i=1}^{\infty} E_i\right) = \lim_{i\to\infty} \mu(E_i). \qquad \square$$

Remark 4.3.3. The hypothesis that $\mu(E_1) < \infty$ in part 2 of Theorem 4.3.2 is essential. In particular, take $E_i = [i, \infty)$ for each i. Then

$$E_1 \supseteq E_2 \supseteq E_3 \supseteq \cdots$$

and $\mu(E_i) = \infty$ for all i, so $\lim_{i\to\infty} \mu(E_i) = \infty$. However,

$$\bigcap_{i=1}^{\infty} E_i = \emptyset,$$

so the intersection has measure zero.

Next, we establish the equivalence between the Carathéodory condition and the definition of measurability given in [WZ77]. As a byproduct, we gain the ability to approximate measurable sets with open, closed, G_δ, and F_σ sets.

Theorem 4.3.4. *Let $E \subseteq \mathbf{R}$. The following are equivalent:*

1. *E is measurable.*
2. *For any $\varepsilon > 0$, there is an open set U containing E such that $\mu^*(U \backslash E) < \varepsilon$.*
3. *There exists a G_δ set A containing E with $\mu^*(A \backslash E) = 0$.*
4. *For any $\varepsilon > 0$, there is a closed set $F \subseteq E$ such that $\mu^*(E \backslash F) < \varepsilon$.*
5. *There exists an F_σ set $B \subseteq E$ with $\mu^*(E \backslash B) = 0$.*

Proof. ($1 \Rightarrow 2$) Suppose E is measurable. Assume first that $\mu(E)$ is finite, and let $\varepsilon > 0$ be given. Then there exists an open set U containing E with

$$\mu^*(U) \leq \mu^*(E) + \varepsilon.$$

Since E is measurable,

$$\mu^*(U) = \mu^*(U \cap E) + \mu^*(U \cap E^c),$$

so

$$\mu^*(U \backslash E) = \mu^*(U) - \mu^*(E) < \varepsilon.$$

Now assume $\mu(E) = \infty$. Then we can express E as a union $E = \bigcup_{i=1}^{\infty} E_i$, where $\mu(E_i) < \infty$ for all i and the E_i are pairwise disjoint. (Intersect E with each interval $[n, n+1)$ for $n \in \mathbf{Z}$, for example.) For each i, find an open set U_i containing E_i such that $\mu^*(U_i \backslash E_i) < \frac{\varepsilon}{2^i}$. Put $U = \bigcup_{i=1}^{\infty} U_i$. Then U is open and contains E, and

$$U \backslash E = \left(\bigcup_{i=1}^{\infty} U_i \right) \cap E^c = \bigcup_{i=1}^{\infty} [U_i \cap E^c] \subseteq \bigcup_{i=1}^{\infty} [U_i \cap E_i^c].$$

Thus

$$\mu^*(U\backslash E) \leq \sum_{i=1}^{\infty} \mu^*(U_i\backslash E_i) = \sum_{i=1}^{\infty} \frac{\varepsilon}{2^i} = \varepsilon,$$

so (2) holds.

($2 \Rightarrow 3$) For each $n \in \mathbf{N}$, we can find an open set U_n containing E such that $\mu^*(U\backslash E) < \frac{1}{n}$. Put $A = \bigcap_{n=1}^{\infty} U_n$. Then A is a G_δ set, $E \subseteq A$, and

$$\mu^*(A\backslash E) \leq \mu^*(U_n\backslash E) < \frac{1}{n}$$

for all $n \in \mathbf{N}$. Thus $\mu^*(A\backslash E) = 0$.

($3 \Rightarrow 1$) Let A be a G_δ set satisfying $E \subseteq A$ and $\mu^*(A\backslash E) = 0$. Then $A\backslash E$ is measurable, as is A. Since we can write

$$E = A\backslash(A\backslash E),$$

it follows that E is measurable.

($1 \Leftrightarrow 4$) Let $\varepsilon > 0$ be given. If E is measurable, then so is E^c. Thus there exists an open set $U \supseteq E^c$ such that $\mu^*(U\backslash E^c) < \varepsilon$ by (2). But $F = U^c$ is closed, $F \subseteq E$, and

$$E\backslash F = E \cap F^c = E \cap U = U \cap (E^c)^c = U\backslash E^c.$$

Thus $\mu^*(E\backslash F) < \varepsilon$. The converse is proved similarly.

($1 \Leftrightarrow 5$) If E is measurable, then so is E^c, so there is a G_δ set A with $E^c \subseteq A$ and $\mu^*(A\backslash E) = 0$. Put $B = A^c$. Then B is an F_σ with $B \subseteq E$. Moreover, $E\backslash B = A\backslash E^c$, so $\mu^*(E\backslash B) = 0$. The converse is similar. □

There is one more crucial property of Lebesgue measure that we have promised, and that we will now establish: Lebesgue measure is translation-invariant. First, we need to introduce some notation: given a set $A \subseteq \mathbf{R}$ and a point $x \in \mathbf{R}$, we let

$$A + x = \{y + x : y \in A\}$$

denote the set obtained by translating A by x.

Proposition 4.3.5. *Outer measure is translation invariant. That is, if $A \subseteq \mathbf{R}$ and $x \in \mathbf{R}$, then*

$$\mu^*(A + x) = \mu^*(A).$$

Proof. Let $A \subseteq \mathbf{R}$ and $x \in \mathbf{R}$. Let $\{I_k\}_{k=1}^{\infty}$ be a covering of A by closed intervals. Then clearly $\{I_k + x\}_{k=1}^{\infty}$ is a covering of $A + x$ by closed intervals, and $v(I_k + x) = v(I_k)$ for all k. Thus

$$\mu^*(A + x) \leq \sum_{k=1}^{\infty} v(I_k + x) = \sum_{k=1}^{\infty} v(I_k).$$

Taking the infimum of the right hand side over all possible coverings of A yields $\mu^*(A+x) \leq \mu^*(A)$. On the other hand, if $\{J_k\}_{k=1}^{\infty}$ is a covering of $A+x$, then $\{J_k - x\}_{k=1}^{\infty}$ is a covering of A, and a similar argument shows that $\mu^*(A) \leq \mu^*(A+x)$. It follows that $\mu^*(A+x) = \mu^*(A)$. $\square$

Proposition 4.3.6. *Suppose $E \subseteq \mathbf{R}$ is measurable. Then $E+x$ is measurable for all $x \in \mathbf{R}$.*

Proof. Let $\varepsilon > 0$ be given, and find an open set U containing E such that $\mu^*(U \backslash E) < \varepsilon$. Let $x \in \mathbf{R}$. Then it is straightforward to check that $U+x$ is open, $E+x \subseteq U+x$, and

$$(U+x)\backslash(E+x) = (U\backslash E) + x.$$

Thus

$$\begin{aligned} \mu^*((U+x)\backslash(E+x)) &= \mu^*((U\backslash E)+x) \\ &= \mu^*(U\backslash E) \\ &< \varepsilon. \end{aligned}$$

Therefore, $E+x$ is measurable. $\square$

Exercises for Section 4.3

Exercise 4.3.1. Suppose $E \subseteq \mathbf{R}$ is measurable with $\mu(E) < \infty$. Show that for any $\varepsilon > 0$, there is a compact set $K \subseteq E$ satisfying

$$\mu(E) \leq \mu(K) + \varepsilon.$$

Conclude that

$$\mu(E) = \sup\{\mu(K) : K \subseteq E \text{ is compact}\}.$$

Exercise 4.3.2. Let $E \subseteq \mathbf{R}$ be a measurable set with $\mu(E) = \infty$. Prove that there is a measurable set $A \subseteq E$ with

$$0 < \mu(A) < \infty.$$

This condition says that Lebesgue measure is **semifinite**.

4.4 Transformations and Non-measurable Sets

When we defined the Lebesgue measure on $\mathbf{R}$, we did so by first defining outer measure, and then restricting the domain of μ^* to so-called *measurable sets*. The fact that certain sets should be deemed "measurable" seems to indicate that some subsets of $\mathbf{R}$ might not satisfy the Carathéodory condition. With the results of the last section, we now all the tools we need to demonstrate the existence of such sets.

We now begin to describe the construction of non-measurable sets, which are often called *Vitali sets*. We will see in the proof that there are actually infinitely many such sets. The idea is to construct a set that is so complicated geometrically that it is impossible to assign a reasonable notion of length to it.

Theorem 4.4.1 (Vitali). *There exist subsets of* $\mathbf{R}$ *that are not Lebesgue measurable.*

Proof. First we define an equivalence relation on the interval $[0,1)$ by

$$x \sim y \text{ if and only if } x - y \in \mathbf{Q}.$$

Now construct a set $N \subseteq [0,1)$ by choosing exactly one element from each of the resulting equivalence classes. For each $r \in \mathbf{Q} \cap [0,1)$, define

$$N_r = \{(x+r) \bmod 1 : x \in N\}.$$

Recall that $(x+r) \bmod 1$ denotes the non-integer part of $x+r$:

$$(x+r) \bmod 1 = \begin{cases} x+r & \text{if } 0 \le x < 1-r \\ x+r-1 & \text{if } 1-r \le x < 1. \end{cases}$$

We claim that the sets N_r partition $[0,1)$.

- Suppose $x \in [0,1)$, and let y be the unique element of N satisfying $x \sim y$. Then we let $r = x - y$ if $x > y$ or $r = x - y + 1$ if $x < y$. Observe that $r \in \mathbf{Q} \cap [0,1)$ and $x = (y+r) \bmod 1$, so $x \in N_r$. It follows that

$$\bigcup_{r \in \mathbf{Q} \cap [0,1)} N_r = [0,1).$$

- Now let $r, s \in \mathbf{Q} \cap [0,1)$, and suppose $x \in N_r \cap N_s$. Then $x = (y+r) \bmod 1$ and $x = (z+s) \bmod 1$ for some $y, z \in N$, meaning that

$$(y+r) - (z+s) = (y-z) + (r-s) = 0 \text{ or } 1,$$

so $y - z \in \mathbf{Q}$. That is, $y \sim z$. Since N contains only one element from each equivalence class, we must have $y = z$. But then it is not hard to check that $r = s$. It follows that $N_r \cap N_s = \emptyset$ if $r \neq s$.

Now we assume that N is measurable. Then N_r is a measurable set for all $r \in \mathbf{Q} \cap [0,1)$, since we can write

$$N_r = \big[(N \cap [0, 1-r)) + r\big] \cup \big[(N \cap [1-r, 1)) + r - 1\big].$$

Furthermore, this decomposition shows that

$$\mu(N_r) = \mu(N \cap [0, 1-r)) + \mu(N \cap [1-r, 1)) = \mu(N).$$

The countable additivity of Lebesgue measure then implies that

$$\mu([0,1)) = \mu\left(\bigcup_r N_r\right) = \sum_r \mu(N_r) = \sum_r \mu(N).$$

Notice that the right hand side is finite if and only if $\mu(N) = 0$. However, in this case we would have $\mu([0,1)) = 0$, which contradicts the fact that $\mu([0,1)) = 1$. Therefore, it must be the case that N is not measurable. $\square$

There is one step in the proof of Vitali's theorem that is fairly subtle. We constructed the non-measurable set N by defining an equivalence relations and then choosing precisely one element from each equivalence class. While it might seem obvious that we should be able to do such a thing, this part of the proof relies heavily on a certain axiom from set theory—the *Axiom of Choice*.

> **Axiom of Choice.** *Given an arbitrary collection $\{X_i\}_{i\in I}$ of nonempty sets, there exists a function $f : X \to \bigcup_{i\in I} X_i$ such that*
>
> $$f(i) \in X_i$$
>
> *for all i.*

Intuitively, the Axiom of Choice (which we abbreviate as AC) guarantees the existence of a **choice function**—it tell us that there is a way to choose exactly one element from each set in our collection. (Note that this is precisely what we did in the proof of Vitali's theorem.) This may seem innocuous (and obvious!), but we can consider a thought experiment (due to Bertrand Russell) to show that there really is something to AC.

Example 4.4.2. Suppose we have an infinite collection of pairs of shoes. That is, we have a collection $\{X_i\}_{i\in I}$, where each X_i denotes a pair of shoes and the index set I is infinite. There is an obvious way to define a choice function—we could simply take the left shoe from each pair, say. Indeed, we do not need to appeal to AC to know that such a choice is possible.

Suppose now that our sets X_i consist of pairs of socks. How do we choose a sock from each pair? It is generally impossible to distinguish left socks from right socks, so our trick for shoes will not work here. We could instead try to construct a choice function by brute force—we choose one sock from the first pair, one sock from the second pair, one sock from the third pair, and so on. However, this process will never terminate. (Note that it also assumes the index set is countable.) We need AC to guarantee that there is a way to choose one sock from each pair.

It is known that AC is independent of the usual Zermelo-Fraenkel (ZF) axioms for set theory. Therefore, it is often taken as an additional axiom, leading to the so-called ZFC axioms for set theory. Some mathematicians express reservations about doing so, since AC is logically equivalent to some strange results. Foremost among them are *Zorn's lemma* and the *well-ordering theorem*. (The latter result is the one that gives many people pause.) However, most mathematicians accept the Axiom of Choice, since it (or one of its variants) is necessary to prove many important theorems in mathematics, including the ones listed below.

- There exist subsets of $\mathbf{R}$ that are not Lebesgue measurable.
- Every vector space has a basis.
- Every nontrivial ring contains a maximal ideal.
- Every field has an algebraic closure (i.e., a field extension over which every polynomial factors completely).
- The Baire Category Theorem.
- Tychonoff's Theorem. (An arbitrary Cartesian product of compact spaces is compact in the product topology.)

4.4.1 Lipschitz Transformations

We now close our discussion of Lebesgue measure with a question: under what conditions does a function $f : \mathbf{R} \to \mathbf{R}$ preserve measurability? Recall the following definition.

Definition 4.4.3. A function $f : \mathbf{R} \to \mathbf{R}$ is **Lipschitz** if there is a constant $\alpha > 0$ such that
$$|f(x) - f(y)| \leq \alpha|x - y|$$
for all $x, y \in \mathbf{R}$.

It is straightforward to check that a Lipschitz function is automatically continuous. Also, the Mean Value Theorem guarantees that a function with bounded derivative is necessarily Lipschitz.

Theorem 4.4.4. *If $f : \mathbf{R} \to \mathbf{R}$ is Lipschitz and $E \subseteq \mathbf{R}$ is measurable, then $f(E)$ is measurable.*

Proof. Suppose f is Lipschitz with Lipschitz constant α. We first show that f takes F_σ sets to F_σ sets. Suppose first that $F \subseteq \mathbf{R}$ is closed. Then we can write $F = \bigcup_{i=1}^\infty K_i$, where each K_i is compact. Since f is continuous, $f(K_i)$ is compact for each i. Therefore,

$$f(F) = \bigcup_{i=1}^\infty f(K_i)$$

is an F_σ set. Now if $A = \bigcup_{i=1}^\infty F_i$ is an F_σ, where each F_i is closed, then

$$f(A) = \bigcup_{i=1}^\infty f(F_i)$$

is a countable union of F_σ sets, hence it is an F_σ.

Now we show that if $Z \subseteq \mathbf{R}$ has measure zero, then $\mu(f(Z)) = 0$ as well. Let $\varepsilon > 0$ be given, and let $\{I_k\}_{k=1}^\infty$ be a covering of Z with $\sum_{k=1}^\infty v(I_k) < \frac{\varepsilon}{\alpha}$. Then clearly we have

$$f(Z) \subseteq \bigcup_{k=1}^\infty f(I_k),$$

and

$$\mu^*(f(I_k)) \leq \alpha \cdot v(I_k) = \alpha \cdot \mu^*(I_k)$$

so

$$\mu^*(f(Z)) \leq \sum_{k=1}^\infty \mu^*(f(I_k)) \leq \sum_{k=1}^\infty \alpha v(I_k) < \varepsilon.$$

This holds for all $\varepsilon > 0$, so $\mu^*(f(Z)) = 0$.

Now suppose E is measurable, and write $E = A \cup Z$ where A is an F_σ set and Z has measure zero. Then

$$f(E) = f(A) \cup f(Z)$$

is the union of an F_σ and a null set, so $f(E)$ is measurable. □

In general, a continuous image of a measurable set need not be measurable. We exhibit a continuous function with this property below.

4.4.2 The Cantor Function

We now set out to construct a quite bizarre function $\varphi : [0,1] \to \mathbf{R}$, which is very closely related to the Cantor set. This approach is closely adapted from the one in [RF10]. We begin by defining

$$f_1(x) = \begin{cases} \frac{1}{2} & \text{if } \frac{1}{3} \le x \le \frac{2}{3} \\ 0 & \text{otherwise,} \end{cases}$$

then

$$f_2(x) = \begin{cases} \frac{1}{4} & \text{if } \frac{1}{9} \le x \le \frac{2}{9} \\ \frac{3}{4} & \text{if } \frac{7}{9} \le x \le \frac{8}{9} \\ 0 & \text{otherwise,} \end{cases}$$

and

$$f_3(x) = \begin{cases} \frac{1}{8} & \text{if } \frac{1}{27} \le x \le \frac{2}{27} \\ \frac{3}{8} & \text{if } \frac{7}{27} \le x \le \frac{8}{27} \\ \frac{5}{8} & \text{if } \frac{19}{27} \le x \le \frac{20}{27} \\ \frac{7}{8} & \text{if } \frac{25}{27} \le x \le \frac{26}{27} \\ 0 & \text{otherwise.} \end{cases}$$

In general we define f_n to be $\frac{2^i-1}{2^n}$ on the ith interval removed at the nth step in the construction of the Cantor set, and zero elsewhere. We then let

$$f(x) = \sum_{n=1}^{\infty} f_n(x)$$

for all $x \in [0,1]$. For each $x \in [0,1]$, there is only one n for which $f_n(x) \neq 0$, so it follows easily that this series converges for all values of x. Notice also that $f(x) = 0$ for all $x \in C$. On the other hand, if we let $U = [0,1] \backslash C$, then f is increasing on U. It turns out that we can actually modify the values of f on C to obtain a function that is increasing on $[0,1]$. If $x \in C$, we define

$$s_x = \sup\bigl\{ f(t) : t < x \text{ and } t \in U \bigr\}.$$

We then set

$$\varphi(x) = \begin{cases} f(x) & \text{if } x \in U \\ s_x & \text{if } x \in C \end{cases}$$

for all $x \in [0,1]$. The function φ is called the **Cantor function**, or sometimes the **Cantor-Lebesgue function**. It is also referred to as the **devil's staircase**, which seems apt when one considers the graph of φ in Figure 4.1.

Remark 4.4.5. There are several other ways of defining the Cantor function. A recursive definition is given in [WZ77]. There is also an algorithm for computing values of the Cantor function that is equivalent to what we have done here:

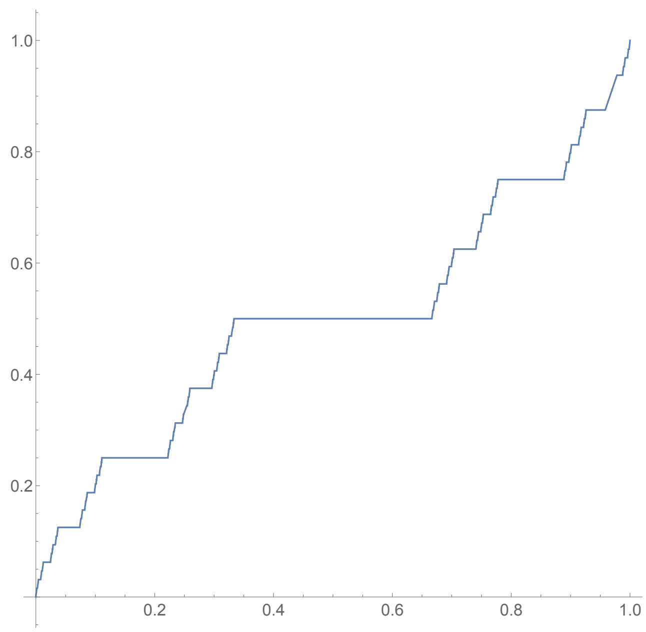

Figure 4.1: The Cantor function. (Image generated in *Mathematica*.)

1. Given a number $x \in [0,1]$, write out the ternary (i.e., base 3) expansion of x.
2. Find the first 1 in the ternary expansion of x (if there is a 1), and change every digit afterward to 0.
3. Change all 2s to 1s, and interpret the result as the binary expansion of some number $\varphi(x)$.

For example, the ternary expansion of $\frac{1}{2}$ is

$$\tfrac{1}{2} = (0.1111\ldots)_3,$$

so we truncate after the first 1 to get 0.1, and then interpret this as a binary number. Thus

$$\varphi(\tfrac{1}{2}) = (0.1)_2 = \tfrac{1}{2}$$

Similarly,

$$\varphi(\tfrac{1}{4}) = \varphi((0.020202\ldots)_3) = (0.010101\ldots)_2 = \tfrac{1}{3}.$$

It is not hard to check that if a point x is deleted at the nth step in the construction of the Cantor set, then $\varphi(x) = f_n(x)$, so this definition agrees with ours.

It is not hard to see that the Cantor function is increasing on $[0,1]$. We already knew it was increasing when restricted to $U = [0,1]\backslash C$, and we have defined φ on

C in such a way that it is increasing everywhere. It follows from a well-known fact in real analysis that if φ has any discontinuities, they must be jump discontinuities. We use this observation to show that φ is actually continuous throughout $[0,1]$.

Proposition 4.4.6. *The Cantor function is continuous on* $[0,1]$.

Proof. First notice that φ is continuous at each point $x_0 \notin C$, since it is piecewise constant on U. By our observation above, we know that φ can only have jump discontinuities. Therefore, to show that f is continuous at each point $x_0 \in C$, it suffices to show that

$$\lim_{x \to x_0^-} \varphi(x) = \lim_{x \to x_0^+} \varphi(x).$$

(Note that if $x_0 = 0$ or $x_0 = 1$, we need to consider a single one-sided limit. The proof in those cases is similar to what follows.) Since φ is increasing, we have

$$\lim_{x \to x_0^-} \varphi(x) = \sup_{x < x_0} \varphi(x)$$

and

$$\lim_{x \to x_0^+} \varphi(x) = \inf_{x > x_0} \varphi(x).$$

Notice that the left-hand limit equals $\varphi(x_0)$ by definition, so φ is continuous from the left by construction. It suffices to show that the right-hand limit equals $\varphi(x_0)$ as well.

Let $\varepsilon > 0$ be given, and choose N such that $\frac{1}{2^N} < \varepsilon$. Let $U_N = [0,1] \setminus \bigcap_{n=1}^{N} C_n$ denote the set of all points removed in the first N steps in the construction of the Cantor set. Then U_N is a union of 2^{N-1} disjoint open intervals, and for each point $x \in U_N$ we have

$$\varphi(x) = \sum_{n=1}^{N} f_n(x).$$

Since at most one term in this partial sum is nonzero, it is easy to check that the possible values of φ on U_N are $\{\frac{k}{2^N} : k = 1, 2, \ldots, 2^N - 1\}$. Since $x_0 \in C$, it must lie between two of the intervals in U_N; choose a point a_N from the first interval and b_N from the second. Then we have $a_N < x_0 < b_N$ and by the above discussion there is some k for which

$$\varphi(b_N) - \varphi(a_N) = \frac{k}{2^N} - \frac{k-1}{2^N} = \frac{1}{2^N} < \varepsilon.$$

It follows that

$$\varphi(b_N) < \varphi(x_0) + \varepsilon.$$

Since we can find such a point for any $\varepsilon > 0$, we have

$$\varphi(x_0) = \inf_{x > x_0} \varphi(x),$$

so φ is continuous from the right. Thus φ is continuous at x_0, and we are done. $\square$

The Cantor function has some other more surprising properties. First observe that since φ is piecewise constant on the open set U, it is differentiable on U with $\varphi'(x) = 0$ for all $x \in U$. On the other hand, φ fails to be differentiable at each point of C. (See [DMRV06, Proposition 9.1].) Thus the Cantor function has the following remarkable property: it is differentiable almost everywhere (i.e., except on a set of measure zero), and its derivative is zero whenever it exists, yet φ manages to increase from 0 to 1 without exhibiting any jump discontinuities. In other words, the Cantor function is an example of a **Lebesgue singular function**.

If we modify the Cantor function slightly, we can obtain function exhibiting other kinds of strange behavior. Define

$$\psi(x) = \varphi(x) + x.$$

Notice that ψ is continuous and strictly increasing, since φ is continuous and increasing and $f(x) = x$ is continuous and strictly increasing. Since $\psi(0) = 0$ and $\psi(1) = 2$, the Intermediate Value Theorem guarantees that ψ maps $[0, 1]$ onto $[0, 2]$. What happens to the Cantor set under this map?

Proposition 4.4.7. *The function ψ maps the Cantor set onto a set of positive Lebesgue measure.*

Proof. We begin by writing $[0, 1] = C \cup U$. For simplicity, let $V = \psi(U)$. Then

$$[0, 2] = \psi(C) \cup V,$$

where $\psi(C) \cap V = \emptyset$. Since ψ is continuous and strictly increasing, it is a homeomorphism. In particular, it is an open map, so V is open and $\psi(C)$ is closed. Consequently, both sets are measurable. We aim to show that $\mu(V) = 1$, which will imply that $\mu(\psi(C)) = 1$ as well. Write

$$U = \bigcup_{n=1}^{\infty} \bigcup_{k=1}^{2^{k-1}} I_{n,k},$$

where $I_{n,k}$ denotes the kth open interval removed at the nth step in the construction of C. Then

$$V = \bigcup_{n=1}^{\infty} \bigcup_{k=1}^{2^{n-1}} \psi(I_{n,k}).$$

Since the intervals $I_{n,k}$ are pairwise disjoint and ψ is one-to-one, the sets $\psi(I_{n,k})$ are pairwise disjoint as well. Moreover, if $x \in I_{n,k}$ we have

$$\psi(x) = \frac{2^k - 1}{2^n} + x,$$

so ψ simply translates each interval $I_{n,k}$ by $\frac{2^k-1}{2^n}$. Thus

$$\mu(V) = \sum_{n=1}^{\infty} \sum_{k=1}^{2^{n-1}} \mu(\psi(I_{n,k})) = \sum_{n=1}^{\infty} \sum_{k=1}^{2^{n-1}} \mu(I_{n,k}) = \mu(U) = 1.$$

It now follows that $\psi(C)$ has positive measure. $\square$

We can milk one more bizarre fact out of this last result. Since $\psi(C)$ has positive measure, a modification of Vitali's theorem shows that it contains a non-measurable subset W. Then $\psi^{-1}(W)$ is a subset of the Cantor set, so it has outer measure zero, and is therefore measurable. But

$$\psi(\psi^{-1}(W)) = W,$$

so we have a measurable that maps onto a non-measurable set. Thus the continuous image of a measurable set need not be measurable. Note that we also obtain a way of constructing a Lebesgue measurable set that is not Borel. In particular, the set $\psi^{-1}(W) \subseteq C$ is Lebesgue measurable, but it cannot be Borel. If it were Borel, then its image under ψ would be be Borel (since ψ is continuous). In other words, W would be Borel, hence Lebesgue measurable.

Exercises for Section 4.4

Exercise 4.4.1. Suppose $Z \subseteq \mathbf{R}$ has measure zero. Show that the set

$$\big\{\arctan(x) : x \in Z\big\}$$

also has measure zero.

Exercise 4.4.2. Show by example that outer measure is not countably additive. That is, show there exists a countable, pairwise disjoint collection of sets $\{E_i\}_{i=1}^{\infty}$ such that

$$\mu^*\left(\bigcup_{i=1}^{\infty} E_i\right) < \sum_{i=1}^{\infty} \mu^*(E_i).$$

4.5 Measurable Functions

Now that we have established all of the essential properties of Lebesgue measure, we can move toward constructing Lebesgue's theory of integration. In order to do so, we first need to single out the sorts of functions that we will be allowed to integrate. Recall that Lebesgue's approach to integration (loosely speaking) was to partition the range of a function f with intervals, and then consider the measure of the preimage of each interval under f. In other words, we might want to find the measure of a set of the form

$$\{x : a < f(x) < b\}$$

for some $a, b \in \mathbf{R}$. For this to make sense, we will need to ensure that such sets are indeed measurable. This leads us to the definition of a *measurable function*.

Let us make two notes before we proceed. First, we will now allow our functions to take values in the *extended real numbers*. That is, given a set $E \subseteq \mathbf{R}$, we will consider functions $f : E \to [-\infty, +\infty]$. Second, we will assume throughout this discussion that the domain E is measurable.

Definition 4.5.1. Let $E \subseteq \mathbf{R}$ be a measurable set. We will say that a function $f : E \to [-\infty, +\infty]$ is **measurable** if for each $a \in \mathbf{R}$, the set

$$\{x \in E : f(x) > a\}$$

is measurable.

To simplify notation, we will often use the abbreviation $\{f > a\}$ to represent the set $\{x : f(x) > a\}$. We will also use other similar abbreviations, such as $\{f < a\}$, $\{f \leq a\}$, and $\{f = a\}$, each of whose meanings should be clear.

We now proceed in proving several important properties of measurable functions. Most follow quickly from facts we have already established about Lebesgue measure, so the proofs are fairly short.

Theorem 4.5.2. *Let $f : E \to [-\infty, +\infty]$ be an extended real-valued function. The following are equivalent:*

1. *f is measurable.*
2. *For all $a \in \mathbf{R}$, the set $\{f \geq a\}$ is measurable.*
3. *For all $a \in \mathbf{R}$, the set $\{f < a\}$ is measurable.*
4. *For all $a \in \mathbf{R}$, the set $\{f \leq a\}$ is measurable.*

Proof. ($1 \Rightarrow 2$) Assume f is measurable, and let $a \in \mathbf{R}$. Then for each $k \in \mathbf{N}$, the set $\{f > a - \frac{1}{k}\}$ is measurable. It is then easy to check that

$$\{f \geq a\} = \bigcap_{k=1}^{\infty} \{f > a - \tfrac{1}{k}\},$$

and this set is measurable since it is an intersection of countably many measurable sets.

($2 \Rightarrow 3$) This follows immediately from the observation that $\{f < a\} = \{f \geq a\}^c$ and that the complement of a measurable set is measurable.

($3 \Rightarrow 4$) This is similar to the proof of the first implication. By assumption, sets of the form $\{f < a + \frac{1}{k}\}$ are measurable for all $k \in \mathbf{N}$, so

$$\{f \leq a\} = \bigcap_{k=1}^{\infty} \{f < a + \tfrac{1}{k}\}$$

is measurable.

($4 \Rightarrow 1$) This again follows by taking complements—we just use the fact that for all $a \in \mathbf{R}$, $\{f > a\} = \{f \leq a\}^c$. □

Theorem 4.5.3. *Let $f : E \to [-\infty, +\infty]$ be a measurable function. Then:*

1. *For all $a \in \mathbf{R}$, the set $\{f = a\}$ is measurable.*
2. *For all $a, b \in \mathbf{R}$ with $a < b$, the sets $\{a < f < b\}$ and $\{a \leq f \leq b\}$ are measurable.*
3. *The sets $\{f = -\infty\}$, $\{f = +\infty\}$, $\{f > -\infty\}$, and $\{f < +\infty\}$ are measurable.*

Proof. For the first assertion, let $a \in \mathbf{R}$ and write

$$\{f = a\} = \{f \geq a\} \cap \{f \leq a\}.$$

Since f is measurable, both sets on the right hand side are measurable by the previous theorem. Thus $\{f = a\}$ is measurable.

Now let $a, b \in \mathbf{R}$ with $a < b$. Then we can write

$$\{a < f < b\} = \{f > a\} \cap \{f < b\},$$

which is measurable by the previous theorem. Similarly, we have

$$\{a \leq f \leq b\} = \{f \geq a\} \cap \{f \leq b\},$$

so $\{a \leq f \leq b\}$ is measurable.

Finally, observe that

$$\{f = +\infty\} = \bigcap_{k=1}^{\infty} \{f > k\},$$

so $\{f = +\infty\}$ is measurable. A similar argument works for $\{f = -\infty\}$. We can then conclude that

$$\{f < +\infty\} = E \backslash \{f = +\infty\}^c$$

is measurable, and similarly for $\{f > -\infty\}$. □

Proposition 4.5.4. *A function* $f : E \to [-\infty, +\infty]$ *is measurable if and only if* $\{a < f < +\infty\}$ *is measurable for all* $a \in \mathbf{R}$ *and* $\{f = +\infty\}$ *is measurable.*

Proof. If f is measurable, then all of the sets listed above are clearly measurable. To establish the converse, let $a \in \mathbf{R}$, and write

$$\{f > a\} = \{a < f < +\infty\} \cup \{f = +\infty\}.$$

The right hand side is measurable by assumption, so f is measurable. $\square$

We will now establish a connection between measurable functions and continuity. To start, we exhibit a characterization of measurability in terms of open sets, which is reminiscent of the topological definition of continuity.

Theorem 4.5.5. *A function* $f : E \to [-\infty, +\infty]$ *is measurable if and only if* $\{f = +\infty\}$ *is measurable and* $f^{-1}(U)$ *is measurable for every open set* $U \subseteq \mathbf{R}$*.*

Proof. Assume f is measurable. Then $\{f = +\infty\}$ is measurable, so suppose $U \subseteq \mathbf{R}$ is open. Recall that we can write U as a union of countably many open intervals:

$$U = \bigcup_{k=1}^{\infty} (a_k, b_k).$$

Therefore,

$$f^{-1}(U) = \bigcup_{k=1}^{\infty} f^{-1}\big((a_k, b_k)\big) = \bigcup_{k=1}^{\infty} \{a_k < f < b_k\},$$

which is measurable.

Now let $a \in \mathbf{R}$, and write $\{f > a\} = \{a < f < +\infty\} \cup \{f = +\infty\}$. Observe that

$$\{a < f < +\infty\} = f^{-1}\big((a, \infty)\big)$$

is the preimage of an open set, so it is measurable by assumption. Since $\{f = +\infty\}$ is assumed to be measurable, $\{f > a\}$ is measurable. Hence f is measurable. $\square$

Corollary 4.5.6. *Any continuous function* $f : E \to \mathbf{R}$ *is measurable.*

Proof. Since f is continuous, it only assumes finite values. Also, if $U \subseteq \mathbf{R}$ is open, then $f^{-1}(U)$ is relatively open in E, hence measurable. $\square$

We will now see that we can often ignore sets of measure zero when working with measurable functions. To that end, we say a property holds **almost everywhere** for a function (abbreviated a.e.) if it holds everywhere except possibly on a set of measure zero. As an example, if f and g are functions, we write $f = g$ a.e. to indicate that

$$\mu\big(\{x : f(x) \neq g(x)\}\big) = 0.$$

Proposition 4.5.7. *If $f = g$ a.e. and f is measurable, then g is also measurable. Moreover,*

$$\mu\big(\{f > a\}\big) = \mu\big(\{g > a\}\big)$$

for all $a \in \mathbf{R}$.

Proof. Let $Z = \{f \neq g\}$ and put $A = E \backslash Z$. Then for any $a \in \mathbf{R}$,

$$\begin{aligned}\{x \in E : g(x) > a\} &= \{x \in A : g(x) > a\} \cup \{x \in Z : g(x) > a\} \\ &= \{x \in A : f(x) > a\} \cup \{x \in Z : g(x) > a\}.\end{aligned}$$

The first set is simply $\{f > a\} \cap A$, so it is measurable since f is measurable. The second set has outer measure zero (since it lies inside Z), hence it is measurable. Thus $\{g > a\}$ is measurable for all a, so g is measurable. For the second assertion, notice that

$$\begin{aligned}\mu\big(\{g > a\}\big) &= \mu\big(\{f > a\} \cap A\big) + \mu\big(\{g > a\} \cap Z\big) \\ &= \mu\big(\{f > a\} \cap A\big) + \mu\big(\{f > a\} \cap Z\big) \\ &= \mu\big(\{f > a\}\big),\end{aligned}$$

where we have again used the fact that any subset of Z has measure 0. ☐

Next we show that linear combinations and products of measurable functions are still measurable. Let us first make a couple of observations. Suppose $f, g : E \to [-\infty, +\infty]$ are measurable functions. We will run into problems if we have $f(x) = +\infty$ and $g(x) = -\infty$ for a particular value of x. In particular, $(f + g)(x)$ is not well-defined in this case. For this reason, we assume f and g are finite a.e., so that the set

$$Z = \{x \in E : f(x) = \pm\infty \text{ or } g(x) = \pm\infty\}$$

has measure zero. We can then define $f + g$ however we want on Z without affecting measurability. We will prove the next theorem under the assumption that f and g are finite, and the case where they are finite a.e. is left as an exercise.

Theorem 4.5.8. *Suppose f and g are finite measurable functions on E, and let $\alpha \in \mathbf{R}$. Then:*

1. *$f + g$ is measurable.*
2. *αf is measurable.*
3. *fg is measurable.*

Proof. Let $a \in \mathbf{R}$. Notice that if $f(x) + g(x) > a$, then $f(x) > a - g(x)$, and we can find a rational number q such that

$$f(x) > q > a - g(x).$$

Thus

$$\{f + g > a\} = \bigcup_{q \in \mathbf{Q}} \{f > q\} \cap \{g > q - a\}.$$

Each set in this union is measurable since f and g are, so we have a countable union of measurable sets. It follows that $f + g$ is measurable.

For the second assertion, notice that the result is immediate if $\alpha = 0$. Therefore, assume $\alpha \neq 0$, and let $a \in \mathbf{R}$. Then

$$\{\alpha f > a\} = \{f > \tfrac{a}{\alpha}\}$$

when $\alpha > 0$, or

$$\{\alpha f > a\} = \{f < \tfrac{a}{\alpha}\}$$

when $\alpha < 0$. In either case, we have a measurable set since f is measurable. Thus αf is measurable.

For the last assertion, we first show that f^2 is measurable. Notice that if $a \geq 0$, then $(f(x))^2 > a$ if and only if $f(x) > \sqrt{a}$ or $f(x) < -\sqrt{a}$. Thus

$$\{f^2 > a\} = \{f > \sqrt{a}\} \cup \{f < -\sqrt{a}\}$$

is measurable. (Note that if $a < 0$, then $\{f^2 > a\} = \{f^2 \geq 0\} = E$.) Now we can write

$$fg = \tfrac{1}{2}\big[(f + g)^2 - f^2 - g^2\big],$$

which is measurable since f, g, and $f + g$ are all measurable. $\square$

As we will see shortly, composition of measurable functions is not so nice. First we obtain a positive result under the assumption that one of the functions is continuous.

Theorem 4.5.9. *Suppose $f : E \to \mathbf{R}$ is measurable, and suppose $g : \mathbf{R} \to \mathbf{R}$ is continuous. Then $g \circ f$ is measurable.*

Proof. Suppose $U \subseteq \mathbf{R}$ is open. Then $g^{-1}(U)$ is open since g is continuous. Hence

$$(g \circ f)^{-1}(U) = f^{-1}\big(g^{-1}(U)\big)$$

is measurable for all open sets U, so $g \circ f$ is measurable. □

Example 4.5.10. It is not the case in general that a composition of two measurable functions is measurable. Let $\varphi : [0,1] \to \mathbf{R}$ denote the Cantor function, and let $\psi(x) = \varphi(x) + x$. Recall that $\psi(C)$ contains a non-measurable subset W. Put $E = \psi^{-1}(W)$, and let χ_E denote the indicator function of E:

$$\chi_E(x) = \begin{cases} 1 & \text{if } x \in E \\ 0 & \text{if } x \notin E. \end{cases}$$

Notice that χ_E is measurable since E is. (We will prove a general result of this sort very soon.) Moreover, ψ^{-1} is continuous, hence measurable. However,

$$(\chi_E \circ \psi^{-1})^{-1}(\{1\}) = \psi(E) = W$$

is not measurable, so $\chi_E \circ \psi^{-1}$ cannot be measurable.

For the rest of this section, we consider results on sequences of measurable functions. Ultimately, we will show that the pointwise limit of a sequence of measurable functions is measurable. As a first step, we begin with the pointwise supremum and infimum.

Theorem 4.5.11. *Let $(f_n)_{n=1}^{\infty}$ be a sequence of measurable functions on a set E. Then $\sup f_n$ and $\inf f_n$ are measurable.*

Proof. Let $a \in \mathbf{R}$. Notice that if $x \in E$, $\sup f_n(x) > a$ if and only if $f_n(x) > a$ for some n. Thus

$$\{\sup f_n > a\} = \bigcup_{n=1}^{\infty} \{f_n > a\}.$$

The right hand side is measurable since each f_n is measurable. Thus $\sup f_n$ is measurable. Similarly, we have

$$\{\inf f_n < a\} = \bigcup_{n=1}^{\infty} \{f_n < a\},$$

so $\inf f_n$ is measurable. □

Remark 4.5.12. As a special case of Theorem 4.5.11 (or more appropriately, as simple adaptation of its proof), notice that if $(f_k)_{k=1}^n$ is a finite collection of functions, then the pointwise maximum and minimum

$$\max_{1\leq k\leq n} f_k, \quad \min_{1\leq k\leq n} f_k$$

are measurable. In particular, the pointwise maximum or minimum of a pair of measurable functions is measurable. As an example,

$$|f| = \max\{f, -f\},$$

is measurable for any measurable function f. Two other important examples that we will encounter are

$$f^+ = \max\{f, 0\},$$

and

$$f^- = -\min\{f, 0\}.$$

Notice that both f^+ and f^- are positive functions, and we can write

$$f = f^+ - f^-.$$

Thus we can write any function as a linear combination of positive functions, which will reduce many of our future arguments to ones about positive functions. Furthermore, observe that f is measurable if and only if f^+ and f^- are both measurable.

Remark 4.5.13. We now recall two important concepts from the study of sequences of real numbers. Notice that if $(x_n)_{n=1}^\infty$ is a bounded sequence in $\mathbf{R}$, the sequence $(s_n)_{n=1}^\infty$ defined by

$$s_n = \sup_{k\geq n} x_k$$

is decreasing and bounded below. Thus it converges, and we can define the **limit superior** of $(x_n)_{n=1}^\infty$ by

$$\limsup_{n\to\infty} x_n = \lim_{n\to\infty} s_n = \lim_{n\to\infty} \sup_{k\geq n} x_k = \inf_n \sup_{k\geq n} x_k.$$

Similarly, the sequence $(l_n)_{n=1}^\infty$ defined by

$$l_n = \inf_{k\geq n} x_k$$

is increasing and bounded above, and we define the **limit inferior** of $(x_n)_{n=1}^\infty$ to be

$$\liminf_{n\to\infty} x_n = \lim_{n\to\infty} l_n = \lim_{n\to\infty} \inf_{k\geq n} x_k = \sup_n \inf_{k\geq n} x_k.$$

It is straightforward to check that

$$\liminf_{n\to\infty} x_n \leq \limsup_{n\to\infty} x_n, \tag{4.1}$$

and that $(x_n)_{n=1}^{\infty}$ converges if and only if

$$\limsup_{n\to\infty} x_n = \liminf_{n\to\infty} x_n = \lim_{n\to\infty} x_n.$$

It is also worth noting that $\limsup x_n$ is the largest possible subsequential limit of $(x_n)_{n=1}^{\infty}$, while $\liminf x_n$ is the smallest one.

If the sequence $(x_n)_{n=1}^{\infty}$ is not bounded, then we may have to modify our definitions. In particular, if $(s_n)_{n=1}^{\infty}$ is not bounded below, we set

$$\limsup_{n\to\infty} x_n = -\infty,$$

and if $(l_n)_{n=1}^{\infty}$ fails to be bounded above, we define

$$\liminf_{n\to\infty} x_n = \infty.$$

However, notice that (4.1) implies that in the first case, $\liminf x_n = -\infty$ as well, so $x_n \to -\infty$. Similarly, in the second case we must have $x_n \to \infty$. Thus the limits superior and inferior will generally be finite, unless the sequence diverges to $\pm\infty$.

Of course if we have a sequence of real-valued functions, it makes sense to take the limsup or liminf of the sequence pointwise. Our next result shows that these operations preserve measurability.

Theorem 4.5.14. *Let $(f_n)_{n=1}^{\infty}$ be a sequence of measurable functions. Then*

$$\limsup_{n\to\infty} f_n, \quad \liminf_{n\to\infty} f_n$$

are measurable. In particular, if $\lim f_n$ *exists a.e., then it is measurable.*

Proof. Recall that

$$\limsup_{n\to\infty} f_n = \inf_n \left\{ \sup_{k\geq n} f_k \right\}.$$

Since each f_k is measurable, $\sup_{k\geq n} f_k$ is measurable for all n by Theorem 4.5.11. Therefore, $\inf_n \sup_{k\geq n} f_k$ is measurable, again by Theorem 4.5.11. Thus $\limsup f_n$ is measurable. The argument for $\liminf f_n$ is done similarly, using the fact that

$$\liminf_{n\to\infty} f_n = \sup_n \left\{ \inf_{k\geq n} f_k \right\}.$$

Finally, if $\lim f_n(x)$ exists for some x, then it equals $\limsup f_n(x)$. Thus if $\lim f_n$ exits a.e., then it is equal a.e. to the measurable function $\limsup f_n$. Therefore, $\lim f_n$ is measurable. □

Exercises for Section 4.5

Exercise 4.5.1 ([WZ77], Exercise 4.6 modified)**.** Let f and g be extended real-valued functions defined on a measurable set E. Assume f and g are measurable and finite a.e. in E, and let

$$Z = \{f = \pm\infty\} \cup \{g = \pm\infty\}.$$

Note that $\mu(Z) = 0$.

(a) Show that $f + g$ is measurable regardless of how we define it on Z.

(b) Do the same for fg.

Exercise 4.5.2 ([WZ77], Exercise 4.12)**.** Suppose $f : [a, b] \to \mathbf{R}$ is continuous a.e. on $[a, b]$. Prove that f is measurable.

Exercise 4.5.3. Suppose $f, g : [a, b] \to \mathbf{R}$ are continuous functions. Show that if $f = g$ a.e., then $f(x) = g(x)$ for all $x \in [a, b]$. Is the same conclusion true if we replace $[a, b]$ with an arbitrary measurable set E?

4.6 The Lebesgue Integral for Nonnegative Functions

We now have essentially all the tools we need to begin constructing Lebesgue's theory of integration. Roughly speaking, Lebesgue's overall approach was the following:

1. Partition the range of a function f using finitely many intervals of the form $[a_k, b_k)$, where $k = 1, 2, \ldots, n$.
2. Partition the domain of f using the preimages $E_k = f^{-1}([a_k, b_k))$.
3. Approximate f with a piecewise-defined function whose restriction to each E_k is constant.

This sort of procedure is illustrated in Figure 4.2. One can also see in that figure that it is presumably easy to define the "area under the graph" of the piecewise function from step 3 using the measures of the sets E_k. As with the Riemann integral, we can then define the Lebesgue integral of f via some sort of limiting process.

Before we can define the Lebesgue integral, we need to investigate the piecewise constant functions that we will use to approximate arbitrary measurable functions. We begin with the fundamental building blocks, which are called *characteristic functions*.

Definition 4.6.1. Let $E \subseteq \mathbf{R}$. The **characteristic** (or **indicator**) **function** of E is the function $\chi_E : \mathbf{R} \to \mathbf{R}$ defined by

$$\chi_E(x) = \begin{cases} 1 & \text{if } x \in E \\ 0 & \text{if } x \notin E. \end{cases}$$

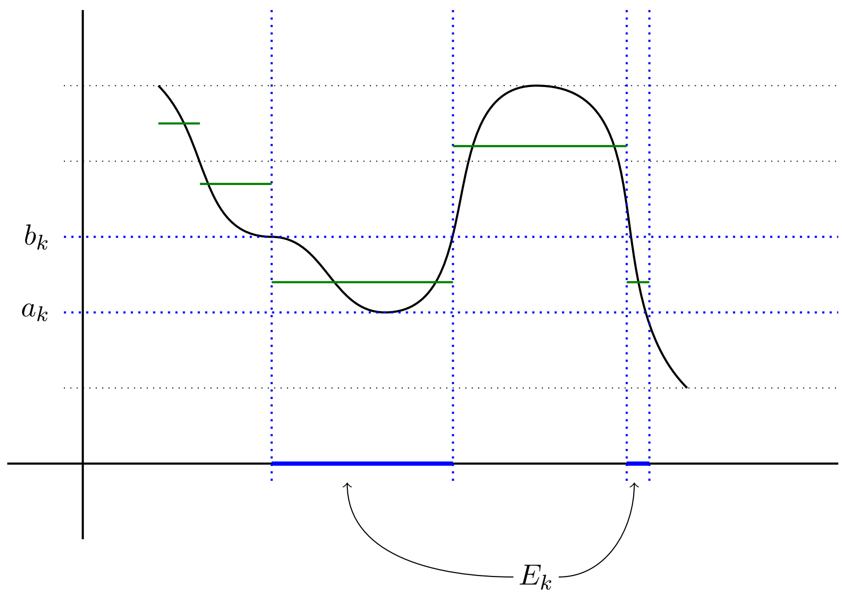

Figure 4.2: An approximation of a function f (shown in black) by a piecewise constant function ϕ (shown in green). The set E_k (blue) is the preimage of the interval $[a_k, b_k)$ under f. Note that ϕ is constant on E_k, i.e., it takes the same value on both connected components of E_k.

Remark 4.6.2. It is easy to check that an indicator function χ_E is measurable if and only if E is measurable. If $a \in \mathbf{R}$, then there are only three cases to consider, since χ_E only takes the values 0 and 1. If $a < 0$, then $\chi_E(x) > a$ for all x, so we have

$$\{\chi_E > a\} = \mathbf{R},$$

which is measurable. Similarly, if $a > 1$, then $\{\chi_E > a\} = \emptyset$ is measurable. Finally, if $0 \le a < 1$, then

$$\{\chi_E > a\} = \{\chi_E > 0\} = E.$$

It is then clear that χ_E is measurable precisely when E is.

Now we assign a term to the kinds of piecewise constant approximating functions that we described above.

Definition 4.6.3. A function $f : E \to \mathbf{R}$ is said to be **simple** if its range is finite.

We can represent any simple function in a canonical way. Let $E \subseteq \mathbf{R}$, and suppose $f : E \to \mathbf{R}$ is simple with range $\{c_1, c_2, \dots, c_n\}$. For $1 \le k \le n$, define

$$E_k = \{x \in E : f(x) = c_k\}.$$

Since the values c_k are distinct, the sets E_k are pairwise disjoint. Furthermore, the function $c_k \cdot \chi_{E_k}$ takes the value c_k on E_k, and it is zero otherwise. Thus we can express f as a linear combination of characteristic functions:

$$f = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}.$$

It is now straightforward to check that f is measurable if and only if each E_k is a measurable set. If f is measurable, then $\{f = c_k\} = E_k$ is measurable for all k. Conversely, if each E_k is measurable, then χ_{E_k} is a measurable function for $1 \le k \le n$, hence f is measurable.

As described above, we plan to define the Lebesgue integral of a function by approximating it with simple functions. This approach requires the following result.

Theorem 4.6.4. *Let $E \subseteq \mathbf{R}$, and suppose $f : E \to [-\infty, +\infty]$ is an extended real-valued function on E.*

1. *There exists a sequence of simple functions that converges to f pointwise on E.*
2. *If $f \ge 0$, then the sequence can be chosen to be increasing.*
3. *If f is measurable, then the simple functions can be chosen to be measurable.*

Proof. We will first assume $f \ge 0$. Given $n \in \mathbf{N}$, we partition the range of f as follows. First divide the interval $[0, n]$ into $n2^n$ intervals of the form

$$\left[\frac{k-1}{2^n}, \frac{k}{2^n}\right), \quad k = 1, 2, \dots, n2^n.$$

Then let

$$E_k = \left\{\frac{k-1}{2^n} \le f < \frac{k}{2^n}\right\}, \quad k = 1, 2, \dots, n2^n,$$

and set

$$E_0 = \{f \ge n\}.$$

Now define $c_k = \frac{k-1}{2^n}$ for $1 \le k \le n2^n$ and $c_0 = n$, and put

$$f_n = \sum_{k=0}^{n2^n} c_k \cdot \chi_{E_k}.$$

Then f_n is a simple function which satisfies $f_n \le f$. Furthermore, if f is measurable, then the sets E_k are measurable for $k = 0, 1, \dots, n2^n$, so f_n is measurable.

Thus far, we have approximated f from below with a sequences of simple functions, each of which is measurable whenever f is. Let us now make two other observations.

- The sequence $(f_n)_{n=1}^\infty$ is increasing pointwise. Fix n and suppose $f(x) \in \left[\frac{k-1}{2^n}, \frac{k}{2^n}\right)$, which means $f_n(x) = \frac{k-1}{2^n}$. If $f(x)$ lies in the left half of the interval, i.e.,
$$\frac{k-1}{2^n} \leq f(x) < \frac{2k-1}{2^{n+1}},$$
then
$$f_{n+1}(x) = \frac{k-1}{2^n} = f_n(x).$$
If x lies in the right half, so
$$\frac{2k-1}{2^{n+1}} \leq f(x) < \frac{k}{2^n},$$
then
$$f_{n+1}(x) = \frac{2k-1}{2^{n+1}} > \frac{2k-2}{2^{n+1}} = \frac{k-1}{2^n} = f_n(x).$$
Thus $f_n(x) \leq f_{n+1}(x)$ for all $x \in E$. It follows that $f_n \leq f_{n+1}$ for all n.

- We have $f_n \to f$ pointwise. Let $x \in E$, and assume first that $f(x)$ is finite. Then there exists $n \in \mathbf{N}$ such that
$$\frac{k-1}{2^n} \leq f(x) < \frac{k}{2^n}$$
for some k, so $f_n(x) = \frac{k-1}{2^n}$ and
$$f(x) - f_n(x) < \frac{k}{2^n} - \frac{k-1}{2^n} = \frac{1}{2^n}.$$
The right hand side tends to zero as $n \to \infty$, so it follows that $f_n(x) \to f(x)$. If $f(x) = +\infty$, then we must have $f_n(x) = n$ for all n, so $f_n(x) \to +\infty$.

We have now established all three conclusions when $f \geq 0$. If f is an arbitrary measurable function, recall that we can write $f = f^+ - f^-$, where $f^+, f^- \geq 0$. Furthermore, f^+ and f^- are measurable if and only if f is. Therefore, we can find sequences $(f_n^+)_{n=1}^\infty$ and $(f_n^-)_{n=1}^\infty$ of simple functions converging to f^+ and f^-, respectively, and the simple functions can be assumed to be measurable if f is. Finally, observe that
$$f_n^+ - f_n^- \to f^+ - f^- = f$$
pointwise, so we are done. □

If we look back to Figure 4.2, it should be clear how to write down the "area under the graph" of a simple function f.

Definition 4.6.5. Let $E \subseteq \mathbf{R}$ be a measurable set and $f : E \to \mathbf{R}$ a nonnegative, measurable simple function with standard representation

$$f = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}.$$

We define the **Lebesgue integral** of f on E to be

$$\int_E f = \sum_{k=1}^{n} c_k \cdot \mu(E_k).$$

We do not require $\mu(E) < \infty$ in the definition of the Lebesgue integral. Therefore, we tacitly allow for the possibility that

$$\int_E f = +\infty.$$

Also, it is possible that for some k we have $c_k = 0$ and $\mu(E_k) = \infty$. In that case, we adopt the convention

$$c_k \cdot \mu(E_k) = 0 \cdot \infty = 0.$$

Finally, observe that if $E \subseteq \mathbf{R}$ is measurable, then

$$\int_{\mathbf{R}} \chi_E = \mu(E).$$

Before defining the integral for more general nonnegative functions, we need to record some properties of the integral for simple functions.

Theorem 4.6.6. *Suppose $f, g : E \to \mathbf{R}$ are nonnegative, measurable simple functions.*

1. *For all $\alpha \geq 0$, $\int_E \alpha f = \alpha \int_E f$.*
2. *$\int_E (f + g) = \int_E f + \int_E g$.*
3. *If $f \leq g$, then $\int_E f \leq \int_E g$.*

Proof. First write f and g in standard form:

$$f = \sum_{k=1}^{n} a_k \cdot \chi_{E_k}, \quad g = \sum_{j=1}^{m} b_j \cdot \chi_{F_j}.$$

Observe that if $\alpha = 0$, then we have

$$\int_E \alpha f = 0 = \alpha \int_E f.$$

If $\alpha > 0$, then

$$\alpha f = \sum_{k=1}^{n} \alpha a_k \cdot \chi_{E_k},$$

so

$$\int_E \alpha f = \sum_{k=1}^{n} \alpha a_k \cdot \mu(E_k) = \alpha \left(\sum_{k=1}^{n} a_k \cdot \mu(E_k) \right) = \alpha \int_E f.$$

Thus 1 holds.

To show that $\int_E (f+g) = \int_E f + \int_E g$, we need to add f and g and express the resulting simple function *in standard form*. To that end, observe that for each $1 \le k \le n$, we can write

$$E_k = \bigcup_{j=1}^{m} (E_k \cap F_j),$$

and the sets $E_k \cap F_j$ are pairwise disjoint. Similarly, we can express

$$F_j = \bigcup_{k=1}^{n} (E_k \cap F_j)$$

for $1 \le j \le m$. Thus

$$f = \sum_{k=1}^{n} \sum_{j=1}^{m} a_k \cdot \chi_{E_k \cap F_j}, \quad g = \sum_{j=1}^{m} \sum_{k=1}^{n} b_j \cdot \chi_{E_k \cap F_j},$$

so we have

$$f + g = \sum_{j,k} (a_k + b_j) \cdot \chi_{E_k \cap F_j}.$$

This is not quite the standard representation of $f+g$, since the coefficients $a_k + b_j$ may not all be distinct. However, the sets $E_k \cap F_j$ are all disjoint, so it is not too hard to see that we can write

$$\begin{aligned} \int_E (f+g) &= \sum_{j,k} (a_k + b_j) \cdot \mu(E_k \cap F_j) \\ &= \sum_{j,k} a_k \cdot \mu(E_k \cap F_j) + \sum_{j,k} b_j \cdot \mu(E_k \cap F_j) \\ &= \int_E f + \int_E g. \end{aligned}$$

Now assume $f \leq g$, and suppose $x \in E_{k_0} \cap F_{j_0}$ for some k_0 and j_0. Then we have

$$f(x) = \sum_{j,k} a_k \cdot \chi_{E_k \cap F_j} = a_{k_0} \cdot \chi_{E_{k_0} \cap F_{j_0}}(x)$$

and

$$g(x) = \sum_{j,k} b_j \cdot \chi_{E_k \cap F_j} = b_{j_0} \cdot \chi_{E_{k_0} \cap F_{j_0}}(x).$$

Hence $f(x) \leq g(x)$ implies $a_k \leq b_j$ for any k and j satisfying $E_k \cap F_j \neq \emptyset$. Therefore,

$$\int_E f = \sum_{j,k} a_k \cdot \mu(E_k \cap F_j) \leq \sum_{j,k} b_j \cdot \mu(E_k \cap F_j) = \int_E g,$$

since $\mu(E_k \cap F_j) = 0$ whenever $E_k \cap F_j$ is empty. □

Now we can define the Lebesgue integral of an arbitrary nonnegative measurable function.

Definition 4.6.7. Suppose $f : E \to [-\infty, +\infty]$ is a nonnegative, measurable function. We define the **Lebesgue integral** of f on E to be

$$\int_E f = \sup \left\{ \int_E g : 0 \leq g \leq f \text{ and } g \text{ is simple} \right\}.$$

For the purposes of computing or analyzing the integral of a function, the supremum defined above could be quite unwieldy. In particular, there could be uncountably many simple functions less than or equal to f. However, we proved in Theorem 4.6.4 that if $f \geq 0$, we can find a sequence $(f_n)_{n=1}^{\infty}$ of simple functions such that $f_n \nearrow f$ pointwise, meaning the sequence is pointwise increasing and $f_n \to f$. Is it necessarily the case that $\int f_n \to \int f$?

Theorem 4.6.8 (Monotone Convergence Theorem)**.** *Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of nonnegative, measurable functions on a set E satisfying $f_n \leq f_{n+1}$ pointwise for all n. Let*

$$f = \lim_{n\to\infty} f_n = \sup_n f_n.$$

Then

$$\int_E f = \lim_{n\to\infty} \int_E f_n.$$

Proof. Notice first that we have $f_n \leq f$ for all n, so $\int_E f_n \leq \int_E f$ for all n by Exercise 4.6.2. Furthermore,

$$\int_E f_n \leq \int_E f_{n+1}$$

for all n, so the integrals $\int_E f_n$ form an increasing sequence of real numbers with

$$\lim_{n\to\infty} \int_E f_n = \sup_n \int_E f_n \le \int_E f.$$

To establish the reverse inequality, let $\alpha \in (0,1)$ and let g be any simple function satisfying $0 \le g \le f$. Put

$$A_n = \big\{x \in E : f_n(x) \ge \alpha g(x)\big\}$$

for each n. Notice that $f_n \le f_{n+1}$ implies $A_n \subseteq A_{n+1}$. The fact that $f_n \to f$ pointwise guarantees that for all $x \in E$, $f_n(x)$ is eventually bigger than $\alpha g(x)$, meaning that

$$E = \bigcup_{n=1}^{\infty} A_n.$$

Now we have

$$\alpha \int_E g \cdot \chi_{A_n} \le \int_E f \cdot \chi_{A_n} \le \int_E f_n,$$

and if we write the standard representation of g as $g = \sum_{k=1}^n c_k \cdot \chi_{E_k}$, then it is easy to check that

$$g \cdot \chi_{A_n} = \sum_{k=1}^n c_k \cdot \chi_{E_k}\chi_{A_n} = \sum_{k=1}^n c_k \cdot \chi_{E_k \cap A_n}.$$

Thus

$$\int_E g \cdot \chi_{A_n} = \sum_{k=1}^n c_k \cdot \mu(E_k \cap A_n),$$

and continuity of measure guarantees that

$$\lim_{n\to\infty} \mu(E_k \cap A_n) = \mu(E_k)$$

for all k. Therefore,

$$\int_E g \cdot \chi_{A_n} \to \int_E g,$$

whence

$$\alpha \int_E g = \lim_{n\to\infty} \alpha \int_E g \cdot \chi_{A_n} \le \lim_{n\to\infty} \int_E f_n.$$

This holds for all $\alpha \in (0,1)$, so taking the supremum over α of the left hand side yields

$$\int_E g \le \lim_{n\to\infty} \int_E f_n.$$

But g was an arbitrary simple function satisfying $0 \le g \le f$, so it follows that

$$\int_E f \le \lim_{n\to\infty} \int_E f_n.$$

Thus $\int_E f_n \to \int_E f$. $\square$

The Monotone Convergence Theorem is the first of three fundamental convergence theorems for the Lebesgue integral. (Actually, there are others, but they are pretty much just variations on the three main ones. For one, we will slightly generalize the Monotone Convergence Theorem once we have defined the Lebesgue integral for arbitrary functions.) It is an extremely useful theorem, due in large part to the fact that it allows us to address questions about the Lebesgue integral using only simple functions. In particular, we can establish some basic properties of the integral by appealing to results about simple functions.

Theorem 4.6.9. *Suppose f and g are nonnegative measurable functions defined on a set $E \subseteq \mathbf{R}$. Then*

$$\int_E (f+g) = \int_E f + \int_E g.$$

Proof. Of course we already know that $f+g$ is measurable. Furthermore, we can find sequences $(f_n)_{n=1}^{\infty}$ and $(g_n)_{n=1}^{\infty}$ of nonnegative, measurable, simple functions such that $f_n \nearrow f$ and $g_n \nearrow g$. Thus $f_n + g_n \nearrow f + g$, and the Monotone Convergence Theorem guarantees that

$$\int_E (f_n + g_n) \to \int_E (f+g).$$

However, we also know that

$$\int_E (f_n + g_n) = \int_E f_n + \int_E g_n \to \int_E f + \int_E g$$

as $n \to \infty$, so it follows that $\int_E (f+g) = \int_E f + \int_E g$. □

By applying the previous result inductively, we can see that if $f_1, f_2, \ldots, f_n$ are nonnegative measurable functions, then

$$\int_E \sum_{k=1}^{n} f_k = \sum_{k=1}^{n} \int_E f_k.$$

Applying the Monotone Convergence Theorem again yields an even stronger result.

Corollary 4.6.10. *Let $(f_n)_{n=1}^{\infty}$ be a sequence of nonnegative measurable functions, and put*

$$f = \sum_{n=1}^{\infty} f_n.$$

Then

$$\int_E f = \sum_{n=1}^{\infty} \int_E f_n.$$

Proof. For each natural number N, put $g_N = \sum_{n=1}^N f_n$. Since each f_n is nonnegative, $g_N \nearrow f$, and the Monotone Convergence Theorem implies that

$$\lim_{N\to\infty} \int_E g_N = \int_E f.$$

However, notice that

$$\int_E g_N = \int_E \sum_{n=1}^N f_n = \sum_{n=1}^N \int_E f_n,$$

so

$$\int_E f = \lim_{N\to\infty} \sum_{n=1}^N \int_E f_n = \sum_{n=1}^\infty \int_E f_n. \qquad \square$$

Next we establish some results to show that we can neglect sets of measure zero when computing integrals. We need two quick and intuitive facts first.

Proposition 4.6.11. *Suppose $f : E \to [0, +\infty]$ is measurable. If we write $E = A \cup B$, where A and B are disjoint, then*

$$\int_E f = \int_A f + \int_B f.$$

Proof. Let $(f_n)_{n=1}^\infty$ be a sequence of simple functions satisfying $f_n \nearrow f$. Given $n \in \mathbf{N}$, write the standard representation of f_n as

$$f_n = \sum_{k=1}^m c_k \cdot \chi_{E_k}.$$

Since $E_k = (E_k \cap A) \cup (E_k \cap B)$, we can write

$$f_n = \sum_{k=1}^m c_k \cdot \chi_{E_k \cap A} + \sum_{k=1}^m c_k \cdot \chi_{E_k \cap B}.$$

Furthermore, A and B are disjoint, so

$$\int_E f_n = \sum_{k=1}^m c_k \cdot \mu(E_k \cap A) + \sum_{k=1}^m c_k \cdot \mu(E_k \cap B) = \int_A f_n + \int_B f_n.$$

Taking the limit of both sides and applying the Monotone Convergence Theorem yields the desired result. $\square$

Proposition 4.6.12. *Suppose $f : E \to [0, +\infty]$ is measurable. If $\mu(E) = 0$, then $\int_E f = 0$.*

Proof. Let g be any simple function satisfying $0 \le g \le f$, and write

$$g = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}$$

in its standard form. Then

$$\int_E g = \sum_{k=1}^{n} c_k \cdot \mu(E_k) = 0$$

since $\mu(E_k) = 0$ for all k. Taking the supremum over all such simple functions then yields $\int_E f = 0$. $\square$

Theorem 4.6.13. *Suppose $f, g : E \to [0, +\infty]$ are measurable functions.*

1. *If $f \le g$ a.e., then $\int_E f \le \int_E g$.*
2. *If $f = g$ a.e., then $\int_E f = \int_E g$.*

Proof. Suppose first that $f \le g$ a.e., and let

$$Z = \{x \in E : f(x) > g(x)\}.$$

Then $\mu(Z) = 0$. Put $A = E \backslash Z$. Then

$$\int_E f = \int_A f + \int_Z f = \int_A f + 0.$$

Since $f \le g$ on A, we have $\int_A f \le \int_A g$ by Exercise 4.6.2, so

$$\int_E f \le \int_A g + 0 = \int_A g + \int_Z g = \int_E f.$$

If $f = g$ a.e., then $f \le g$ a.e. and $g \le f$ a.e., so $\int_E f = \int_E g$. $\square$

Next we prove a very crude estimation result for the Lebesgue integral, which will yield one final "almost everywhere" theorem.

Theorem 4.6.14 (Chebyshev's Inequality)**.** *Suppose $f : E \to [0, +\infty]$ is measurable. For any $\alpha > 0$ we have*

$$\mu\big(\{f > \alpha\}\big) \leq \frac{1}{\alpha} \int_E f.$$

Proof. Let $\alpha > 0$ and put $A = \{f > \alpha\}$. Then

$$\int_E f \geq \int_A f \geq \int_A \alpha = \alpha \cdot \mu(A).$$

Rearranging, we get

$$\mu(A) \leq \frac{1}{\alpha} \int_E f,$$

whence the result. □

Theorem 4.6.15. *Suppose $f : E \to [0, +\infty]$ is measurable. Then $\int_E f = 0$ if and only if $f = 0$ a.e. on E.*

Proof. Notice that $f = 0$ a.e. clearly implies $\int_E f = 0$. Conversely, for all $\alpha > 0$ we have

$$\mu\big(\{f > \alpha\}\big) \leq \frac{1}{\alpha} \int_E f = 0$$

by Chebyshev's inequality. Therefore,

$$\{f > 0\} = \bigcup_{n=1}^{\infty} \{f > \tfrac{1}{n}\}$$

has measure zero, meaning that $f = 0$ almost everywhere. □

Before closing out this section, let us revisit the Monotone Convergence Theorem. What could go wrong if $f_n \to f$ a.e., but we do not assume that $f_n \nearrow f$?

Example 4.6.16. For each $n \in \mathbf{N}$, define $f_n : [0, \infty) \to \mathbf{R}$ by $f_n = \chi_{[n,n+1]}$. Notice that $f_n \to 0$ pointwise on $[0, \infty)$, and that

$$\int_{[0,\infty)} f_n = \mu([n, n+1]) = 1$$

for all n. Thus

$$\lim_{n\to\infty} \int_{[0,\infty)} f_n = 1 \neq 0 = \int_{[0,\infty)} \lim_{n\to\infty} f_n.$$

Our finding does not violate the Monotone Convergence Theorem, since the sequence $(f_n)_{n=1}^{\infty}$ is not increasing pointwise. The real issue here is that the areas under the graphs of the functions "escape to infinity" as $n \to \infty$.

Example 4.6.17. In Example 3.1.4, we considered the sequence of functions $f_n : [0,1] \to \mathbf{R}$ defined by

$$f_n(x) = \begin{cases} 2^{2n}x & \text{if } 0 \le x \le \frac{1}{2^n} \\ 2^{2n}\left(\frac{1}{2^{n-1}} - x\right) & \text{if } \frac{1}{2^n} < x \le \frac{1}{2^{n-1}} \\ 0 & \text{if } \frac{1}{2^{n-1}} < x \le 1, \end{cases}$$

whose graphs look as follows:

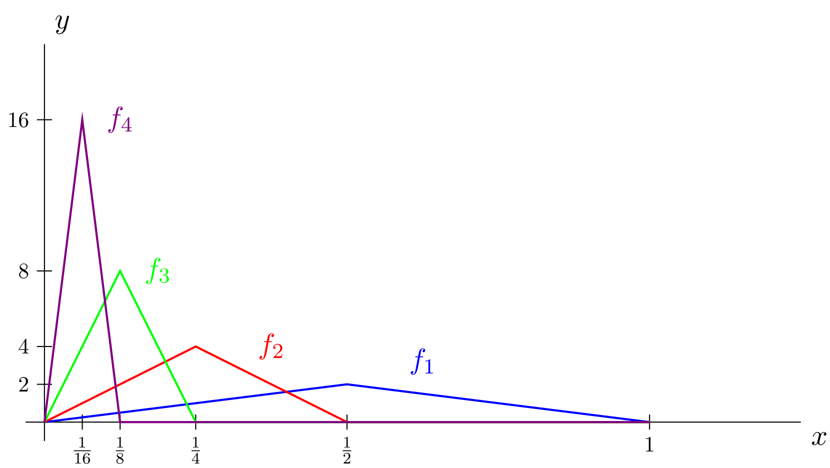

For each n, the area under the graph of f_n is 1. However, $f_n \to 0$ pointwise, so $\lim \int f_n \neq \int \lim f_n$. Again, the sequence is not pointwise monotone, so we have not violated the Monotone Convergence Theorem.

Even in situations where $f_n \to f$ a.e. but $\int f_n \not\to \int f$, we still have a useful inequality at our disposal. In fact, it is not necessary to assume that a sequence $(f_n)_{n=1}^{\infty}$ converges at all.

Theorem 4.6.18 (Fatou's Lemma)**.** *Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of nonnegative measurable functions defined on a set E. Then*

$$\int_E \liminf_{n\to\infty} f_n \le \liminf_{n\to\infty} \int_E f_n.$$

Proof. For each n, put $g_n = \inf_{k\ge n} f_k$. Then $g_n \le g_{n+1}$ for all n, and

$$\lim_{n\to\infty} g_n = \liminf_{n\to\infty} f_n$$

by definition. Thus $g_n \nearrow \liminf f_n$, so the Monotone Convergence Theorem implies

$$\int_E \liminf_{n\to\infty} f_n = \lim_{n\to\infty} \int_E g_n.$$

Now fix n, and observe that $g_n \leq f_k$ for all $k \geq n$ by construction. Thus

$$\int_E g_n \leq \int_E f_k$$

for all $k \geq n$, which implies that

$$\int_E g_n \leq \inf_{k \geq n} \int_E f_k.$$

Hence

$$\lim_{n\to\infty} \int_E g_n \leq \lim_{n\to\infty} \inf_{k \geq n} \int_E f_k = \liminf_{n\to\infty} \int_E f_n$$

by the definition of the limit inferior. It then follows that

$$\int_E \liminf_{n\to\infty} f_n \leq \liminf_{n\to\infty} \int_E f_n,$$

and we are done. □

In the case that the sequence $(f_n)_{n=1}^{\infty}$ converges a.e., Fatou's lemma takes the following form.

Corollary 4.6.19. *Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of nonnegative measurable functions defined on a set E, and $(f_n)_{n=1}^{\infty}$ converges a.e. to a function on E. Then*

$$\int_E \lim_{n\to\infty} f_n \leq \liminf_{n\to\infty} \int_E f_n.$$

Exercises for Section 4.6

Exercise 4.6.1. Suppose f is a nonnegative measurable function on a set E, and $\alpha \geq 0$. Prove that

$$\int_E \alpha f = \alpha \int_E f.$$

Exercise 4.6.2. Suppose f and g are nonnegative measurable functions defined on a measurable set $E \subseteq \mathbf{R}$, and $f \leq g$ on E. Prove that

$$\int_E f \leq \int_E g.$$

Exercise 4.6.3. Let f be a nonnegative measurable function on $\mathbf{R}$. Prove that

$$\lim_{n\to\infty} \int_{[-n,n]} f = \int_{\mathbf{R}} f.$$

Exercise 4.6.4 ([WZ77], Exercise 5.3)**.** Let $(f_n)_{n=1}^{\infty}$ be a sequence of nonnegative measurable functions defined on a set E. If $f_n \to f$ a.e. and $f_n \leq f$ for all n, show that $\int_E f_n \to \int_E f$.

4.7 Integrals of Arbitrary Measurable Functions

Using what we already know about nonnegative functions, we now extend the definition of the Lebesgue integral to arbitrary Lebesgue measurable functions. Given a measurable function $f : E \to [-\infty, +\infty]$, recall that we can decompose f into its positive and negative parts:

$$f = f^+ - f^-,$$

where f^+ and f^- are measurable and nonnegative. Since we know how to integrate nonnegative functions, our natural inclination is to define

$$\int_E f = \int_E f^+ - \int_E f^-.$$

Of course we could run into issues if $\int_E f^+$ and $\int_E f^-$ are both infinite. In this case, we can't assign a meaningful value to $\int_E f$. Therefore, we need to require that at least one of the integrals on the right hand side is finite; in fact, things are much easier if we assume they are both finite.

Definition 4.7.1. A measurable function $f : E \to [-\infty, +\infty]$ is said to be **Lebesgue integrable** on E if $\int_E f^+ < \infty$ and $\int_E f^- < \infty$. If f is an integrable function on E, we define its **Lebesgue integral** over E to be

$$\int_E f = \int_E f^+ - \int_E f^-.$$

Given a set $E \subseteq \mathbf{R}$, we denote the set of all Lebesgue integrable functions on E by $L^1(E)$. This notation may be reminiscent of ℓ^1. Indeed, we will see later that there is a connection between these two spaces.

Since the integral of a function f is defined in terms of the integrals of f^+ and f^-, we can quickly prove results about the general Lebesgue integral using facts regarding nonnegative functions. Before doing that, we establish three quick estimation results.

Proposition 4.7.2. *A measurable function $f : E \to [-\infty, +\infty]$ is integrable if and only if $\int_E |f| < +\infty$.*

Proof. Both implications follow from the observation that

$$|f| = f^+ + f^-.$$

If f is integrable, then we have

$$\int_E |f| = \int_E f^+ + \int_E f^- < \infty.$$

Conversely, if $\int_E |f| < \infty$, then it follows that f is integrable since $\int_E f^+ \le \int_E |f|$ and $\int_E f^- \le \int_E |f|$. $\square$

Proposition 4.7.3. *For all $f \in L^1(E)$, we have $\left|\int_E f\right| \le \int_E |f|$.*

Proof. Let $f \in L^1(E)$. Notice that

$$\left|\int_E f\right| = \left|\int_E f^+ - \int_E f^-\right| \le \int_E f^+ + \int_E f^- = \int_E |f|,$$

whence the result. $\square$

Proposition 4.7.4. *If $f \in L^1(E)$, then f is finite a.e. on E.*

Proof. Let $f \in L^1(E)$. Then $\int_E |f| < \infty$. Put $C = \int_E |f|$. The for each $n \in \mathbf{N}$, Chebyshev's inequality implies

$$\mu\big(\{|f| > n\}\big) \le \frac{1}{n}\int_E |f| = \frac{C}{n},$$

and the right hand side tends to zero as $n \to \infty$. Since $\{|f| > n\} \supseteq \{|f| > n+1\}$ for all n, and

$$\{|f| = \infty\} = \bigcap_{n=1}^{\infty} \{|f| > n\},$$

continuity of measure implies that

$$\mu\big(\{|f| = \infty\}\big) = \lim_{n\to\infty} \mu\big(\{|f| > n\}\big) = 0.$$

Thus $|f|$ is finite a.e., so f is finite a.e. on E. $\square$

Now we can use what we already know about nonnegative functions to establish basic results about the Lebesgue integral for arbitrary integrable functions.

Proposition 4.7.5. *Let $f, g \in L^1(E)$.*

1. *The function $f + g \in L^1(E)$, and $\int_E (f+g) = \int_E f + \int_E g$.*
2. *For all $\alpha \in \mathbf{R}$, $\alpha f \in L^1(E)$, and $\int_E \alpha f = \alpha \int_E f$.*
3. *If we write $E = A \cup B$ with A and B disjoint, then $\int_E f = \int_A f + \int_B f$.*
4. *If $f \le g$ a.e., then $\int_E f \le \int_E g$.*
5. *If $f = g$ a.e., then $\int_E f = \int_E g$.*

Proof. Notice that if $f, g \in L^1(E)$, then

$$\int_E |f+g| \le \int_E |f| + \int_E |g| < \infty,$$

so $f + g \in L^1(E)$. Put $h = f + g$. Then we have

$$h^+ - h^- = f^+ - f^- + g^+ - g^-,$$

and rearranging yields

$$h^+ + f^- + g^- = f^+ + g^+ + h^-.$$

Both sides consist of sums of nonnegative functions, so we have

$$\int_E h^+ + \int_E f^- + \int_E g^- = \int_E f^+ + \int_E g^+ + \int_E h^-$$

by our previous results for nonnegative functions. Rearranging again, we get

$$\int_E h^+ - \int_E h^- = \int_E f^+ - \int_E f^- + \int_E g^+ - \int_E g^-,$$

or

$$\int_E h = \int_E f + \int_E g,$$

so (1) holds.

Now let $\alpha \in \mathbf{R}$. If $\alpha \ge 0$, then we have $\alpha f = \alpha f^+ - \alpha f^-$ with αf^+ and αf^- nonnegative. Thus it is easy to see that $(\alpha f)^+ = \alpha f^+$ and $(\alpha f)^- = \alpha f^-$, so

$$\int_E \alpha f = \int_E \alpha f^+ - \int_E \alpha f^- = \alpha \int_E f^+ - \alpha \int_E f^- = \alpha \int_E f.$$

Now suppose $\alpha < 0$. Since we can write $\alpha = -|\alpha|$ with $|\alpha| \ge 0$, it suffices to show that (2) holds with $\alpha = -1$. Notice that

$$(-f)^+ = \max\{-f, 0\} = -\min\{f, 0\} = f^-$$

and

$$(-f)^- = \max\{-(-f), 0\} = \max\{f, 0\} = f^+,$$

so

$$\int_E (-f) = \int_E f^- - \int_E f^+ = -\left(\int_E f^+ - \int_E f^-\right) = -\int_E f.$$

Therefore, for any $\alpha < 0$ we have

$$\int_E \alpha f = -\int_E |\alpha| f = -|\alpha| \int_E f = \alpha \int_E f,$$

so (2) holds for all $\alpha \in \mathbf{R}$.

The proof of (3) is straightforward. Again, using previous results for nonnegative functions we get

$$\begin{aligned}\int_E f &= \int_E f^+ - \int_E f^- \\ &= \left(\int_A f^+ + \int_B f^+\right) - \left(\int_A f^- + \int_B f^-\right) \\ &= \left(\int_A f^+ - \int_A f^-\right) + \left(\int_B f^+ - \int_B f^-\right) \\ &= \int_A f + \int_B f.\end{aligned}$$

Now suppose $f \leq g$ a.e., and let $Z = \{x \in E : f(x) > g(x)\}$. Then $\mu(Z) = 0$, and the function $g - f$ is nonnegative everywhere on $E \backslash Z$. Thus $\int_{E \backslash Z} (g - f) \geq 0$, and

$$\left| \int_Z (g - f) \right| \leq \int_Z |g - f| = 0$$

since $\mu(Z) = 0$. It follows that

$$\int_E (g - f) = \int_{E \backslash Z} (g - f) + \int_Z (g - f) \geq 0,$$

or $\int_E f \leq \int_E g$. Thus (4) holds. Property (5) is then immediate from (4) since $f = g$ a.e. precisely when $f \leq g$ a.e. and $g \leq f$ a.e. on E. □

Now we arrive at the convergence theorems for the general Lebesgue integral. The *pièce de résistance*[2] is Lebesgue's Dominated Convergence Theorem.

Theorem 4.7.6 (Dominated Convergence Theorem)**.** *Let $(f_n)_{n=1}^\infty$ be a sequence of functions in $L^1(E)$. Suppose f_n converges a.e. to a function f and that there is a function $g \in L^1(E)$ such that $|f_n| \leq g$ a.e. for all n. Then $f \in L^1(E)$ and*

$$\lim_{n \to \infty} \int_E f_n = \int_E f.$$

Proof. Notice first that $|f_n| \leq g$ a.e. for all n implies $|f| \leq g$ a.e., so

$$\int_E |f| \leq \int_E g < \infty.$$

Therefore $f \in L^1(E)$.

[2]Thank you, Jody Trout.

Now we consider the sequence $(g+f_n)_{n=1}^{\infty}$. Since $|f_n| \le g$ a.e., we have $g+f_n \ge 0$ a.e., and Fatou's Lemma implies[3]

$$\int_E \liminf_{n\to\infty}(g+f_n) \le \liminf_{n\to\infty} \int_E (g+f_n) = \int_E g + \liminf_{n\to\infty} \int_E f_n.$$

But $g+f_n \to g+f$ a.e., so

$$\int_E g + \int_E f = \int_E (g+f) \le \int_E g + \liminf_{n\to\infty} \int_E f_n,$$

Since $\int_E g$ is finite, it follows that

$$\int_E f \le \liminf_{n\to\infty} \int_E f_n.$$

Now we apply similar reasoning to $(g-f_n)_{n=1}^{\infty}$. We have $g-f_n \ge 0$ a.e. and $g-f_n \to g-f$ a.e., so Fatou's Lemma gives

$$\begin{aligned}\int_E g - \int_E f &\le \liminf_{n\to\infty} \int_E (g-f_n) \\ &= \int_E g + \liminf_{n\to\infty}\left(-\int_E f_n\right) \\ &= \int_E g - \limsup_{n\to\infty} \int_E f_n.\end{aligned}$$

Thus $\int_E f \ge \limsup \int_E f_n$, so we have

$$\limsup_{n\to\infty} \int_E f_n \le \int_E f \le \liminf_{n\to\infty} \int_E f_n,$$

which forces $\liminf \int_E f_n = \limsup \int_E f_n$. Thus the sequence of integrals $\int_E f_n$ converges, and

$$\lim_{n\to\infty} \int_E f_n = \int_E f. \qquad \square$$

As a quick application of the Dominated Convergence Theorem, we have a condition that tells us when one can interchange a limit with an infinite summation. This result will be crucial later for showing that $L^1(E)$ is a Banach space.

[3]To avoid any possible confusion, it's worth noting here that the limit inferior is *not* additive. In general, we can only say

$$\liminf_{n\to\infty}(a_n+b_n) \ge \liminf_{n\to\infty} a_n + \liminf_{n\to\infty} b_n.$$

However, since we are simply adding the constant $\int_E g$ here, we really do have equality.

Theorem 4.7.7. *Let $(f_n)_{n=1}^\infty$ be a sequence in $L^1(E)$ with $\sum_{n=1}^\infty \int_E |f_n| < \infty$. Then $\sum_{n=1}^\infty f_n$ converges a.e. to a function $f \in L^1(E)$, and*

$$\int_E \sum_{n=1}^\infty f_n = \sum_{n=1}^\infty \int_E f_n.$$

Proof. First observe that

$$\int_E \sum_{n=1}^\infty |f_n| = \sum_{n=1}^\infty \int_E |f_n| < \infty$$

by Corollary 4.6.10. Hence the function $g = \sum_{n=1}^\infty |f_n|$ belongs to $L^1(E)$. Since g is finite a.e., the series $\sum_{n=1}^\infty f_n(x)$ is absolutely convergent for almost all $x \in E$. Thus $\sum_{n=1}^\infty f_n$ converges a.e. to a function f. Now observe that for each $N \in \mathbf{N}$ we have

$$\left|\sum_{n=1}^N f_n\right| \le \sum_{n=1}^N |f_n| \le g,$$

so the Dominated Convergence Theorem implies that $f \in L^1(E)$ and

$$\int_E f = \lim_{N\to\infty} \int_E \sum_{n=1}^N f_n = \lim_{N\to\infty} \sum_{n=1}^N \int_E f_n = \sum_{n=1}^\infty \int_E f_n. \qquad \square$$

We will see another application of the Dominated Convergence Theorem momentarily. First we discuss extensions of the Monotone Convergence Theorem and Fatou's Lemma to arbitrary measurable functions. In both cases, we need some sort of "bounding function" as in the statement of the Dominated Convergence Theorem.

Theorem 4.7.8 (Monotone Convergence Theorem)**.** *Let $(f_n)_{n=1}^\infty$ be a sequence of integrable functions on a set E.*

1. *Suppose $f_n \nearrow f$ a.e. for some measurable function f, and there exists $g \in L^1(E)$ such that $f_n \ge g$ a.e. for all n. Then $\int_E f_n \to \int_E f$.*

2. *Suppose $f_n \searrow f$ a.e. for some measurable function f, and there exists $h \in L^1(E)$ such that $f_n \le g$ a.e. for all n. Then $\int_E f_n \to \int_E f$.*

Proof. To prove (1), we consider the sequence $(f_n - g)_{n=1}^\infty$. Notice that $f_n - g \to f - g$ a.e. and $f_n - g \ge 0$ a.e. for all n. In other words, $f_n - g \nearrow f - g$, and the original version of the Monotone Convergence Theorem implies that

$$\lim_{n\to\infty} \int_E (f_n - g) = \int_E (f - g).$$

Equivalently,

$$\lim_{n\to\infty} \int_E f_n - \int_E g = \int_E f - \int_E g,$$

and since $\int_E g$ is finite, we can cancel to obtain

$$\lim_{n\to\infty} \int_E f_n = \int_E f.$$

To prove (2), we simply consider the sequence $(-f_n)_{n=1}^{\infty}$. Then $-f_n \nearrow -f$ a.e., and $-f_n \geq -h$ a.e. for all n, so (1) implies that $\int_E(-f_n) \to \int_E(-f)$. The result follows. □

The extension of Fatou's Lemma is proved in an analogous manner using our original version of Fatou's Lemma for nonnegative functions.

Theorem 4.7.9 (Fatou's Lemma). *Let $(f_n)_{n=1}^{\infty}$ be a sequence of integrable functions on a set E.*

1. *Suppose there is a function $g \in L^1(E)$ such that $f_n \geq g$ a.e. for all n. Then*

$$\int_E \liminf_{n\to\infty} f_n \leq \liminf_{n\to\infty} \int_E f_n.$$

2. *Suppose there is a function $h \in L^1(E)$ such that $f_n \leq h$ a.e. for all n. Then*

$$\int_E \limsup_{n\to\infty} f_n \geq \limsup_{n\to\infty} \int_E f_n.$$

Remark 4.7.10. We can further relax the hypotheses in these more general versions of the Monotone Convergence Theorem and Fatou's Lemma. Notice that we did not actually use the hypothesis that the functions f_n to lie in $L^1(E)$ in either of the proofs. Indeed, we can remove the assumption that the functions f_n are integrable, though some care is needed. We did not define the Lebesgue integral for functions outside of $L^1(E)$, though we can do so under the right assumptions. Notice that if f is a measurable function on E, $g \in L^1(E)$, and $f \geq g$ a.e., then we have

$$f^- = -\min\{f, 0\} \leq -\min\{g, 0\} = g^-.$$

Thus

$$\int_E f^- \leq \int_E g^- < \infty,$$

so

$$\int_E f = \int_E f^+ - \int_E f^-$$

is well-defined—either $f \in L^1(E)$, or the integral is infinite. Similarly, if f is measurable and there exists $g \in L^1(E)$ with $f \le g$ a.e., then $\int_E f^+$ is finite and $\int_E f$ is well-defined. In other words, the hypotheses of the Monotone Convergence Theorem and Fatou's Lemma can be relaxed (namely to just the requirement that each f_n is measurable) while still guaranteeing that all the relevant integrals make sense.

4.7.1 The Relationship with the Riemann Integral

As one final application of the Dominated Convergence Theorem, we investigate the relationship between the Riemann and Lebesgue integrals. In particular, we show that Lebesgue's theory of integration subsumes the Riemann integral, at least when considering integrals of bounded functions on bounded intervals. (We will not discuss the concept of an improper Riemann integral right now.)

Before stating and proving the theorem, we need to introduce some notation and recall some facts about the Riemann integral. Recall first that the Riemann integral is most easily defined via Darboux sums. Given a partition

$$P = \{x_0 = a < x_1 < x_2 < \cdots < x_{n-1} < x_n = b\}$$

of the interval $[a,b]$, we let $\Delta x_i = x_i - x_{i-1}$, and we call $\|P\| = \max_{1\le i\le n} \Delta x_i$ the **mesh** of P. Now we define the **upper and lower Darboux sums** by

$$U(f,P) = \sum_{i=1}^{n} M_i \Delta x_i, \quad L(f,P) = \sum_{i=1}^{n} m_i \Delta x_i,$$

where M_i and m_i denote the supremum and infimum, respectively, of f on the interval $[x_{i-1}, x_i]$. For any partition P we have $L(f,P) \le U(f,P)$, and we declare f to be **Riemann integrable** on $[a,b]$ if given any $\varepsilon > 0$ there exists a partition P_ε satisfying

$$U(f,P_\varepsilon) - L(f,P_\varepsilon) < \varepsilon.$$

If f is Riemann integrable, we define its **Riemann integral** to be

$$(R)\int_a^b f = \inf_P U(f,P) = \sup_P L(f,P).$$

We use the notation $(R)\int_a^b f$ to distinguish the Riemann integral of f from the Lebesgue integral, which we denote with the customary $\int_{[a,b]} f$.

Finally, to any partition P, we can associate two simple functions u_P and l_P, which will be useful below:

$$u_P = \sum_{i=1}^{n} M_i \chi_{[x_{i-1},x_i)}, \quad l_P = \sum_{i=1}^{n} m_i \chi_{[x_{i-1},x_i)}. \tag{4.2}$$

Notice that u_P and l_P are both measurable, and their Lebesgue integrals are precisely the upper and lower Darboux sums of f:

$$\int_{[a,b]} u_P = \sum_{i=1}^{n} M_i \mu([x_{i-1}, x_i)) = \sum_{i=1}^{n} M_i \Delta x_i$$

and likewise,

$$\int_{[a,b]} l_P = \sum_{i=1}^{n} m_i \Delta x_i.$$

Also, it is straightforward to check that if P' is a refinement of P (meaning that $P' \subseteq P$), then $u_{P'} \leq u_P$ and $l_{P'} \geq l_P$.

Theorem 4.7.11. *Suppose $f : [a,b] \to \mathbf{R}$ is bounded. If f is Riemann integrable on $[a,b]$, then it is Lebesgue measurable (hence integrable), and its Riemann and Lebesgue integrals agree:*

$$(R) \int_a^b f = \int_{[a,b]} f.$$

Proof. Assume f is Riemann integrable on $[a,b]$. Then for each $n \in \mathbf{N}$, we can find a partition P_n of $[a,b]$ satisfying

$$U(f, P_n) - L(f, P_n) < \frac{1}{n}.$$

Thus we obtain a sequence of partitions $(P_n)_{n=1}^{\infty}$ such that

$$\lim_{n\to\infty} U(f, P_n) = \lim_{n\to\infty} L(f, P_n) = (R) \int_a^b f.$$

By adding points to each P_n if necessary, we can assume that P_{n+1} is a refinement of P_n for all n, and that $\|P_n\| \to 0$ as $n \to \infty$.

For each n, let $u_n = u_{P_n}$ and $l_n = l_{P_n}$ denote the simple functions associated to P_n as in (4.2). Since P_{n+1} is a refinement of P_n for all n, it follows that $u_{n+1} \leq u_n$ for all n. Thus the sequence $(u_n)_{n=1}^{\infty}$ is pointwise decreasing and bounded, so it converges pointwise to a function g on $[a,b]$. Similarly, $(l_n)_{n=1}^{\infty}$ is pointwise decreasing and bounded, so it converges pointwise to a function h. Notice that g and h are both measurable, since they are pointwise limits of measurable functions, and

$$h \leq f \leq g.$$

Moreover, the Dominated Convergence Theorem implies that

$$\int_{[a,b]} g = \lim_{n\to\infty} \int_{[a,b]} u_n = \lim_{n\to\infty} U(f, P_n) = (R) \int_a^b f$$

and

$$\int_{[a,b]} h = \lim_{n\to\infty} \int_{[a,b]} l_n = \lim_{n\to\infty} L(f, P_n) = (R)\int_a^b f.$$

It follows that

$$\int_{[a,b]} (g-h) = 0,$$

so $g - h = 0$ a.e. on $[a, b]$. Hence $f = g$ a.e. (and $f = h$ a.e.), so f is measurable and

$$\int_{[a,b]} f = \int_{[a,b]} g = (R)\int_a^b f.$$

Thus the Riemann and Lebesgue integrals of f agree. □

Exercises for Section 4.7

Exercise 4.7.1. Use the Dominated Convergence Theorem to compute

$$\lim_{n\to\infty} \int_0^1 \frac{n\sqrt{x}}{1+n^2 x}\, dx.$$

Exercise 4.7.2. Let $E \subseteq \mathbf{R}$ be a measurable set with $\mu(E) < \infty$. Suppose f_n is Lebesgue integrable for all n, and that $f_n \to f$ uniformly on E. Prove that $\int_E f_n \to \int_E f$.

Exercise 4.7.3. Suppose f and g are extended real-valued functions on $E \subseteq \mathbf{R}$. Assume $f \in L^1(E)$ and g is a.e. bounded, meaning there exists a constant $C > 0$ such that

$$\mu\bigl(\{x \in E : |g(x)| > C\}\bigr) = 0.$$

Prove that $fg \in L^1(E)$.

Exercise 4.7.4 ([WZ77], Exercise 5.2)**.** Show that the conclusion of the general Monotone Convergence Theorem can fail if the bounding function is not assumed to lie in $L^1(E)$. In particular, exhibit a sequence $(f_n)_{n=1}^\infty$ of functions on a set E such that $f_n \searrow f$ and $f_n \le g$ for some measurable function g (which necessarily does not belong to $L^1(E)$), but

$$\int_E f_n \not\to \int_E f.$$

Exercise 4.7.5 ([WZ77], Exercise 5.21)**.** Let f be a measurable function defined on a set E. If $\int_A f = 0$ for every measurable set $A \subseteq E$, prove that $f = 0$ a.e. on E.

Exercise 4.7.6. Let $E \subseteq \mathbf{R}$, and let $(f_n)_{n=1}^\infty$ be a sequence of functions in $L^1(E)$.

(a) Given a function $f \in L^1(E)$, we say $f_n \to f$ in L^1 if

$$\|f_n - f\|_1 = \int_E |f_n - f| \to 0$$

as $n \to \infty$. Prove that if $f_n \to f$ a.e. and there exists $g \in L^1(E)$ such that $|f_n| \le g$ for all n, then $f_n \to f$ in L^1.

(b) If $f_n \to f$ in L^1, must it be the case that $f_n \to f$ a.e. on E? Give either a proof or a counterexample.

(c) If $f_n \ge 0$ a.e. for all n and $f_n \to f$ in L^1, must it be the case that $f \ge 0$ a.e. as well?

Exercise 4.7.7. Prove the **Generalized Dominated Convergence Theorem**:

> **Theorem 4.7.12.** *Let $(f_n)_{n=1}^\infty$ be a sequence of measurable functions on a set E that converges a.e. to a function f. Suppose there is a sequence $(g_n)_{n=1}^\infty$ of nonnegative functions in $L^1(E)$ satisfying*
>
> $$|f_n| \le g$$
>
> *for all $n \in \mathbf{N}$, and there exists $g \in L^1(E)$ such that $g_n \to g$ a.e. and $\int_E g_n \to \int_E g$. Then*
>
> $$\int_E f_n \to \int_E f.$$

Exercise 4.7.8. (a) Let $f : [0,\infty) \to \mathbf{R}$ be a nonnegative measurable function. Prove that

$$\lim_{n\to\infty} \int_{[0,n]} f = \int_{[0,\infty)} f.$$

(b) Suppose $f : [0,\infty) \to \mathbf{R}$ is improperly Riemann integrable on $[0,\infty)$, meaning that

$$\int_0^\infty f(x)\,dx = \lim_{t\to\infty} \int_0^t f(x)\,dx$$

exists and is finite. Assume moreover that the improper integral converges *absolutely*, i.e.

$$\int_0^\infty |f(x)|\,dx = \lim_{t\to\infty} \int_0^t |f(x)|\,dx < \infty.$$

Prove that f is Lebesgue integrable on $[0,\infty)$, and that its Lebesgue integral agrees with its improper Riemann integral.

(c) Let $f(x) = \frac{\sin x}{x}$. Show that f is improperly Riemann integrable on $[0,\infty)$, but not Lebesgue integrable on $[0,\infty)$. (Hint: Show that $\int_{[0,\infty)} f^+ = \infty$.)

4.8 The Theorems of Egorov and Lusin

We now arrive at two (somewhat) surprising theorems regarding measurable functions. These theorems are precise realizations of heuristic statements laid out by J. E. Littlewood in his three principles of real analysis. Paraphrasing from his original statements, Littlewood's principles are the following:

1. Every measurable set is nearly an open set.
2. Every pointwise convergent sequence of measurable functions is nearly uniformly convergent.
3. Every measurable function is nearly continuous.

In each of these statements, Littlewood's use of "nearly" should be interpreted as "except for a set of arbitrarily small measure". Indeed, we have already seen a precise statement of Littlewood's first principle—given a Lebesgue measurable set $E \subseteq \mathbf{R}$ and $\varepsilon > 0$, there exists an open set $U \supseteq E$ such that

$$\mu(U \backslash E) < \varepsilon.$$

That is, E differs from an open set by a set whose measure is as small as we like.

Littlewood's other two principles can also be made precise, and we will now look at two theorems that do just that. First we have Egorov's theorem, which quantifies Littlewood's second principle.

Theorem 4.8.1 (Egorov)**.** *Let $E \subseteq \mathbf{R}$ be a measurable set with $\mu(E) < \infty$. Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of measurable functions on E that converges pointwise to a finite function f. Then given $\varepsilon > 0$, there exists a closed set $F \subseteq E$ such that $\mu(E \backslash F) < \varepsilon$ and $f_n \to f$ uniformly on F.*

The proof of Egorov's theorem is not all that difficult, but we do need to prove a quick lemma first.

Lemma 4.8.2. *Given $\varepsilon, \eta > 0$, there exists a closed set $F \subseteq E$ and $M \in \mathbf{N}$ such that*

$$\mu(E \backslash F) < \eta \quad \textit{and} \quad |f(x) - f_n(x)| < \varepsilon$$

for all $x \in F$ and all $n \geq M$.

Proof. Let $\varepsilon, \eta > 0$ be given. Since f is finite, the function $|f - f_n|$ is well-defined and measurable, so for each $n \in \mathbf{N}$ the set

$$E_n = \big\{x \in E : |f(x) - f_k(x)| < \varepsilon \text{ for all } k > n\big\} = \bigcap_{k>n} \{|f - f_n| < \varepsilon\}$$

is measurable. Also, clearly $E_n \subseteq E_{n+1}$ for all n, and the assumption that $f_n \to f$ pointwise implies that each $x \in E$ belongs to some E_n for sufficiently large n. Therefore,

$$E = \bigcup_{n=1}^{\infty} E_n,$$

and continuity of measure implies that

$$\mu(E) = \lim_{n\to\infty} \mu(E_n).$$

Since $\mu(E)$ is finite, we can choose M such that $\mu(E_M) > \mu(E) - \frac{\eta}{2}$, or

$$\mu(E \backslash E_M) = \mu(E) - \mu(E_M) < \frac{\eta}{2}.$$

Now choose a closed set $F \subseteq E_M$ with $\mu(E_M \backslash F) < \frac{\eta}{2}$. Then

$$\mu(E \backslash F) = \mu(E \backslash E_M) + \mu(E_M \backslash F) < \eta,$$

and

$$|f(x) - f_n(x)| < \varepsilon$$

for all $x \in F$ when $n \geq M$. □

Proof of Egorov's Theorem. Let $\varepsilon > 0$ be given. By the previous lemma, for each $k \geq 1$ we can find a closed set $F_k \subseteq E$ and $M_k \in \mathbf{N}$ such that

$$\mu(E \backslash F_k) < \frac{\varepsilon}{2^k}$$

and

$$|f(x) - f_n(x)| < \frac{1}{k}$$

for all $x \in F_k$ when $n \geq M_k$. Set

$$F = \bigcap_{k=1}^{\infty} F_k.$$

Then F is closed, and since $F \subseteq F_k$ for all k, it follows that $f_n \to f$ uniformly on F. Finally, we have

$$E \backslash F = E \cap \bigcup_{k=1}^{\infty} F_k^c = \bigcup_{k=1}^{\infty} (E \backslash F_k),$$

so

$$\mu(E \backslash F) \leq \sum_{k=1}^{\infty} \mu(E \backslash F_k) < \sum_{k=1}^{\infty} \frac{\varepsilon}{2^k} = \varepsilon. \qquad \square$$

Finally, Littlewood's third principle is made precise by Lusin's theorem. Again, we will have to prove the theorem in two steps—first for simple functions, and then for arbitrary measurable functions.

Theorem 4.8.3 (Lusin). *Suppose $E \subseteq \mathbf{R}$ is a measurable set with $\mu(E) < \infty$, and $f : E \to \mathbf{R}$ is measurable. Given $\varepsilon > 0$, there is a closed set $F \subseteq E$ such that $\mu(E \backslash F) < \varepsilon$ and f is continuous on F.*

Lemma 4.8.4. *The conclusion of Lusin's theorem holds for measurable simple functions.*

Proof. Let f be a measurable simple function on E, and write f in its standard form:

$$f = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}.$$

Let $\varepsilon > 0$ be given, and for $1 \leq k \leq n$, find a closed set $F_k \subseteq E_k$ such that $\mu(E_k \backslash F_k) < \frac{\varepsilon}{n}$. Then

$$F = \bigcup_{k=1}^{n} F_k$$

is closed, and

$$E \backslash F = \bigcup_{k=1}^{n} (E_k \backslash F_k),$$

so

$$\mu(E \backslash F) = \sum_{k=1}^{n} \mu(E_k \backslash F_k) < n \cdot \frac{\varepsilon}{n} = \varepsilon.$$

Notice that f is constant on each F_k. Moreover, the sets $F_1, F_2, \ldots, F_n$ are pairwise disjoint, so they are relatively open in E. (In particular, F_i and F_j are separated whenever $i \neq j$.) Thus f is locally constant on F, hence it is continuous on F. This completes the proof. $\square$

Proof of Lusin's Theorem. Assume $f : E \to \mathbf{R}$ is measurable. Then there is a sequence $(f_n)_{n=1}^{\infty}$ of measurable simple functions such that $f_n \to f$ pointwise on E. For each n, the previous lemma implies that there is a closed set $F_n \subseteq E$ such that

$$\mu(E \backslash F_n) < \frac{\varepsilon}{2^{n+1}}$$

and f_n is continuous on F_n. By Egorov's theorem, there is a closed set $F_0 \subseteq E$ such that $\mu(E \backslash F_0) < \frac{\varepsilon}{2}$ and $f_n \to f$ uniformly on F_0. Let

$$F = F_0 \cap \left(\bigcap_{n=1}^{\infty} F_n \right).$$

Then F is closed, each f_n is continuous on F, and $f_n \to f$ uniformly on F. Therefore, f is continuous on F. Moreover,

$$\mu(E \backslash F) \leq \mu(E \backslash F_0) + \sum_{n=1}^{\infty} \mu(E \backslash F_n) < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon. \qquad \square$$

It is also possible to prove Lusin's theorem when $\mu(E) = \infty$, but we omit the details.

Exercises for Section 4.8

Exercise 4.8.1 ([WZ77], Exercise 4.20)**.** Suppose $f : [a,b] \to \mathbf{R}$ is measurable. Show that for any $\varepsilon > 0$, there is a continuous function $g : [a,b] \to \mathbf{R}$ such that

$$\mu\big(\{x \in [a,b] : f(x) \neq g(x)\}\big) < \varepsilon.$$

Chapter 5

Abstract Measure and Integration

So far we have studied Lebesgue's theory of measure and integration on the real line. However, some careful thought should reveal that the definition of the Lebesgue integral relies only on the presence of a *measure*, i.e., a function that allows one to determine the sizes of sets and has several crucial properties similar to those of Lebesgue measure (such as countable additivity). In other words, we have unwittingly laid the groundwork for integration theories on other, more general spaces. Consequently, this chapter is devoted to generalizing the Lebesgue theory to abstract *measure spaces*.

5.1 Measure Spaces

If we think back to the construction of Lebesgue measure on $\mathbf{R}$, we were able to establish many desirable properties using only two crucial facts.

1. The collection of Lebesgue measurable sets forms a σ-algebra.
2. Lebesgue measure is countably additive on disjoint sets.

Consequently, many of the nice properties of Lebesgue measure (as well as the Lebesgue integral) would generalize to other settings if we simply require these two conditions to hold. Thus we will build them directly into our abstract definition of a measure.

Definition 5.1.1. Let X be a set, and suppose $\mathcal{M} \subseteq \mathcal{P}(X)$ is a σ-algebra on X. A **measure** on $\mathcal{M}$ is a function $\mu : \mathcal{M} \to [0, \infty]$ satisfying the following two conditions:

1. $\mu(\emptyset) = 0$

2. if $\{E_j\}_{j=1}^{\infty}$ is a countable collection of disjoint sets in $\mathcal{M}$, then

$$\mu\left(\bigcup_{j=1}^{\infty} E_j\right) = \sum_{j=1}^{\infty} \mu(E_j).$$

The triple $(X, \mathcal{M}, \mu)$ is called a **measure space**, and the elements of $\mathcal{M}$ are called the μ-**measurable** subsets of X.

Example 5.1.2. Recall that $\mathcal{L}$ denotes the σ-algebra of Lebesgue measurable sets on $\mathbf{R}$. If we let $\mu : \mathcal{L} \to [0, \infty]$ denote the usual Lebesgue measure, then we already know that $(\mathbf{R}, \mathcal{L}, \mu)$ satisfies the definition of a measure space.

Though the Lebesgue measure on $\mathbf{R}$ is the prototypical example of a measure, it is far from the only one. There are several other commonly-studied examples, and in many cases they look quite different from what we have seen so far.

Example 5.1.3. Let X be any set, put $\mathcal{M} = \mathcal{P}(X)$, and define

$$\mu(E) = \operatorname{card}(E)$$

for all $E \in \mathcal{M}$. It is straightforward to check that $(X, \mathcal{M}, \mu)$ is a measure space. The measure μ is referred to as the **counting measure** on X.

Example 5.1.4. Let X be a set, and take $\mathcal{M} = \mathcal{P}(X)$ again. Now fix an element $x_0 \in X$ and define $\mu : \mathcal{M} \to [0, \infty]$ by

$$\mu(E) = \begin{cases} 1 & \text{if } x_0 \in E \\ 0 & \text{if } x_0 \notin E. \end{cases}$$

Then $\mu(\emptyset) = 0$ and μ is clearly countably additive. Thus $(X, \mathcal{M}, \mu)$ is a measure space. The measure μ is called the **Dirac measure** (or **point mass**) concentrated at x_0.

Example 5.1.5 (Lebesgue-Stieltjes measures on $\mathbf{R}$)**.** Let $F : \mathbf{R} \to \mathbf{R}$ be a function that is increasing and right-continuous. We can use F to define a measure μ_F on $\mathbf{R}$ as follows. Begin by defining μ on half-open intervals (more precisely, intervals that are open on the left and closed on the right) by

$$\mu_F((a, b]) = F(b) - F(a),$$

and set $\mu_F(\emptyset) = 0$. Next, extend μ_F additively to *finite* disjoint unions of half-open intervals:

$$\mu_F\left(\bigcup_{j=1}^{n} (a_j, b_j]\right) = \sum_{j=1}^{n} \big[F(b_j) - F(a_j)\big]$$

At this point, μ_F is an example of a **premeasure**, which we will define later. It is possible to extend μ_F to a measure defined on the σ-algebra generated by the half-open intervals, which is precisely the Borel σ-algebra $\mathcal{B}$. (The ability to extend μ_F requires a considerable amount of proof, and we will not deal with it just yet.) It follows that μ_F is an example of what we call a **Borel measure** on $\mathbf{R}$. It is worth noting that if we take $F(x) = x$, then we recover the usual Lebesgue measure on $\mathbf{R}$.

The function F is sometimes called the **distribution function** for the measure μ_F. This construction should be reminiscent of the idea of a cumulative distribution function in probability theory. Indeed, if we begin with a Borel probability measure μ on $\mathbf{R}$ (meaning that μ is defined on $\mathcal{B}$ and $\mu(\mathbf{R}) = 1$), we can define a function $F : \mathbf{R} \to \mathbf{R}$ by

$$F(x) = \mu((-\infty, x]).$$

It can then be checked that F is increasing and continuous from the right, and that $\mu_F = \mu$. Thus F really is the cumulative distribution function for μ.

As we mentioned earlier, many properties of Lebesgue measure carry over easily to abstract measure spaces. Moreover, the proofs often look identical to those from the simpler Lebesgue setting, so we can omit them in many cases. Indeed, this is true of some of the statements in the following theorem.

Theorem 5.1.6. *Let $(X, \mathcal{M}, \mu)$ be a measure space.*

1. (Monotonicity) *If $E_1, E_2 \in \mathcal{M}$ with $E_1 \subseteq E_2$, then $\mu(E_1) \leq \mu(E_2)$.*
2. (Subadditivity) *Let $\{E_j\}_{j=1}^{\infty}$ be a countable collection of sets in $\mathcal{M}$. Then*
$$\mu\left(\bigcup_{j=1}^{\infty} E_j\right) \leq \sum_{j=1}^{\infty} \mu(E_j).$$
3. (Continuity from below) *Suppose $E_1 \subseteq E_2 \subseteq E_3 \subseteq \cdots$ is an increasing sequence of sets in $\mathcal{M}$, and put $E = \bigcup_{j=1}^{\infty} E_j$. Then*
$$\mu(E) = \lim_{n\to\infty} \mu(E_j).$$
4. (Continuity from above) *Suppose $E_1 \supseteq E_2 \supseteq E_3 \supseteq \cdots$ is a decreasing sequence of sets in $\mathcal{M}$, and assume $\mu(E_1) < \infty$. If we let $E = \bigcap_{j=1}^{\infty} E_j$, then*
$$\mu(E) = \lim_{j\to\infty} \mu(E_j).$$

Proof. To establish (1) and (2), we cannot simply generalize the corresponding proofs for Lebesgue measure—both proofs relied on the definition of the Lebesgue outer measure in terms of intervals. However, we really only need countable additivity for in both cases. First observe that if $E_1 \subseteq E_2$, then we can write $E_2 = E_1 \cup E_2 \backslash E_1$, so

$$\mu(E_2) = \mu(E_1 \cup E_2 \backslash E_1) = \mu(E_1) + \mu(E_2 \backslash E_1)$$

since E_1 and $E_2 \backslash E_1$ are disjoint. Now $\mu(E_2 \backslash E_1) \geq 0$, so $\mu(E_2) - \mu(E_1) \geq 0$, and (1) follows.

Now suppose $\{E_j\}_{j=1}^{\infty} \subseteq \mathcal{M}$. Recall that we can replace this collection with a sequence $\{A_j\}_{j=1}^{\infty}$ of *pairwise disjoint* sets in $\mathcal{M}$ such that $\bigcup_{j=1}^{\infty} A_j = \bigcup_{j=1}^{\infty} E_j$ by taking

$$\begin{aligned}
A_1 &= E_1 \\
A_2 &= E_2 \backslash E_1 \\
A_3 &= E_3 \backslash (E_1 \cup E_2) \\
&\vdots \\
A_j &= E_j \backslash (E_1 \cup E_2 \cup \cdots \cup E_{j-1}).
\end{aligned}$$

Then the countable additivity of μ guarantees that

$$\mu\left(\bigcup_{j=1}^{\infty} E_j\right) = \mu\left(\bigcup_{j=1}^{\infty} A_j\right) = \sum_{j=1}^{\infty} \mu(A_j).$$

But for each j we have $A_j \subseteq E_j$, so $\mu(A_j) \leq \mu(E_j)$ by the monotonicity of μ. Hence

$$\mu\left(\bigcup_{j=1}^{\infty} E_j\right) = \sum_{j=1}^{\infty} \mu(A_j) \leq \sum_{j=1}^{\infty} \mu(E_j)$$

and (2) holds.

The proofs of (3) and (4) are now identical to those contained in the proof of Theorem 4.3.2, since that proof only relied on the countable additivity of Lebesgue measure. □

Just as many statements about the Lebesgue integral followed from the properties of measure described above, we will be able to use Theorem 5.1.6 to prove analogous results for integrals on other measure spaces with relative ease. We will tackle this goal in the next section.

Before moving on to integrals, we need to lay out a few more facts about measures in general. Unfortunately, some things that we take for granted when working with Lebesgue measure no longer hold when we consider abstract measure spaces. Therefore, we need to exhibit extra care in certain situations. We begin with some definitions.

Definition 5.1.7. Let $(X, \mathcal{M}, \mu)$ be a measure space.

1. We say μ is **finite** if $\mu(X) < \infty$.
2. If $\mu(X) = 1$, then μ is called a **probability measure**.
3. We say μ is **σ-finite** if X can be written as a union

$$X = \bigcup_{j=1}^{\infty} E_j,$$

where $E_j \in \mathcal{M}$ and $\mu(E_j) < \infty$ for all j.

Many reasonable measures (but not all!) are σ-finite. This is fortunate, since it is necessary to assume σ-finiteness in the statements of some of the more technical theorems that we will encounter. (In particular, Fubini's theorem has a hypothesis requiring the relevant measures to be σ-finite.)

Example 5.1.8. The usual Lebesgue measure on $\mathbf{R}$ is not finite, though it is σ-finite. Note that $\mathbf{R}$ can be expressed as a countable union of bounded intervals, say

$$\mathbf{R} = \bigcup_{n \in \mathbf{Z}} [n, n+1],$$

and each such interval necessarily has finite measure.

Example 5.1.9. Let X be a set. It is easy to check that the counting measure X is finite if and only if X is a finite set, and it is σ-finite if and only if X is countable.

Example 5.1.10. If X is a set and $x_0 \in X$, then the Dirac measure concentrated at x_0 is a finite measure.

Example 5.1.11. A Lebesgue-Stieltjes measure μ_F on $\mathbf{R}$ is always σ-finite, for the same reason that Lebesgue measure is σ-finite. Notice that μ_F is finite if and only if the distribution function F is bounded.

Another tool we have used often is the ability to ignore sets of Lebesgue measure zero, which is tantamount to the requirement that certain conditions hold *almost everywhere*.

Definition 5.1.12. Let $(X, \mathcal{M}, \mu)$ be a measure space. A set $Z \in \mathcal{M}$ is said to be **μ-null** if $\mu(Z) = 0$. A statement holds **μ-almost everywhere** (abbreviated **μ-a.e.**) if it holds everywhere except possibly on a μ-null set.

We now arrive at one unfortunate issue regarding null sets for arbitrary measures. When working with Lebesgue measure, we repeatedly used the fact that any subset of a null set has outer measure zero, and is therefore measurable. This is not always the case for other measures.

Example 5.1.13. Take $F : \mathbf{R} \to \mathbf{R}$ to be the function $F(x) = x$, and let μ_F be the associated Lebesgue-Stieltjes measure. Notice that if $(a, b]$ is a half-open interval, then

$$\mu_F((a, b]) = F(b) - F(a) = b - a.$$

It is not hard to see that μ_F behaves similarly on open and closed intervals. In other words, μ_F agrees with the usual Lebesgue measure on intervals, hence on Borel sets. However, our construction of μ_F yields a measure that is only defined on the Borel σ-algebra. If we let C denote the Cantor set, then C is μ_F-measurable with $\mu_F(C) = 0$. However, we have previously seen that the Cantor set contains non-Borel subsets, which are not μ_F-measurable. However, any such set is Lebesgue measurable.

Definition 5.1.14. A measure space $(X, \mathcal{M}, \mu)$ is said to be **complete** if $E \subseteq Z$ with $Z \in \mathcal{M}$ and $\mu(Z) = 0$ implies $E \in \mathcal{M}$.

The measure space in Example 5.1.13 is not complete, but it really only fails due to a technicality. It can clearly be extended to a complete measure—we can enrich its domain in such a way that subsets of null sets are always measurable. Fortunately, this is something that we can always do.

Theorem 5.1.15. *Suppose $(X, \mathcal{M}, \mu)$ is a measure space, let*

$$\mathcal{N} = \{Z \in \mathcal{M} : \mu(Z) = 0\},$$

and set

$$\overline{\mathcal{M}} = \{A \cup B : A \in \mathcal{M} \text{ and } B \subseteq Z \text{ for some } Z \in \mathcal{N}\}.$$

Then $\overline{\mathcal{M}}$ is a σ-algebra containing $\mathcal{M}$, and there is a unique extension of μ to a complete measure $\overline{\mu} : \overline{\mathcal{M}} \to [0, \infty]$.

Proof. It is straightforward to check that $\overline{\mathcal{M}}$ is a σ-algebra. If $\{E_j\}_{j=1}^{\infty}$ is a countable collection of sets in $\overline{\mathcal{M}}$, we write $E_j = A_j \cup B_j$ with $A_j \in \mathcal{M}$ and $B_j \subseteq Z$ for some $Z \in \mathcal{N}$. Then observe that

$$\bigcup_{j=1}^{\infty} E_j = \bigcup_{j=1}^{\infty} (A_j \cup B_j) = \bigcup_{j=1}^{\infty} A_j \cup \bigcup_{j=1}^{\infty} B_j$$

belongs to $\overline{\mathcal{M}}$, since $\mathcal{M}$ is a σ-algebra and $\mathcal{N}$ is closed under countable unions. Now suppose $E \in \overline{\mathcal{M}}$, and write $E = A \cup B$ with $A \in \mathcal{M}$ and $B \subseteq Z$ for some $Z \in \mathcal{N}$. Then

$$E^c = (E^c \cap Z) \cup (E^c \cap Z^c) = Z \backslash E \cup (E \cup Z)^c.$$

Notice that $E \cup Z \in \mathcal{M}$ and $Z \backslash E \subseteq Z$, so $E^c \in \overline{\mathcal{M}}$. Thus $\overline{\mathcal{M}}$ is a σ-algebra.

Now define $\overline{\mu} : \overline{\mathcal{M}} \to [0, \infty]$ by

$$\overline{\mu}(A \cup B) = \mu(A).$$

We need to check that this is well-defined: suppose $A_1 \cup B_1 = A_2 \cup B_2$, where $A_1, A_2 \in \mathcal{M}$ and $B_1 \subseteq Z_1$ and $B_2 \subseteq Z_2$ for some null sets Z_1, Z_2. Then

$$\mu(A_1) \le \mu(A_2) + \mu(Z_2) = \mu(A_2),$$

and similarly $\mu(A_2) \le \mu(A_1)$, so

$$\overline{\mu}(A_1 \cup B_1) = \mu(A_1) = \mu(A_2) = \overline{\mu}(A_2 \cup B_2).$$

Thus $\overline{\mu}$ is well-defined. It is now straightforward to check that $\overline{\mu}$ is a complete measure. First, if $\{A_j \cup B_j\}_{j=1}^{\infty}$ is a collection of disjoint sets in $\overline{M}$, then we have

$$\begin{aligned} \overline{\mu}\left(\bigcup_{j=1}^{\infty}(A_j \cup B_j)\right) &= \overline{\mu}\left(\bigcup_{j=1}^{\infty} A_j \cup \bigcup_{j=1}^{\infty} B_j\right) \\ &= \mu\left(\bigcup_{j=1}^{\infty} A_j\right) \\ &= \sum_{j=1}^{\infty} \mu(A_j) \\ &= \sum_{j=1}^{\infty} \overline{\mu}(A_j \cup B_j). \end{aligned}$$

Thus $\overline{\mu}$ is countably additive. Since $\overline{\mu}(\emptyset) = \mu(\emptyset) = 0$, it follows that $\overline{\mu}$ is a measure. It is immediate that the restriction of μ to $\mathcal{M}$ agrees with μ. Finally, it is easy to check that $\overline{\mu}$ is the unique extension of μ to $\overline{\mathcal{M}}$. Let $A \cup B \in \overline{\mathcal{M}}$ with $B \subseteq Z$ for some $Z \in \mathcal{N}$, and suppose ν is another measure on $\overline{\mathcal{M}}$ extending μ. Then

$$\nu(A \cup B) \le \nu(A \cup Z) = \mu(A \cup Z) = \mu(A) = \overline{\mu}(A \cup B)$$

and

$$\overline{\mu}(A \cup B) = \mu(A) = \nu(A) \le \nu(A \cup B),$$

so $\nu = \overline{\mu}$. □

Exercises for Section 5.1

Exercise 5.1.1. Let X be a set, and let $\mu : \mathcal{P}(X) \to [0, \infty]$ denote the counting measure on X. Prove that μ is a measure.

Exercise 5.1.2. Let X be a set. Fix $x_0 \in X$ and let $\mu : \mathcal{P}(X) \to [0, \infty]$ denote the Diract measure concentrated at x_0. Prove that μ is a measure.

Exercise 5.1.3. Let X be an uncountable set, and define

$$\mathcal{M} = \big\{E \subseteq X : E \text{ is countable or } E^c \text{ is countable}\big\}.$$

(a) Prove that $\mathcal{M}$ is a σ-algebra.

(b) Define $\mu : \mathcal{M} \to [0, \infty]$ by

$$\mu(E) = \begin{cases} 0 & \text{if } E \text{ is countable} \\ 1 & \text{if } E^c \text{ is countable.} \end{cases}$$

Show that μ defines a measure on $\mathcal{M}$.

Exercise 5.1.4. Suppose $F : \mathbf{R} \to \mathbf{R}$ is an increasing, right-continuous function, and let μ_F denote the associated Lebesgue-Stieltjes measure characterized by

$$\mu_F((a, b]) = F(b) - F(a)$$

for all $a, b \in \mathbf{R}$ with $a < b$.

(a) Show that for all $a, b \in \mathbf{R}$ $(a < b)$ we have

(i) $\mu_F(\{a\}) = F(a) - \lim\limits_{x \to a^-} F(x)$

(ii) $\mu_F([a, b)) = \lim\limits_{x \to b^-} F(x) - \lim\limits_{x \to a^-} F(x)$

(iii) $\mu_F([a, b]) = F(b) - \lim\limits_{x \to a^-} F(x)$

(iv) $\mu_F((a, b)) = \lim\limits_{x \to b^-} F(x) - F(a)$

(b) Suppose F is given by the **Heaviside function**:

$$F(x) = \begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x \geq 0. \end{cases}$$

What is the associated Lebesgue-Stieltjes measure μ_F?

Exercise 5.1.5. Let $(X, \mathcal{M}, \mu)$ be a measure space. We say μ is **semifinite** if given any set $E \in \mathcal{M}$ with $\mu(E) = \infty$, there exists $A \subseteq E$ such that $A \in \mathcal{M}$ and

$$0 < \mu(A) < \infty.$$

(a) Show that if μ is σ-finite, then it is semifinite. Is the converse true?

(b) Define $\mu_0 : \mathcal{M} \to [0, \infty]$ by

$$\mu_0(E) = \sup\big\{\mu(A) : A \subseteq E, A \in \mathcal{M}, \text{ and } \mu(A) < \infty\big\}.$$

Show that μ_0 is a semifinite measure.

(c) Suppose μ is already semifinite. Prove that $\mu_0 = \mu$.

5.2 Measurable Functions

Once we have a measure space $(X, \mathcal{M}, \mu)$ on which to work, the definition and properties of the Lebesgue integral carry over more or less verbatim from the basic case that we have already studied. Of course we must begin by defining what it means for a function to be *measurable*, since these are the functions that we will ultimately be allowed to integrate. The definition we will give is slightly fancier than our previous one, in part because it generalizes more easily to functions that are necessarily $\mathbf{R}$-valued. However, we will see immediately afterward that all of our old criteria for measurability still apply.

Definition 5.2.1. Suppose X and Y are sets, and let $\mathcal{M} \subseteq \mathcal{P}(X)$ and $\mathcal{N} \subseteq \mathcal{P}(Y)$ be σ-algebras. A function $f : X \to Y$ is $(\mathcal{M}, \mathcal{N})$**-measurable** if $f^{-1}(E) \in \mathcal{M}$ for all $E \in \mathcal{N}$.

This definition obviously gives us the ability to discuss measurability for functions between two measure spaces. Obviously, we are still mainly interested in the case of real-valued functions. To fit into this general definition, we need to single out a particular σ-algebra on $\mathbf{R}$ that we will use to test measurability. It turns out that the Lebesgue σ-algebra is too restrictive, so we will use the Borel σ-algebra $\mathcal{B}$.

Definition 5.2.2. Let $(X, \mathcal{M}, \mu)$ be a measure space. We say a function $f : X \to \mathbf{R}$ is $\mathcal{M}$**-measurable** if whenever $E \subseteq \mathbf{R}$ is a Borel set, $f^{-1}(E) \in \mathcal{M}$.

Despite it being different from our original definition of measurability, this new definition does align with one of our old criteria. Recall that if $E \subseteq \mathbf{R}$ is a Lebesgue measurable set, then a function $f : E \to \mathbf{R}$ is measurable if and only if $f^{-1}(U)$ is Lebesgue measurable for every open set $U \subseteq \mathbf{R}$. The same condition holds here as well, since the Borel σ-algebra is generated by the open sets in $\mathbf{R}$.

Proposition 5.2.3. *Let $(X, \mathcal{M}, \mu)$ be a measure space. A function $f : X \to \mathbf{R}$ is measurable if and only if $f^{-1}(U) \in \mathcal{M}$ for every open set $U \subseteq \mathbf{R}$.*

Proof. Suppose first that $f : X \to \mathbf{R}$ is measurable. Since any open set is Borel, it is immediate that $f^{-1}(U) \in \mathcal{M}$ for any open set $U \subseteq \mathbf{R}$.

Conversely, suppose that preimages of open sets under f belong to the σ-algebra $\mathcal{M}$. Define

$$\mathcal{N} = \big\{E \subseteq \mathbf{R} : f^{-1}(E) \in \mathcal{M}\big\}.$$

It is not hard to verify that $\mathcal{N}$ is a σ-algebra. If $\{E_j\}_{j=1}^{\infty} \subseteq \mathcal{N}$, then

$$f^{-1}\left(\bigcup_{j=1}^{\infty} E_j\right) = \bigcup_{j=1}^{\infty} f^{-1}(E_j)$$

belongs to $\mathcal{M}$, since $f^{-1}(E_j) \in \mathcal{M}$ for all j and $\mathcal{M}$ is a σ-algebra. Hence $\bigcup_{j=1}^{\infty} E_j \in \mathcal{N}$. By similar reasoning, if $E \in \mathcal{N}$, then

$$f^{-1}(E^c) = \big(f^{-1}(E)\big)^c$$

belongs to $\mathcal{M}$ as well. Hence $E^c \in \mathcal{N}$, and it follows that $\mathcal{N}$ is a σ-algebra. Moreover, $\mathcal{N}$ contains the open sets in $\mathbf{R}$ by assumption, so it must contain the σ-algebra generated by them. In other words, $\mathcal{B} \subseteq \mathcal{N}$, and it follows that f is measurable. $\square$

In Chapter 4 we had some other useful criteria for determining when a function is measurable (one of which was our actual definition of measurability). As in that case, given $f : X \to \mathbf{R}$ we will define

$$\{f > a\} = \{x \in X : f(x) > a\},$$

and the sets $\{f \geq a\}$, $\{f < a\}$, $\{f \leq a\}$, and $\{f = a\}$ are all defined similarly.

Theorem 5.2.4. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $f : X \to \mathbf{R}$. The following are equivalent:*

1. *f is $\mathcal{M}$-measurable.*
2. *For all $a \in \mathbf{R}$, the set $\{f > a\}$ belongs to $\mathcal{M}$.*
3. *For all $a \in \mathbf{R}$, the set $\{f \geq a\}$ belongs to $\mathcal{M}$.*
4. *For all $a \in \mathbf{R}$, the set $\{f < a\}$ belongs to $\mathcal{M}$.*
5. *For all $a \in \mathbf{R}$, the set $\{f \leq a\}$ belongs to $\mathcal{M}$.*

Proof. The proof is identical to the one for Theorem 4.5.2, since that proof relied only upon the fact that the Lebesgue measurable sets form a σ-algebra. $\square$

Several other facts about measurable functions from Chapter 4 translate to the general setting in a similar fashion. For example, if f is measurable then sets of the form $\{f = a\}$ are measurable for all $a \in \mathbf{R}$. In addition, we can (and will!) discuss measurability for extended real-valued functions, though doing so takes some care given our working definition of measurability. While it is possible to define a notion of Borel subsets of the extended real line, it is much easier to single out the sets where f takes infinite values.

Definition 5.2.5. Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $f : X \to [-\infty, +\infty]$ is an extended real-valued function on X. We say f is $\mathcal{M}$**-measurable** if the sets

$$\{f = +\infty\} = \big\{x \in X : f(x) = +\infty\big\}, \quad \{f = -\infty\} = \big\{x \in X : f(x) = -\infty\big\}$$

belong to $\mathcal{M}$, and $f^{-1}(E) \in \mathcal{M}$ for every Borel set $E \subseteq \mathbf{R}$.

It is also quite easy to generalize the notion of measurability to complex-valued functions. We may as well do so now, since we will gain some additional generality with little extra work. (In particular, we will ultimately have the ability to integrate complex-valued functions.)

Definition 5.2.6. Let $(X, \mathcal{M}, \mu)$ be a measure space. A function $f : X \to \mathbf{C}$ is said to be $\mathcal{M}$-**measurable** if $\operatorname{Re} f$ and $\operatorname{Im} f$ are $\mathcal{M}$-measurable.

Now we quickly run through some facts regarding both complex-valued and extended real-valued measurable functions. Most of these results generalize ones from Chapter 4, and we therefore omit many of the proofs.

Proposition 5.2.7. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $f, g : X \to \mathbf{R}$ are $\mathcal{M}$-measurable. Then*

1. *$f + g$ is $\mathcal{M}$-measurable.*
2. *αf is $\mathcal{M}$-measurable for any $\alpha \in \mathbf{R}$.*
3. *fg is $\mathcal{M}$-measurable.*

The same results hold if $f, g : X \to \mathbf{C}$ and $\alpha \in \mathbf{C}$.

Proof. In the real-valued case, the proof is identical to that of Theorem 4.5.8. The results for complex-valued functions then follow from the real case by considering real and imaginary parts. □

Proposition 5.2.8. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $(f_n)_{n=1}^{\infty}$ is a sequence of $\mathcal{M}$-measurable, extended real-valued functions on X. Then the functions*

$$\sup f_n, \quad \inf f_n, \quad \limsup f_n, \quad \liminf f_n$$

are all $\mathcal{M}$-measurable. Consequently, if the function

$$f = \lim f_n$$

exists everywhere, then f is $\mathcal{M}$-measurable.

Proof. See the proofs of Theorems 4.5.11 and 4.5.14. □

Corollary 5.2.9. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $(f_n)_{n=1}^{\infty}$ is a sequence of $\mathcal{M}$-measurable functions $f_n : X \to \mathbf{C}$. If $f = \lim f_n$ exists everywhere, then it is measurable.*

Proof. If $\lim f_n$ exists, then so do $\lim \operatorname{Re} f_n$ and $\lim \operatorname{Im} f_n$. Since $\operatorname{Re} f_n$ and $\operatorname{Im} f_n$ are $\mathcal{M}$-measurable real-valued functions for all n, it follows that $\operatorname{Re} f = \lim \operatorname{Re} f_n$ and $\operatorname{Im} f = \lim \operatorname{Im} f_n$ are $\mathcal{M}$-measurable. Thus f is $\mathcal{M}$-measurable. □

Recall that in Theorem 4.5.14, we only had to assume that $\lim f_n$ existed almost everywhere (since a function that is equal a.e. to a Lebesgue measurable function is itself measurable). In general, things may be complicated if μ is not a complete measure. This issue is rectified if we assume completeness, as the next proposition shows. The proof is left as an exercise (see Exercise 5.2.1).

Proposition 5.2.10. *Let $(X, \mathcal{M}, \mu)$ be a measure space. The following statements hold if and only μ is a complete measure.*

1. *Suppose $f = g$ μ-a.e. and f is $\mathcal{M}$-measurable. Then g is also $\mathcal{M}$-measurable.*

2. *If $(f_n)_{n=1}^{\infty}$ is a sequence of $\mathcal{M}$-measurable functions and $f_n \to f$ μ-a.e., then f is $\mathcal{M}$-measurable.*

Now we return to other straightforward consequences of Proposition 5.2.8 for measurable functions.

Corollary 5.2.11. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose f and g are μ-measurable extended real-valued functions on X. Then the functions*

$$\max\{f, g\}, \quad \min\{f, g\}$$

are $\mathcal{M}$-measurable.

Corollary 5.2.12. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose f is a μ-measurable extended real-valued function on X. Then the functions*

$$f^+ = \max\{f, 0\}, \quad f^- = -\min\{f, 0\},$$

and

$$|f| = f^+ + f^-$$

are all $\mathcal{M}$-measurable.

Finally, we finish our discussion of measurable functions with one new result that applies specifically to complex-valued functions. Recall that if $\alpha = a + ib \in \mathbf{C}$, the **modulus** of α is defined to be

$$|\alpha| = \sqrt{a^2 + b^2}.$$

Moreover, we can write

$$|\alpha|^2 = \alpha\overline{\alpha},$$

where $\overline{\alpha}$ denotes the **complex conjugate**:

$$\overline{\alpha} = a - ib.$$

Proposition 5.2.13. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $f : X \to \mathbf{C}$ is μ-measurable. Then the function $|f|$ is $\mathcal{M}$-measurable.*

Proof. Since f is $\mathcal{M}$-measurable, so are $\operatorname{Re} f$ and $\operatorname{Im} f$. Therefore, the function

$$\overline{f} = \operatorname{Re} f - i(\operatorname{Im} f)$$

is $\mathcal{M}$-measurable. Consequently, $|f|^2 = f\overline{f}$ is $\mathcal{M}$-measurable. Since the function $x \mapsto \sqrt{x}$ is continuous from $[0, \infty)$ to $\mathbf{R}$ (which implies that the preimage of any Borel set is Borel), it follows that the composite function

$$z \mapsto \sqrt{|f(z)|^2} = |f(z)|$$

is $\mathcal{M}$-measurable. □

Exercises for Section 5.2

Exercise 5.2.1. Prove Proposition 5.2.10.

5.3 Integration on Measure Spaces

It is now time to begin generalizing the construction of the Lebesgue integral to abstract measure spaces. As promised, the construction is identical to the usual one on the real line. The celebrated convergence theorems also hold in the general setting, with very similar proofs.

Recall that we originally built up the Lebesgue integral in stages. We began with characteristic functions and simple functions, for which the definition of the integral is obvious. Then we considered nonnegative functions, which we handled by approximating with simple functions. Finally, we integrated arbitrary measurable functions by writing them as linear combinations of nonnegative functions. We can mimic this procedure in our current setting, and also add one final step for dealing with complex-valued functions.

Fix a measure space $(X, \mathcal{M}, \mu)$. Recall that if $E \subseteq X$, the **characteristic function** of E is the function $\chi_E : X \to \mathbf{R}$ defined by

$$\chi_E(x) = \begin{cases} 1 & \text{if } x \in E \\ 0 & \text{if } x \notin E. \end{cases}$$

As before, it is easy to check that χ_E is $\mathcal{M}$-measurable if and only if $E \in \mathcal{M}$. More generally, we say a function $f : X \to \mathbf{R}$ is **simple** if it has finite range. Any simple function has a standard representation via characteristic functions: if f is simple with range $\{c_1, c_2, \ldots, c_n\}$, then we set define

$$E_k = \{x \in E : f(x) = c_k\}$$

for $1 \leq k \leq n$. Notice that the sets E_k are pairwise disjoint, and that the function $c_k \cdot \chi_{E_k}$ takes the value c_k on E_k, and it is zero otherwise. Consequently, we have a canonical way of representing f as a linear combination of characteristic functions:

$$f = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}.$$

Of course it is straightforward to check that f is measurable if and only if each E_k is a measurable set. Furthermore, it is obvious how to define the integral of a nonnegative simple function. If we write $f = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}$ in its standard form, then it makes sense to declare the **integral** of f with respect to μ to be

$$\int_X f \, d\mu = \sum_{k=1}^{n} c_k \cdot \mu(E_k).$$

As in Chapter 4, it is not necessary to assume the sets E_k have finite measure. It is perfectly fine if $\int_X f\,d\mu = \infty$, and we use the convention $0 \cdot \infty = 0$ to deal with the situation where $c_k = 0$ and $\mu(E_k) = \infty$.

The next step is to integrate nonnegative measurable functions by approximating with simple functions. Of course the following two results were crucial in Chapter 4, and they generalize easily to our new setting.

Theorem 5.3.1. *Suppose $f : X \to [0, +\infty]$ is a measurable extended real-valued function. Then there exists a sequence $\{f_n\}_{n=1}^{\infty}$ of measurable, simple functions such that*

$$f_n \leq f_{n+1}$$

for all n and $f_n \to f$ pointwise on X.

Proof. See the proof of Theorem 4.6.4. □

Theorem 5.3.2. *Suppose $f, g : X \to \mathbf{R}$ are nonnegative, measurable simple functions.*

1. *For all $\alpha \geq 0$, $\int_X \alpha f\,d\mu = \alpha \int_X f\,d\mu$.*
2. *$\int_X (f + g)\,d\mu = \int_X f\,d\mu + \int_X g\,d\mu$.*
3. *If $f \leq g$, then $\int_X f\,d\mu \leq \int_X g\,d\mu$.*

Proof. The proof is identical to that of Theorem 4.6.6. □

Now we can easily define the integral of an arbitrary nonnegative measurable function. If $f : X \to [0, +\infty]$ is measurable, we define the **integral** of f over X with respect to μ to be

$$\int_X f\,d\mu = \sup\left\{\int_X g\,d\mu : 0 \leq g \leq f \text{ and } g \text{ is simple}\right\}.$$

Next, we extend the integral to measurable functions $f : X \to [-\infty, +\infty]$ by writing

$$f = f^+ - f^-.$$

Recall that we need to make the stipulation that both $\int_X f^+\,d\mu$ and $\int_X f^-\,d\mu$ are finite, which is equivalent to requiring

$$\int_X |f|\,d\mu < \infty.$$

Any such function is said to be **integrable**. For arbitrary integrable functions, we define

$$\int_X f\,d\mu = \int_X f^+\,d\mu - \int_X f^-\,d\mu.$$

Finally, suppose $f : X \to \mathbf{C}$ is measurable. We declare f to be **integrable** if

$$\int_X |f|\,d\mu < \infty.$$

(It is a good exercise to check that f is integrable if and only if $\operatorname{Re} f$ and $\operatorname{Im} f$ are integrable.) We let

$$L^1(X,\mu) = \left\{ f : X \to \mathbf{C} : \int_X |f|\,d\mu < \infty \right\}$$

denote the set of all integrable, complex-valued functions on X. Given an integrable function f, we define

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

We will investigate some examples of integrals for specific measure spaces shortly. However, we first need to generalize some properties of the integral from Chapter 4. We begin with results on nonnegative functions, specifically the Monotone Convergence Theorem and its consequences.

Theorem 5.3.3 (Monotone Convergence Theorem)**.** *Suppose* $(f_n)_{n=1}^{\infty}$ *is a sequence of* $\mathcal{M}$*-measurable functions* $f_n : X \to [0,+\infty]$ *satisfying* $f_n \leq f_{n+1}$ *pointwise for all* n*. Let*

$$f = \lim_{n\to\infty} f_n = \sup_n f_n.$$

Then

$$\int_X f\,d\mu = \lim_{n\to\infty} \int_X f_n\,d\mu.$$

Proof. See the proof of Theorem 4.6.8. □

In particular, the Monotone Convergence Theorem provides a computational tool for calculating integrals, as opposed to the unwieldy supremum definition. In addition, the Monotone Convergence Theorem has many useful applications. Recall that additivity for both finite and countable collections of measurable functions follows directly from it.

Corollary 5.3.4. *Suppose $f, g : X \to [0, +\infty]$ are $\mathcal{M}$-measurable. Then*

$$\int_X (f+g)\, d\mu = \int_X f\, d\mu + \int_X g\, d\mu.$$

Proof. See the proof of Theorem 4.6.9. $\square$

Corollary 5.3.5. *Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of $\mathcal{M}$-measurable functions $f_n : X \to [0, +\infty]$, and put*

$$f = \sum_{n=1}^{\infty} f_n.$$

Then

$$\int_X f\, d\mu = \sum_{n=1}^{\infty} \int_X f_n\, d\mu.$$

Proof. See the proof of Corollary 4.6.10. $\square$

With all of these properties in hand, let's look at the integral for some of the specific examples of measures that we have encountered.

Example 5.3.6. Let $X = \mathbf{N}$ equipped with its counting measure μ, and suppose $f : \mathbf{N} \to \mathbf{R}$ is nonnegative. Notice that f is automatically measurable, since every subset of $\mathbf{N}$ is measurable. The previous corollary implies

$$\int_{\mathbf{N}} f\, d\mu = \sum_{n=1}^{\infty} \int_{\mathbf{N}} f \cdot \chi_{\{n\}}\, d\mu,$$

where

$$\int_{\mathbf{N}} f \cdot \chi_{\{n\}}\, d\mu = f(n) \cdot \mu(\{n\}) = f(n)$$

for all n. Thus

$$\int_{\mathbf{N}} f\, d\mu = \sum_{n=1}^{\infty} f(n), \tag{5.1}$$

and we have $\int_{\mathbf{N}} f\, d\mu < \infty$ precisely when f is summable (i.e., when the series on the right converges). Thus a function $f : X \to \mathbf{C}$ is integrable precisely when

$$\sum_{n=1}^{\infty} |f(n)| < \infty.$$

Once we establish the general version of the Dominated Convergence Theorem, it will be easy to check that the integral of f is again given by (5.1).

Example 5.3.7. Suppose X is a set, fix $x_0 \in X$, and let μ be the Dirac measure concentrated at x_0. Let $f : X \to [0, +\infty]$, and write

$$f = f \cdot \chi_{\{x_0\}} + f \cdot \chi_{X \setminus \{x_0\}}.$$

Notice that

$$f \cdot \chi_{\{x_0\}} = f(x_0) \cdot \chi_{\{x_0\}},$$

so

$$\int_X f \cdot \chi_{\{x_0\}} \, d\mu = f(x_0).$$

Now suppose $g : X \to [0, +\infty]$ is a simple function satisfying $g \leq f \cdot \chi_{X \setminus \{x_0\}}$, and write g in its standard form as $g = \sum_{k=1}^{n} c_k \cdot \chi_{E_k}$. Notice that if $c_k \neq 0$ for some k, then we must have $E_k \subseteq X \setminus \{x_0\}$, hence $\mu(E_k) = 0$. It follows that $\int_X g \, d\mu = 0$, whence

$$\int_X f \cdot \chi_{X \setminus \{x_0\}} \, d\mu = 0.$$

Thus

$$\int_X f \, d\mu = f(x_0)$$

for all nonnegative measurable functions, and it should be clear that the same holds for all integrable functions on X.

Now we summarize some of the other useful results that are specific to nonnegative measurable functions.

Proposition 5.3.8. *Suppose $f, g : X \to [0, +\infty]$ are measurable functions.*

1. *If $f \leq g$ μ-a.e., then $\int_X f \, d\mu \leq \int_X g \, d\mu$.*
2. *If $f = g$ μ-a.e., then $\int_X f \, d\mu = \int_X g \, d\mu$.*
3. (Chebyshev's Inequality) *For any $\alpha > 0$ we have*

$$\mu\big(\{f > \alpha\}\big) \leq \frac{1}{\alpha} \int_X f \, d\mu.$$

4. *We have $\int_X f \, d\mu = 0$ if and only if $f = 0$ μ-almost everywhere.*

Proof. See the proofs of Theorems 4.6.13, 4.6.14, and 4.6.15. □

Of course we can easily obtain a version of Fatou's lemma for the generalized integral as well.

Theorem 5.3.9 (Fatou's Lemma)**.** *Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of $\mathcal{M}$-measurable functions $f_n : X \to [0, +\infty]$. Then*

$$\int_X \liminf_{n\to\infty} f_n \, d\mu \leq \liminf_{n\to\infty} \int_X f_n \, d\mu.$$

Proof. See the proof of Theorem 4.6.18. □

In the event that the sequence $(f_n)_{n=1}^{\infty}$ converges, we get the following special case of Fatou's lemma.

Corollary 5.3.10. *Let $(f_n)_{n=1}^{\infty}$ be a sequence of $\mathcal{M}$-measurable functions $f_n : X \to [0, +\infty]$ that converges μ-almost everywhere. Then*

$$\int_X \lim_{n\to\infty} f_n \, d\mu \leq \liminf_{n\to\infty} \int_X f_n \, d\mu.$$

Now we consider some results for integrable complex-valued functions, i.e. functions in $L^1(X, \mu)$.

Proposition 5.3.11. *Let $f, g \in L^1(X, \mu)$. Then:*

1. $\left|\int_X f \, d\mu\right| \leq \int_X |f| \, d\mu$.
2. *f is μ-a.e. finite.*
3. *$f + g \in L^1(X, \mu)$, and $\int_X (f + g) \, d\mu = \int_X f \, d\mu + \int_X g \, d\mu$.*
4. *For all $\alpha \in \mathbf{C}$, $\alpha f \in L^1(X, \mu)$, and $\int_X \alpha f \, d\mu = \alpha \int_X f \, d\mu$.*
5. *If $f = g$ a.e., then $\int_X f \, d\mu = \int_X g \, d\mu$.*

Last but certainly not least is the *pièce de résistance*—the Dominated Convergence Theorem.

Theorem 5.3.12 (Dominated Convergence Theorem)**.** *Let $(f_n)_{n=1}^{\infty}$ be a sequence of functions in $L^1(X, \mu)$. Suppose f_n converges a.e. to a function f and that there is a function $g \in L^1(X, \mu)$ such that $|f_n| \leq g$ a.e. for all n. Then $f \in L^1(X, \mu)$ and*

$$\lim_{n\to\infty} \int_X f_n \, d\mu = \int_X f \, d\mu.$$

Proof. See the proof of Theorem 4.7.6. $\square$

As a consequence of the Dominated Convergence Theorem, we have the following useful corollary.

Corollary 5.3.13. *Suppose $(f_n)_{n=1}^\infty$ is a sequence in $L^1(X,\mu)$ satisfying $\sum_{n=1}^\infty \int_X |f_n|\,d\mu < \infty$. Then $\sum_{n=1}^\infty f_n$ converges a.e. to a function $f \in L^1(X,\mu)$, and*

$$\int_X \sum_{n=1}^\infty f_n\,d\mu = \sum_{n=1}^\infty \int_X f_n\,d\mu.$$

Proof. See the proof of Theorem 4.7.7. $\square$

Example 5.3.14. Let μ denote the counting measure on $\mathbf{N}$. Given an integrable function $f : \mathbf{N} \to \mathbf{C}$, define

$$f_n = f \cdot \chi_{\{n\}}$$

for each $n \in \mathbf{N}$. Observe that

$$f = \sum_{n=1}^\infty f_n\,d\mu$$

and for each n we have

$$\int_{\mathbf{N}} f_n\,d\mu = f(n)$$

so

$$\int_{\mathbf{N}} f\,d\mu = \int_{\mathbf{N}} \sum_{n=1}^\infty f_n\,d\mu = \sum_{n=1}^\infty \int_{\mathbf{N}} f_n\,d\mu = \sum_{n=1}^\infty f(n)$$

by Corollary 5.3.13.

Exercises for Section 5.2

Exercise 5.3.1. Let $(X, \mathcal{M}, \mu)$ be a measure space, $f : X \to [0,\infty]$ a measurable function, and define a set function $\lambda : \mathcal{M} \to [0,\infty]$ by

$$\lambda(E) = \int_E f\,d\mu = \int_X f \cdot \chi_E\,d\mu.$$

Prove that λ is a measure on $\mathcal{M}$, and that

$$\int_X g\,d\lambda = \int_X fg\,d\mu$$

for any measurable function $g : X \to [0,\infty]$.

5.4 Modes of Convergence

Throughout our studies so far, we have seen several ways in which a sequence of functions on a measure space $(X, \mathcal{M}, \mu)$ can be deemed convergent. If we have a sequence $(f_n)_{n=1}^{\infty}$ of functions $f_n : X \to \mathbf{C}$, we say that

- $f_n \to f$ **pointwise** if $f_n(x) \to f(x)$ in $\mathbf{C}$ for all $x \in X$,
- $f_n \to f$ **pointwise almost everywhere** (or **pointwise a.e.**) if $f_n(x) \to f(x)$ μ-a.e. on X, or
- $f_n \to f$ **uniformly** if given any $\varepsilon > 0$, there exists N such that for all $n \geq N$ we have $|f_n(x) - f(x)| < \varepsilon$ for all $x \in X$.

Of course uniform convergence implies pointwise convergence, which in turn guarantees pointwise a.e. convergence. The converse of neither implication is true in general.

There are actually several other ways that one can talk about convergence for sequences of measurable functions, each with its own uses. We will focus on three such notions here, along with their relationships to one another and our previous modes of convergence. Throughout the discussion, we let $(X, \mathcal{M}, \mu)$ denote a measure space.

Definition 5.4.1. Let $(f_n)_{n=1}^{\infty}$ be a sequence of $\mathcal{M}$-measurable functions on X. We say that f_n **converges in** L^1 to a function f if

$$\int_X |f_n - f| \, d\mu \to 0$$

as $n \to \infty$.

In the context of probability theory, convergence in L^1 is often called *convergence in mean.* It will be important later on, since we will eventually define a metric on $L^1(X, \mu)$ via

$$d(f, g) = \int_X |f - g| \, d\mu,$$

and convergence in L^1 will correspond precisely to convergence with respect to this metric.[1]

[1] The function defined above is not actually a metric, but a *pseudometric*—it is not positive definite, since $d(f, g) = 0$ whenever $f = g$ μ-almost everywhere. We will need to correct for this by instead considering equivalence classes of measurable functions that are μ-a.e. equal.

Definition 5.4.2. A sequence $(f_n)_{n=1}^\infty$ of $\mathcal{M}$-measurable functions on X is said to **converge in measure** if for all $\varepsilon > 0$,

$$\mu\big(\{x \in X : |f_n(x) - f(x)| \geq \varepsilon\}\big) \to 0$$

as $n \to \infty$.

Convergence in measure is meant to generalize the notion of *convergence in probability* from probability theory.

It is natural to ask how convergence in L^1 and convergence in measure relate to uniform, pointwise, and pointwise a.e. convergence, as well as how they relate to one another. Unfortunately, the answer is complicated. We begin with a positive result.

Theorem 5.4.3. *Let $(f_n)_{n=1}^\infty$ be a sequence of measurable functions on X. If f_n converges to a function f in L^1, then $f_n \to f$ in measure.*

Proof. Suppose $f_n \to f$ in L^1, and let $\varepsilon > 0$ be given. For each n, define

$$E_n = \{x \in X : |f_n(x) - f(x)| \geq \varepsilon\}.$$

Observe that

$$\varepsilon \cdot \mu(E_n) \leq \int_X |f_n - f| \cdot \chi_{E_n}\, d\mu \leq \int_X |f_n - f|\, d\mu,$$

so

$$\mu(E_n) \leq \frac{1}{\varepsilon} \int_X |f_n - f|\, d\mu$$

for all n. The right hand side tends to zero as $n \to \infty$ since $f_n \to f$ in L^1. Hence $\mu(E_n) \to 0$, so $f_n \to f$ in measure. □

The converse is not true in general, though it does hold if we assume that the f_n are all dominated by an integrable function g. (See Exercise 5.4.1.)

Example 5.4.4. Convergence in measure does not imply convergence in L^1. Let $(\mathbf{R}, \mathcal{L}, \mu)$ be the usual Lebesgue measure space on $\mathbf{R}$, and define

$$f_n = \frac{1}{n}\chi_{[0,n]}$$

for all $n \in \mathbf{N}$. It is straightforward to prove that $f_n \to 0$ in measure. Let $\varepsilon > 0$ be given. For any $n > \frac{1}{\varepsilon}$ we have

$$|f_n(x)| \leq \frac{1}{n} < \varepsilon$$

for all $x \in \mathbf{R}$. Thus
$$\mu\big(\{x \in X : |f_n(x)| \geq \varepsilon\}\big) = 0$$
for sufficiently large n. Hence $f_n \to 0$ in measure. However, the sequence does not converge in L^1 since
$$\int_{\mathbf{R}} |f_n| \, d\mu = \frac{1}{n} \cdot n = 1$$
for all n.

Remark 5.4.5. The fact that the sequence in Example 5.4.4 converges in measure also follows from the fact that uniform convergence guarantees convergence in measure (Exercise 5.4.2). Consequently, Example 5.4.4 also shows that uniform convergence does not imply convergence in L^1.

While uniform convergence implies convergence in measure, pointwise a.e. convergence does not suffice.

Example 5.4.6. Let $(\mathbf{R}, \mathcal{M}, \mu)$ be the Lebesgue measure space, and define
$$f_n = \chi_{[n,n+1]}$$
for each $n \in \mathbf{N}$. Notice that $f_n \to 0$ pointwise, but
$$\mu\big(\{x \in X : |f_n(x)| \geq \tfrac{1}{2}\}\big) = 1$$
for all n. Thus $f_n \not\to 0$ in measure.

Finally, we show that neither convergence in measure nor convergence in L^1 need imply pointwise a.e. convergence.

Example 5.4.7. Again, we consider the Lebesgue measure space $(\mathbf{R}, \mathcal{M}, \mu)$. Define $f_n : \mathbf{R} \to \mathbf{R}$ by
$$f_n = \chi_{\left[\frac{n-2^k}{2^k}, \frac{n-2^k+1}{2^k}\right]},$$
where $k \geq 0$ satisfies $2^k \leq n < 2^{k+1}$. (Terry Tao calls $(f_n)_{n=1}^\infty$ the *typewriter sequence*, since the intervals shift from left to right across $[0,1]$, returning to the left end after each traversal of the unit interval.) Notice that the measures of the intervals are shrinking as well, so it follows that
$$\int_{\mathbf{R}} |f_n| \, d\mu \to 0$$
as $n \to \infty$. Thus $f_n \to 0$ in L^1, hence in measure. However, this sequence does not converge to 0 pointwise a.e.—given any $x \in [0,1]$ and any $n_0 \in \mathbf{N}$, there always exists $n \geq n_0$ such that $f_n(x) = 1$.

Though convergence in measure does not imply pointwise a.e. convergence, there is always a subsequence that converges almost everywhere.

Theorem 5.4.8. *Let* $(f_n)_{n=1}^{\infty}$ *be a sequence of measurable functions on* X*, and suppose* f_n *converges in measure to a function* f*. Then there is a subsequence* $(f_{n_k})_{k=1}^{\infty}$ *that converges pointwise a.e. to* f*.*

Proof. Since $f_n \to f$ in measure, for each $k \in \mathbf{N}$ there exists n_k such that

$$\mu\big(\big\{x \in X : |f_n(x) - f(x)| \geq \tfrac{1}{k}\big\}\big) < \tfrac{1}{2^k}$$

for all $n \geq n_k$. For each $k \in \mathbf{N}$, define

$$E_k = \big\{x \in X : |f_{n_k}(x) - f(x)| \geq \tfrac{1}{k}\big\}.$$

For each $m \in \mathbf{N}$, set

$$A_m = \bigcup_{k=m}^{\infty} E_k.$$

Then

$$\mu(A_m) \leq \sum_{k=m}^{\infty} \mu(E_k) \leq \sum_{k=m}^{\infty} \frac{1}{2^k} = \frac{1}{2^{m-1}}.$$

Furthermore, if $k \geq m$, we have

$$|f_{n_k}(x) - f(x)| < \frac{1}{k}$$

for all $x \notin A_m$. Therefore, the sequence $(f_{n_k}(x))_{k=1}^{\infty}$ converges to $f(x)$ whenever $x \in A_m^c$. Now set $A = \bigcap_{m=1}^{\infty} A_m$. Then

$$\mu(A) = \lim_{m\to\infty} \mu(A_m) = \lim_{m\to\infty} \frac{1}{2^{m-1}} = 0$$

by continuity of measure. For all $x \in A^c$, the sequence $(f_{n_k}(x))_{k=1}^{\infty}$ converges to $f(x)$, so $f_{n_k} \to f$ pointwise almost everywhere. $\square$

Corollary 5.4.9. *If* $f_n \to f$ *in* L^1*, then there is a subsequence* $(f_{n_k})_{k=1}^{\infty}$ *that converges pointwise a.e. to* f*.*

Proof. This follows from the previous theorem and the fact that L^1 convergence implies convergence in measure. $\square$

There is now one last type of convergence that we will discuss. Along with it comes a major theorem in analysis, which we have already encountered in a special case.

Definition 5.4.10. Let $(f_n)_{n=1}^{\infty}$ be a sequence of measurable functions on X. We say f_n converges **almost uniformly** to a function f if given any $\varepsilon > 0$ there exists a set $E \subseteq X$ such that $\mu(E) < \varepsilon$ and $f_n \to f$ uniformly on E^c.

Clearly uniform convergence implies almost uniform convergence, which in turn implies pointwise a.e. convergence. Surprisingly enough, a.e. convergence is enough to guarantee almost uniform convergence when μ is a finite measure. We already saw a preliminary version of this result, known as Egorov's theorem, in Chapter 4. The proof is more or less the same, so we will omit it.

Theorem 5.4.11 (Egorov)**.** *Let $(X, \mathcal{M}, \mu)$ be a measure space with $\mu(E) < \infty$. Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of measurable functions $f_n : X \to \mathbf{C}$, and assume f_n converges pointwise a.e. to a function $f : X \to \mathbf{C}$. Then $f_n \to f$ almost uniformly.*

It is a straightforward exercise to show that almost uniform convergence implies convergence in measure. Consequently, pointwise a.e. convergence implies convergence in measure on *finite* measure spaces.

We end this section with a diagram illustrating the various implications that we have just explored. A solid arrow indicates the implication always holds, while a dashed arrow requires the assumption that $\mu(X) < \infty$.

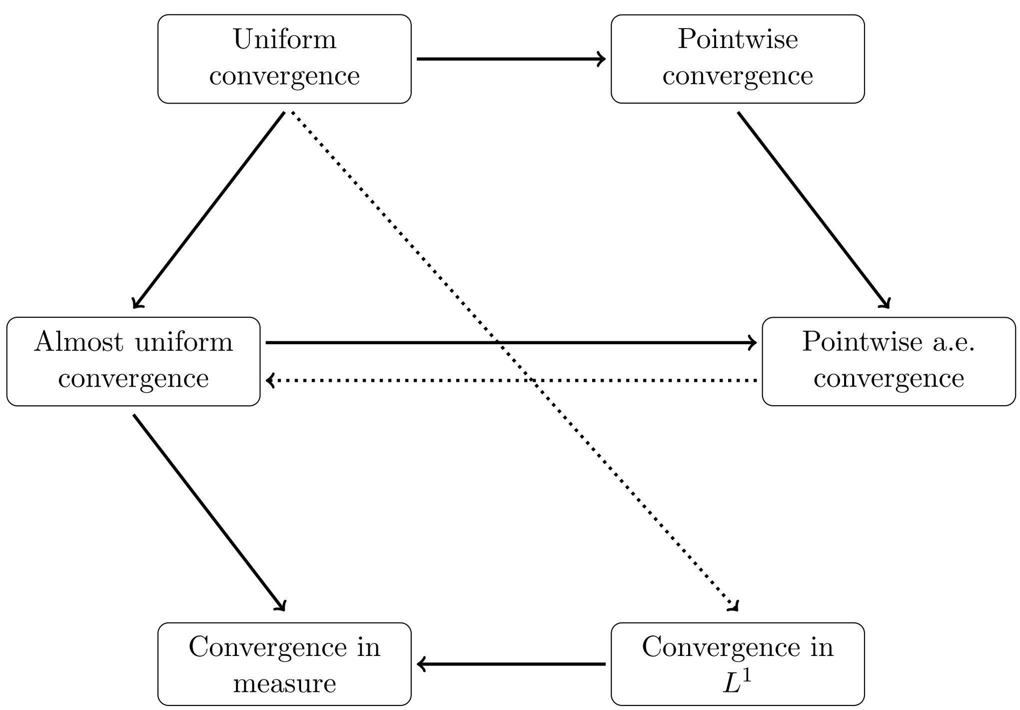

It is worth noting that some implications become true if one assumes the functions in the sequence are all dominated by an integrable function.

Exercises for Section 5.4

Exercise 5.4.1. Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of measurable functions on a measure space $(X, \mathcal{M}, \mu)$, and that there exists a function $g \in L^1(X, \mu)$ such that $|f_n| \leq g$ for all n. Show that if $f_n \to f$ in measure, then $f_n \to f$ in L^1.

Exercise 5.4.2. Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of measurable functions on a measure space $(X, \mathcal{M}, \mu)$. Show that if $f_n \to f$ almost uniformly, then $f_n \to f$ in measure. In particular, uniform convergence implies convergence in measure.

Exercise 5.4.3. Suppose $(f_n)_{n=1}^{\infty}$ is a sequence of measurable functions on a *finite* measure space $(X, \mathcal{M}, \mu)$. Show that if $f_n \to f$ uniformly, then $f_n \to f$ in L^1.

Exercise 5.4.4. Let X be a set, and for each $x \in X$ let δ_x denote the Dirac measure concentrated at x. Prove that a sequence of functions $f_n : X \to \mathbf{C}$ converges pointwise to a function f if and only if $f_n \to f$ in measure with respect to δ_x for all $x \in X$.

5.5 Product Measures and Fubini's Theorem

We now arrive at a result that you have already used, and likely taken for granted, in multivariable calculus. Given a sufficiently nice (usually meaning continuous) function $f : \mathbf{R}^2 \to \mathbf{R}$, one can compute the integral of f via iterated integration:

$$\int_{\mathbf{R}^2} f(x,y)\, d(x,y) = \int_{\mathbf{R}} \left[\int_{\mathbf{R}} f(x,y)\, dx \right] dy = \int_{\mathbf{R}} \left[\int_{\mathbf{R}} f(x,y)\, dy \right] dx.$$

Furthermore, the order of integration does not matter, as we have intimated above. Geometrically, this result amounts to the fact that we can compute the volume under the graph of f by first computing the areas of several cross-sections via integration, then adding up those areas with another integral.

This section is devoted to adapting the above discussion to general measure spaces, and then proving two similar theorems—the theorems of Fubini and Tonelli—that allow one to compute double integrals via iterated integration. Therefore, we begin with two measure spaces $(X, \mathcal{M}, \mu)$ and $(Y, \mathcal{N}, \nu)$, and we attempt to determine how to integrate functions over the Cartesian product $X \times Y$. Obviously we first need to construct an appropriate measure $\mu \times \nu$ on $X \times Y$. We first single out the subsets of $X \times Y$ for which the definition of measure should be obvious.

Definition 5.5.1. Given two sets $A \in \mathcal{M}$ and $B \in \mathcal{N}$, we call their product $A \times B \subseteq X \times Y$ a **measurable rectangle**.

For a measurable rectangle $A \times B \subseteq X \times Y$, it seems natural to define the measure of $A \times B$ by

$$(\mu \times \nu)(A \times B) = \mu(A)\nu(B).$$

Of course this map does not define a measure on $X \times Y$. The domain of our product measure needs to be a σ-algebra of sets on $X \times Y$, and the measurable rectangles do not form a σ-algebra. The obvious thing to do would be to consider the σ-algebra generated by the measurable rectangles, which we denote by $\mathcal{M} \otimes \mathcal{N}$. However, it is still not clear how to extend the definition of $\mu \times \nu$ to the full σ-algebra $\mathcal{M} \otimes \mathcal{N}$.

We will now outline a general method for constructing measures, which will allow us to properly construct our product measure space. We begin with a definition.

Definition 5.5.2. Let X be a set. We say $\mathcal{A} \subseteq \mathcal{P}(X)$ is an **algebra** of sets on X if $\mathcal{A}$ is closed under complements and finite unions.

A common method of constructing measures is to first define a set function with certain desirable properties on an algebra, and then extend the domain to an appropriate σ-algebra. This preliminary set function is called a *premeasure*.

Definition 5.5.3. Let X be a set and suppose $\mathcal{A} \subseteq \mathcal{P}(X)$ is an algebra. A function $\mu_0 : \mathcal{A} \to [0, \infty]$ is called a **premeasure** if

1. $\mu_0(\emptyset) = 0$
2. If $\{E_j\}_{j=1}^{\infty}$ is a collection of pairwise disjoint sets such that $\bigcup_{j=1}^{\infty} E_j \in \mathcal{A}$, then
$$\mu_0\left(\bigcup_{j=1}^{\infty} E_n\right) = \sum_{j=1}^{\infty} \mu_0(E_j).$$

Notice that the second condition in the above definition implies that a premeasure is finitely additive on families of pairwise disjoint sets.

Example 5.5.4. Let $\mathcal{A}$ be the algebra on $\mathbf{R}$ generated by all half-open intervals of the form $(a, b]$ for $a, b \in \mathbf{R}$. It is possible to show that every element of $\mathcal{A}$ can be written as a finite disjoint union of half-open intervals and rays of the form (a, ∞) or $(-\infty, b]$.

Now suppose $F : \mathbf{R} \to \mathbf{R}$ is increasing and right-continuous, and define a set function $\mu_0 : \mathcal{A} \to [0, \infty]$ as follows: put $\mu_0(\emptyset) = 0$, and define

$$\mu_0((a, b]) = F(b) - F(a)$$

for all $a, b \in \mathbf{R}$. Then declare $\mu_0((a,\infty)) = \infty$ and $\mu_0((-\infty,b]) = \infty$, and extend the definition of μ_0 additively to elements of $\mathcal{A}$. In particular, if $a_j, b_j \in \mathbf{R}$ for $1 \leq j \leq n$, then

$$\mu_0\left(\bigcup_{j=1}^{n}(a_j, b_j]\right) = \sum_{j=1}^{n}\big[F(b_j) - F(a_j)\big].$$

One can then show that μ_0 is a premeasure on $\mathcal{A}$. (The proof is not difficult, though it is quite technical due to the fact that there are many different ways of writing an element of $\mathcal{A}$ as a disjoint union of half-open intervals and rays.) Indeed, this premeasure is the precursor to the Lebesgue-Stieltjes measure associated to F, as we have already seen.

Once one has a premeasure defined on an algebra, the next step is to define an associated *outer measure* for subsets of X.

Definition 5.5.5. Let X be a set. An **outer measure** on X is a set function $\mu^* : \mathcal{P}(X) \to [0,\infty]$ satisfying the following conditions:

1. $\mu^*(\emptyset) = 0$.
2. (Monotonicity) If $E_1 \subseteq E_2$, then $\mu^*(E_1) \leq \mu^*(E_2)$.
3. (Subadditivity) If $\{E_j\}_{j=1}^{\infty}$ is a countable collection of subsets of X, then

$$\mu^*\left(\bigcup_{j=1}^{\infty} E_j\right) \leq \sum_{j=1}^{\infty}\mu^*(E_j).$$

We can use a premeasure on an algebra $\mathcal{A} \subseteq \mathcal{P}(X)$ to define an outer measure as follows. If $E \subseteq X$, we set

$$\mu^*(E) = \inf\left\{\sum_{j=1}^{\infty}\mu_0(A_j) : A_j \in \mathcal{A} \text{ for all } j \text{ and } E \subseteq \bigcup_{j=1}^{\infty} A_j\right\}.$$

Note the similarity to the construction of Lebesgue outer measure on $\mathbf{R}$. (In that situation, the algebra $\mathcal{A}$ was the one generated by the closed, bounded intervals.) Indeed, the proof that μ^* has the required properties is quite similar to the one for Lebesgue outer measure.

In a similar fashion to the Lebesgue case, we declare a set $E \subseteq X$ to be μ^*-**measurable** if

$$\mu^*(A) = \mu^*(A \cap E) + \mu^*(A \cap E^c)$$

for all $A \subseteq X$. Based on our experience with Lebesgue measure, one might expect that the μ^*-measurable sets form a σ-algebra, and that μ^* restricts to a measure on those sets. This is the content of Carathéodory's theorem, the proof of which is nearly identical to the one for Lebesgue measure.

Theorem 5.5.6 (Carathéodory)**.** *Let X be a set, and suppose μ^* is an outer measure on X. Let $\mathcal{M}$ denote the set of all μ^*-measurable subsets of X. Then $\mathcal{M}$ is a σ-algebra, and μ^* restricts to a complete measure $\mu : \mathcal{M} \to [0, \infty]$.*

Proof. The proof follows the same progression that we used in Chapter 4 to show that both the Lebesgue measurable sets form a σ-algebra and Lebesgue measure is countably additive.

- Show $\mathcal{M}$ is an algebra.
- Show μ^* is finitely additive on $\mathcal{M}$.
- Use finite additivity to argue that $\mathcal{M}$ is a σ-algebra.
- Show μ^* is countably additive on $\mathcal{M}$.
- Show that the measure $\mu = \mu^*|_{\mathcal{M}}$ is complete.

As in the Lebesgue case, it is obvious that the condition for μ^*-measurability is symmetric in E and E^c. Thus $\mathcal{M}$ is closed under complements. Now suppose $E_1, E_2 \in \mathcal{M}$ and let $A \subseteq X$. First observe that

$$\mu^*(A \cap (E_1 \cup E_2)) \leq \mu^*(A \cap (E_1 \cap E_2)) + \mu^*(A \cap (E_1 \cap E_2^c)) + \mu^*(A \cap (E_1^c \cap E_2))$$

by countable subadditivity. Since E_1, is μ^*-measurable,

$$\mu^*(A \cap E_1) = \mu^*(A \cap E_1 \cap E_2) + \mu^*(A \cap E_1 \cap E_2^c)$$

so the previous inequality simplifies to

$$\mu^*(A \cap (E_1 \cup E_2)) \leq \mu^*(A \cap E_1) + \mu^*(A \cap (E_1^c \cap E_2)).$$

Furthermore,

$$\mu^*(A \cap E_1^c) = \mu^*(A \cap E_1^c \cap E_2) + \mu^*(A \cap E_1^c \cap E_2^c),$$

so

$$\begin{aligned} \mu^*(A \cap (E_1 \cup E_2)) + \mu^*(A \cap (E_1 \cup E_2)^c) &\leq \mu^*(A \cap E_1) + \mu^*(A \cap E_1^c) \\ &= \mu^*(A), \end{aligned}$$

where we have again used the fact that E_1 is measurable. Since we always have

$$\mu^*(A) \leq \mu^*(A \cap (E_1 \cup E_2)) + \mu^*(A \cap (E_1 \cup E_2)^c)$$

by countable subadditivity, it follows that $E_1 \cup E_2 \in \mathcal{M}$. One can then argue via induction that $\mathcal{M}$ is closed under finite unions, hence it is an algebra.

Suppose $E_1, E_2 \in \mathcal{M}$ are disjoint. Since E_1 is μ^*-measurable and $E_1 \cap E_2 = \emptyset$, we have

$$\mu^*(E_1 \cup E_2) = \mu^*((E_1 \cup E_2) \cap E_1) + \mu^*((E_1 \cup E_2) \cap E_1^c) = \mu^*(E_1) + \mu^*(E_2).$$

It now follows by induction that μ^* is finitely additive on $\mathcal{M}$.

Now suppose $\{E_j\}_{j=1}^\infty$ is a countable collection of pairwise disjoint sets in $\mathcal{M}$, and put $E = \bigcup_{j=1}^\infty E_j$. Let $A \subseteq X$. We first claim that for each n we have

$$\mu^*\left(A \cap \bigcup_{j=1}^n E_j\right) = \sum_{j=1}^n \mu^*(A \cap E_j).$$

This claim is obvious for $n = 1$. Assuming it holds for $n-1$, we have

$$\begin{aligned}
\mu^*\left(A \cap \bigcup_{j=1}^n E_j\right) &= \mu^*\left(\left(A \cap \bigcup_{j=1}^n E_j\right) \cap E_n\right) + \mu^*\left(\left(A \cap \bigcup_{j=1}^n E_j\right) \cap E_n^c\right) \\
&= \mu^*(A \cap E_n) + \mu^*\left(A \cap \bigcup_{j=1}^{n-1} E_j\right) \\
&= \mu^*(A \cap E_n) + \sum_{j=1}^{n-1} \mu^*(A \cap E_j) \\
&= \sum_{j=1}^n \mu^*(A \cap E_j),
\end{aligned}$$

and the claim is proven. Now observe that

$$\begin{aligned}
\mu^*(A) &= \mu^*\left(A \cap \bigcup_{j=1}^n E_j\right) + \mu^*\left(A \cap \left(\bigcup_{j=1}^n E_j\right)^c\right) \\
&= \sum_{j=1}^n \mu^*(A \cap E_j) + \mu^*\left(A \cap \left(\bigcup_{j=1}^n E_j\right)^c\right) \\
&\geq \sum_{j=1}^n \mu^*(A \cap E_j) + \mu^*(A \cap E^c)
\end{aligned}$$

for all n. Therefore, letting n tend to ∞, we get

$$\begin{aligned}\mu^*(A) &\geq \sum_{j=1}^{\infty} \mu^*(A \cap E_j) + \mu^*(A \cap E^c) \\ &\geq \mu^*(A \cap E) + \mu^*(A \cap E^c) \\ &\geq \mu^*(A),\end{aligned}$$

so

$$\mu^*(A) = \mu^*(A \cap E) + \mu^*(A \cap E^c),$$

whence E is μ^*-measurable. Therefore, $\mathcal{M}$ is a σ-algebra. Letting $A = E$ above, we obtain

$$\mu^*(E) = \sum_{j=1}^{\infty} \mu^*(E_j),$$

so μ^* is countably additive on $\mathcal{M}$.

Everything we have done so far shows that μ^* restricts to a measure on $\mathcal{M}$. It remains to see that $\mu = \mu^*|_{\mathcal{M}}$ is complete. It suffices to show that if $E \subseteq X$ with $\mu^*(E) = 0$, then $E \in \mathcal{M}$. Well, if $\mu^*(E) = 0$, then for all $A \subseteq X$ we have

$$\mu^*(A) \leq \mu^*(A \cap E) + \mu^*(A \cap E^c) \leq \mu^*(E) + \mu^*(A) = \mu^*(A),$$

so

$$\mu^*(A) = \mu^*(A \cap E) + \mu^*(A \cap E^c),$$

which shows that E is μ^*-measurable. It then follows from monotonicity that any subset of a μ-null set is measurable, hence μ is complete. $\square$

By applying Carathédory's theorem to the outer measure induced by a premeasure, we obtain the following result almost immediately.

Theorem 5.5.7 (Carathéodory-Hahn)**.** *Let X be a set, $\mathcal{A} \subseteq \mathcal{P}(X)$ an algebra, and $\mu_0 : \mathcal{A} \to [0, \infty]$ a premeasure on $\mathcal{A}$. Let $\mathcal{M}$ denote the σ-algebra generated by $\mathcal{A}$. There exists a measure $\mu : \mathcal{M} \to [0, \infty]$ such that $\mu|_{\mathcal{A}} = \mu_0$. If μ_0 is σ-finite, then so is μ, and μ is the unique extension of μ_0 to $\mathcal{M}$.*

Proof. The existence of μ is more or less guaranteed by Carathéodory's theorem—simply build the outer measure μ^* associated to μ_0, and then restrict the resulting measure μ to $\mathcal{M}$. However, we need to know this last step makes sense. That is, we need to check that every set in $\mathcal{A}$ (and hence every set in $\mathcal{M}$) is μ^*-measurable.

Suppose $E \in \mathcal{A}$ and let $A \subseteq X$. Given $\varepsilon > 0$, there exists a collection of sets $\{A_j\}_{j=1}^{\infty} \subseteq \mathcal{A}$ such that $A \subseteq \bigcup_{j=1}^{\infty} A_j$ and

$$\sum_{j=1}^{\infty} \mu_0(A_j) \leq \mu^*(A) + \varepsilon$$

by the definition of μ^*. Since μ_0 is finitely additive on $\mathcal{A}$, we have

$$\mu_0(A_j) = \mu_0(A_j \cap E) + \mu_0(A_j \cap E^c)$$

for all j, so

$$\begin{aligned} \mu^*(A) + \varepsilon &\geq \sum_{j=1}^{\infty} \mu_0(A_j \cap E) + \sum_{j=1}^{\infty} \mu_0(A_j \cap E^c) \\ &\geq \mu^*(A \cap E) + \mu^*(A \cap E^c), \end{aligned}$$

where the last inequality comes from the definition of μ^* and the fact that the families $\{A_j \cap E\}_{j=1}^{\infty}$ and $\{A_j \cap E^c\}_{j=1}^{\infty}$ are coverings of $A \cap E$ and $A \cap E^c$, respectively, by elements of $\mathcal{A}$. This inequality holds for all $\varepsilon > 0$, so

$$\mu^*(A) \geq \mu^*(A \cap E) + \mu^*(A \cap E^c).$$

Hence E is μ^*-measurable.

Carathéodory's theorem now guarantees that $\mu = \mu^*|_{\mathcal{M}}$ is a measure defined on $\mathcal{M}$. It is not too difficult to show that $\mu^*|_{\mathcal{A}} = \mu_0$ (see Exercise 5.5.1), which implies that $\mu|_{\mathcal{A}} = \mu_0$.

Since $\mu|_{\mathcal{A}} = \mu_0$, it is clear that μ is σ-finite whenever μ_0 is. We leave the proof of uniqueness in this case as an exercise. □

Remark 5.5.8. Suppose $\mathcal{A}$ is an algebra on a set X, μ_0 is a premeasure on $\mathcal{A}$, and μ^* is the associated outer measure. It is possible (in fact, likely) that the σ-algebra $\mathcal{M}$ generated by $\mathcal{A}$ is not the full σ-algebra of μ^*-measurable sets. In other words, the measure μ afforded by the Carathéodory-Hahn theorem may turn out to not be a complete measure. This statement might seem contradictory, since Carathéodory's theorem guarantees that every outer measure induces a *complete* measure. However, the issue here is not the measure itself, but the σ-algebra upon which it is defined.

As an example, let $\mathcal{A}$ be the algebra on $\mathbf{R}$ generated by the half-open intervals, and let μ_0 be the Lebesgue-Stieltjes premeasure defined in Example 5.5.4. Then the Lebesgue-Stieltjes measure that we obtain by applying the Carathéodory-Hahn theorem is defined on the σ-algebra generated by $\mathcal{A}$, which is precisely the Borel σ-algebra $\mathcal{B}_{\mathbf{R}}$. However, we know the special case of Lebesgue measure that $\mathcal{B}_{\mathbf{R}}$ is not generally the full σ-algebra of μ^*-measurable sets. We will soon discuss another specific example that exhibits this phenomenon in the context of product measures.

We know that any measure can be extended to a complete measure, which seems to make this discussion moot. Indeed, this issue is easy to fix when μ_0 is σ-finite. Suppose μ_0 is σ-finite, and let $\mathcal{M}^*$ denote the σ-algebra consisting of all μ^*-measurable sets. Then $\mu^*|_{\mathcal{M}^*} = \overline{\mu}$, where $\overline{\mu}$ is the completion of μ. In other words, $\overline{\mathcal{M}} = \mathcal{M}^*$—when we enlarge the σ-algebra $\mathcal{M}$ in order to complete μ, we simply throw in all of the remaining μ^*-measurable sets. If μ_0 is *not* σ-finite, then $\overline{\mathcal{M}} \neq \mathcal{M}^*$, and we need to enlarge the σ-algebra even further. In this case, $\mu^*|_{\mathcal{M}^*}$ is not $\overline{\mu}$, but an extension known as the **saturation** of $\overline{\mu}$.

Now we return to the problem of constructing product measures. Let $(X, \mathcal{M}, \mu)$ and $(Y, \mathcal{N}, \nu)$ be measure spaces, and let $\mathcal{A}$ denote the collection of all subsets of $X \times Y$ that can be written as finite disjoint unions of measurable rectangles. The set $\mathcal{A}$ still does not form a σ-algebra, though it does form an algebra.

Proposition 5.5.9. *The collection $\mathcal{A}$ of finite disjoint unions of measurable rectangles on $X \times Y$ is an algebra.*

Proof. We begin with two initial observations about measurable rectangles. Notice first that if $A \times B$ is a measurable rectangle, then

$$(A \times B)^c = (A^c \times Y) \cup (A \times B^c),$$

which is a disjoint union of measurable rectangles, hence a member of $\mathcal{A}$. Now let $A_1 \times B_1$ and $A_2 \times B_2$ be measurable rectangles. Notice that

$$(A_1 \times B_1) \cap (A_2 \times B_2) = (A_1 \cap A_2) \times (B_1 \cap B_2),$$

and an inductive proof then shows that any finite intersection of measurable rectangles is again a measurable rectangle.

To verify that $\mathcal{A}$ is an algebra, it is enough to show that any finite union of measurable rectangles belongs to $\mathcal{A}$, and that $\mathcal{A}$ is closed under complements. To this end, let $A_1 \times B_1$ and $A_2 \times B_2$ be measurable rectangles. Observe that

$$\begin{aligned}(A_1 \times B_1) \cap (A_2 \times B_2)^c &= (A_1 \times B_1) \cap \big[(A_2^c \times Y) \cup (A_2 \times B_2^c)\big] \\ &= \big[(A_1 \cap A_2^c) \times B_1\big] \cup \big[(A_1 \cap A_2) \times (B_1 \cap B_2^c)\big].\end{aligned}$$

It is easy to see that this union is disjoint, and it follows that $(A_1 \times B_1) \backslash (A_2 \times B_2)$ belongs to $\mathcal{A}$. Therefore,

$$(A_1 \times B_1) \cup (A_2 \times B_2) = (A_1 \times B_1) \backslash (A_2 \times B_2) \cup (A_2 \times B_2)$$

is a finite union of disjoint measurable rectangles, so it lies in $\mathcal{A}$. Applying this proof inductively shows that $\mathcal{A}$ is closed under finite unions. The proof that $\mathcal{A}$ is closed under complements is straightforward—if $\{A_j \times B_j\}_{j=1}^n$ is a finite collection of disjoint measurable rectangles, then

$$\left(\bigcup_{j=1}^{n} (A_j \times B_j)\right)^c = \bigcap_{j=1}^{n} \big[(A_j^c \times B_j) \cup (A_j \times B_j^c)\big].$$

Each term in this intersection is a disjoint union of two measurable rectangles, and it is not too difficult to see from our earlier observations that the intersection can itself be expressed as a finite disjoint union of measurable rectangles. Hence $\mathcal{A}$ is an algebra. $\square$

Proposition 5.5.10. *The function* $(\mu \times \nu)_0 : \mathcal{A} \to [0, \infty]$ *defined by*

$$(\mu \times \nu)_0 \left(\bigcup_{j=1}^{n} (A_j \times B_j) \right) = \sum_{j=1}^{n} \mu(A_j)\nu(B_j)$$

is a premeasure on $\mathcal{A}$.

Proof. Clearly $\mu_0(\emptyset) = 0$. Suppose $\{E_j\}_{j=1}^{\infty}$ is a collection of disjoint sets in $\mathcal{A}$ such that $\bigcup_{j=1}^{\infty} E_j \in \mathcal{A}$. Since each E_j and $\bigcup_{j=1}^{\infty} E_j$ are each finite disjoint unions of measurable rectangles, it really suffices to assume $E_j = A_j \times B_j$ is a measurable rectangle for all j, and that $\bigcup_{j=1}^{\infty} E_j = A \times B$ is a measurable rectangle. Therefore, to establish countable additivity we simply need to show that

$$\mu(A)\nu(B) = \sum_{j=1}^{\infty} \mu(A_j)\nu(B_j).$$

Fix a point $x \in A$. Notice that if $y \in B$, then the pair (x, y) belongs to precisely one of the rectangles $A_j \times B_j$. Therefore, we can write B as a disjoint union

$$B = \bigcup_{j : x \in A_j} B_j.$$

Thus

$$\nu(B) = \sum_{j : x \in A_j} \nu(B_j).$$

Now for each $x \in A$ we have

$$\nu(B)\chi_A(x) = \sum_{j=1}^{\infty} \nu(B_j)\chi_{A_j}(x).$$

The same equality holds trivially if $x \notin A$, so we can write

$$\nu(B)\chi_A = \sum_{j=1}^{\infty} \nu(B_j)\chi_{A_j}.$$

Finally we can integrate over X and apply Corollary 5.3.5 to conclude that

$$\begin{aligned} \mu(A)\nu(B) &= \nu(B) \int_X \chi_A \, d\mu \\ &= \int_X \sum_{j=1}^{\infty} \nu(B_j)\chi_{A_j} \, d\mu \end{aligned}$$

$$= \sum_{j=1}^{\infty} \nu(B_j) \int_X \chi_{A_j}\, d\mu$$
$$= \sum_{j=1}^{\infty} \mu(A_j)\nu(B_j).$$

Hence μ_0 is a premeasure. □

Now we can apply the Carathéodory-Hahn theorem to obtain an extension $\mu \times \nu$ of $(\mu \times \nu)_0$ to the σ-algebra $\mathcal{M} \otimes \mathcal{N}$. We call the measure $\mu \times \nu$ the **product** of μ and ν, and the resulting measure space $(X \times Y, \mathcal{M} \otimes \mathcal{N}, \mu \times \nu)$ is called a **product measure space**.

Remark 5.5.11. It is worth noting that the σ-algebra $\mathcal{M} \otimes \mathcal{N}$ might not contain all of the $(\mu \times \nu)^*$-measurable sets. In particular, $\mu \times \nu$ might not be a complete measure, even if μ and ν are. To see this, let $X = Y = \mathbf{R}$, take $\mathcal{M} = \mathcal{N} = \mathcal{L}$, and let $\mu = \nu$ both be Lebesgue measure. If $A \subseteq \mathbf{R}$ is a non-Lebesgue measurable set, then the set $\{0\} \times A \subseteq X \times Y$ cannot belong to $\mathcal{M} \otimes \mathcal{N}$, since the cross-section corresponding to $x = 0$ is not ν-measurable. However, $\{0\} \times A \subseteq \{0\} \times \mathbf{R}$, and the latter set is a $\mu \times \nu$-null set.

5.5.1 Integration on Product Spaces

Now that we have determined how to properly define the product of two measure spaces $(X, \mathcal{M}, \mu)$ and $(Y, \mathcal{N}, \nu)$, we turn to the problem of integrating a function $f : X \times Y \to \mathbf{C}$ with respect to the product measure $\mu \times \nu$. As mentioned in the introduction, we would like to show that the integral can be done via iterated integration, and that the order of the iterated integral does not matter:

$$\int_{X \times Y} f\, d(\mu \times \nu) = \int_Y \left[\int_X f\, d\mu\right] d\nu = \int_X \left[\int_Y f\, d\nu\right] d\mu. \tag{5.2}$$

Of course the equation we have just written does not quite make sense. Since f is defined on $X \times Y$, we cannot really integrate f over X. As in multivariable calculus, we need to hold one variable fixed, and then integrate the resulting function over X. Thus we could write (5.2) more precisely as

$$\int_{X \times Y} f(x,y)\, d(\mu \times \nu) = \int_Y \left[\int_X f(x,y)\, d\mu(x)\right] d\nu(y) = \int_X \left[\int_Y f(x,y)\, d\nu(y)\right] d\mu(x).$$

This problem leads us to define a notion of "cross-section" for both sets and functions on $X \times Y$.

Definition 5.5.12. Let $E \subseteq X \times Y$. For each $x \in X$, we define the x**-section** of E to be

$$E_x = \{y \in Y : (x, y) \in E\}.$$

Similarly, for $y \in Y$ we define the y**-section** of E by

$$E^y = \{x \in X : (x, y) \in E\}.$$

Definition 5.5.13. Let $f : X \times Y \to \mathbf{C}$. For each $x \in X$, we define the x**-section** of f to be the function $f_x : Y \to \mathbf{C}$ given by

$$f_x(y) = f(x, y).$$

Similarly, for $y \in Y$ we define the y**-section** $f_y : X \to \mathbf{C}$ by

$$f^y(x) = f(x, y).$$

Remark 5.5.14. Of course we can also define the x- and y-sections of an extended real-valued function accordingly.

By using the x- and y-sections of a function f, we can restate (5.2) more precisely as follows:

$$\int_{X\times Y} f \, d(\mu \times \nu) = \int_Y \left[\int_X f^y \, d\mu \right] d\nu = \int_X \left[\int_Y f_x \, d\nu \right] d\mu. \tag{5.3}$$

With the appropriate hypotheses, (5.3) is more or less the content of Fubini's theorem. Before we can state and prove that theorem, however, there are some natural questions regarding measurability that we need to answer.

Proposition 5.5.15. *Suppose $E \in \mathcal{M} \otimes \mathcal{N}$. Then $E_x \in \mathcal{N}$ for all $x \in X$, and $E^y \in \mathcal{M}$ for all $y \in Y$.*

Proof. We begin by defining

$$\mathcal{S} = \{E \subseteq X \times Y : E_x \in \mathcal{N} \text{ for all } x \in X \text{ and } E^y \in \mathcal{M} \text{ for all } y \in Y\}.$$

Obviously the goal is to show that $\mathcal{M} \otimes \mathcal{N} \subseteq \mathcal{S}$. It is straightforward to check that if $\{E_j\}_{j=1}^\infty$ is a collection of subsets of $X \times Y$, then

$$\left(\bigcup_{j=1}^\infty E_j \right)_x = \bigcup_{j=1}^\infty (E_j)_x, \quad \left(\bigcup_{j=1}^\infty E_j \right)^y = \bigcup_{j=1}^\infty (E_j)^y.$$

for all $x \in X$ and all $y \in Y$. Also, if $E \subseteq X \times Y$,

$$(E^c)_x = (E_x)^c, \quad (E^c)^y = (E^y)^c$$

for all $x \in X$ and all $y \in Y$. It is then easy to see that $\mathcal{S}$ is closed under countable unions and intersections, so it is a σ-algebra. Now notice that if $A \in \mathcal{M}$ and $B \in \mathcal{N}$, then $(A \times B)_x = B$ whenever $x \in A$ (and it is empty otherwise), and likewise $(A \times B)^y = A$ for all $y \in B$ (or it equals $\emptyset$ when $y \notin B$). Thus $\mathcal{S}$ contains all the measurable rectangles, so it must contain the σ-algebra they generate. Hence $\mathcal{M} \otimes \mathcal{N} \subseteq \mathcal{S}$, and the result follows. □

Proposition 5.5.16. *Suppose $f : X \times Y \to \mathbf{C}$ is $\mathcal{M} \otimes \mathcal{N}$-measurable. Then f_x is $\mathcal{N}$-measurable for all $x \in X$, and f^y is $\mathcal{M}$-measurable for all $y \in Y$.*

Proof. By considering real and imaginary parts, it suffices to assume f is real-valued. Let $B \subseteq \mathbf{R}$ be a Borel set. Then it is not hard to check that

$$(f_x)^{-1}(B) = \big(f^{-1}(B)\big)_x$$

for all $x \in X$. Since f is $\mathcal{M} \otimes \mathcal{N}$-measurable, $f^{-1}(B) \in \mathcal{M} \otimes \mathcal{N}$, and the previous proposition guarantees that $(f^{-1}(B))_x \in \mathcal{N}$ for all $x \in X$. Thus f_x is $\mathcal{N}$-measurable for all x. Similarly, we have

$$(f^y)^{-1}(B) = \big(f^{-1}(B)\big)^y$$

for all $y \in Y$, and the same argument shows that f^y is $\mathcal{M}$-measurable for all y. □

Before we can start in on Fubini's theorem, we need a technical result on σ-algebras. Let X be a set. We say a collection of sets $\mathcal{C} \subseteq \mathcal{P}(X)$ is a **monotone class** if it is closed under countable *increasing* unions and countable *decreasing* intersections. That is, if $\{E_j\}_{j=1}^{\infty} \subseteq \mathcal{C}$ with

$$E_1 \subseteq E_2 \subseteq E_3 \subseteq \cdots,$$

then we require $\bigcup_{j=1}^{\infty} E_j \in \mathcal{C}$, and if $E_1 \supseteq E_2 \supseteq E_3 \supseteq \cdots$, then $\bigcap_{j=1}^{\infty} E_j \in \mathcal{C}$. It is easy to see that any σ-algebra is automatically a monotone class.

Proposition 5.5.17 (Monotone Class Lemma)**.** *Let X be a set, and suppose $\mathcal{A} \subseteq \mathcal{P}(X)$ is an algebra of sets. Let $\mathcal{M}$ be the σ-algebra generated by $\mathcal{A}$. Then $\mathcal{M}$ is the smallest monotone class containing $\mathcal{A}$.*

Proof. Exercise 5.5.3. □

Naturally, the first step in proving Fubini's theorem is to establish the result for simple functions. The next theorem allows us to handle characteristic functions, and then simple functions follow by linearity.

Theorem 5.5.18. *Suppose $(X, \mathcal{M}, \mu)$ and $(Y, \mathcal{N}, \nu)$ are σ-finite measure spaces. For each set $E \in \mathcal{M} \otimes \mathcal{N}$, the function $x \mapsto \nu(E_x)$ is $\mathcal{M}$-measurable on X and the function $y \mapsto \mu(E^y)$ is $\mathcal{N}$-measurable on Y. Moreover,*

$$(\mu \times \nu)(E) = \int_X \nu(E_x)\, d\mu = \int_Y \mu(E^y)\, d\nu.$$

Proof. First assume that μ and ν are finite. Let $\mathcal{C}$ denote collection of all sets in $\mathcal{M} \otimes \mathcal{N}$ for which the conclusion of the theorem holds. We aim to show that $\mathcal{C}$ is a monotone class and that it contains finite disjoint unions of measurable rectangles, which will imply that $\mathcal{M} \otimes \mathcal{N} \subseteq \mathcal{C}$ by the Monotone Class Lemma.

First suppose that $E = A \times B$ is a measurable rectangle in $X \times Y$. Then for each $x \in X$,

$$\nu(E_x) = \begin{cases} \nu(B) & \text{if } x \in A \\ 0 & \text{if } x \notin A, \end{cases}$$

so we can write $\nu(E_x) = \chi_A(x)\nu(B)$. Thus $x \mapsto \nu(E_x)$ is measurable and

$$\int_X \nu(E_x)\, d\mu = \nu(B) \int_X \chi_A\, d\mu = \mu(A)\nu(B).$$

Similarly, we have $\mu(E^y) = \chi_B(y)\mu(A)$ for all $y \in Y$, so $y \mapsto \mu(E^y)$ is measurable and

$$\int_Y \mu(E^y)\, d\nu = \mu(A) \int_Y \chi_B\, d\nu = \mu(A)\nu(B).$$

Thus $E \in \mathcal{C}$. By linearity, these conditions hold for finite disjoint unions of rectangles as well. Hence the algebra $\mathcal{A}$ consisting of finite disjoint unions of measurable rectangles is contained in $\mathcal{C}$.

Now suppose $\{E_j\}_{j=1}^\infty$ is an increasing sequence of sets in $\mathcal{C}$, and put $E = \bigcup_{j=1}^\infty E_j$. For each j, define

$$f_j(y) = \mu\big((E_j)^y\big)$$

for all $y \in Y$. Then each f_j is measurable, the sequence $(f_j)_{j=1}^\infty$ is pointwise increasing, and $f_j \to f$ pointwise, where $f(y) = \mu(E^y)$ for all $y \in Y$. Thus f is measurable, and the Monotone Convergence Theorem implies that

$$\int_Y \mu(E^y)\, d\nu = \int_Y f\, d\nu = \lim_{j\to\infty} \int_Y f_j\, d\nu = \lim_{j\to\infty} (\mu \times \nu)(E_j) = (\mu \times \nu)(E),$$

with the last equality following from continuity of measure. Of course a similar argument shows that

$$\int_X \nu(E_x)\,d\mu = (\mu \times \nu)(E),$$

so $E \in \mathcal{C}$. Thus $\mathcal{C}$ is closed under countable increasing unions. Now suppose $\{E_j\}_{j=1}^{\infty}$ is a decreasing sequence of sets, and put $E = \bigcap_{j=1}^{\infty} E_j$. For each j, define a function f_j by

$$f_j(y) = \mu\big((E_j)^y\big).$$

Then the sequence $(f_j)_{j=1}^{\infty}$ is pointwise decreasing, and $f_j \to f$ pointwise, where $f(y) = \mu(E^y)$ for all $y \in Y$. Notice that

$$\int_Y f_1\,d\nu \le \int_Y \mu(X)\,d\nu = \mu(X)\nu(Y) < \infty,$$

so $f_1 \in L^1(Y, \nu)$. Thus the Dominated Convergence Theorem implies that

$$\int_Y \mu(E^y)\,d\nu = \int_Y f\,d\nu = \lim_{j\to\infty} \int_Y f_j\,d\nu = \lim_{j\to\infty} (\mu \times \nu)(E_j) = (\mu \times \nu)(E).$$

Similarly, $\int_X \nu(E_x)\,d\mu = (\mu \times \nu)(E)$, so $E \in \mathcal{C}$. This shows that $\mathcal{C}$ is closed under countable decreasing intersections, hence it is a monotone class. Therefore, $\mathcal{C}$ contains $\mathcal{M} \otimes \mathcal{N}$ by the Monotone Class Lemma.

Now we assume the two measure spaces are σ-finite. Write $X = \bigcup_{j=1}^{\infty} X_j$ and $Y = \bigcup_{j=1}^{\infty} Y_j$, where the unions are increasing and the measures of X_j and Y_j are finite for all j. Then we have

$$X \times Y = \bigcup_{j=1}^{\infty} X_j \times Y_j.$$

Let $E \in \mathcal{M} \otimes \mathcal{N}$. Then $E \cap (X_j \times Y_j)$ has finite measure for each j, and

$$\nu\big((E \cap (X_j \times Y_j))_x\big) = \begin{cases} \nu(E_x \cap Y_j) & \text{if } x \in X_j \\ 0 & \text{if } x \notin X_j. \end{cases}$$

A similar result holds for the y-sections. Thus our previous work shows that

$$(\mu \times \nu)(E \cap (X_j \times Y_j)) = \int_X \nu(E_x \cap Y_j)\chi_{X_j}\,d\mu = \int_Y \mu(E^y \cap X_j)\chi_{Y_j}\,d\nu.$$

Since the sequences $(\nu(E_x \cap Y_j)\chi_{X_j})_{j=1}^{\infty}$ and $(\mu(E^y \cap X_j)\chi_{Y_j})_{j=1}^{\infty}$ increase pointwise to $\nu(E_x)$ and $\mu(E^y)$, respectively, the Monotone Convergence Theorem implies that

$$(\mu \times \nu)(E) = \int_X \nu(E_x)\,d\mu = \int_Y \mu(E^y)\,d\nu.$$

Thus the result holds for σ-finite measures as well. $\square$

Now we arrive at the general statement of Fubini's theorem. The theorem is actually called the Fubini-Tonelli theorem, since it comes in two parts—a statement for nonnegative functions on $X \times Y$ (Tonelli) and an analogous statement for $\mathcal{M} \otimes \mathcal{N}$-integrable functions (Fubini).

Theorem 5.5.19 (Fubini-Tonelli)**.** *Suppose* $(X, \mathcal{M}, \mu)$ *and* $(Y, \mathcal{N}, \nu)$ *are* σ*-finite measure spaces.*

1. (Tonelli's theorem) *If* $f : X \times Y \to [0, +\infty]$ *is* $\mathcal{M} \otimes \mathcal{N}$*-measurable, then the functions* $g : X \to [0, +\infty]$ *and* $h : Y \to [0, +\infty]$ *defined by*

$$g(x) = \int_Y f_x \, d\nu, \quad h(y) = \int_X f^y \, d\mu$$

are $\mathcal{M}$*- and* $\mathcal{N}$*-measurable, respectively. Moreover,*

$$\int_{X \times Y} f \, d(\mu \times \nu) = \int_X g \, d\mu = \int_Y h \, d\nu.$$

2. (Fubini's theorem) *Suppose* $f \in L^1(X \times Y, \mu \times \nu)$*. Then* $f_x \in L^1(Y, \nu)$ *for* μ*-a.e.* $x \in X$*,* $f^y \in L^1(X, \mu)$ *for* ν*-a.e.* $y \in Y$*, and the functions* $g : X \to \mathbf{C}$ *and* $h : Y \to \mathbf{C}$ *defined a.e. by*

$$g(x) = \int_Y f_x \, d\nu, \quad h(y) = \int_X f^y \, d\mu$$

belong to $L^1(X, \mu)$ *and* $L^1(Y, \nu)$*, respectively. Moreover,*

$$\int_{X \times Y} f \, d(\mu \times \nu) = \int_X g \, d\mu = \int_Y h \, d\nu.$$

Proof. As we mentioned earlier, we have already established Tonelli's theorem for characteristic functions in Theorem 5.5.18. Of course it would then hold for nonnegative simple functions by the linearity of the integral. Suppose $f : X \times Y \to [0, +\infty]$ is $\mathcal{M} \otimes \mathcal{N}$-measurable, and let $(f_n)_{n=1}^{\infty}$ be a sequence of nonnegative simple functions that increase pointwise to f. For each n, set

$$g_n(x) = \int_Y (f_n)_x \, d\nu, \quad h_n(x) = \int_X (f_n)^y \, d\mu.$$

Certainly $(f_n)_x$ increases pointwise to f_x for each $x \in X$, and similarly $(f_n)^y$ increases pointwise to f^y for each $y \in Y$. Thus the Monotone Convergence Theorem implies that g_n and h_n increase pointwise to g and h, respectively. In particular, g and h are measurable, and the Monotone Convergence Theorem again guarantees that

$$\int_X g\,d\mu = \lim_{n\to\infty} \int_X g_n\,d\mu = \lim_{n\to\infty} \int_{X\times Y} f_n\,d(\mu\times\nu) = \int_{X\times Y} f\,d(\mu\times\nu)$$

and

$$\int_Y h\,d\nu = \lim_{n\to\infty} \int_X h_n\,d\nu = \lim_{n\to\infty} \int_{X\times Y} f_n\,d(\mu\times\nu) = \int_{X\times Y} f\,d(\mu\times\nu).$$

Thus Tonelli's theorem is proved.

Now suppose $f \in L^1(X\times Y, \mu\times\nu)$. By considering real and imaginary parts, it suffices to assume $f : X\times Y \to \mathbf{R}$. By further reducing to f^+ and f^-, we can assume f is nonnegative and measurable and $\int_{X\times Y} f\,d(\mu\times\nu) < \infty$. Then by Tonelli's theorem we know that g and h are finite a.e., or equivalently, $f_x \in L^1(Y,\nu)$ for μ-a.e. x and $f^y \in L^1(X,\mu)$ for ν-a.e. y. The rest of the assertions in Fubini's theorem then clearly follow from Tonelli's theorem and the linearity of the integral. □

Remark 5.5.20. The assumption in the Fubini-Tonelli theorem that both measures are σ-finite is essential. In fact, one can find two measure spaces $(X, \mathcal{M}, \mu)$ and $(Y, \mathcal{N}, \nu)$ where μ is σ-finite and ν is not, and a function $f : X\times Y \to \mathbf{R}$ such that the integrals $\int_{X\times Y} f\,d(\mu\times\nu)$, $\int_X\int_Y f\,d\nu\,d\mu$, and $\int_Y\int_X f\,d\mu\,d\nu$ all give different values. (See Exercise 5.5.4.)

Remark 5.5.21. The hypothesis that f is either nonnegative or integrable is essential in the Fubini-Tonelli theorem. Here are some things that might happen if these hypotheses are omitted:

- Even if $f : X\times Y \to \mathbf{C}$ is not measurable with respect to the σ-algebra $\mathcal{M}\otimes\mathcal{N}$, the sections f_x and f^y may turn out to be measurable for all x and y. However, the iterated integrals need not be equal.

- If f is not a nonnegative function that is not integrable, the sections f_x and f^y may still turn out to be integrable for all x and y, and the iterated integrals $\int_X\int_Y f\,d\nu\,d\mu$ and $\int_Y\int_X f\,d\mu\,d\nu$ might still exist. However, the two iterated integrals need not be equal. (Exercise 5.5.5.)

Remark 5.5.22. It is common practice for one to use the Fubini and Tonelli theorems in succession. Given a function $f : X\times Y \to \mathbf{C}$, one can use the Tonelli theorem to compute $\int_{X\times Y} |f|\,d(\mu\times\nu)$ via an iterated integral, and thus show that $f \in L^1(X\times Y, \mu\times\nu)$. With that done, Fubini's theorem can then be applied to compute $\int_{X\times Y} f\,d(\mu\times\nu)$.

If the reader has a keen eye, they may have noticed that we have made no mention of completeness in the statement of the Fubini-Tonelli theorem. However, we did observe earlier that the product measure will not turn out to be complete in general. Of course we can complete it, and then the Fubini-Tonelli theorem still holds, albeit with some slight modifications. The proof is not all that different from the one we have already given.

Theorem 5.5.23 (Fubini-Tonelli for Complete Measures). *Suppose* $(X, \mathcal{M}, \mu)$ *and* $(Y, \mathcal{N}, \nu)$ *are complete* σ*-finite measure spaces, and let* $(X \times Y, \overline{\mathcal{M} \otimes \mathcal{N}}, \overline{\mu \times \nu})$ *denote the completion of the product space* $(X \times Y, \mathcal{M} \otimes \mathcal{N}, \mu \times \nu)$*.*

1. (Tonelli's theorem) *If* $f : X \times Y \to [0, +\infty]$ *is* $\overline{\mathcal{M} \otimes \mathcal{N}}$*-measurable, then* f_x *is* $\mathcal{N}$*-measurable for* μ*-a.e.* x*, and* f^y *is* $\mathcal{M}$*-measurable for* ν*-a.e.* y*. Moreover, the functions* $g : X \to [0, +\infty]$ *and* $h : Y \to [0, +\infty]$ *defined by*

$$g(x) = \int_Y f_x \, d\nu, \quad h(y) = \int_X f^y \, d\mu$$

   are $\mathcal{M}$*- and* $\mathcal{N}$*-measurable, respectively, and*

$$\int_{X \times Y} f \, d(\mu \times \nu) = \int_X g \, d\mu = \int_Y h \, d\nu.$$

2. (Fubini's theorem) *Suppose* $f \in L^1(X \times Y, \overline{\mu \times \nu})$*. Then* $f_x \in L^1(Y, \nu)$ *for* μ*-a.e.* $x \in X$*,* $f^y \in L^1(X, \mu)$ *for* ν*-a.e.* $y \in Y$*, and the functions* g *and* h *defined a.e. as above belong to* $L^1(X, \mu)$ *and* $L^1(Y, \nu)$*, respectively. Moreover,*

$$\int_{X \times Y} f \, d(\mu \times \nu) = \int_X g \, d\mu = \int_Y h \, d\nu.$$

We close out this section with some elementary examples and applications of the Fubini-Tonelli theorem.

Example 5.5.24 (Integration on $\mathbf{R}^n$). Let $X = Y = \mathbf{R}$ and $\mathcal{M} = \mathcal{N} = \mathcal{L}$, both equipped with the Lebesgue measure μ. The completion $\overline{\mu \times \mu}$ is the two-dimensional Lebesgue measure on $\mathbf{R}^2$, which generalizes the notion of *area* in the plane in the same way that the Lebesgue measure on $\mathbf{R}$ generalizes *length*. Indeed, one can define the Lebesgue outer measure of a set $E \subseteq \mathbf{R}^2$ by covering E with rectangles whose sides are intervals, adding up the areas of the rectangles, and then taking the infimum over all such coverings. The resulting measure agrees with $\overline{\mu \times \mu}$.

Suppose $E = [a, b] \times [c, d]$ is a rectangle in $\mathbf{R}^2$ and $f : E \to \mathbf{R}$ is continuous. Then f is both Riemann integrable and Lebesgue integrable on E, and we have

$$\iint_E f(x, y) \, dA = \int_E f \, d(\overline{\mu \times \mu}),$$

where the integral on the left denotes the usual two-dimensional Riemann integral.

Moreover, the version of Fubini's theorem for complete measures gives

$$\begin{aligned}\int_E f(x,y)\,d(\overline{\mu\times\mu})(x,y) &= \int_{[a,b]}\left[\int_{[c,d]} f(x,y)\,d\mu(y)\right]\,d\mu(x)\\ &= \int_{[c,d]}\left[\int_{[a,b]} f(x,y)\,d\mu(x)\right]\,d\mu(y).\end{aligned}$$

However, it is not hard to check that the sections of f are all continuous, hence Riemann integrable, so we get

$$\int_{[a,b]} f(x,y)\,d\mu(x) = \int_a^b f(x,y)\,dx$$

for all $y \in \mathbf{R}$ and

$$\int_{[c,d]} f(x,y)\,d\mu(y) = \int_c^d f(x,y)\,dy$$

for all $x \in \mathbf{R}$. Furthermore, we can easily verify that the functions $g : [a,b] \to \mathbf{R}$ and $h : [c,d] \to \mathbf{R}$ defined by

$$g(x) = \int_c^d f(x,y)\,dy, \quad h(x) = \int_a^b f(x,y)\,dx$$

are continuous. For example, let $\varepsilon > 0$ be given, and choose $\delta > 0$ such that $\|(x_1,y_1)-(x_2,y_2)\|_2 < \delta$ implies

$$|f(x_1,y_1) - f(x_2,y_2)| < \frac{\varepsilon}{d-c}$$

for $(x_1,y_1),(x_2,y_2) \in E$. (We can do this since f is continuous on the compact set E, hence it is uniformly continuous.) In particular, for each $y \in [c,d]$ we have

$$|f(x_1,y) - f(x_2,y)| < \frac{\varepsilon}{d-c}$$

whenever $x_1, x_2 \in [a,b]$ satisfy $|x_1 - x_2| < \delta$. Therefore, $|x_1 - x_2| < \delta$ implies

$$\begin{aligned}|g(x_1) - g(x_2)| &= \left|\int_c^d f(x_1,y)\,dy - \int_c^d f(x_2,y)\,dy\right|\\ &\le \int_c^d |f(x_1,y) - f(x_2,y)|\,dy\\ &< \int_c^d \frac{\varepsilon}{d-c}\,dy\\ &= \varepsilon,\end{aligned}$$

so g is (uniformly) continuous on E. Hence g is Riemann integrable on $[a, b]$, and we have

$$\int_{[a,b]} \left[\int_{[c,d]} f(x, y)\, d\mu(y) \right] d\mu(x) = \int_a^b \left[\int_c^d f(x, y)\, dy \right] dx.$$

A similar argument works to show that h is continuous, hence Riemann integrable, and we get

$$\int_{[c,d]} \left[\int_{[a,b]} f(x, y)\, d\mu(x) \right] d\mu(y) = \int_c^d \left[\int_a^b f(x, y)\, dx \right] dy.$$

Therefore,

$$\int_E f(x, y)\, dA = \int_a^b \int_c^d f(x, y)\, dy\, dx = \int_c^d \int_a^b f(x, y)\, dx\, dy,$$

which is the version of Fubini's theorem that one usually encounters in multivariable calculus.

Example 5.5.25 (Counting measure)**.** Let $X = Y = \mathbf{N}$, both equipped with the counting measure μ. It is not hard to see that the product measure $\mu \times \mu$ is nothing more than the counting measure on $\mathbf{N}^2 = \mathbf{N} \times \mathbf{N}$, so a function $f : \mathbf{N}^2 \to \mathbf{C}$ belongs to $L^1(\mathbf{N}^2, \mu \times \mu)$ if and only if

$$\int_{\mathbf{N}^2} |f|\, d(\mu \times \mu) = \sum_{(n,m) \in \mathbf{N}^2} |f(n, m)| < \infty.$$

That is, f is integrable if and only if the doubly-indexed series $\sum_{n,m=1}^{\infty} f(n, m)$ converges absolutely. In this case, Fubini's theorem applies, and we have

$$\sum_{n,m=1}^{\infty} f(n, m) = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} f(n, m) = \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} f(n, m).$$

In other words, if a doubly-indexed series converges absolutely, then we may interchange the order of summation.

One might wonder how we could check that a doubly-indexed series converges absolutely without already knowing that we can interchange the order of summation. Well, this is where Tonelli's theorem comes to the rescue: we can write

$$\sum_{(n,m) \in \mathbf{N}^2} |f(n, m)| = \sum_{n=1}^{\infty} \sum_{m=1}^{\infty} |f(n, m)| = \sum_{m=1}^{\infty} \sum_{n=1}^{\infty} |f(n, m)|,$$

so the original sum converges absolutely if and only if one of the sums on the right hand side converges.

Our last application of the Fubini-Tonelli theorem is a little more complicated, though it is of great importance in harmonic analysis. Given two Lebesgue measurable functions $f, g : \mathbf{R} \to \mathbf{C}$, we define the **convolution** of f and g to be the function $f * g : \mathbf{R} \to \mathbf{C}$ defined by

$$(f * g)(x) = \int_{\mathbf{R}} f(x-t)g(t)\, d\mu(t),$$

whenever this integral makes sense. By making a change of variables and using the translation invariance of Lebesgue measure, it is not hard to check that convolution is commutative. We will consider another property of it, however: if $f, g \in L^1(\mathbf{R}, \mu)$, then $f * g \in L^1(\mathbf{R}, \mu)$ as well. Consequently, $L^1(\mathbf{R}, \mu)$ forms a commutative algebra with multiplication given by convolution. Once we have defined the L^1-norm of a function later on, we will see that the our theorem also shows that convolution is *submultiplicative* with respect to the L^1-norm. Consequently, $L^1(\mathbf{R}, \mu)$ is an example of what we will call a *Banach algebra.*

Before we can prove the aforementioned result, we need a quick lemma about measurability. The proof is due to Dana Williams [Wil07].

Lemma 5.5.26. *Suppose $f : \mathbf{R} \to \mathbf{C}$ is Lebesgue measurable. Then the function $F : \mathbf{R}^2 \to \mathbf{C}$ defined by*

$$F(x,t) = f(x-t)$$

is measurable with respect to the completion of $\mu \times \mu$.

Proof. Since f is measurable and can be expressed as a pointwise limit of measurable simple functions, it suffices to assume $f = \chi_E$ for some Lebesgue measurable set $E \subseteq \mathbf{R}$. Checking that F is measurable amounts to verifying that the set

$$E' = \big\{(x,t) \in \mathbf{R}^2 : x - t \in E\big\}$$

belongs to the product σ-algebra $\mathcal{L} \otimes \mathcal{L}$. By Theorem 4.3.4, we can find an F_σ-set A and a null set Z such that $E = A \cup Z$. If we define

$$A' = \big\{(x,t) \in \mathbf{R}^2 : x - t \in A\big\},$$

then A' is the preimage of the Borel set A under the continuous map $(x,t) \mapsto x - t$. Hence A' is a Borel set, so $A' \in \mathcal{L} \otimes \mathcal{L}$. Now consider

$$Z' = \big\{(x,t) \in \mathbf{R}^2 : x - t \in Z\big\}.$$

Again, Theorem 4.3.4 guarantees the existence of a Borel null set $W \supseteq Z$, and we define

$$W' = \big\{(x,t) \in \mathbf{R}^2 : x - t \in W\big\}.$$

The same arguments as above show that $W' \in \mathcal{L} \otimes \mathcal{L}$. Now notice that for all $x \in \mathbf{R}$,

$$W'_x = \{t \in \mathbf{R} : (x,t) \in W'\} = \{t \in \mathbf{R} : x - t \in W\} = x - W,$$

which is a null set. Tonelli's theorem then implies that

$$(\mu \times \mu)(W') = \int_{\mathbf{R}} \nu(W'_x)\, d\mu(x) = 0,$$

so W' is $\mu \times \mu$-null. Since $Z' \subseteq W'$, it follows that Z' is measurable. (This is where we need to be working with the complete product measure.) It is easy to see that $E' = A' \cup Z'$, so E' is $\overline{\mu \times \mu}$-measurable. Hence F is measurable. □

Theorem 5.5.27. *Let $f, g \in L^1(\mathbf{R}, \mu)$. Then the convolution $f * g$ is a.e.-defined on $\mathbf{R}$. Moreover, $f * g \in L^1(\mathbf{R}, \mu)$, and*

$$\int_{\mathbf{R}} |f * g|\, d\mu \le \left(\int_{\mathbf{R}} |f|\, d\mu\right)\left(\int_{\mathbf{R}} |g|\, d\mu\right).$$

Proof. Assume first that f and g are nonnegative. Notice that the function $(x,t) \mapsto f(x-t)g(t)$ is $\overline{\mu \times \mu}$-measurable on $\mathbf{R}^2$, since it can be thought of as a product of two measurable functions. By Tonelli's theorem for complete measures, we have

$$\begin{aligned}
\int_{\mathbf{R}} (f * g)(x)\, d\mu(x) &= \int_{\mathbf{R}} \int_{\mathbf{R}} f(x-t)g(t)\, d\mu(t)\, d\mu(x) \\
&= \int_{\mathbf{R}} \int_{\mathbf{R}} f(x-t)g(t)\, d\mu(x)\, d\mu(t) \\
&= \int_{\mathbf{R}} g(t) \left[\int_{\mathbf{R}} f(x-t)\, d\mu(x)\right] d\mu(t) \\
&= \left(\int_{\mathbf{R}} g(t)\, d\mu(t)\right)\left(\int_{\mathbf{R}} f(x-t)\, d\mu(x)\right).
\end{aligned}$$

Since Lebesgue measure is translation-invariant, it follows that

$$\int_{\mathbf{R}} f * g\, d\mu = \left(\int_{\mathbf{R}} f\, d\mu\right)\left(\int_{\mathbf{R}} g\, d\mu\right)$$

when f and g are nonnegative. In general, we have

$$|(f * g)(x)| \le \int_{\mathbf{R}} |f(x-t)g(t)|\, d\mu(t) = (|f| * |g|)(x),$$

so $|f * g| \le |f| * |g|$. Therefore,

$$\int_{\mathbf{R}} |f * g|\, d\mu \le \int_{\mathbf{R}} |f| * |g|\, d\mu = \left(\int_{\mathbf{R}} |f|\, d\mu\right)\left(\int_{\mathbf{R}} |g|\, d\mu\right).$$ □

Exercises for Section 5.5

Exercise 5.5.1. Let X be a set and $\mathcal{A} \subseteq \mathcal{P}(X)$ an algebra on X. Suppose μ_0 is a premeasure on $\mathcal{A}$ and μ^* is the associated outer measure. Show that $\mu^*(E) = \mu_0(E)$ for every $E \in \mathcal{A}$. (Hint: Mimic the proof of Proposition 4.1.5.)

Exercise 5.5.2. Complete the proof of the Carathéodory-Hahn theorem by showing that if μ_0 is a σ-finite premeasure on an algebra $\mathcal{A} \subseteq \mathcal{P}(X)$, then the associated measure μ is σ-finite and μ is the unique extension of μ_0 to the σ-algebra generated by $\mathcal{A}$.

Exercise 5.5.3. This exercise will lead to a proof of the Monotone Class Lemma.

Let X be a set, and suppose $\mathcal{A} \subseteq \mathcal{P}(X)$ is an algebra of sets. Let $\mathcal{C}$ denote the monotone class generated by $\mathcal{A}$, and let $\mathcal{M}$ denote the σ-algebra generated by $\mathcal{A}$. Given a set $E \in \mathcal{C}$, define

$$\mathcal{C}(E) = \{F \in \mathcal{C} : E\backslash F,\ F\backslash E,\ E \cap F \in \mathcal{C}\}.$$

(a) Show that for all $E \in \mathcal{C}$, $\emptyset \in \mathcal{C}(E)$ and $E \in \mathcal{C}(E)$.

(b) Prove that if $E, F \in \mathcal{C}$, then $F \in \mathcal{C}(E)$ if and only if $E \in \mathcal{C}(F)$.

(c) Prove that $\mathcal{C}(E)$ is a monotone class for all $E \in \mathcal{C}$.

(d) Let $E \in \mathcal{A}$. Show that $\mathcal{A} \subseteq \mathcal{C}(E)$. Conclude that $\mathcal{C} \subseteq \mathcal{C}(E)$ for all $E \in \mathcal{C}$.

(e) Prove that $\mathcal{C}$ is a σ-algebra. Conclude that $\mathcal{C} = \mathcal{M}$.

Exercise 5.5.4. Define two measure spaces $(X, \mathcal{M}, \mu)$ and $(Y, \mathcal{N}, \nu)$ as follows. Let $X = Y = [0,1]$, take $\mathcal{M}$ and $\mathcal{N}$ to both be the Borel σ-algebra on $[0,1]$, let μ be the restriction of Lebesgue measure to $[0,1]$, and let ν be the counting measure. Let

$$D = \big\{(x,x) \in X \times Y : x \in [0,1]\big\}$$

denote the diagonal in $X \times Y$. Show that the integrals

$$\int_{X\times Y} \chi_D \, d(\mu \times \nu), \qquad \int_X \int_Y \chi_D \, d\nu \, d\mu, \qquad \int_Y \int_X \chi_D \, d\mu \, d\nu$$

all yield different values. Conclude that the Fubini-Tonelli theorem fails if the measures are not both assumed to be σ-finite.

Exercise 5.5.5. Let $X = Y = \mathbf{N}$, and set $\mathcal{M} = \mathcal{N} = \mathcal{P}(\mathbf{N})$. Let μ and ν both denote the counting measure on $\mathbf{N}$. Define $f : X \times Y \to \mathbf{C}$ by

$$f(m,n) = \begin{cases} 1 & \text{if } m = n \\ -1 & \text{if } m = n+1 \\ 0 & \text{otherwise.} \end{cases}$$

Show that $\int_{X\times Y} |f| \, d(\mu \times \nu) = \infty$, and that the iterated integrals $\int_X \int_Y f \, d\nu \, d\mu$ and $\int_Y \int_X f \, d\mu \, d\nu$ both exist, but are unequal.

Exercise 5.5.6. Let μ denote the Lebesgue measure on $\mathbf{R}$. For a function $f \in L^1(\mathbf{R}, \mu)$, we define the **Fourier transform** $\hat{f} : \mathbf{R} \to \mathbf{C}$ by

$$\hat{f}(x) = \int_{\mathbf{R}} f(t) e^{-ixt} \, d\mu(t).$$

Use Fubini's theorem to show that for all $f, g \in L^1(\mathbf{R}, \mu)$,

$$(f * g)\hat{}\,(x) = \hat{f}(x)\hat{g}(x)$$

for all $x \in \mathbf{R}$.

5.6 Signed and Complex Measures

In the remainder of this chapter, we will investigate the situation of a space (and σ-algebra) equipped with multiple different measures. In particular, we are interested in the ways in which such measures can relate to one another. For example, if two measures μ and ν are closely related in an appropriate sense, perhaps we can write down a "change-of-variables" formula for their respective integrals. The ability to do so is a fundamental consequence of a major result known as the Radon-Nikodym theorem.

Before we can discuss the Radon-Nikodym theorem, we need to determine what it means for two measures to be "closely related". In doing so, we will introduce a bit more generality to our measure spaces. Throughout this section, we will study multiple measures defined on the same space X and the same σ-algebra $\mathcal{M} \subseteq \mathcal{P}(X)$. For brevity, we will refer to such a pair $(X, \mathcal{M})$ as a **measurable space**.

Definition 5.6.1. Let $(X, \mathcal{M})$ be a measurable space. A **signed measure** on $\mathcal{M}$ is a function $\nu : \mathcal{M} \to [-\infty, \infty]$ satisfying the following conditions:

1. $\nu(\emptyset) = 0$.
2. ν takes at most one of the values ∞ or $-\infty$.
3. If $\{E_j\}_{j=1}^{\infty}$ is a collection of pairwise disjoint sets in $\mathcal{M}$ and $E = \bigcup_{j=1}^{\infty} E_j$,

$$\nu(E) = \sum_{j=1}^{\infty} \nu(E_j),$$

where we require the sum to converge absolutely whenever $\nu(E)$ is finite.

While we cannot interpret a signed measure as representing the volume of a set *per se*, there are several reasons for considering them.

1. Rather than studying volume, one can study "signed volume" or, more precisely, certain signed quantities like electric charge density.

2. Given a measure space $(X, \mathcal{M}, \mu)$ and a measurable function $f : X \to [0, \infty]$, we can define a new measure λ on X by
$$\lambda(E) = \int_E f \, d\mu$$
for all $E \in \mathcal{M}$, which we abbreviate by writing $\lambda = f \, d\mu$. (See Exercise 5.3.1.) If we simply require f to be real-valued rather than nonnegative, we no longer obtain a measure. However, $\lambda = f \, d\mu$ is a signed measure.

3. In certain situations, it can be useful to form linear combinations of measures on a space X. As a simple case, one may want to take two positive measures μ_1 and μ_2 and form
$$\nu = \mu_1 - \mu_2,$$
which will generally turn out to be a signed measure.

4. This application of signed measures is much deeper, and we will not encounter it for quite some time. Suppose X is a compact metric space and μ is a Borel measure on X. Define $\varphi : C(X) \to \mathbf{C}$ by
$$\varphi(f) = \int_X f \, d\mu.$$
The basic properties of the integral show that φ is a linear map from $C(X)$ to the complex numbers; it is an example of a *continuous linear functional.* Thus every Borel measure on X yields a continuous linear functional on $C(X)$. We can actually go the other way—*every* continuous linear functional on $C(X)$ is given by integration against a Radon measure (i.e., a Borel measure with certain regularity properties). This result is known as the Riesz Representation Theorem. In general, the Radon measure associated to a linear functional will need to be a signed measure (or perhaps even a complex-valued measure).

The second and third points above are the ones on which we will focus presently. Indeed, we will see that every signed measure can be expressed in one of these two ways. This is fortunate, since we can then use either characterization to easily define the integral of a function with respect to a signed measure.

Before moving on to our first major result, let us make a couple more comments about the third point. It is straightforward to check that if μ_1 and μ_2 are signed measures on $\mathcal{M} \subseteq \mathcal{P}(X)$, then the set function $\mu_1 + \mu_2$ defined by
$$(\mu_1 + \mu_2)(E) = \mu_1(E) + \mu_2(E)$$

for $E \in \mathcal{M}$ is again a signed measure. Likewise, given $\alpha \in \mathbf{R}$ we can define

$$(\alpha\mu_1)(E) = \alpha \cdot \mu_1(E),$$

and $\alpha\mu_1$ is a signed measure. In other words, we have observed the following result:

Proposition 5.6.2. *Let X be a set equipped with a σ-algebra $\mathcal{M} \subseteq \mathcal{P}(X)$. The set of all signed measures on $\mathcal{M}$ is a vector space over $\mathbf{R}$ with respect to the operations described above.*

As our first step toward proving the Radon-Nikodym theorem, we are about to establish that every signed measure can be written as the difference of two positive measures. This result is known as the Jordan Decomposition Theorem. First we discuss a version of this result at the level of measurable sets.

Definition 5.6.3. Let $(X, \mathcal{M})$ be a measurable space, and suppose $\nu : \mathcal{M} \to [-\infty, \infty]$ is a signed measure. We say a set $E \in \mathcal{M}$ is

- **positive** if $\nu(A) \geq 0$ for all $A \in \mathcal{M}$ such that $A \subseteq E$;
- **negative** if $\nu(A) \leq 0$ for all $A \in \mathcal{M}$ such that $A \subseteq E$;
- **null** if $\nu(A) = 0$ for all $A \in \mathcal{M}$ such that $A \subseteq E$.

Of course our goal its to decompose X as the disjoint union of a positive set and a negative set, which will in turn yield the desired decomposition of ν into a difference of positive measures. First recall that if $A, B \subseteq X$, the **symmetric difference** of A and B is defined by

$$A \triangle B = (A \backslash B) \cup (B \backslash A).$$

Theorem 5.6.4 (Hahn Decomposition Theorem)**.** *Let $(X, \mathcal{M})$ be a measurable space and $\nu : \mathcal{M} \to [-\infty, \infty]$ a signed measure on X. Then there exist disjoint sets $P, N \in \mathcal{M}$ such that P is positive, N is negative, and $X = P \cup N$. Moreover, if $X = P' \cup N'$ is any other such decomposition, then $P \triangle P' = N \triangle N'$ is a ν-null set.*

In order to prove the theorem, we need two short lemmas. The first one provides a version of continuity of measure for signed measures, and the proof is nearly identical to the one for positive measures.

Lemma 5.6.5. *Let $(X, \mathcal{M})$ be a measurable space and ν a signed measure on $\mathcal{M}$.*

1. *Suppose $\{E_j\}_{j=1}^{\infty}$ is an increasing sequence of sets in $\mathcal{M}$. If we put $E = \bigcup_{j=1}^{\infty} E_j$, then*
$$\nu(E) = \lim_{j\to\infty} \nu(E_j).$$

2. *Suppose $\{E_j\}_{j=1}^{\infty}$ is a decreasing sequence of sets in $\mathcal{M}$ such that $\nu(E_1) \neq \pm\infty$, and let $E = \bigcap_{j=1}^{\infty} E_j$. Then*
$$\nu(E) = \lim_{j\to\infty} \nu(E_j).$$

The second result establishes two seemingly obvious facts about positive sets. Indeed, both follow quickly from the definition.

Lemma 5.6.6. *Let $(X, \mathcal{M})$ be a measurable space, and suppose ν a signed measure on $\mathcal{M}$.*

1. *Let $E \in \mathcal{M}$ be a positive set. Then any set $A \in \mathcal{M}$ satisfying $A \subseteq E$ is a positive set.*

2. *Suppose $\{E_j\}_{j=1}^{\infty}$ is a countable family of positive sets in $\mathcal{M}$. Then $\bigcup_{j=1}^{\infty} E_j$ is a positive set.*

Proof. The first assertion is obvious from the definition of a positive set. For the second, suppose $A \subseteq \bigcup_{j=1}^{\infty} E_j$ with $A \in \mathcal{M}$. For each $m \in \mathbf{N}$, define

$$A_m = \bigcup_{j=1}^{m} A \cap E_j.$$

Since each E_j is a positive set, $\nu(A \cap E_j) \geq 0$ for all j. Moreover, $A = \bigcup_{m=1}^{\infty} A_m$ and this union is increasing, so

$$\nu(A) = \lim_{m\to\infty} \nu(A_m) \geq 0$$

by continuity of measure. It follows that $\bigcup_{j=1}^{\infty} E_j$ is a positive set. $\square$

Proof of Theorem 5.6.4. By replacing ν with $-\nu$ if necessary, we may assume that ν omits the value ∞. Set

$$\alpha = \sup\{\nu(E) : E \in \mathcal{M} \text{ is positive}\}.$$

Then there is a sequence $\{E_j\}_{j=1}^{\infty}$ of positive sets for which $\nu(E_j) \to \alpha$ as $j \to \infty$. If we let $P = \bigcup_{j=1}^{\infty} E_j$, then P is positive and $\nu(P) = \alpha$ by our previous two lemmas. Notice that we have shown $\alpha < \infty$.

Now let $N = P^c$. We claim that N is a negative set. First we will show that if $Q \in \mathcal{M}$ is a positive set with $Q \subseteq N$, then $\nu(Q) = 0$. Well, if $\nu(Q) > 0$, then $P \cup Q$ is a positive set and

$$\nu(P \cup Q) = \nu(P) + \nu(Q) > \alpha,$$

which contradicts the definition of α from above. Thus Q must be a null set.

Next we claim that if $A \subseteq N$ and $\nu(A) > 0$, then there exists a set $B \subseteq A$ with $\nu(B) > \nu(A)$. Since A is not a positive set, it must contain a subset C with $\nu(C) < 0$. If we then set $B = A \backslash C$, we have

$$\nu(A) = \nu(B) + \nu(C),$$

so $\nu(B) > \nu(A)$.

Now we assume N is not a negative set and derive a contradiction. Since N contains a set of positive measure, there is a smallest integer n_1 and a set $A_1 \subseteq N$ such that

$$\nu(A_1) > \frac{1}{n_1}.$$

Now let n_2 be the smallest integer such that A_1 contains a set of measure at least $\nu(A_1) + \frac{1}{n_2}$ (which we can do by the previous claim), and let A_2 be such a set. Thus we can proceed inductively and find a sequence of sets

$$A_1 \supseteq A_2 \supseteq A_3 \supseteq \cdots$$

such that

$$\nu(A_j) > \nu(A_{j-1}) + \frac{1}{n_j}$$

for all j. Put $A = \bigcap_{j=1}^{\infty} A_j$. Then

$$\nu(A) = \lim_{j \to \infty} \nu(A_j) > \sum_{j=1}^{\infty} \frac{1}{n_j}.$$

Since $\nu(A) < \infty$, the sum on the right hand side converges, meaning $n_j \to \infty$ as $j \to \infty$. But $\nu(A) > 0$, so we can find a natural number n and a set $B \subseteq A$ with $\nu(B) > \nu(A) + \frac{1}{n}$. Since $n_j \to \infty$, we have $n_j > n$ for sufficiently large j. Then notice that $B \subseteq A_{j-1}$ and

$$\nu(B) > \nu(A) + \frac{1}{n} \geq \nu(A_{j-1}) + \frac{1}{n},$$

contradicting the definition of n_j. Therefore, our assumption that N is not a negative set must not be valid. Thus we can express $X = P \cup N$ with P positive and N negative.

Finally, suppose $X = P' \cup N'$ is another such decomposition of X. Then $P \backslash P'$ is a subset of P, hence it is a positive set. However, we also have $P \backslash P' \subseteq N'$, since $P' \cap N' = \emptyset$. Thus $P \backslash P'$ is also a negative set, so it must be the case that $P \backslash P'$ is ν-null. A similar argument shows that $P' \backslash P$ is ν-null, whence $P \triangle P'$ is ν-null. $\square$

Any such decomposition of X into the disjoint union of a positive set and a negative set is aptly called a **Hahn decomposition** of X. As we saw at the end of the proof, such a decomposition is not generally unique, though differences can only arise through ν-null sets.

We can now use the Hahn Decomposition Theorem to establish the promised decomposition result for signed measures. In order to state that result, we need a couple more definitions.

Definition 5.6.7. Let $(X, \mathcal{M})$ be a measurable space, and suppose μ and ν are signed measures on $\mathcal{M}$.

1. Let $A \in \mathcal{M}$. We say μ is **concentrated** on A if we have

$$\mu(E) = \mu(E \cap A)$$

for all $E \in \mathcal{M}$. Equivalently, we have $\mu(E) = 0$ whenever $E \cap A = \emptyset$.

2. We say μ and ν are **mutually singular**, written

$$\mu \perp \nu,$$

if there exist disjoint sets $A, B \in \mathcal{M}$ such that μ is concentrated on A and ν is concentrated on B.

Loosely speaking, two signed measures μ and ν are mutually singular if we can decompose X into two corresponding disjoint sets that do not interact in a measure-theoretic sense. This is exactly the same sort of phenomenon that we observed in the Hahn Decomposition Theorem.

Theorem 5.6.8 (Jordan Decomposition Theorem)**.** *Let $(X, \mathcal{M})$ be a measurable space, and suppose ν is a signed measure on $\mathcal{M}$. Then there exist unique positive measures ν^+ and ν^- on $\mathcal{M}$ such that $\nu = \nu^+ - \nu^-$ and $\nu^+ \perp \nu^-$.*

Proof. Let $X = P \cup N$ be a Hahn decomposition of X associated to ν. For each $E \in \mathcal{M}$, we set

$$\nu^+(E) = \nu(E \cap P), \quad \nu^-(E) = -\nu(E \cap N).$$

Then clearly ν^+ and ν^- are both positive measures, since they arise from restricting ν to the positive and negative sets P and N, respectively. Furthermore, we have

$$\nu(E) = \nu^+(E) - \nu^-(E)$$

for all $E \in \mathcal{M}$ since $P \cap N = \emptyset$. Finally, it is easy to see that ν^+ is concentrated on P and ν^- is concentrated on N, so $\nu^+ \perp \nu^-$.

Now suppose $\nu = \mu^+ - \mu^-$ is another such decomposition. Since $\mu^+ \perp \mu^-$, we can find disjoint sets $A, B \in \mathcal{M}$ such that $X = A \cup B$ and μ^+ and μ^- are concentrated on A and B, respectively. But then A is positive and B is negative for ν, so we have found another Hahn decomposition for ν. In particular, $P \triangle A$ is ν-null, so

$$\mu^+(E) = \mu^+(E \cap A) = \nu^+(E \cap A) = \nu^+(E \cap P) = \nu^+(E)$$

for all $E \in \mathcal{M}$. Similarly, $\mu^- = \nu^-$. □

The decomposition of a signed measure ν guaranteed by the previous theorem is called the **Jordan decomposition** of ν. The positive measures ν^+ and ν^- appearing in the Jordan decomposition are respectively called the **positive** and **negative variations** of ν.

The Jordan decomposition gives one a natural way of defining integration with respect to a signed measure. In particular, if ν is a signed measure on a σ-algebra $\mathcal{M} \subseteq \mathcal{P}(X)$ and $f : X \to \mathbf{C}$ is measurable with respect to $\mathcal{M}$, we define

$$\int_X f \, d\nu = \int_X f \, d\nu^+ - \int_X f \, d\nu^-.$$

Notice that this integral is well-defined whenever f belongs to both $L^1(X, \nu^+)$ and $L^1(X, \nu^-)$. Consequently, we define the ν-integrable functions by setting

$$L^1(X, \nu) = L^1(X, \nu^+) \cap L^1(X, \nu^-).$$

Eventually we will see that the Radon-Nikodym theorem yields another way of defining integration with respect to a signed measure.

Before proceeding any further, we need to observe that to any signed measure ν we can naturally associate a positive measure $|\nu|$ that plays the role of the absolute value for measures.

Definition 5.6.9. Suppose ν is a signed measure. The positive measure $|\nu|$ defined by

$$|\nu| = \nu^+ + \nu^-$$

is called the **total variation** of ν.

Remark 5.6.10. It is not generally the case that $|\nu|(E) = |\nu(E)|$ for all $E \in \mathcal{M}$. For example, let $X = [-1, 1]$, let $\mathcal{M}$ be the Borel σ-algebra on X, and consider the signed measure ν defined on $\mathcal{M}$ by

$$\nu(E) = \int_X x \, d\mu(x),$$

where μ denotes the Lebesgue measure on $[-1, 1]$. Then

$$\nu([-\tfrac{1}{2}, 1]) = \tfrac{3}{8},$$

while

$$|\nu|([-\tfrac{1}{2}, 1]) = \nu^+([0, 1]) + \nu^-([-\tfrac{1}{2}, 0]) = \tfrac{5}{8}.$$

However, it is easy to check that the inequality $|\nu(E)| \le |\nu|(E)$ always holds.

By examining the proof of the Hahn Decomposition Theorem, one can see that if ν omits the value ∞, then ν^+ is actually bounded. Similarly, if ν does not take the value $-\infty$, then ν^- is bounded. Therefore, a signed measure that omits infinite values is necessarily bounded. In particular, $|\nu|(X)$ is finite. It is then a straightforward exercise to verify the following result.

Proposition 5.6.11. *Let $(X, \mathcal{M})$ be a measurable space. The set of all finite signed measures on $\mathcal{M}$ is a normed vector space, with the norm given by*

$$\|\nu\| = |\nu|(X).$$

We close out this section with a discussion of measures taking complex values. This will allow us to achieve even further generality in the upcoming Radon-Nikodym theorem.

Definition 5.6.12. Let $(X, \mathcal{M})$ be a measurable space. A set function $\nu : \mathcal{M} \to \mathbf{C}$ is called a **complex measure** if it satisfies the following conditions:

1. $\nu(\emptyset) = 0$.
2. If $\{E_j\}_{j=1}^{\infty}$ is any countable collection of pairwise disjoint sets in $\mathcal{M}$ and we let $E = \bigcup_{j=1}^{\infty} E_j$, then

$$\nu(E) = \sum_{j=1}^{\infty} \nu(E_j),$$

where we require that the sum on the right hand side to converge absolutely.

Notice that a complex measure is forbidden to assume infinite values by definition. In particular, any *finite* signed (or simply positive) measure is a complex measure. We also have two natural ways of constructing complex measures. First, if ν_1 and ν_2 are finite signed measures, then

$$\nu = \nu_1 + i\nu_2$$

is a complex measure. It should not be at all shocking that any complex measure can be written in terms of signed measures in such a way, nor that integration against such a measure is defined by

$$\int_X f\, d\nu = \int_X f\, d\nu_1 + i \int_X f\, d\nu_2$$

whenever $f \in L^1(X, \nu_1) \cap L^1(X, \nu_2)$. Another more interesting way of constructing a complex measure is by starting with a positive measure μ and a function $f \in L^1(X, \mu)$, and defining

$$\nu(E) = \int_E f\, d\mu,$$

just as we did for signed measures. We will soon see that every complex measure arises this way, thanks to the Radon-Nikodym theorem.

Some of the other concepts we have for signed measures also generalize to complex measures. For one, we declare two complex measures μ and ν to be **mutually singular** (again written $\mu \perp \nu$) if there exist disjoint sets $A, B \in \mathcal{M}$ such that μ is concentrated on A and ν is concentrated on B. Equivalently:

Proposition 5.6.13. *Suppose $\mu = \mu_1 + i\mu_2$ and $\nu = \nu_1 + i\nu_2$, where μ_1, μ_2, ν_1, and ν_2 are signed measures. We have $\mu \perp \nu$ if and only if $\mu_1 \perp \nu_1$, $\mu_2 \perp \nu_2$, $\mu_1 \perp \nu_2$, and $\mu_2 \perp \nu_1$.*

Proof. Exercise 5.6.4. □

It is also possible to extend the definition of the total variation to a complex measure. Unfortunately, we do not have the Jordan decomposition at our disposal, so the definition is more complicated than the one for signed measures.

Definition 5.6.14. Let $(X, \mathcal{M})$ be a measurable space and $\nu : \mathcal{M} \to \mathbf{C}$ a complex measure. Given $E \in \mathcal{M}$, we define

$$|\nu|(E) = \sup \left\{ \sum_{j=1}^{\infty} |\nu(E_j)| : E_1, E_2, \ldots \in \mathcal{M} \text{ are disjoint, } E = \bigcup_{j=1}^{\infty} E_j \right\}.$$

The set function $|\nu| : \mathcal{M} \to \mathbf{R}$ is called the **total variation** of ν.

Notice that it is not immediately clear that $|\nu|$ even defines a measure on $\mathcal{M}$. We will prove this fact now.

Theorem 5.6.15. *The total variation $|\nu|$ is a positive measure on $\mathcal{M}$.*

Proof. Clearly $|\nu|$ is nonnegative and $|\nu|(\emptyset) = 0$, so the only issue is countable additivity. Suppose $\{E_j\}_{j=1}^{\infty}$ is a collection of pairwise disjoint sets in $\mathcal{M}$, and put $E = \bigcup_{j=1}^{\infty} E_j$. The by the definition of $|\nu|$, we clearly have

$$\sum_{j=1}^{\infty} |\nu(E_j)| \le |\nu|(E).$$

Now let $\{A_k\}_{k=1}^{\infty}$ be another collection of disjoint sets in $\mathcal{M}$ with $E = \bigcup_{k=1}^{\infty} A_k$. Then for each k we have

$$|\nu(A_k)| = \left|\sum_{j=1}^{\infty} \nu(A_k \cap E_j)\right| \le \sum_{j=1}^{\infty} |\nu(A_k \cap E_j)|,$$

so

$$\sum_{k=1}^{\infty} |\nu(A_k)| \le \sum_{k=1}^{\infty}\sum_{j=1}^{\infty} |\nu(A_k \cap E_j)| = \sum_{j=1}^{\infty}\sum_{k=1}^{\infty} |\nu(A_k \cap E_j)|.$$

Since the family $\{A_k \cap E_j\}_{k=1}^{\infty}$ is a partition of E_j, we have

$$\sum_{k=1}^{\infty} |\nu(A_k \cap E_j)| \le |\nu|(E_j).$$

Therefore,

$$\sum_{k=1}^{\infty} |\nu(A_k)| \le \sum_{j=1}^{\infty} |\nu|(E_j).$$

Since $\{A_k\}_{k=1}^{\infty}$ is an arbitrary partition of E, it follows that

$$|\nu|(E) \le \sum_{j=1}^{\infty} |\nu|(E_j).$$

To prove the reverse inequality, we first choose $t_j \in \mathbf{R}$ such that $|\nu|(E_j) > t_j$ for each j. Then for each j we can find a partition $\{A_{j,k}\}_{k=1}^{\infty}$ of E_j such that

$$\sum_{k=1}^{\infty} |\nu(A_{j,k})| > t_j.$$

But then $E = \bigcup_{j,k} A_{j,k}$ so we have

$$|\nu|(E) \geq \sum_{j=1}^{\infty}\sum_{k=1}^{\infty} |\nu(A_{j,k})| \geq \sum_{j=1}^{\infty} t_j.$$

Now take the supremum over all possible choices of t_j to obtain

$$\sum_{j=1}^{\infty} |\nu|(E_j) \leq |\nu|(E).$$

Thus $|\nu|$ is a measure. □

We close this section by checking that the two definitions of total variation agree when ν is a finite signed measure.

Proposition 5.6.16. *Suppose ν is a signed measure. Then the total variation $|\nu|$ defined as above is equal to $\nu^+ + \nu^-$.*

Proof. Let $E \in \mathcal{M}$, and let $\{E_j\}_{j=1}^{\infty} \subseteq \mathcal{M}$ be a partition of E. Then

$$\sum_{j=1}^{\infty} |\nu(E_j)| \leq \sum_{j=1}^{\infty} \big[\nu^+(E_j) + \nu^-(E_j)\big] = \nu^+(E) + \nu^-(E),$$

and taking the supremum over all partitions of E yields

$$|\nu|(E) \leq \nu^+(E) + \nu^-(E).$$

Now let $X = P \cup N$ be a Hahn decomposition for ν. Then $E \cap P$ and $E \cap N$ partition E, so

$$\nu^+(E) + \nu^-(E) = |\nu(E \cap P)| + |\nu(E \cap N)| \leq |\nu|(E).$$

Thus $|\nu|(E) = \nu^+(E) + \nu^-(E)$. □

Exercises for Section 5.6

Exercise 5.6.1. Let X be a set, $\mathcal{M} \subseteq \mathcal{P}(X)$ a σ-algebra, and suppose ν, μ, and λ are signed measures on $\mathcal{M}$. Prove the following assertions.

(a) If ν is concentrated on a set $E \in \mathcal{M}$, then so is $|\nu|$.

(b) If $\nu \perp \mu$, then $|\nu| \perp |\mu|$.

(c) If $\nu \perp \lambda$ and $\mu \perp \lambda$, then $\nu + \mu \perp \lambda$.

Exercise 5.6.2. Prove Proposition 5.6.11.

Exercise 5.6.3. Let ν be a signed measure on a σ-algebra $\mathcal{M} \subseteq \mathcal{P}(X)$. Prove that $L^1(X, \nu) = L^1(X, |\nu|)$.

Exercise 5.6.4. Prove Proposition 5.6.13.

Exercise 5.6.5. Let ν be a complex measure on a σ-algebra $\mathcal{M} \subseteq \mathcal{P}(X)$. Prove that $|\nu(E)| \leq |\nu|(E)$ for all $E \in \mathcal{M}$.

5.7 The Radon-Nikodym Theorem

As we foreshadowed in the last section, we intend to prove a major theorem regarding integration with respect to two closely related measures on the same space. In particular, we would like to transform integration with respect to a signed or complex measure into integration against some related *positive* measure. Before we can even state this theorem, we need to make the term "closely related" much more precise.

Definition 5.7.1. Let $(X, \mathcal{M})$ be a measurable space, and suppose ν is a signed or complex measure on $\mathcal{M}$ and μ is a positive measure on $\mathcal{M}$. We say that ν is **absolutely continuous** with respect to μ, written

$$\nu \ll \mu,$$

if we have $\nu(E) = 0$ for all $E \in \mathcal{M}$ such that $\mu(E) = 0$.

One might wonder how common absolutely continuous measures are. As it turns out, it is quite easy to construct signed or complex measures that are absolutely continuous with respect to a prescribed positive measure μ.

Example 5.7.2. Let $(X, \mathcal{M}, \mu)$ be a measure space, and let $f \in L^1(X, \mu)$. Recall that we can define a complex measure ν on $\mathcal{M}$ by setting

$$\nu(E) = \int_E f \, d\mu$$

for all $E \in \mathcal{M}$. It is easy to see that ν is absolutely continuous with respect to μ: if $\mu(E) = 0$, then $\int_E f \, d\mu = 0$, so $\nu(E) = 0$. More surprisingly, the Radon-Nikodym theorem will show that *every* complex measure on $\mathcal{M}$ that is absolutely continuous with respect to μ arises in this way.

Since we will have much occasion to talk about such measures in the near future, we will introduce some notation for measures defined like the one from the previous example. Given a positive measure μ and a function $f \in L^1(X, \mu)$, we write

$$d\nu = f \, d\mu$$

for the measure defined by
$$\nu(E) = \int_E f \, d\mu$$
for all $E \in \mathcal{M}$.

Now we will look at an example of two measures that are not absolutely continuous with respect to one another. We will see that the notion of mutual singularity ends up playing a role.

Example 5.7.3. Let $(\mathbf{R}, \mathcal{L}, \mu)$ denote the usual Lebesgue measure space, and let ν denote the Dirac measure on $\mathbf{R}$ concentrated at 0. Then notice that $\mu(\{0\}) = 0$, while $\nu(\{0\}) = 1$. Thus ν is not absolutely continuous with respect to μ. In fact, $\nu \perp \mu$ since ν is concentrated on the singleton $\{0\}$, while μ can be thought of as being concentrated on $\mathbf{R}\backslash\{0\}$.

It is not hard to check (see Exercise 5.7.1) that the notion of absolute continuity is roughly the antithesis of mutual singularity.

Proposition 5.7.4. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose ν is a signed or complex measure on $\mathcal{M}$. If $\nu \ll \mu$ and $\nu \perp \mu$, then $\nu = 0$.*

It is also a worthwhile exercise (see Exercise 5.7.2) to check the assertions in the following proposition regarding absolute continuity.

Proposition 5.7.5. *Let $(X, \mathcal{M})$ be a measurable space, and suppose μ is a positive measure on $\mathcal{M}$.*

1. *If ν is a signed measure on $\mathcal{M}$, then $\nu \ll \mu$ if and only if $\nu^+ \ll \mu$ and $\nu^- \ll \mu$.*
2. *If ν is a complex measure on $\mathcal{M}$, then $\nu \ll \mu$ if and only if its real and imaginary parts are absolutely continuous with respect to μ.*
3. *If ν is a signed or complex measure on $\mathcal{M}$, then $\nu \ll \mu$ if and only if $|\nu| \ll \mu$.*

It is worth noting that the term "absolutely continuous" is not unwarranted—there is an "epsilon-delta" definition of absolute continuity that makes the name seem quite apt.

Theorem 5.7.6. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose ν is a complex measure on $\mathcal{M}$. Then $\nu \ll \mu$ if and only if for every $\varepsilon > 0$ there exists $\delta > 0$ such that $\mu(E) < \delta$ implies $|\nu(E)| < \varepsilon$.*

Proof. Without loss of generality, we may assume that ν is positive. (Otherwise, we could instead consider $|\nu|$, since $\nu \ll \mu$ if and only if $|\nu| \ll \mu$ and $|\nu(E)| \leq |\nu|(E)$ for all $E \in \mathcal{M}$.) Suppose the ε-δ assertion holds, and let $E \in \mathcal{M}$ with $\mu(E) = 0$. Then given $\varepsilon > 0$, it is immediate that $\nu(E) < \varepsilon$, whence $\nu(E) = 0$. Thus $\nu \ll \mu$.

On the other hand, suppose the ε-δ condition fails to hold. Then there exists an $\varepsilon > 0$ such that for all $n \in \mathbf{N}$ there is a set $E_n \in \mathcal{M}$ satisfying $\mu(E_n) < \frac{1}{2^n}$ but $\nu(E_n) \geq \varepsilon$. For each $m \in \mathbf{N}$, set $F_m = \bigcup_{n=m}^{\infty} E_n$, and define $F = \bigcap_{m=1}^{\infty} F_m$. Then observe that

$$\mu(F_m) \leq \sum_{n=m}^{\infty} \mu(E_n) \leq \sum_{m=1}^{\infty} \frac{1}{2^n} = \frac{1}{2^{m-1}},$$

so $\mu(F_m) \to 0$ as $m \to \infty$. It then follows from continuity of measure that $\mu(F) = 0$. However, $\nu(F_m) \geq \varepsilon$ for all m, so $\nu(F) \geq \varepsilon$ as well. Thus ν is not absolutely continuous with respect to μ. □

Now we digress for a moment to say a little more in the situation where ν and μ are both positive measures.

Definition 5.7.7. Suppose μ and ν are both positive measures defined on the same σ-algebra $\mathcal{M}$. We say that μ and ν are **mutually absolutely continuous** if $\nu \ll \mu$ and $\mu \ll \nu$.

Mutually absolutely continuous measures are sometimes said to be **equivalent**. As an instructive example, we show that any σ-finite measure is equivalent to a *finite* measure.

Theorem 5.7.8. *Suppose $(X, \mathcal{M}, \mu)$ is a measure space, where μ is σ-finite. Then there is a function $f \in L^1(X, \mu)$ such that $0 < f(x) < 1$ for all $x \in X$. Consequently, the finite measure*

$$d\nu = f\, d\mu$$

is mutually absolutely continuous with μ.

Proof. Since μ is σ-finite, we can write $X = \bigcup_{j=1}^{\infty} E_j$ with the E_j pairwise disjoint and $\mu(E_j) < \infty$ for all j. Define a sequence of functions $(f_n)_{n=1}^{\infty}$ on X by

$$f_n = \frac{1}{2^n(1 + \mu(E_n))} \cdot \chi_{E_n},$$

and put $f = \sum_{n=1}^{\infty} f_n$. Then for each n, $0 < f_n(x) < 1$ for all $x \in X$. Since the E_n are pairwise disjoint, it follows that $0 < f(x) < 1$ for all $x \in X$. Furthermore, the Monotone Convergence Theorem guarantees that

$$\int_X f\, d\mu = \sum_{n=1}^{\infty} \int_X f_n\, d\mu = \sum_{n=1}^{\infty} \frac{1}{2^n(1 + \mu(E_n))} \cdot \mu(E_n) \leq \sum_{n=1}^{\infty} \frac{1}{2^n} = 1,$$

so $f \in L^1(X, \mu)$.

If ν is the measure defined by $d\nu = f\,d\mu$, then we have already argued that $\nu \ll \mu$. On the other hand, if

$$\nu(E) = \int_X f \cdot \chi_E \, d\mu = 0$$

for some $E \in \mathcal{M}$, then it must be the case that $f \cdot \chi_E = 0$ μ-almost everywhere. But $f(x) \neq 0$ for all x, so the only possibility is that $\mu(E) = 0$. Hence ν and μ are mutually absolutely continuous. □

Now we arrive at the promised Radon-Nikodym theorem. The theorem actually comes in two parts—the first is sometimes called the "Lebesgue-Radon-Nikodym" theorem, while the second is the more specialized result that we have been referencing throughout our discussion.

Theorem 5.7.9 (Lebesgue-Radon-Nikodym). *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space, and suppose ν is a complex measure on $\mathcal{M}$.*

1. *There exist unique complex measures ν_a and ν_s on $\mathcal{M}$ such that $\nu_a \ll \mu$, $\nu_s \perp \mu$, and*
$$\nu = \nu_a + \nu_s.$$
Furthermore, if ν is a finite signed (or positive) measure, then so are ν_a and ν_s.

2. *There is a function $f \in L^1(X, \mu)$ such that*
$$\nu_a(E) = \int_E f \, d\mu$$
for all $E \in \mathcal{M}$. (In other words, $d\nu_a = f\,d\mu$.) Furthermore, any other such function agrees with f μ-almost everywhere.

Before we can start the proof, we need a small lemma regarding mutually singular measures.

Lemma 5.7.10. *Suppose μ and ν are finite positive measures on a σ-algebra $\mathcal{M}$. Either $\nu \perp \mu$, or there exists an $\varepsilon > 0$ and a set $E \in \mathcal{M}$ such that $\mu(E) > 0$ and*

$$\nu(A) \geq \varepsilon \mu(A)$$

for all sets $A \in \mathcal{M}$ with $A \subseteq E$.

Proof. For each $n \in \mathbf{N}$, we consider the signed measure $\nu - \frac{1}{n}\mu$, and we let $X = P_n \cup N_n$ be a Hahn decomposition for this measure. Set

$$P = \bigcup_{n=1}^{\infty} P_n$$

and $N = P^c$. Then $N \subseteq N_n$ for all n, so N is a negative set for the measure $\nu - \frac{1}{n}\mu$ for all n. Consequently,

$$0 \le \nu(N) \le \tfrac{1}{n}\mu(N)$$

for all n, meaning that $\nu(N) = 0$. If we also have $\mu(P) = 0$, then $\nu \perp \mu$. On the other hand, $\mu(P) > 0$ implies that $\mu(P_n) > 0$ for some n. Since P_n is a positive set for $\nu - \frac{1}{n}\mu$, we have

$$\nu(P_n) \ge \tfrac{1}{n}\mu(P_n),$$

so we can take $E = P_n$ and $\varepsilon = \frac{1}{n}$. □

Proof of the Lebesgue-Radon-Nikodym theorem. We first assume that μ is finite and ν is positive. We first focus on constructing the "absolutely continuous part" ν_a. We prove this and assertion (2) in one fell swoop by finding an appropriate function $f \in L^1(X, \mu)$ and defining $d\nu_a = f\, d\mu$. To this end, let

$$\mathcal{F} = \left\{ f : X \to [0, \infty] : \int_E f\, d\mu \leq \nu(E) \text{ for all } E \in \mathcal{M} \right\}.$$

Notice that $0 \in \mathcal{F}$, so $\mathcal{F} \neq \emptyset$. Now suppose $f, g \in \mathcal{F}$, and let

$$A = \{x \in X : f(x) > g(x)\}.$$

If we let $h = \max\{f, g\}$, then for all $E \in \mathcal{M}$ we have

$$\begin{aligned}
\int_E h\, d\mu &= \int_{E \cap A} h\, d\mu + \int_{E \backslash A} h\, d\mu \\
&= \int_{E \cap A} f\, d\mu + \int_{E \backslash A} g\, d\mu \\
&\leq \nu(E \cap A) + \nu(E \backslash A) \\
&= \nu(E).
\end{aligned}$$

Hence $h \in \mathcal{F}$. Now put

$$\alpha = \sup \left\{ \int_X f\, d\mu : f \in \mathcal{F} \right\}.$$

Notice that $\alpha \leq \nu(X) < \infty$ by the definition of $\mathcal{F}$. Now find a sequence of functions $(f_n)_{n=1}^{\infty}$ in $\mathcal{F}$ satisfying $\int_X f_n\, d\mu \to \alpha$ as $n \to \infty$. For each $m \in \mathbf{N}$, define

$$g_m = \max\{f_1, f_2, \ldots, f_m\},$$

and set $f = \sup_n f_n$. Then $g_m \in \mathcal{F}$ for all m, clearly g_m increases pointwise to f, and

$$\int_X f_m\, d\mu \leq \int_X g_m\, d\mu$$

for all m. Consequently,

$$\alpha = \lim_{n \to \infty} \int_X f_n\, d\mu \leq \lim_{n \to \infty} \int_X g_n\, d\mu \leq \alpha,$$

so $\int_X g_n\, d\mu \to \alpha$ as $n \to \infty$. It follows from the Monotone Convergence Theorem that $\int_X f\, d\mu = \alpha$, and that

$$\int_E f\, d\mu = \lim_{n \to \infty} \int_E g_n\, d\mu \leq \nu(E)$$

for all $E \in \mathcal{M}$. Thus $f \in \mathcal{F}$ and $f \in L^1(X, \mu)$.

Define ν_a by $d\nu_a = f\,d\mu$. Then $\nu_a \ll \mu$ by previous arguments. If we set $\nu_s = \nu - \nu_a$, then we claim that $\nu_s \perp \mu$. Suppose, to the contrary, that ν_s is not singular with respect to μ. Then by the lemma there is an $\varepsilon > 0$ and a set $E \in \mathcal{M}$ such that $\mu(E) > 0$ and $\nu_s(A) \geq \varepsilon\mu(A)$ for all measurable sets $A \subseteq E$. If $A \in \mathcal{M}$ is any measurable set, then

$$\begin{aligned}\int_A f + \varepsilon\chi_E\,d\mu &= \int_A f\,d\mu + \varepsilon\int_A \chi_E\,d\mu \\ &= \nu_a(A) + \varepsilon\mu(A\cap E) \\ &\leq \nu_a(A) + \nu_s(A\cap E) \\ &\leq \nu_a(A) + \nu_s(A) \\ &= \nu(A),\end{aligned}$$

so the function $f + \varepsilon\chi_E$ belongs to $\mathcal{F}$. Moreover,

$$\int_X f + \varepsilon\chi_E\,d\mu = \int_X f\,d\mu + \varepsilon\int_X \chi_E\,d\mu = \alpha + \varepsilon\mu(E) > \alpha,$$

which contradicts the definition of α. Therefore, it must be the case that $\nu_s \perp \mu$.

Before generalizing the proof, we show that ν_a and ν_s are unique. Suppose $\nu = \nu_a' + \nu_s'$ is another decomposition, with $\nu_a' \ll \mu$ and $\nu_s' \perp \mu$. Then

$$\nu_a - \nu_a' = \nu_s' - \nu_s,$$

where $\nu_a - \nu_a' \ll \mu$ and $\nu_s' - \nu_s \perp \mu$. Therefore, $\nu_a - \nu_a' = \nu_s' - \nu_s = 0$. Similarly, if there is another function g satisfying $d\nu_a = g\,d\mu$, then we have

$$\int_E f\,d\mu = \int_E g\,d\mu$$

for all $E \in \mathcal{M}$, which implies that $f = g$ μ-almost everywhere.

Now suppose ν is a finite signed measure, and let $\nu = \nu^+ - \nu^-$ be its Jordan decomposition. Then by our previous work, we have decompositions $\nu^+ = \nu_a^+ + \nu_s^+$ and $\nu^- = \nu_a^- + \nu_s^-$ with $\nu_a^+, \nu_a^- \ll \mu$ and $\nu_s^+, \nu_s^- \perp \mu$. Thus

$$\nu_a^+ - \nu_a^- \ll \mu, \quad \nu_s^+ - \nu_s^- \perp \mu,$$

and we have the appropriate decomposition of ν:

$$\nu_a = \nu_a^+ - \nu_a^-, \quad \nu_s = \nu_s^+ - \nu_s^-.$$

The function f is found similarly—we can find $f^+, f^- \in L^1(X,\mu)$ such that $d\nu_a^+ = f^+\,d\mu$ and $d\nu_a^- = f^-\,d\mu$, and putting $f = f^+ - f^-$ gives

$$d\nu_a = f\,d\mu.$$

Indeed, if $E \in \mathcal{M}$, then we have

$$\nu_a(E) = \nu_a^+(E) - \nu_a^-(E) = \int_E f^+ \, d\mu - \int_E f^- \, d\mu = \int_E f^+ - f^- \, d\mu = \int_E f \, d\mu.$$

We can work similarly to pass to the case where ν is a complex measure—simply decompose the real and imaginary parts of ν and combine the results together in the appropriate way.

Finally, we assume that μ is σ-finite. By Theorem 5.7.8, there exists a *finite* measure μ' on $\mathcal{M}$ that is mutually absolutely continuous with μ. Applying what we have already proven, there exist complex measures ν_a and ν_s such that $\nu_a \ll \mu'$, $\nu_s \perp \mu'$, and $\nu = \nu_a + \nu_s$. However, it then follows that $\nu_a \ll \mu$ and $\nu_s \perp \mu$. Furthermore, there exists $f \in L^1(X, \mu')$ such that

$$d\nu_a = f \, d\mu'.$$

Let $g \in L^1(X, \mu)$ be a function satisfying $0 < g(x) < 1$ for all $x \in X$ and $d\mu' = g \, d\mu$. Then $fg \in L^1(X, \mu)$ and

$$\nu_a(E) = \int_E f \, d\mu' = \int_E fg \, d\mu$$

for all $E \in \mathcal{M}$. Thus $d\nu_a = fg \, d\mu$. □

It is worth noting that the Lebesgue-Radon-Nikodym theorem still holds for more general signed measures, provided that we also assume ν is σ-finite. The only changes are that ν_a and ν_s are now σ-finite signed measures, and the function f satisfying $d\nu_a = f \, d\mu$ does not necessarily belong to $L^1(X, \mu)$. Instead, we can only say that f is an **extended μ-integrable** function, meaning either $\int_X f^+ \, d\mu$ or $\int_X f^- \, d\mu$ is finite. In fact, the proof shows that f is "locally" integrable.

Theorem 5.7.11 (Lebesgue-Radon-Nikodym for σ-finite signed measures)**.** *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space, and suppose ν is a σ-finite signed measure on $\mathcal{M}$.*

1. *There exist unique σ-finite signed measures ν_a and ν_s on $\mathcal{M}$ such that $\nu_a \ll \mu$, $\nu_s \perp \mu$, and*
$$\nu = \nu_a + \nu_s.$$

2. *There is an extended μ-integrable function $f : X \to \mathbf{R}$ such that*
$$\nu_a(E) = \int_E f \, d\mu$$
for all $E \in \mathcal{M}$. (In other words, $d\nu_a = f \, d\mu$.) Furthermore, any other such function agrees with f μ-almost everywhere.

Proof. Assume first that ν is positive. Write $X = \bigcup_{j=1}^{\infty} E_j$ with $E_j \in \mathcal{M}$ and $|\nu(E_j)| < \infty$ for all j. Let ν_j denote the restriction of ν to E_j; that is, ν_j is the measure on $\mathcal{M}$ defined by

$$\nu_j(A) = \nu_j(A \cap E_j).$$

Notice that ν_j is a finite measure that is concentrated on E_j by definition. Then by the previous version of the Radon-Nikodym theorem, there exist finite measures $\nu_{a,j}$ and $\nu_{s,j}$ on $\mathcal{M}$ such that $\nu_{a,j} \ll \mu$, $\nu_{s,j} \perp \mu$, and

$$\nu_j = \nu_{a,j} + \nu_{s,j}.$$

Furthermore, there exist functions $f_j \in L^1(X, \mu)$ such that $d\nu_{a,j} = f_j \, d\mu$. We may assume f_j vanishes outside E_j.

Set $\nu_a = \sum_{j=1}^{\infty} \nu_{a,j}$ and $\nu_s = \sum_{j=1}^{\infty} \nu_{s,j}$. Then by countable additivity, it is not hard to see that $\nu = \nu_a + \nu_s$: given $A \in \mathcal{M}$, we have

$$\nu(A) = \sum_{j=1}^{\infty} \nu(A \cap E_j) = \sum_{j=1}^{\infty} \nu_j(A) = \sum_{j=1}^{\infty} \nu_{a,j}(A) + \sum_{j=1}^{\infty} \nu_{s,j}(A).$$

Furthermore, we have $\nu_a \ll \mu$, since $\mu(A) = 0$ implies $\nu_{a,j}(A) = 0$ for all j, whence $\nu_a(A) = 0$. That $\nu_s \perp \mu$ follows from Exercise 5.7.5.

Now we set $f = \sum_{j=1}^{\infty} f_j$. Then the Monotone Convergence Theorem guarantees that for all $E \in \mathcal{M}$ we have

$$\int_E f \, d\mu = \sum_{j=1}^{\infty} \int_E f_j \, d\mu = \sum_{j=1}^{\infty} \nu_{a,j}(E) = \nu_a(E).$$

Thus $d\nu_a = f \, d\mu$.

Now suppose ν is a σ-finite signed measure. Then as in the proof of the Lebesgue-Radon-Nikodym theorem, we can work separately with the positive measures ν^+ and ν^- and add the resulting decompositions. The only thing worth verifying is the fact that f is extended μ-integrable. Let $X = P \cup N$ be a Hahn decomposition for ν. Then

$$\int_X f \, d\mu = \int_X f^+ \, d\mu - \int_X f^- \, d\mu = \nu^+(P) - \nu^-(N).$$

Since either $\nu^+(P)$ or $\nu^-(N)$ is finite, it follows that one of $\int_X f^+ \, d\mu$ or $\int_X f^- \, d\mu$ is finite. $\square$

The decomposition of a signed or complex measures into absolutely continuous and singular parts (with respect to a specified positive measure) is called the **Lebesgue decomposition**. Before moving on, we will look at one specific example of a Lebesgue decomposition.

Example 5.7.12. Let μ denote Lebesgue measure on $\mathbf{R}$, and define $F : \mathbf{R} \to \mathbf{R}$ by

$$F(x) = \begin{cases} x & \text{if } x < 0 \\ x+1 & \text{if } x \geq 0. \end{cases}$$

Let ν denote the Lebesgue-Stieltjes measure associated to F. Then ν is not absolutely continuous with respect to μ, since $\mu(\{0\}) = 0$ but $\nu(\{0\}) = 1$. These measures are not mutually singular, either. It should not be hard to convince oneself that

$$\nu_a = \mu, \quad \nu_s = \delta_0,$$

where δ_0 denotes the Dirac measure concentrated at 0. Indeed, we can check that $\nu = \mu + \delta_0$. Given a half-open interval $(a, b]$, we have

$$\nu((a,b]) = F(b) - F(a) = \begin{cases} b-a & \text{if } 0 \notin (a,b] \\ b+1-a & \text{if } 0 \in (a,b], \end{cases}$$

which agrees with $\mu((a,b]) + \delta_0((a,b])$. Since ν and $\mu + \delta_0$ agree on half-open intervals and they are σ-finite, they agree on Borel sets. Furthermore, we clearly have $\mu \ll \mu$ and $\delta_0 \perp \mu$, so we have found the Lebesgue decomposition of ν. It is also obvious that the constant function $f = 1$ is the function satisfying $d\nu_a = f\,d\mu$.

What if μ is a positive σ-finite measure, ν is either a complex measure or a σ-finite signed measure, and $\nu \ll \mu$? Since Lebesgue decompositions are unique, we could conclude that $\nu_a = \nu$ and $\nu_s = 0$. Thus we have the following corollary, which is what analysts often think of as the Radon-Nikodym theorem.

Corollary 5.7.13. *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space, and let ν be either a σ-finite signed measure or a complex measure on $\mathcal{M}$. If $\nu \ll \mu$, then there exists a function f such that*

$$\nu(E) = \int_E f\,d\mu$$

for all $E \in \mathcal{M}$.

Definition 5.7.14. Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space, and let ν be either a σ-finite signed measure or a complex measure on $\mathcal{M}$ satisfying $\nu \ll \mu$. Then any function f satisfying $d\nu = f\,d\mu$ is called a **Radon-Nikodym derivative** of ν with respect to μ, denoted

$$f = \frac{d\nu}{d\mu}.$$

Notice that the Radon-Nikodym derivative is not unique, though any two derivatives are necessarily equal μ-almost everywhere.

As we promised earlier, the Radon-Nikodym theorem provides us with a "change-of-variables" formula between measures, which is akin to the usual u-substitution formula from calculus.

Theorem 5.7.15. *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space, and let ν be either a σ-finite signed measure or a complex measure on $\mathcal{M}$ satisfying $\nu \ll \mu$. If $f \in L^1(X, \mu)$, then $f \cdot \frac{d\nu}{d\mu} \in L^1(X, \mu)$ and*

$$\int_X f\, d\nu = \int_X f \cdot \frac{d\nu}{d\mu}\, d\mu. \tag{5.4}$$

Proof. Assume first that ν is positive. If f is the characteristic function of a measurable set, then (5.4) holds simply by the construction of $\frac{d\nu}{d\mu}$. It then holds for simple functions by linearity, and our usual mantra regarding the Monotone Convergence Theorem implies that (5.4) holds for nonnegative functions. If $f \in L^1(X, \nu)$ is real-valued, then $\int_X f^+\, d\nu$ and $\int_X f^-\, d\nu$ are both finite, hence

$$\int_X f^+ \cdot \frac{d\nu}{d\mu}\, d\mu, \quad \int_X f^- \cdot \frac{d\nu}{d\mu}\, d\mu < \infty.$$

It follows that $f \cdot \frac{d\nu}{d\mu} \in L^1(X, \mu)$, and linearity again implies that (5.4) holds. The result then follows for all functions in $L^1(X, \nu)$ by considering real and imaginary parts.

Now suppose ν is a signed measure. Then the result follows by applying the previous work to ν^+ and ν^-, together with the observation that

$$\frac{d\nu}{d\mu} = \frac{d\nu^+}{d\mu} - \frac{d\nu^-}{d\mu}.$$

Similarly, we extend to the situation where ν is complex by considering the real and imaginary parts and invoking the fact that

$$\frac{d\nu}{d\mu} = \frac{d\nu_{\mathrm{Re}}}{d\mu} + i\frac{d\nu_{\mathrm{Im}}}{d\mu}. \qquad \square$$

Notice that if ν is a signed or complex measure and $|\nu|(E) = 0$ for some $E \in \mathcal{M}$, then $\nu(E) = 0$ since $|\nu(E)| \le |\nu|(E)$. Thus $\nu \ll |\nu|$. As a result, we obtain the following characterization of complex measures or σ-finite signed measures.

Theorem 5.7.16. *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space, and let ν be either a σ-finite signed measure or a complex measure on $\mathcal{M}$. Then $\nu \ll |\nu|$, and for all $f \in L^1(X, \nu) = L^1(X, |\nu|)$, we have*

$$\int_X f \, d\nu = \int_X f \cdot \frac{d\nu}{d|\nu|} \, d|\nu|.$$

Finally, we have the following corollary when ν and μ are mutually absolutely continuous positive measures.

Corollary 5.7.17. *Let X be a set, $\mathcal{M}$ a σ-algebra on X, and suppose μ and ν are σ-finite positive measures on $\mathcal{M}$. If μ and ν are mutually absolutely continuous, then*

$$\frac{d\mu}{d\nu} \cdot \frac{d\nu}{d\mu} = 1$$

almost everywhere.

Proof. For all $E \in \mathcal{M}$, we have

$$\int_E 1 \, d\nu = \nu(E) = \int_E \frac{d\nu}{d\mu} \, d\mu = \int_E \frac{d\nu}{d\mu} \cdot \frac{d\mu}{d\nu} \, d\nu$$

by the proof of Theorem 5.7.15. (In particular, (5.4) holds for all nonnegative measurable functions.) Since this holds for all measurable sets E, it must be the case that the integrands agree μ-almost everywhere (or equivalently, ν-almost everywhere). $\square$

5.7.1 An Application to Probability Theory

We now discuss an application of the Radon-Nikodym theorem to probability theory. In a graduate-level probability course, one learns to think of probability in terms of measure spaces. That is, we model a random process with a **probability space**, which is simply a measure space $(X, \mathcal{M}, \mu)$ satisfying $\mu(X) = 1$. The set X is called the **sample space**, while the σ-algebra $\mathcal{M}$ consists of the **events** that are allowed to occur. The **probability** that an event $E \in \mathcal{M}$ occurs is precisely $\mu(E)$.

Once one views a probability space as a special kind of measure space, many familiar concepts from probability can be couched in measure-theoretic terms. For example, a **random variable** is simply a measurable function $f : X \to \mathbf{R}$, and the **expected value** of a random variable is computed by averaging the function—that is, integrating it—over X:

$$\mathbf{E}(f) = \int_X f \, d\mu.$$

In probability theory, one would also like to "condition" on certain events, meaning that we want to compute the probability of an event under the assumption that some event (or collection of events) has occurred. We know from classical probability theory that the **conditional probability** of A given B is simply

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)} = \frac{\mu(A \cap B)}{\mu(B)},$$

which makes sense whenever $\mu(B) \neq 0$. Likewise, one defines the **conditional expectation** of a random variable f given an event E by averaging f over all possible outcomes in E:

$$\mathbf{E}(f \mid E) = \frac{1}{\mu(E)} \int_E f \, d\mu.$$

Again, this definition is meaningless if $\mu(E) = 0$. (Go look up the Borel-Kolmogorov paradox to see what can go wrong when one conditions on events of probability zero.)

We can get around these issues by conditioning on a σ-algebra instead of a particular event. Suppose $\mathcal{N}$ is a sub-σ-algebra of $\mathcal{M}$, and let ν denote the restriction of μ to $\mathcal{N}$. Given an integrable random variable $f : X \to \mathbf{R}$, we are tempted to define $\mathbf{E}(f \mid \mathcal{N})$ by "locally averaging" f over the events in $\mathcal{N}$:

$$\mathbf{E}(f \mid \mathcal{N}) = \frac{1}{\nu(E)} \int_E f \, d\nu,$$

or to avoid problems with null sets,

$$\mathbf{E}(f \mid \mathcal{N}) \cdot \nu(E) = \int_E f \, d\nu \tag{5.5}$$

for all $E \in \mathcal{N}$. However, there are two problems with this definition. First, f is only assumed to be $\mathcal{M}$-measurable—it is unlikely that f is measurable with respect to the smaller σ-algebra $\mathcal{N}$, so the integral $\int_E f \, d\nu$ technically does not make sense. Second, there is no reason to expect the existence of a universal constant $\mathbf{E}(f \mid \mathcal{N})$ satisfying (5.5) for all events E unless f is constant. To correct these problems, we seek a new *random variable* $\mathbf{E}(f \mid \mathcal{N})$ whose "local averages" over events in $\mathcal{N}$ agree with those of f. That is, we want the equation

$$\int_E \mathbf{E}(f \mid \mathcal{N}) \, d\nu = \int_E f \, d\mu$$

to hold for all $E \in \mathcal{N}$. Observe that the measure λ defined on $\mathcal{M}$ by $d\lambda = f \, d\mu$ is absolutely continuous with respect to μ, and the restriction of λ to $\mathcal{N}$ is absolutely continuous with respect to ν. Therefore, the Radon-Nikodym theorem guarantees that there is an $\mathcal{N}$-measurable function $\mathbf{E}(f \mid \mathcal{N}) : X \to \mathbf{R}$ satisfying

$$\lambda(E) = \int_E \mathbf{E}(f \mid \mathcal{N}) \, d\nu$$

for all $E \in \mathcal{N}$. Indeed, $\mathbf{E}(f \mid \mathcal{N}) = \frac{d\lambda}{d\nu}$ is nothing more than the Radon-Nikodym derivative of $\lambda|_{\mathcal{N}}$ with respect to ν.

5.7.2 The Lebesgue Differentiation Theorem

Now we turn to another discussion related to the Radon-Nikodym theorem. Here we will focus on the situation where $X = \mathbf{R}$ (or $\mathbf{R}^n$) equipped with its Lebesgue measure μ, and ν is either a σ-finite signed measure or a complex measure that is absolutely continuous with respect to μ. The notation $\frac{d\nu}{d\mu}$ and the term "Radon-Nikodym derivative" seem to suggest that we are somehow measuring the "change" in ν relative to the "change" in μ. But how can one quantify the rate of change of one measure with respect to another? We can make sense of this concept by thinking in terms of open balls. For each $x \in \mathbf{R}^n$, we will define

$$F(x) = \lim_{r \to 0} \frac{\nu(B(x;r))}{\mu(B(x;r))}.$$

In light of the Radon-Nikodym theorem, we know there is a measurable function f satisfying $d\nu = f\, d\mu$, so we can rewrite F as

$$F(x) = \lim_{r \to 0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} f\, d\mu.$$

There are two ways to interpret the right hand side of this equation. For one, we are simply computing the average value of f over the ball $B(x;r)$, and then taking the limit as $r \to 0$. Consequently, one might expect that $F(x) = f(x)$, provided f does not oscillate too wildly at x. On the other hand, one could also think of F as some sort of derivative. Putting these two together, it seems that we are saying that the derivative of the indefinite integral of f (i.e., the measure ν) equals f μ-a.e., which sounds an awful lot like a version of the fundamental theorem of calculus for the Lebesgue integral.

In order to analyze the Radon-Nikodym derivative from this perspective, we will consider a specific function associated to f that "maximizes" the average value of f over all balls centered at x. First, we need to make one observation. Recall that $\frac{d\nu}{d\mu}$ belongs to $L^1(\mathbf{R}^n, \mu)$ when ν is finite (i.e., a complex measure), but it is only guaranteed to be an extended integrable function when ν is a σ-finite signed measure. In fact, if we look back at the proof of the Radon-Nikodym theorem, we can see that f is constructed in such a way that it is integrable over certain sets of finite measure.

Definition 5.7.18. A function $f : \mathbf{R}^n \to \mathbf{C}$ is said to be **locally integrable** if $\int_K |f|\, d\mu < \infty$ for every compact set $K \subseteq \mathbf{R}^n$. The set of all locally integrable functions is denoted by $L^1_{\text{loc}}(\mathbf{R}^n, \mu)$.

Notice that we certainly have $L^1(\mathbf{R}^n, \mu) \subseteq L^1_{\text{loc}}(\mathbf{R}^n, \mu)$, but locally integrable functions need not be integrable on $\mathbf{R}^n$. (For example, the constant function $f = 1$ is locally integrable.) To each locally integrable function f, we can define a particular function that encodes the maximum possible average value of f on an open ball centered at each point $x \in \mathbf{R}^n$.

Definition 5.7.19. Given a function $f \in L^1_{\text{loc}}(\mathbf{R}^n, \mu)$, we define $Hf : \mathbf{R}^n \to \mathbf{R}$ by

$$Hf(x) = \sup_{r>0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} |f| \, d\mu.$$

We call Hf the **Hardy-Littlewood maximal function** of f.

For each $x \in \mathbf{R}^n$ and $r > 0$, we let $A_r f(x)$ denote the average value of f over the ball $B(x;r)$:

$$A_r f(x) = \frac{1}{\mu(B(x;r))} \int_{B(x;r)} f \, d\mu.$$

For a fixed f, it is not hard to check that $A_r f(x)$ varies continuously in both r and x. Consequently, we can write

$$Hf(x) = \sup_{r>0} A_r|f|(x),$$

so Hf is the pointwise supremum of a family of continuous functions. Therefore, it is lower semicontinuous[2], hence measurable. Can we say anything about its integrability? Unfortunately, Hf need not belong to $L^1(\mathbf{R}^n, \mu)$ even when f is integrable. However, we do have a certain boundedness result for the maximal function, which says that Hf is a "weak L^1" function whenever $f \in L^1(\mathbf{R}^n, \mu)$.

Theorem 5.7.20 (Hardy-Littlewood Maximal Inequality)**.** *For all $f \in L^1(\mathbf{R}^n, \mu)$ and all $\alpha > 0$, we have*

$$\mu\big(\{x \in \mathbf{R}^n : Hf(x) > \alpha\}\big) \le \frac{3^n}{\alpha} \int_{\mathbf{R}^n} |f| \, d\mu.$$

The proof of the maximal inequality is not difficult, but it requires a technical lemma (due to Wiener) regarding unions of open balls in $\mathbf{R}^n$. We omit both proofs and forge ahead to the main result on differentation.

[2] A function $f : \mathbf{R}^n \to \mathbf{R}$ is **lower semicontinuous** if the set

$$\{x \in \mathbf{R}^n : f(x) > \alpha\}$$

is open for all $\alpha \in \mathbf{R}$. It is evident that any lower semicontinuous function is measurable.

Theorem 5.7.21 (Lebesgue Differentiation Theorem)**.** *For all $f \in L^1_{\text{loc}}(\mathbf{R}^n, \mu)$, we have*

$$\lim_{r\to 0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} f\, d\mu = f(x)$$

for μ-a.e. $x \in \mathbf{R}^n$. In fact,

$$\lim_{r\to 0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} |f(x) - f(y)|\, d\mu(y) = 0$$

for μ-a.e. $x \in \mathbf{R}^n$.

Proof. Since we are only trying to establish pointwise convergence, it suffices to show the result holds for all x in a ball of arbitrarily large radius centered at 0. That is, we may assume $\|x\| < R$ for some $R > 0$ and that f vanishes outside $B(0; R)$. In other words, we can assume $f \in L^1(\mathbf{R}^n, \mu)$.

Now we need a fact that we have not encountered yet, but that we will see in the relatively near future. Given $\varepsilon > 0$, we can find a compactly supported[3] continuous function $g : \mathbf{R}^n \to \mathbf{C}$ satisfying

$$\int_{\mathbf{R}^n} |f - g|\, d\mu < \varepsilon.$$

Since g is compactly supported, it belongs to $L^1(\mathbf{R}^n, \mu)$. Furthermore, the continuity of g guarantees that for each $x \in \mathbf{R}^n$ and $\delta > 0$ there exists an $r > 0$ such that $\|x - y\| < r$ implies

$$|g(x) - g(y)| < \delta,$$

so

$$\begin{aligned}
|A_r g(x) - g(x)| &= \left| \frac{1}{\mu(B(x;r))} \int_{B(x;r)} g(y)\, d\mu(y) - g(x) \right| \\
&= \left| \frac{1}{\mu(B(x;r))} \int_{B(x;r)} g(y)\, d\mu(y) - \frac{1}{\mu(B(x;r))} \int_{B(x;r)} g(x)\, d\mu(y) \right| \\
&\le \frac{1}{\mu(B(x;r))} \int_{B(x;r)} |g(y) - g(x)|\, d\mu(y) \\
&< \delta.
\end{aligned}$$

[3] A function $f : \mathbf{R}^n \to \mathbf{C}$ is **compactly supported** if the closure of the set

$$\{x \in \mathbf{R}^n : f(x) \neq 0\}$$

is compact.

It follows that

$$\lim_{r\to 0} A_r g(x) = g(x)$$

for all $x \in \mathbf{R}^n$. Now observe that[4]

$$\begin{aligned}\limsup_{r\to 0} |A_r f(x) - f(x)| &= \limsup_{r\to 0} |A_r(f-g)(x) + A_r g(x) - g(x) + (f-g)(x)| \\ &\le \sup_{r>0} |A_r(f-g)(x)| + 0 + |f(x) - g(x)| \\ &\le H(f-g)(x) + |f(x) - g(x)|.\end{aligned}$$

To finish the proof, it suffices to show that the last line can be made arbitrarily small. To this end, for each $\alpha > 0$ we define

$$E_\alpha = \left\{ x \in \mathbf{R}^n : \limsup_{r\to 0} |A_r f(x) - f(x)| > \alpha \right\}$$

and

$$F_\alpha = \big\{ x \in \mathbf{R}^n : |f(x) - g(x)| > \alpha \big\}.$$

By the work above, we see that

$$E_\alpha \subseteq F_{\alpha/2} \cup \big\{ x \in \mathbf{R}^n : H(f-g)(x) > \tfrac{\alpha}{2} \big\}.$$

The Hardy-Littlewood maximal inequality guarantees that

$$\mu\big(\big\{x \in \mathbf{R}^n : H(f-g)(x) > \tfrac{\alpha}{2}\big\}\big) \le \frac{3^n}{\alpha/2} \int_{\mathbf{R}^n} |f-g| \, d\mu < \frac{2 \cdot 3^n \varepsilon}{\alpha}.$$

On the other hand,

$$\int_{F_{\alpha/2}} |f-g| \, d\mu \ge \frac{\alpha}{2} \cdot \mu(F_{\alpha/2}),$$

whence

$$\mu(F_{\alpha/2}) < \frac{2\varepsilon}{\alpha}.$$

Therefore,

$$\mu(E_\alpha) < \frac{2 \cdot 3^n \varepsilon}{\alpha} + \frac{2\varepsilon}{\alpha} = \frac{2(1+3^n)\varepsilon}{\alpha}$$

for all $\alpha > 0$. Since $\varepsilon > 0$ was arbitrary, we must have $\mu(E_\alpha) = 0$ for all $\alpha > 0$. If we let $E = \bigcup_{n=1}^\infty E_{1/n}$, then $\mu(E) = 0$ and

$$\lim_{r\to 0} A_r f(x) = f(x)$$

[4] For a function $f : \mathbf{R}^n \to \mathbf{R}$, we define the **limit superior** of f at $a \in \mathbf{R}^n$ to be

$$\limsup_{x\to a} f(x) = \lim_{\varepsilon\to 0} \sup_{0<\|x-a\|<\varepsilon} f(x) = \inf_{\varepsilon>0} \sup_{0<\|x-a\|<\varepsilon} f(x).$$

for all $x \in E^c$. Thus we have proven our first claim.

To prove the second estimate, which is seemingly stronger than the first, we begin by defining

$$L_f = \left\{ x \in \mathbf{R}^n : \lim_{r \to 0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} |f(x) - f(y)| \, d\mu(y) = 0 \right\}.$$

We claim that $\mu(L_f^c) = 0$. For each $z \in \mathbf{C}$, define $g_z(x) = |f(x) - z|$. Then applying the estimate that we have already established to g_z, we have

$$\lim_{r \to 0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} g_z(y) \, d\mu(y) = g_z(x) = |f(x) - c| \tag{5.6}$$

for μ-a.e. x, i.e., for all x outside a null set E_z. Let $D \subseteq \mathbf{C}$ be a countable dense subset. Then the set

$$E = \bigcup_{d \in D} E_d$$

is μ-null. Furthermore, given any $x \in E^c$ and $\varepsilon > 0$, there is a $d \in D$ satisfying $|f(x) - d| < \frac{\varepsilon}{2}$. Therefore,

$$|f(x) - f(y)| \leq |f(x) - d| + |f(y) - d| < |f(y) - d| + \frac{\varepsilon}{2}$$

for all $y \in \mathbf{R}^n$, so

$$\begin{aligned}
\limsup_{r \to 0} \frac{1}{\mu(B(x;r))} & \int_{B(x;r)} |f(x) - f(y)| \, d\mu(y) \\
&\leq \lim_{r \to 0} \frac{1}{\mu(B(x;r))} \int_{B(x;r)} |f(y) - d| + \frac{\varepsilon}{2} \, d\mu(y) \\
&= |f(x) - d| + \frac{\varepsilon}{2} \\
&< \varepsilon,
\end{aligned}$$

where the last equality follows from (5.6). Since $\varepsilon > 0$ was arbitrary, we are done. $\square$

The set L_f that we defined in the last proof is called the **Lebesgue set** of f, and its elements are called **Lebesgue points**. Thus the latter part of the proof establishes that x is a Lebesgue point of f for μ-a.e. $x \in \mathbf{R}^n$. Notice that the Lebesgue points of f are precisely the points at which f does not oscillate very much, at least on average. This condition is a weakened form of continuity—indeed, f is continuous at x if and only if

$$\lim_{r \to 0} \sup_{y \in B_r(x)} |f(x) - f(y)| = 0.$$

Thus the second estimate in the Lebesgue differentiation theorem can be thought of as another manifestation of Littlewood's third principle—a measurable function is almost continuous.

Finally, let us make one more observation regarding the Radon-Nikodym derivative. Coming back to our original discussion, the Lebesgue differentiation theorem has the following immediate corollary.

Corollary 5.7.22. *Let ν be a σ-finite signed measure or a complex measure on $\mathbf{R}^n$ that is absolutely continuous with respect to the Lebesgue measure μ. Then we have*

$$\lim_{r\to 0} \frac{\nu(B_r(x))}{\mu(B_r(x))} = \frac{d\nu}{d\mu}(x)$$

for μ-a.e. $x \in \mathbf{R}^n$.

Exercises for Section 5.7

Exercise 5.7.1. Suppose $(X, \mathcal{M}, \mu)$ is a measure space and ν is a signed measure on $\mathcal{M}$. Prove that $\nu \ll \mu$ and $\nu \perp \mu$ implies $\nu = 0$.

Exercise 5.7.2. Prove Proposition 5.7.5.

Exercise 5.7.3. Let $(X, \mathcal{M}, \mu)$ be a measure space, and let ν be a complex measure on $\mathcal{M}$.

(a) Suppose λ is another positive measure on $\mathcal{M}$. Show that if $\nu \ll \lambda$ and $\lambda \ll \mu$, then $\nu \ll \mu$.

(b) Show that if $\lambda \ll \mu$ and $\nu \perp \mu$, then $\nu \perp \lambda$.

Exercise 5.7.4. Let X be a set and $\mathcal{M}$ a σ-algebra on X. Define a relation $\approx$ on the set of positive measures on $\mathcal{M}$ by $\mu \approx \nu$ if and only if μ and ν are mutually absolutely continuous. Prove that $\approx$ is an equivalence relation.

Exercise 5.7.5. Suppose $(X, \mathcal{M}, \mu)$ is a measure space and $\{\nu_j\}_{j=1}^{\infty}$ is a collection of positive measures on $\mathcal{M}$ satisfying $\nu_j \perp \mu$ for all j. If we define $\nu = \sum_{j=1}^{\infty} \nu_j$, prove that $\nu \perp \mu$.

Exercise 5.7.6. Suppose $F : \mathbf{R} \to \mathbf{R}$ is bounded, increasing, and differentiable, and let ν denote the associated Lebesgue-Stieltjes measure, which is characterized by

$$\nu((a, b]) = F(b) - F(a)$$

for all $a, b \in \mathbf{R}$ with $a < b$. Let μ denote Lebesgue measure on $\mathbf{R}$, and assume $\nu \ll \mu$. Prove that

$$\frac{d\nu}{d\mu} = F'$$

μ-almost everywhere.

Exercise 5.7.7. Let $E \subseteq \mathbf{R}^n$ be a Borel set, and let μ denote the Lebesgue measure on $\mathbf{R}^n$. Given a point $x \in \mathbf{R}^n$, we define the **density** of E at x to be

$$D_E(x) = \lim_{r \to 0} \frac{\mu(E \cap B(x;r))}{\mu(B(x;r))},$$

provided the limit exists. Show that $D_E(x) = 1$ for μ-a.e. $x \in E$, and that $D_E(x) = 0$ for μ-a.e. $x \in E^c$. (For an added challenge, find an example of a Borel set E and a point $x \in \mathbf{R}^n$ such that $0 < D_E(x) < 1$.)

Chapter 6

Banach Spaces

With abstract measure theory under our belts, we now turn our attention toward some topics in functional analysis. We will specifically study Banach spaces and Hilbert spaces, with an eye toward operators and linear functionals on such spaces. We begin with a preliminary discussion on L^p spaces, which provide some informative examples from the realm of measure theory.

6.1 The L^p Spaces

Throughout this section we let $(X, \mathcal{M}, \mu)$ denote a fixed measure space. Recall that we previously defined the space $L^1(X, \mu)$ to be the set of all complex-valued μ-integrable functions on X:

$$L^1(X,\mu) = \left\{ f : X \to \mathbf{C} \;\middle|\; \int_X |f| \, d\mu < \infty \right\}.$$

We will now study $L^1(X, \mu)$ from a different perspective than we have considered before—we will try to analyze the structure of $L^1(X, \mu)$ as a whole, much like we did with the function space $C(X)$ in Chapter 3.

Let us begin with some minor observations regarding $L^1(X, \mu)$. First, observe that if $f, g \in L^1(X, \mu)$, then

$$\int_X |f+g| \, d\mu \le \int_X |f| + |g| \, d\mu = \int_X |f| \, d\mu + \int_X |g| \, d\mu < \infty,$$

so $f + g \in L^1(X, \mu)$. Also, for any $\alpha \in \mathbf{C}$ we have

$$\int_X |\alpha f| \, d\mu = \int_X |\alpha||f| \, d\mu = |\alpha| \int_X |f| \, d\mu < \infty,$$

so $\alpha f \in L^1(X, \mu)$. It is then evident that $L^1(X, \mu)$ satisfies the axioms for a vector space. In fact, we have all the necessary tools to turn $L^1(X, \mu)$ into a *normed* vector space in a natural way. Given $f \in L^1(X,\mu)$, we define

$$\|f\|_1 = \int_X |f| \, d\mu.$$

Notice that the computations we performed above guarantee that $\|\alpha f\|_1 = |\alpha| \|f\|_1$ and $\|f+g\|_1 \leq \|f\|_1 + \|g\|_1$ for all $f, g \in L^1(X,\mu)$ and all $\alpha \in \mathbf{C}$. Furthermore, it is clear that we always have $\|f\|_1 \geq 0$. However, notice that $\|f\|_1 = 0$, i.e.,

$$\int_X |f| \, d\mu = 0,$$

only implies that $f = 0$ μ-almost everywhere. In other words, $\|\cdot\|_1$ is not positive definite. (Such a function is called a **seminorm**.) In order to get a true norm on $L^1(X,\mu)$, we need to identify functions that agree μ-almost everywhere. This is accomplished by defining an equivalence relation $\sim$ on $L^1(X,\mu)$, where we declare $f \sim g$ if and only if $f = g$ μ-almost everywhere. For the more algebraically-minded reader, this is equivalent to defining

$$V = \big\{ f \in L^1(X,\mu) \mid f = 0 \ \mu\text{-a.e.} \big\}$$

and then forming the quotient of $L^1(X,\mu)$ by the subspace V. In any event, $\|\cdot\|_1$ becomes a full-fledged norm on the resulting set of equivalence classes. However, we will continue to write $L^1(X,\mu)$, with the tacit understanding that we are thinking of an element $f \in L^1(X,\mu)$ not as a single function, but as an equivalence class of functions that are μ-a.e. equal.

We can define an entire family of related function spaces, called the **L^p-spaces**, in a similar fashion. For each $1 \leq p < \infty$, we set

$$L^p(X,\mu) = \left\{ f : X \to \mathbf{C} \;\middle|\; \int_X |f|^p \, d\mu < \infty \right\},$$

where we again identify functions that agree almost everywhere. Clearly $L^p(X,\mu)$ is closed under multiplication by complex scalars, and closure under addition follows from the observation that

$$|f+g|^p \leq \big(|f| + |g|\big)^p \leq \big(2\max\{|f|, |g|\}\big)^p \leq 2^p\big(|f|^p + |g|^p\big).$$

We also define the corresponding **L^p-norm** by

$$\|f\|_p = \left(\int_X |f|^p \, d\mu \right)^{1/p}.$$

It should not be clear at all that $\|\cdot\|_p$ defines a norm. To prove it, we need to establish two important inequalities due to Hölder and Minkowski.

Theorem 6.1.1 (Hölder's Inequality). *Suppose $1 \leq p, q < \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$. Then given two measurable functions $f, g : X \to \mathbf{C}$, we have*

$$\|fg\|_1 \leq \|f\|_p \|g\|_q.$$

In particular, if $f \in L^p(X, \mu)$ and $g \in L^q(X, \mu)$, then $fg \in L^1(X, \mu)$.

Proof. Notice that the inequality holds trivially if one of $\|f\|_p$ and $\|g\|_q$ is ∞. Similarly, if $\|f\|_p = 0$ or $\|g\|_q = 0$, then $fg = 0$ μ-a.e., so $\|fg\|_1 = 0$. Therefore, we may assume that $0 < \|f\|_p, \|g\|_q < \infty$. In fact, we can divide both sides by $\|f\|_p \|g\|_q$ and assume that $\|f\|_p = \|g\|_q = 1$. The rest of the proof relies on the fact that the exponential function is convex: for all $a, b \in \mathbf{R}$ and $0 \leq \lambda \leq 1$, we have

$$e^{\lambda a + (1-\lambda) b} \leq \lambda e^a + (1-\lambda) e^b.$$

Given $x \in X$ with $0 < |f(x)|, |g(x)| < \infty$, take $a = \log |f(x)|^p$, $b = \log |g(x)|^q$, $\lambda = \frac{1}{p}$, and $1 - \lambda = \frac{1}{q}$. Then

$$|f(x)g(x)| \leq \tfrac{1}{p}|f(x)|^p + \tfrac{1}{q}|g(x)|^q.$$

If $f(x) = 0$ or $g(x) = 0$, then the same inequality holds trivially. Now if we integrate both sides, we get

$$\|fg\|_1 = \int_X |fg| \, d\mu \leq \frac{1}{p} \int_X |f|^p \, d\mu + \frac{1}{q} \int_X |g|^q \, d\mu = \frac{1}{p} + \frac{1}{q} = 1 = \|f\|_p \|g\|_q,$$

so the desired inequality holds. □

The exponents p and q satisfying $\frac{1}{p} + \frac{1}{q} = 1$ in Hölder's inequality are called **conjugate exponents**. Notice that the case $p = q = 2$ is special, and the corresponding case of Hölder's inequality is worth singling out.

Theorem 6.1.2 (Cauchy-Schwarz Inequality). *If $f, g \in L^2(X, \mu)$, then $fg \in L^1(X, \mu)$ and*

$$\int_X |fg| \, d\mu \leq \left(\int_X |f|^2 \, d\mu \right)^{1/2} \left(\int_X |g|^2 \, d\mu \right)^{1/2}.$$

Now we can use Hölder's inequality to establish Minkowski's inequality, which is really just the statement that the L^p-norms satisfy the triangle inequality.

Theorem 6.1.3 (Minkowski's Inequality)**.** *Let* $1 \leq p < \infty$*. For all* $f, g \in L^p(X, \mu)$ *we have*

$$\|f+g\|_p \leq \|f\|_p + \|g\|_p.$$

Proof. We begin with the observation that

$$|f+g|^p \leq |f||f+g|^{p-1} + |g||f+g|^{p-1}.$$

Now let q be the conjugate exponent to p, and notice that

$$(p-1)q = (p-1)\left(\frac{p-1}{p}\right) = p.$$

Thus

$$\int_X |f+g|^{(p-1)q}\, d\mu = \int_X |f+g|^p\, d\mu < \infty$$

since $f + g \in L^p(X, \mu)$. Now Hölder's inequality gives

$$\int_X |f||f+g|^{p-1}\, d\mu \leq \left(\int_X |f|^p\, d\mu\right)^{1/p} \left(\int_X |f+g|^{(p-1)q}\, d\mu\right)^{1/q},$$

and similarly,

$$\int_X |g||f+g|^{p-1}\, d\mu \leq \left(\int_X |g|^p\, d\mu\right)^{1/p} \left(\int_X |f+g|^{(p-1)q}\, d\mu\right)^{1/q}.$$

Thus

$$\int_X |f+g|^p\, d\mu \leq \left(\|f\|_p + \|g\|_p\right) \left(\int_X |f+g|^{(p-1)q}\, d\mu\right)^{1/q},$$

and writing $(p-1)q = p$ and $\frac{1}{q} = 1 - \frac{1}{p}$ gives

$$\int_X |f+g|^p\, d\mu \leq \left(\|f\|_p + \|g\|_p\right) \left(\int_X |f+g|^p\, d\mu\right)^{1-\frac{1}{p}}.$$

If $f + g = 0$ μ-a.e., then Minkowski's inequality holds trivially. Otherwise, it follows from the previous inequality that

$$\left(\int_X |f+g|^p\, d\mu\right)^{1/p} \leq \|f\|_p + \|g\|_p,$$

whence the result. □

Corollary 6.1.4. *For all $1 \le p < \infty$, $L^p(X,\mu)$ is a normed vector space.*

Before investigating L^p spaces further, we discuss a specific example that we have actually seen once before.

Example 6.1.5. Let $X = \mathbf{N}$ equipped with counting measure. We can think of a function $f : \mathbf{N} \to \mathbf{C}$ as a sequence $(x_n)_{n=1}^\infty$, and the condition for f to be in $L^p(X,\mu)$ translates to the requirement that

$$\sum_{n=1}^\infty |x_n|^p < \infty.$$

Furthermore, the L^p-norm takes the form

$$\|(x_n)_{n=1}^\infty\|_p = \left(\sum_{n=1}^\infty |x_n|^p\right)^{1/p}$$

In particular, $L^1(\mathbf{N},\mu)$ consists of the absolutely summable sequences,

$$\sum_{n=1}^\infty |x_n| < \infty,$$

while $L^2(\mathbf{N},\mu)$ is the space of all square summable sequences,

$$\sum_{n=1}^\infty |x_n|^2 < \infty.$$

In other words, the L^1 and L^2 spaces in this case are nothing more than the spaces ℓ^1 and ℓ^2 that we encountered in Chapter 1. (Actually, there is a slight difference—we are now dealing with complexified versions of those spaces.) Therefore, we denote the L^p spaces in this particular situation by $\ell^p(\mathbf{N})$, or simply ℓ^p.

Example 6.1.6. Let $X = \{1, 2, \ldots, n\}$, and let μ denote the counting measure on X. Notice that the set of all functions from X to $\mathbf{C}$ is precisely $\mathbf{C}^n$, and every such function is μ-measurable. Furthermore, for all $1 \le p < \infty$ we have

$$\sum_{j=1}^n |x_j|^p < \infty,$$

for all n-tuples $(x_j)_{j=1}^n \in \mathbf{C}^n$. Therefore, $L^p(X,\mu)$ is simply $\mathbf{C}^n$ endowed with the corresponding p-norm. If we restrict to the subspace of real-valued functions, we recover $\mathbf{R}^n$ with its various p-norms.

There is also an analog of ℓ^∞ for general measure spaces. Recall that the norm on ℓ^∞ was defined by

$$\|(x_n)_{n=1}^\infty\|_\infty = \sup_n |x_n|.$$

We would like to do something similar for measurable functions on $(X, \mathcal{M}, \mu)$, though our definition should allow for the possibility that a function is unbounded (or even takes infinite values) on a μ-null set. Therefore, we replace the usual supremum norm with the *essential supremum*.

Definition 6.1.7. Suppose $f : X \to \mathbf{C}$ is measurable. We define the **essential supremum** of f to be

$$\operatorname*{ess\,sup}_{x\in X} |f(x)| = \inf\left\{\alpha \geq 0 : \mu\big(\{x \in X : |f(x)| > \alpha\}\big) = 0\right\}.$$

Now we define $L^\infty(X, \mu)$ to be the set of all (equivalence classes of μ-a.e. equal) measurable functions having a finite essential supremum:

$$L^\infty(X,\mu) = \left\{ f : X \to \mathbf{C} \;\middle|\; \operatorname*{ess\,sup}_{x\in X} |f(x)| < \infty \right\}.$$

Such functions are sometimes called *essentially bounded* functions. We then set the norm on $L^\infty(X, \mu)$ to be

$$\|f\|_\infty = \operatorname*{ess\,sup}_{x\in X} |f(x)|.$$

It is then straightforward to verify that $L^\infty(X, \mu)$ is a normed vector space. (See Exercise 6.1.1.) We also have a version of Hölder's inequality for $\|\cdot\|_\infty$, where we formally interpret 1 and ∞ as conjugate exponents.

Theorem 6.1.8 (Hölder's Inequality for $p = 1$ and $q = \infty$)**.** *Suppose $f \in L^1(X, \mu)$ and $g \in L^\infty(X, \mu)$. Then $fg \in L^1(X, \mu)$, and*

$$\|fg\|_1 \leq \|f\|_1 \|g\|_\infty.$$

Proof. By the definition of the essential supremum, we have $|g(x)| \leq \|g\|_\infty$ for μ-a.e. $x \in X$. Thus $|fg| \leq |f|\|g\|_\infty$ μ-a.e., and it follows that

$$\int_X |fg|\, d\mu \leq \int_X |f|\|g\|_\infty\, d\mu = \|g\|_\infty \int_X |f|\, d\mu = \|f\|_1 \|g\|_\infty. \qquad \square$$

Example 6.1.9. Let $X = \mathbf{N}$ equipped with the counting measure μ. Notice that $\mu(E) = 0$ if and only if $E = \emptyset$, so it is easy to see that

$$\operatorname*{ess\,sup}_{n\in\mathbf{N}} |f(n)| = \sup_{n\in\mathbf{N}} |f(n)|$$

for any function $f : \mathbf{N} \to \mathbf{C}$. Thus $\|\cdot\|_\infty$ is the usual supremum norm, and we see that $L^\infty(\mathbf{N}, \mu)$ is nothing more than ℓ^∞.

Example 6.1.10. If $X = \{1, 2, \ldots, n\}$ endowed with counting measure μ, then $L^\infty(X, \mu)$ is just $\mathbf{C}^n$ equipped with the norm

$$\left\|(x_j)_{j=1}^n\right\| = \max_{1 \le j \le n} |x_j|.$$

Remark 6.1.11. Suppose X is a compact metric space and μ is a Borel measure on X. (For example, we could take $X = [0, 1]$ and μ to be Lebesgue measure.) Then any continuous function $f : X \to \mathbf{C}$ is measurable, and it is bounded. Therefore,

$$\operatorname*{ess\,sup}_{x \in X} |f(x)| = \sup_{x \in X} |f(x)|,$$

so we recover the usual supremum norm on continuous functions.

We now close our preliminary discussion of L^p spaces with a few small yet interesting results. The first states that the infinity norm can be thought of as a "limiting case" for the norms $\|\cdot\|_p$ when $\mu(X)$ is finite.

Theorem 6.1.12. *Suppose $\mu(X) < \infty$. Then for all $f \in L^\infty(X, \mu)$,*

$$\|f\|_\infty = \lim_{p \to \infty} \|f\|_p.$$

Proof. Notice first that the given statement tacitly assumes that if $f \in L^\infty(X, \mu)$, then $f \in L^p(X, \mu)$ for all $1 \le p < \infty$. This is actually straightforward to check:

$$\|f\|_p^p = \int_X |f|^p \, d\mu \le \int_X \|f\|_\infty^p \, d\mu = \|f\|_\infty^p \cdot \mu(X) < \infty.$$

Now let $\alpha < \|f\|_\infty$ and set

$$E = \{x \in X : |f(x)| > \alpha\}.$$

Then $\mu(E) > 0$. Furthermore, we have

$$\|f\|_p \ge \left(\int_E |f|^p \, d\mu \right)^{1/p} \ge \alpha \cdot \mu(E)^{1/p}.$$

Therefore,

$$\liminf_{p \to \infty} \|f\|_p \ge \alpha \cdot \lim_{p \to \infty} \mu(E)^{1/p} = \alpha,$$

since $\mu(E)^{1/p} \to 1$ as $p \to \infty$. This holds for all $\alpha < \|f\|_\infty$, so we can conclude that

$$\liminf_{p \to \infty} \|f\|_p \ge \|f\|_\infty.$$

On the other hand,

$$\|f\|_p \leq \left(\int_X \|f\|_\infty^p \, d\mu\right)^{1/p} = \|f\|_\infty \cdot \mu(X)^{1/p},$$

whence

$$\limsup_{p\to\infty} \|f\|_p \leq \|f\|_\infty.$$

Hence $\|f\|_p \to \|f\|_\infty$ as $p \to \infty$. $\square$

Our first step in the last proof was to observe that $L^\infty(X,\mu) \subseteq L^p(X,\mu)$ for all $p \geq 1$ when $\mu(X)$ is finite. In fact, we can make a more general statement about L^p spaces in this situation.

Theorem 6.1.13. *Suppose $\mu(X) < \infty$ and $1 \leq p < q \leq \infty$. Then $L^q(X,\mu) \subseteq L^p(X,\mu)$. Furthermore,*

$$\|f\|_p \leq \|f\|_q \cdot \mu(X)^{\frac{1}{p}-\frac{1}{q}}.$$

Proof. We have already established the result when $q = \infty$. Suppose $1 \leq p < q < \infty$ and let $f \in L^q(X,\mu)$. Then observe that $\frac{q}{p}$ and $\frac{q}{q-p}$ are conjugate exponents, so Hölder's inequality yields

$$\|f\|_p^p = \int_X |f|^p \cdot 1 \, d\mu \leq \||f|^p\|_{q/p}\|1\|_{q/(q-p)} = \left(\int_X |f|^{p\cdot\frac{q}{p}} \, d\mu\right)^{p/q} \mu(X)^{\frac{q-p}{q}},$$

or

$$\|f\|_p^p \leq \left(\int_X |f|^q \, d\mu\right)^{p/q} \mu(X)^{1-\frac{p}{q}}.$$

Taking pth roots of both sides, we get

$$\|f\|_p \leq \left(\int_X |f|^q \, d\mu\right)^{1/q} \mu(X)^{\frac{1}{p}-\frac{1}{q}} < \infty.$$

Thus $f \in L^p(X,\mu)$, and we have the desired bound on its norm. $\square$

Exercises for Section 6.1

Exercise 6.1.1. Prove that $\|\cdot\|_\infty$ defines a norm on $L^\infty(X,\mu)$.

6.2 Fundamentals of Banach Spaces

The L^p spaces introduced in the last section will serve as informative examples of Banach spaces, which are the foundation for much of functional analysis. We have already introduced the definition of a Banach space—and we have seen several examples—but we will still go through some preliminary discussion here. We will then prove that the L^p spaces are indeed Banach spaces before moving on to more general Banach space theory.

Let $\mathbf{F}$ be a field. We will almost always take $\mathbf{F} = \mathbf{C}$, but we may occasionally consider situations where $\mathbf{F} = \mathbf{R}$.[1] Recall that a **vector space** over $\mathbf{F}$ is an abelian group $(V, +)$ together with a map $(\alpha, x) \mapsto \alpha x$ from $\mathbf{F} \times V \to V$ satisfying the following axioms:

1. $\alpha(v + w) = \alpha x + \alpha y$ for all $\alpha \in \mathbf{F}$ and all $x, y \in V$;
2. $(\alpha_1 + \alpha_2)x = \alpha_1 x + \alpha_2 x$ for all $\alpha_1, \alpha_2 \in \mathbf{F}$ and all $x \in V$;
3. $(\alpha_1 \alpha_2)x = \alpha_1(\alpha_2 x)$ for all $\alpha_1, \alpha_2 \in \mathbf{F}$ and all $x \in V$;
4. $1 \cdot x = x$ for all $x \in V$, where $1 \in \mathbf{F}$ denotes the multiplicative identity.

Now assume $\mathbf{F} = \mathbf{C}$ or $\mathbf{F} = \mathbf{R}$, and let V be an $\mathbf{F}$-vector space. As we have defined before, a **norm** on V is a function $\|\cdot\| : V \to [0, \infty)$ satisfying the following conditions:

1. (Positive definite) $\|x\| = 0$ if and only if $x = 0$;
2. (Homogeneous) $\|\alpha x\| = |\alpha|\|x\|$ for all $\alpha \in \mathbf{F}$ and all $x \in V$;
3. (Triangle inequality) $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in V$.

Recall that if V is a normed vector space, then it automatically comes equipped with a metric: for all $x, y \in V$, we set

$$d(x, y) = \|x - y\|.$$

Definition 6.2.1. Let V be a normed vector space. If V is complete in with respect to the metric induced by its norm, we say that V is a **Banach space**.

We have already seen several examples of Banach spaces in previous chapters. We review some of them now, while also setting the stage for some new examples.

[1] Much of what we will do relies on the existence of a norm on the underlying field $\mathbf{F}$, as well as the completeness of $\mathbf{F}$. Therefore, one could make some of the theory work for complete subfields of $\mathbf{C}$, or for fields of p-adic numbers. We will often need results specific to the structure of the complex numbers including the ability to take conjugates, as well as the fact that $\mathbf{C}$ is algebraically closed.

Example 6.2.2. For all n, $\mathbf{C}^n$ and $\mathbf{R}^n$ are both Banach spaces with respect to their Euclidean norms as a consequence of Exercise 2.3.3 and the completeness of $\mathbf{C}$ and $\mathbf{R}$. (We shall see shortly that $\mathbf{C}^n$ and $\mathbf{R}^n$ are both complete with respect to *any* norm.) In particular, $\mathbf{R}$ and $\mathbf{C}$ are both Banach spaces with respect to the norm

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

Example 6.2.3. As shown in Exercise 2.3.6, the sequence space ℓ^2 is a Banach space. We will soon extend this result to other L^p spaces.

Example 6.2.4. Let X be a compact metric space. Then the space $C(X)$ of continuous complex-valued functions on X is a Banach space with respect to the norm

$$\|f\|_\infty = \sup_{x\in X} |f(x)|$$

by Corollary 3.1.13.

Example 6.2.5. We can generalize the last example to obtain new kinds of Banach spaces. We say a metric space X is **locally compact** if every point has a compact neighborhood. (Equivalently, for each $x \in X$ there is an $r > 0$ such that the closed ball $B[x;r]$ is compact.) Define

$$C_b(X) = \left\{ f : X \to \mathbf{C} \;\middle|\; f \text{ is continuous and } \|f\|_\infty = \sup_{x\in X} |f(x)| < \infty \right\}.$$

It is straightforward to check that $C_b(X)$ is a Banach space with respect to the norm $\|\cdot\|_\infty$. Indeed, it is immediate from Theorems 3.1.10 and 3.1.11 that if $(f_n)_{n=1}^\infty$ is a Cauchy sequence in $C_b(X)$, then it converges to a continuous function $f : X \to \mathbf{C}$. Thus one only needs to check that f is bounded. Well, given $\varepsilon > 0$ we can find an $n \in \mathbf{N}$ such that

$$|f_n(x) - f(x)| < \varepsilon$$

for all $x \in X$, whence

$$|f(x)| \le |f_n(x)| + \varepsilon \le \|f_n\|_\infty + \varepsilon$$

for all $x \in X$. Hence $f \in C_b(X)$.

We can also define a useful subspace of $C_b(X)$ as follows. We say a function $f : X \to \mathbf{C}$ **vanishes at infinity** if the set

$$\{x \in X : |f(x)| \ge \varepsilon\}$$

is compact for all $\varepsilon > 0$. We then define

$$C_0(X) = \{ f : X \to \mathbf{C} \mid f \text{ is continuous and vanishes at infinity} \}.$$

It is easy to see that any function in $C_0(X)$ is necessarily bounded, so $C_0(X) \subseteq C_b(X)$. In fact, $C_0(X)$ is closed in $C_b(X)$ (see Exercise 6.2.1), so $C_0(X)$ is a Banach space under the supremum norm.

Finally, for a function $f : X \to \mathbf{C}$ we define the **support** of f, denoted $\operatorname{supp}(f)$, to be the closure of the set

$$\{x \in X : f(x) \neq 0\}.$$

Let

$$C_c(X) = \bigl\{f : X \to \mathbf{C} \bigm| f \text{ is continuous and } \operatorname{supp}(f) \text{ is compact}\bigr\},$$

which is called the set of **compactly supported** continuous functions. Clearly $C_c(X) \subseteq C_0(X)$, but $C_c(X)$ is not a Banach space under the supremum norm. In fact, $C_c(X)$ is dense in $C_0(X)$. (The proof requires Urysohn's lemma.)

Next we show that the L^p spaces are always complete. Before we can prove it, we need a lemma regarding convergent series in a Banach space.

Definition 6.2.6. Let V be a normed vector space, and let $(x_j)_{j=1}^\infty$ be a sequence of elements in V. We say the series $\sum_{j=1}^\infty x_j$ **converges absolutely** if

$$\sum_{j=1}^\infty \|x_j\| < \infty.$$

Recall that an absolutely convergent series of real or complex numbers is convergent. This fact relies heavily on the completeness of $\mathbf{R}$ and $\mathbf{C}$, so it seems plausible that an absolutely convergent series in a normed space V might not converge if V is not complete. In fact, the convergence of such series can be used to characterize when V is a Banach space.

Proposition 6.2.7. *A normed vector space V is a Banach space if and only if every absolutely convergent series in V converges to an element of V.*

Proof. Assume first that V is a Banach space, and suppose $(x_j)_{j=1}^\infty$ is a sequence in V such that $\sum_{j=1}^\infty x_j$ converges absolutely. Let $\varepsilon > 0$ be given. Then there exists a natural number N such that

$$\sum_{j=n+1}^{m} \|x_j\| < \varepsilon$$

for all $m > n \geq N$. Hence

$$\left\| \sum_{j=n+1}^{m} x_j \right\| \leq \sum_{j=n+1}^{m} \|x_j\| < \varepsilon,$$

so the partial sums of $\sum_{j=1}^{\infty} x_j$ form a Cauchy sequence. Since V is complete, it follows that $\sum_{j=1}^{\infty} x_j$ converges.

Now assume every absolutely convergent series converges in V, and let $(x_j)_{j=1}^{\infty}$ be a Cauchy sequence in V. Then we can construct a subsequence $(x_{j_k})_{k=1}^{\infty}$ satisfying

$$\left\|x_{j_{k+1}} - x_{j_k}\right\| < \frac{1}{2^k}$$

for all k. Thus

$$\sum_{k=1}^{\infty} \left\|x_{j_{k+1}} - x_{j_k}\right\| \le \sum_{k=1}^{\infty} \frac{1}{2^k} < \infty,$$

so the series $\sum_{k=1}^{\infty}(x_{j_{k+1}} - x_{j_k})$ is absolutely convergent. By assumption, this series converges to an element $x \in V$. But the series telescopes, so the partial sums have the form

$$\sum_{k=1}^{n}(x_{j_{k+1}} - x_{j_k}) = x_{j_{n+1}} - x_{j_1}.$$

Hence $x_{j_{n+1}} \to x + x_{j_1}$. Since the Cauchy sequence $(x_j)_{j=1}^{\infty}$ has a convergent subsequence, it converges as well. It follows that V is complete. $\square$

This result about convergent series provides one way of proving the Riesz-Fischer theorem, which says that the L^p spaces are always complete.

Theorem 6.2.8 (Riesz-Fischer). *Let $(X, \mathcal{M}, \mu)$ be a measure space. Then $L^p(X, \mu)$ is a Banach space for all $1 \le p \le \infty$.*

Proof. Assume first that $1 \le p < \infty$. Let $(f_j)_{j=1}^{\infty}$ be a sequence in $L^p(X, \mu)$ such that $\sum_{j=1}^{\infty} f_j$ converges absolutely. Notice that for each $n \in \mathbf{N}$ we have

$$\left\|\sum_{j=1}^{n} |f_j|\right\|_p \le \sum_{j=1}^{n} \|f_j\|_p$$

by Minkowski's inequality. It follows that

$$\int_X \left(\sum_{j=1}^{n} |f_j|\right)^p d\mu \le \left(\sum_{j=1}^{\infty} \|f_j\|_p\right)^p < \infty$$

for all n. Applying the Monotone Convergence Theorem to the left hand side, we see that

$$\int_X \left(\sum_{j=1}^{\infty} |f_j|\right)^p d\mu = \lim_{n\to\infty} \int_X \left(\sum_{j=1}^{n} |f_j|\right)^p d\mu < \infty.$$

Thus if we set $f = \sum_{j=1}^{\infty} f_j$, then

$$\int_X |f|^p \, d\mu \le \int_X \left(\sum_{j=1}^{\infty} |f_j| \right)^p d\mu < \infty,$$

so $f \in L^p(X,\mu)$. That is, the series converges pointwise to a function in $L^p(X,\mu)$. It remains to see that the partial sums of the series converge with respect to the L^p-norm. To this end, observe that for all n,

$$\left| f - \sum_{j=1}^{n} f_j \right| = \left| \sum_{j=n+1}^{\infty} f_j \right| \le \sum_{j=n+1}^{\infty} |f_j|$$

Since $\left(\sum_{j=n+1}^{\infty} |f_j|\right)^p \in L^1(X,\mu)$ and $\sum_{j=n+1}^{\infty} |f_j| \to 0$ as $n \to \infty$, the Generalized Dominated Convergence Theorem implies that

$$\lim_{n\to\infty} \left| f - \sum_{j=1}^{n} f_j \right|^p d\mu = 0.$$

It follows that $\|f - \sum_{j=n+1}^{\infty} f_j\|_p \to 0$ as $n \to \infty$, so the series $\sum_{j=1}^{\infty} f_j$ converges in the L^p-norm. Hence $L^p(X,\mu)$ is a Banach space by the previous proposition.

Now we consider the case $p = \infty$. Suppose $(f_j)_{j=1}^{\infty}$ is a Cauchy sequence in $L^{\infty}(X,\mu)$. By the definition of the essential supremum, for each pair $j,k \in \mathbf{N}$ with $j > k$, there is a μ-null set $E_{j,k}$ such that

$$|f_j(x) - f_k(x)| \le \|f_j - f_k\|_{\infty}$$

for all $x \in E_{j,k}$. Put $E = \bigcup_{j,k} E_{j,k}$. Then $\mu(E) = 0$, and $(f_j)_{j=1}^{\infty}$ is uniformly Cauchy on E^c. It then follows from Theorem 3.1.10 that $f_j \to f$ uniformly on E^c for some function $f : X \to \mathbf{C}$, and that f is necessarily bounded on E^c. Thus $f \in L^{\infty}(X,\mu)$, and we clearly have $\|f_j - f\|_{\infty} \to 0$. Hence $L^{\infty}(X,\mu)$ is a Banach space. □

Now we present two approximation results for L^p spaces. The first shows that certain kinds of simple functions are dense in $L^p(X,\mu)$; the second result does the same for continuous functions.

Theorem 6.2.9. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and let $\mathcal{S}$ denote the set of all simple functions $g : X \to \mathbf{C}$ with the property that*

$$\mu\big(\{x \in X : f(x) \ne 0\}\big) < \infty.$$

Then $\mathcal{S}$ is dense in $L^p(X,\mu)$ for $1 \le p < \infty$.

Proof. Since any simple function is bounded and the elements of $\mathcal{S}$ are nonzero only on a set of finite measure, it is clear that $\mathcal{S} \subseteq L^p(X,\mu)$. In fact, any simple function in $L^p(X,\mu)$ must belong to $\mathcal{S}$. If $g = \sum_{k=1}^n c_k \chi_{E_k}$ is a simple function that does not lie in $\mathcal{S}$, then we must have $c_k \neq 0$ and $\mu(E_k) = \infty$ for some k. Then

$$\int_X |g|^p \, d\mu \geq \int_{E_k} |c_k|^p \, d\mu = \infty,$$

so $g \notin L^p(X,\mu)$.

Suppose $f \in L^p(X,\mu)$ is a nonnegative function. Then there is a sequence $(g_n)_{n=1}^\infty$ of simple functions that increases pointwise to f. Since $g_n \leq f$ for all n, we clearly have

$$\int_X g_n^p \, d\mu \leq \int_X f^p \, d\mu < \infty,$$

so $g_n \in L^p(X,\mu)$, hence in $\mathcal{S}$, for all n. Furthermore, we have $|f - g_n|^p \leq f^p$ for all n and f^p is integrable, so the Dominated Convergence Theorem guarantees that

$$\lim_{n\to\infty} \int_X |f - g_n|^p \, d\mu = 0.$$

Thus $\|f - g_n\|_p \to 0$ as $n \to \infty$, so $g_n \to f$ in $L^p(X,\mu)$. If $f \in L^p(X,\mu)$ is real-valued, we can find sequences $(g_n^+)_{n=1}^\infty$ and $(g_n^-)_{n=1}^\infty$ of simple functions converging to f^+ and f^-, respectively, in $L^p(X,\mu)$, and it follows that

$$\left\| f - (g_n^+ - g_n^-) \right\|_p \leq \left\| f^+ - g_n^+ \right\|_p + \left\| f^- - g_n^- \right\|_p \to 0$$

as $n \to \infty$. The case where f is complex-valued is handled similarly. □

For the next result, we assume $X = \mathbf{R}$ and μ is Lebesgue measure. The theorem actually holds for any Radon measure (which we have not defined yet) on a locally compact Hausdorff space, though it requires more technology.

Theorem 6.2.10. *The set of continuous, compactly-supported functions $C_c(\mathbf{R})$ is dense in $L^p(\mathbf{R},\mu)$ for $1 \leq p < \infty$.*

Proof. By the previous theorem, we know that the set $\mathcal{S}$ of simple functions that vanish outside a set of finite measure are dense in $L^p(X,\mu)$. Therefore, it suffices to show that we can approximate such simple functions in the L^p norm with elements of $C_c(\mathbf{R})$. First assume that $E \subseteq \mathbf{R}$ is a measurable set with $\mu(E) < \infty$, and let $\varepsilon > 0$ be given. Then by the regularity properties of Lebesgue measure, there is an open set $U \supseteq E$ satisfying $\mu(U \backslash E) < \frac{\varepsilon^p}{2}$, and there is a compact set $K \subseteq E$ with $\mu(E \backslash K) < \frac{\varepsilon^p}{2}$. By Urysohn's lemma[2], there is a function $\varphi \in C_c(\mathbf{R})$ satisfying the following three conditions:

[2] We actually do not need the full power of Urysohn's lemma here. It is possible to prove a weaker version for locally compact metric spaces, and then use it to construct the function we need. See Exercise 6.2.3.

1. $0 \le \varphi(x) \le 1$ for all $x \in \mathbf{R}$;
2. $\varphi(x) = 1$ for all $x \in K$; and
3. $\varphi(x) = 0$ for all $x \in U^c$.

Notice that φ agrees with χ_E on K and on U^c, so

$$\mu\big(\{x \in \mathbf{R} \mid \varphi(x) \neq \chi_E(x)\}\big) < \varepsilon.$$

Thus

$$\|\varphi - \chi_E\|_p^p = \int_{\mathbf{R}} |\varphi - \chi_E|^p \, d\mu \le \int_{U \setminus K} |\varphi - \chi_E|^p \, d\mu \le \mu(U \setminus K) < \varepsilon^p,$$

or $\|\varphi - \chi_E\|_p < \varepsilon$. Now we can approximate an arbitrary element of $\mathcal{S}$ by writing it as a linear combination of characteristic functions, and then approximating each of those functions sufficiently closely. □

Remark 6.2.11. The previous two results do not hold for $L^\infty(X, \mu)$. The simple functions vanishing outside a set of finite measure are not dense in $L^\infty(X, \mu)$. In fact, notice that *every* simple function belongs to $L^\infty(X, \mu)$, and the set of all simple functions is dense in $L^\infty(X, \mu)$. (This follows from the fact that any bounded measurable function can be approximated *uniformly* by a sequence of simple functions.) Also, the set $C_c(\mathbf{R})$ is not dense in $L^\infty(\mathbf{R}, \mu)$, since the completion of $C_c(\mathbf{R})$ with respect to the norm $\|\cdot\|_\infty$ is $C_0(X)$.

We now close this section with some remarks about finite-dimensional vector spaces. First, we need a definition.

Definition 6.2.12. Let V be a vector space. We say that two norms $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ on V are **equivalent** if there exist constants $C_1, C_2 > 0$ such that

$$C_1 \|x\|_\alpha \le \|x\|_\beta \le C_2 \|x\|_\alpha.$$

It follows quickly from the definition that equivalent norms yield the same open sets and the same convergent sequences. In other words, they generate the same *topology* on V. It is also easy to check that this notion of equivalence defines an equivalence relation on the set of norms on V.

Theorem 6.2.13. *Let V be a finite-dimensional vector space over $\mathbf{R}$ or $\mathbf{C}$. Any two norms on V are equivalent.*

Proof. Let $n = \dim(V)$, and choose a basis $\{x_1, x_2, \ldots, x_n\}$ for V. Given a vector $x = \sum_{i=1}^n \alpha_i x_i \in V$, we define

$$\|x\|_\infty = \max\{|\alpha_1|, |\alpha_2|, \ldots, |\alpha_n|\}.$$

Then $\|\cdot\|_\infty$ is clearly a norm on V. Suppose $\|\cdot\|$ is another norm. Since equivalence of norms defines an equivalence relation, it suffices to show that $\|\cdot\|$ is equivalent to $\|\cdot\|_\infty$. To this end, observe that for all $x \in V$ we have

$$\|x\| = \left\| \sum_{i=1}^n \alpha_i x_i \right\| \leq \sum_{i=1}^n |\alpha_i| \|x_i\| \leq \|x\|_\infty \sum_{i=1}^n \|x_i\|.$$

If we set $C_2 = \sum_{i=1}^n \|x_i\|$, then we have $\|x\| \leq C_2 \|x\|_\infty$ for all $x \in V$. Now consider the set

$$A = \{x \in V : \|x\|_\infty = 1\}.$$

Then A is compact.[3] The map $\|\cdot\| : V \to [0, \infty)$ is continuous, so the image of A under this map is compact in $[0, \infty)$. The image cannot contain 0 (since $\|\cdot\|$ is a norm), so there is a constant C_1 such that $C_1 \leq \|x\|$ for all $x \in A$. Now given any $x \in V$, notice that $x/\|x\|_\infty \in A$, so we have

$$C_1 \leq \left\| \frac{x}{\|x\|_\infty} \right\| = \frac{\|x\|}{\|x\|_\infty},$$

or $C_1 \|x\|_\infty \leq \|x\|$. Therefore,

$$C_1 \|x\|_\infty \leq \|x\| \leq C_2 \|x\|_\infty,$$

so $\|\cdot\|$ and $\|\cdot\|_\infty$ are equivalent. □

When we choose a basis for a vector space V, we are really setting up an isomorphism of V with $\mathbf{F}^n$, where $\mathbf{F} = \mathbf{R}$ or $\mathbf{F} = \mathbf{C}$. We are then tempted to use this isomorphism to call on some of the topological properties of $\mathbf{F}^n$. However, we do not know whether the isomorphism is continuous or not. We will address this question in the next section, where we consider continuity properties of linear operators on Banach spaces. We do have the following immediate corollary for $\mathbf{F}^n$, however.

Corollary 6.2.14. *Any two norms on $\mathbf{R}^n$ or $\mathbf{C}^n$ are equivalent. As a result, $\mathbf{R}^n$ and $\mathbf{C}^n$ are Banach spaces with respect to any norm.*

[3] We are tempted to invoke the Heine-Borel theorem here, and conclude that A is compact because it is closed and bounded. However, we should be careful since we only know that V is isomorphic to $\mathbf{R}^n$ or $\mathbf{C}^n$. It is not hard to check (using the definition of $\|\cdot\|_\infty$) that A is sequentially compact, hence compact.

Proof. The first assertion is evident. We know $\mathbf{R}^n$ and $\mathbf{C}^n$ are complete with respect to their Euclidean norms, hence they are complete with respect to any norm that is equivalent to the Euclidean norm. But this is true of all norms, so any norm makes $\mathbf{R}^n$ or $\mathbf{C}^n$ into Banach a space. $\square$

Exercises for Section 6.2

Exercise 6.2.1. Let X be a locally compact metric space. Prove that $C_0(X)$ is a subspace of $C_b(X)$, and that $C_0(X)$ is closed in $C_b(X)$. Conclude that $C_0(X)$ is a Banach space.

Exercise 6.2.2. Let V be a normed vector space, and let $W \subseteq V$ be a subspace. Show that the closure $\overline{W}$ is also a subspace of V.

Exercise 6.2.3. This exercise will guide you to a proof of Urysohn's lemma for locally compact metric spaces.

(a) Let X be a locally compact metric space, and fix $x_0 \in X$ and $r > 0$. Define $f : \mathbf{R} \to [0,1]$ to be the piecewise linear function

$$f(t) = \begin{cases} 1 & \text{if } t \le r \\ 1 - \frac{2}{t}(t-r) & \text{if } r < t \le \frac{3r}{2} \\ 0 & \text{if } t > \frac{3r}{2}. \end{cases}$$

Now define $\varphi : X \to [0,1]$ by $\varphi(x) = f(d(x,x_0))$. Check that φ is continuous, $\varphi(x) = 1$ for all $x \in B(x_0;r)$, and $\operatorname{supp}(\varphi) \subseteq B(x_0;2r)$.

(b) Suppose $K \subseteq X$ is compact and $U \subseteq X$ is open with $K \subseteq U$. Show that there is a function $\varphi \in C_c(X)$ satisfying the following conditions:

- $0 \le \varphi(x) \le 1$ for all $x \in X$;
- $\varphi(x) = 1$ for all $x \in K$;
- $\varphi(x) = 0$ for all $x \in U^c$.

Exercise 6.2.4. Let V be a vector space, and suppose $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ are two equivalent norms on V.

(a) Fix $x_0 \in V$ and $r > 0$. Prove that the ball

$$B_\alpha(x_0;r) = \{x \in V \mid \|x - x_0\|_\alpha < r\}$$

is open with respect to the norm $\|\cdot\|_\beta$, and that the ball

$$B_\beta(x_0;r) = \{x \in V \mid \|x - x_0\|_\beta < r\}$$

is open with respect to the norm $\|\cdot\|_\alpha$. Conclude that a set $U \subseteq V$ is open with respect to $\|\cdot\|_\alpha$ if and only if it is open with respect to $\|\cdot\|_\beta$.

(b) Let $(x_n)_{n=1}^{\infty}$ be a sequence in V. Prove that $(x_n)_{n=1}^{\infty}$ is Cauchy (respectively, convergent) with respect to $\|\cdot\|_{\alpha}$ if and only if it is Cauchy (respectively, convergent) with respect to $\|\cdot\|_{\beta}$. Conclude that V is complete with respect to the norm $\|\cdot\|_{\alpha}$ if and only if it is complete with respect to $\|\cdot\|_{\beta}$.

6.3 Bounded Operators

As one learns in a course on linear algebra, it is important to study not just vector spaces, but maps between them that respect the algebraic structure. Recall that if V and W are vector spaces over a field $\mathbf{F}$, we say a map $T : V \to W$ is *linear* if

$$T(x + y) = T(x) + T(y)$$

and

$$T(\alpha x) = \alpha T(x)$$

for all $x, y \in V$ and all $\alpha \in \mathbf{F}$. In linear algebra, we usually call such maps *linear transformations*. In functional analysis, another term is commonly used.

Definition 6.3.1. Let V and W be vector spaces over a field $\mathbf{F}$. A linear map $T : V \to W$ is called a **(linear) operator**.

We will usually refer to a linear operator as simply an "operator", where it is understood that the map is assumed to be linear.

Remark 6.3.2. It is customary in functional analysis to express the action of an operator on a vector via juxtaposition: if $T : V \to W$ is an operator and $x \in V$, we write Tx in place of the usual function notation $T(x)$. This notation should be reminiscent of the finite-dimensional setting, where every linear operator corresponds to multiplication by an appropriate matrix.

Since we are currently in the business of studying *normed* vector spaces, it would be a good idea to single out the operators that play nicely with the relevant norms. Throughout the rest of this section, we assume the field of scalars for any vector space is $\mathbf{C}$.

Definition 6.3.3. Let V and W be normed vector spaces. We say an operator $T : V \to W$ is **bounded** if there exists a constant $K \geq 0$ such that

$$\|Tx\| \leq K\|x\|$$

for all $x \in V$.

Example 6.3.4. Let $(X, \mathcal{M}, \mu)$ be a measure space, and fix a function $f \in L^\infty(X, \mu)$. Given $1 \le p \le \infty$, we define $T_f : L^p(X, \mu) \to L^p(X, \mu)$ by

$$T_f g = fg.$$

Of course we should check that T_f actually maps into $L^p(X, \mu)$. If $1 \le p < \infty$, then we have

$$\int_X |fg|^p \, d\mu \le \int_X \|f\|_\infty^p |g|^p \, d\mu = \|f\|_\infty^p \int_X |g|^p \, d\mu < \infty$$

for all $g \in L^p(X, \mu)$. Notice that this computation also shows that

$$\|T_f g\|_p = \|fg\|_p \le \|f\|_\infty \|g\|_p,$$

so taking $K = \|f\|_\infty$ shows that T_f is bounded. If $p = \infty$, then it is not hard to check that

$$\|fg\|_\infty \le \|f\|_\infty \|g\|_\infty,$$

so $T_f : L^\infty(X, \mu) \to L^\infty(X, \mu)$ is bounded.

Example 6.3.5. Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $1 \le p, q \le \infty$ are conjugate exponents. Given a function $f \in L^p(X, \mu)$, define $T_f : L^q(X, \mu) \to L^1(X, \mu)$ by

$$T_f g = fg.$$

Notice that T_f does map into $L^1(X, \mu)$, since Hölder's inequality guarantees that $fg \in L^1(X, \mu)$ and

$$\|fg\|_1 \le \|f\|_p \|g\|_q.$$

Hence T_f is bounded as well, and we could take K to be $\|f\|_p$.

Example 6.3.6. Define $S : \ell^2 \to \ell^2$ by

$$S(x_1, x_2, \dots) = (0, x_1, x_2, \dots).$$

It is easy to see that S is bounded—in fact,

$$\|S(x_1, x_2, \dots)\|_2 = \|(0, x_1, x_2, \dots)\|_2 = \left(\sum_{j=1}^\infty |x_j|^2 \right)^{1/2} = \|(x_1, x_2, \dots)\|_2,$$

so $\|Sx\|_2 = \|x\|_2$ for all $x \in \ell^2$. (In other words, S is an *isometry*.) The operator S has a special name—it is called the **unilateral shift** on ℓ^2.

We can define a related operator $S^* : \ell^2 \to \ell^2$ by shifting our sequences to the left instead:

$$S^*(x_1, x_2, \dots) = (x_2, x_3, \dots).$$

(The reason for using the notation S^* will become clear later.) Clearly

$$\|S^*(x_1, x_2, \ldots)\|_2 = \left(\sum_{j=2}^{\infty} |x_j|^2\right)^{1/2} \leq \left(\sum_{j=1}^{\infty} |x_j|^2\right)^{1/2} = \|(x_1, x_2, \ldots)\|_2,$$

so S^* is bounded.

Example 6.3.7. Let X be a compact metric space, and suppose μ is a finite Borel measure on X. Define $T : C(X) \to \mathbf{C}$ by

$$Tf = \int_X f \, d\mu.$$

Then T is bounded, since

$$|Tf| = \left|\int_X f \, d\mu\right| \leq \int_X |f| \, d\mu \leq \int_X \|f\|_\infty \, d\mu = \|f\|_\infty \cdot \mu(X).$$

As we will see later, the operator T is an example of a **bounded linear functional** on $C(X)$.

Example 6.3.8. Consider the vector space $C^\infty([0,1])$ of functions $f : [0,1] \to \mathbf{C}$ that have continuous derivatives of all orders, equipped with the supremum norm $\|\cdot\|_\infty$. Define an operator $D : C^\infty([0,1]) \to C^\infty([0,1])$ by

$$(Df)(x) = f'(x).$$

Then D is a linear operator, but it is not bounded. Notice that if we take $f(x) = x^n$ for some $n \in \mathbf{N}$, we have $\|f\|_\infty = 1$. However,

$$(Df)(x) = nx^{n-1},$$

so $\|Df\|_\infty = n\|f\|_\infty$. By taking n to be arbitrarily large, we can see that D cannot possibly be bounded.

So far it seems that we have ignored the simple examples of operators between finite dimensional spaces. Well, there is little need to single out specific examples in this case—*every* operator defined on a finite-dimensional space is bounded.

Theorem 6.3.9. *Let V and W be normed vector spaces, and assume V is finite-dimensional. If $T : V \to W$ is an operator, then T is bounded.*

Proof. It should not be hard to convince oneself that if $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ are equivalent norms on V, then an operator $T : V \to W$ is bounded with respect to $\|\cdot\|_\alpha$ if and only if it is bounded with respect to $\|\cdot\|_\beta$. Therefore, we begin as in the proof of

Theorem 6.2.13: let $n = \dim(V)$, choose a basis $\{x_1, x_2, \ldots, x_n\}$ for V, and define a norm $\|\cdot\|_\infty$ on V by

$$\|\alpha_1 x_1 + \alpha_2 x_2 + \cdots + \alpha_n x_n\|_\infty = \max\{|\alpha_1|, |\alpha_2|, \ldots, |\alpha_n|\}.$$

Since all norms on V are equivalent by Theorem 6.2.13, we may as well work with $\|\cdot\|_\infty$. Suppose $T : V \to W$ is an operator, and let $x = \sum_{j=1}^n \alpha_j x_j \in V$. Then

$$Tx = \sum_{j=1}^n \alpha_j T x_j,$$

so

$$\|Tx\| = \left\| \sum_{j=1}^n \alpha_j T x_j \right\| \leq \sum_{j=1}^n |\alpha_j| \|Tx_j\| \leq \sum_{j=1}^\infty \|x\|_\infty \|Tx_j\|.$$

Set $K = \sum_{j=1}^n \|Tx_j\|$. Then we have $\|Tx\| \leq K\|x\|_\infty$ for all $x \in V$, so T is bounded. $\square$

As mentioned in the last section, we plan to show that every finite-dimensional normed vector space over **C** is a Banach space. The previous theorem is a big step toward this result, though we still need to connect the notion of boundedness to continuity. Based on the definition, it should not be hard to see that a bounded operator is automatically continuous. In fact, it turns out that boundedness for linear operators is equivalent to continuity.

Proposition 6.3.10. *Let V and W be normed vector spaces, and suppose $T : V \to W$ is an operator. The following are equivalent.*

1. *T is bounded.*
2. *T is continuous.*
3. *T is continuous at a single point.*

Proof. Assume first that T is bounded, and choose $K > 0$ such that $\|Tx\| \leq K\|x\|$ for all $x \in V$. Let $\varepsilon > 0$ be given, and set $\delta = \frac{\varepsilon}{K}$. Then $\|x - y\| < \delta$ implies

$$\|Tx - Ty\| \leq K\|x - y\| < K \cdot \frac{\varepsilon}{K} = \varepsilon.$$

Hence T is continuous on V. (Indeed, it is uniformly continuous.)

It is clear that (2) implies (3), so we simply need to show that (3) implies (1). Suppose T is continuous at a point $y \in V$. Then there is a $\delta > 0$ such that $\|x - y\| \leq \delta$ implies $\|Tx - Ty\| \leq 1$. Let $x \in V$ with $\|x\| = 1$. Then $\|\delta x\| = \delta$, so

$$\|(\delta x + y) - y\| \leq \delta.$$

and it follows that

$$\big\|T(\delta x + y) - Ty\big\| \le 1.$$

But the left hand side is simply $\|T(\delta x)\| = \delta\|Tx\|$, so

$$\|Tx\| \le \tfrac{1}{\delta}.$$

Now if $x \in V$ is any nonzero vector, then $z = x/\|x\|$ is a unit vector. Then by our previous work we have $\|Tz\| \le \frac{1}{\delta}$, whence

$$\|Tx\| \le \tfrac{1}{\delta}\|x\|.$$

Note that if $x = 0$, it is obvious that $\|Tx\| \le \frac{1}{\delta}\|x\|$. Thus $\|Tx\| \le \frac{1}{\delta}\|x\|$ for all $x \in V$, so T is bounded. □

Putting together the last two theorems, we can now prove the following result about finite-dimensional spaces.

Theorem 6.3.11. *Let V be a finite-dimensional normed vector space. Then V is a Banach space.*

Proof. Let $n = \dim(V)$, and choose a basis $\{x_1, x_2, \dots, x_n\}$ for V. Define a map $T : V \to \mathbf{C}^n$ by

$$T(\alpha_1 x_1 + \alpha_2 x_2 + \cdots + \alpha_n x_n) = (\alpha_1, \alpha_2, \dots, \alpha_n).$$

It is easy to see that T defines a linear isomorphism of V onto $\mathbf{C}^n$. Then T is bounded, so there is a constant $K > 0$ such that $\|Tx\| \le K\|x\|$ for all $x \in V$. Let $(y_j)_{j=1}^{\infty}$ be a Cauchy sequence in V. Let $\varepsilon > 0$ be given, and find an N such that

$$\|y_j - y_k\| < \frac{\varepsilon}{K}$$

for all $j, k \ge N$. Then

$$\|Ty_j - Ty_k\| \le K\|y_j - y_k\| < \varepsilon$$

for all $j, k \ge N$, so the sequence $(Ty_j)_{j=1}^{\infty}$ in $\mathbf{C}^n$ is Cauchy. Since $\mathbf{C}^n$ is complete, there is a point $z \in \mathbf{C}^n$ such that $Ty_j \to z$. By a well-known fact from linear algebra, the inverse map $T^{-1} : \mathbf{C}^n \to V$ is linear, so it is automatically continuous. Therefore,

$$y_j = T^{-1}Ty_j \to T^{-1}z.$$

Thus $(y_j)_{j=1}^{\infty}$ converges in V, so V is complete. □

While it is interesting to study individual operators between normed vector spaces, one can gain quite a bit of insight by collectively considering *all* bounded operators between two Banach spaces. To this end, given two normed vector spaces V and W, we let

$$B(V,W) = \{\text{bounded operators } T : V \to W\}.$$

In the special case where $V = W$, we simply write $B(V)$ in place of $B(V,V)$. It is not hard to check that $B(V,W)$ is a vector space under the "pointwise" operations

$$(T+S)x = Tx + Sx$$

and

$$(\alpha T)x = \alpha(Tx)$$

for $x, y \in V$ and $\alpha \in \mathbf{C}$. In fact, it is possible to equip $B(V,W)$ with a norm in a natural way.

Definition 6.3.12. Let V and W be normed vector spaces. Given $T \in B(V,W)$, we define the **operator norm** of T to be

$$\|T\| = \sup_{\|x\| \leq 1} \|Tx\|.$$

Intuitively, the operator norm measures the maximum amount by which an operator can stretch the unit ball in any given direction.

Theorem 6.3.13. *Let V and W be normed vector spaces. The operator norm defines a norm on $B(V,W)$.*

Proof. It is fairly clear that $\|\cdot\| : B(V,W) \to [0,\infty)$ is positive definite and homogeneous, so we need only check that it satisfies the triangle inequality. Let $S, T \in B(V,W)$, and observe that if $\|x\| \leq 1$, then

$$\|(S+T)x\| \leq \|Sx\| + \|Tx\| \leq \|S\| + \|T\|.$$

The right side is therefore an upper bound for the set $\{\|(S+T)x\|\}_{\|x\|\leq 1}$, so it must be at least as large of the supremum of this set. Hence

$$\sup_{\|x\|\leq 1} \|(S+T)x\| \leq \|S\| + \|T\|.$$

Therefore, $\|S+T\| \leq \|S\| + \|T\|$, and $\|\cdot\|$ is in fact a norm. $\square$

We have very nearly encountered the operator norm already in our study of bounded operators. In particular, it can be taken as the constant K in the definition of boundedness.

Theorem 6.3.14. *If $T \in B(V, W)$, then*

$$\|Tx\| \leq \|T\|\|x\|$$

for all $x \in V$. In fact, $\|T\| = \inf\{K \geq 0 : \|Tx\| \leq K\|x\| \text{ for all } x \in V\}$.

Proof. Notice that the first assertion clearly holds when $x = 0$. Observe that if $x \in V$ and $\|x\| = 1$, then

$$\|Tx\| \leq \|T\|$$

by definition. Given an arbitrary $x \neq 0$, set $z = x/\|x\|$. Then $\|z\| = 1$, and

$$\frac{\|Tx\|}{\|x\|} = \|Tz\| \leq \|T\|,$$

so $\|Tx\| \leq \|T\|\|x\|$.

To prove the second assertion, we begin by defining

$$A = \{K \geq 0 : \|Tx\| \leq K\|x\| \text{ for all } x \in V\}.$$

We have already shown that $\|Tx\| \leq \|T\|\|x\|$ for all $x \in V$, so we see that $\|T\| \in A$. Therefore, $\|T\| \geq \inf A$. Now suppose $K \in A$. Then for all $x \in V$ with $\|x\| \leq 1$, we have

$$\|Tx\| \leq K\|x\| \leq K,$$

so K is an upper bound for the set $\{\|Tx\| : \|x\| \leq 1\}$. As a result, K is greater than or equal to the supremum of this set, which is precisely $\|T\|$. In other words, $\|T\| \leq K$. This holds for all $K \in A$, so $\|T\|$ is a lower bound for A. It must then be the case that $\|T\| \leq \inf A$. Therefore, $\|T\| = \inf A$. □

Theorem 6.3.15. *Let V, W, and Z be normed vector spaces. If $T \in B(V, W)$ and $S \in B(W, Z)$, then $S \circ T \in B(V, Z)$ and $\|S \circ T\| \leq \|S\|\|T\|$.*

Proof. Observe that if $x \in V$ and $\|x\| \leq 1$, then

$$\|(S \circ T)x\| \leq \|S\|\|Tx\| \leq \|S\|\|T\|\|x\| \leq \|S\|\|T\|$$

by the previous theorem. Hence

$$\sup_{\|v\| \leq 1} \|(S \circ T)(v)\| \leq \|S\|\|T\|,$$

so $\|S \circ T\| \leq \|S\|\|T\|$. This computation also shows that $S \circ T \in B(V, Z)$. □

Theorem 6.3.16. *Suppose V is a normed vector space and W is a Banach space. Then $B(V,W)$ is a Banach space with respect to the operator norm.*

Proof. The only interesting thing to check is that $B(V,W)$ is complete. Assume $(T_n)_{n=1}^{\infty}$ is a Cauchy sequence in $B(V,W)$ and fix $x \in V$. Note that

$$\|T_n x - T_m x\| = \|(T_n - T_m)x\| \leq \|T_n - T_m\| \|x\|,$$

which then implies that $(T_n x)_{n=1}^{\infty}$ is a Cauchy sequence in W. Since W is complete, there exists a vector $Tx \in W$ such that $T_n x \to Tx$. Thus we get a function $T : V \to W$ which is easily checked to be linear. Furthermore, we claim that the set $\{\|T_n\|\}$ is bounded. Since $(T_n)_{n=1}^{\infty}$ is a Cauchy sequence, for any $\varepsilon > 0$ and n, m sufficiently large, we have

$$|\|T_n\| - \|T_m\|| \leq \|T_n - T_m\| < \varepsilon.$$

Therefore, $(\|T_n\|)_{n=1}^{\infty}$ is a Cauchy sequence of real numbers, hence convergent. In particular, $(T_n)_{n=1}^{\infty}$ is bounded, so there exists $K \geq 0$ such that $\|T_n\| \leq K$ for all n. Then

$$\|Tx\| = \lim \|T_n x\| \leq \limsup \|T_n\| \|x\| \leq K\|x\|,$$

so T is bounded. Hence $T \in B(V,W)$. It remains to show that $T_n \to T$ in $B(V,W)$. Given $\varepsilon > 0$, there is an N such that $n, m \geq N$ implies that

$$\|T_n - T_m\| \leq \frac{\varepsilon}{2}.$$

Observe that given $x \in V$,

$$\|T_n x - Tx\| = \lim_{n \to \infty} \|T_n x - T_m x\|.$$

Since $\|T_n x - T_m x\| \leq \|T_n - T_m\| \|x\|$, for each $n \in \mathbf{N}$ we have

$$\|T_n x - Tx\| \leq \limsup_{m \to \infty} \|T_n - T_m\| \|x\|.$$

Thus if $n \geq N$,

$$\|T_n - T\| = \sup_{\|x\| \geq 1} \|T_n x - Tx\| \leq \limsup_{m \to \infty} \|T_n - T_m\| \leq \frac{\varepsilon}{2} < \varepsilon.$$

Therefore, $T_n \to T$ in $B(V,W)$, and $B(V,W)$ is complete. $\square$

The last theorem implies that if V is a Banach space, then $B(V)$ is itself a Banach space with respect to the operator norm. We have also shown that the composition of any two operators $S, T \in B(V)$ is again an element of $B(V)$, and that

$$\|S \circ T\| \leq \|S\|\|T\|.$$

It is common in this case to write ST in place of $S \circ T$, and think of ST as the "product" of two operators. Thus we have endowed $B(V)$ with a multiplication operation, so $B(V)$ has additional structure beyond that of a Banach space.

Definition 6.3.17. Let A be an algebra over $\mathbf{C}$. We say that A is a **Banach algebra** if A is a Banach space with the property that

$$\|ab\| \leq \|a\|\|b\|$$

for all $a, b \in A$. In this case we say that the norm is **submultiplicative**.

Observe that our previous work shows that if V is a Banach space, then $B(V)$ is a Banach algebra.

Exercises for Section 6.3

Exercise 6.3.1. Prove that each function defines a bounded operator between the specified normed vector spaces, and compute its operator norm.

(a) Define $T : C_0(\mathbf{R}) \to \mathbf{C}$ by $Tf = f(0)$, where $C_0(\mathbf{R})$ is equipped with the supremum norm.

(b) Fix $\theta \in [0, 2\pi)$, and define $T : \mathbf{C}^2 \to \mathbf{C}^2$ by

$$T\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix},$$

where $\mathbf{C}^2$ is equipped with the Euclidean norm $\|(x_1, x_2)\| = \sqrt{|x_1|^2 + |x_2|^2}$.

(c) Let μ denote Lebesgue measure on $\mathbf{R}$, fix $K \in L^2(\mathbf{R}^2, \mu \times \mu)$, and define $T : L^2(\mathbf{R}, \mu) \to L^2(\mathbf{R}, \mu)$ by

$$(Tf)(x) = \int_{\mathbf{R}} K(x, y) f(y)\, d\mu(y).$$

Exercise 6.3.2. Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $1 < p, q < \infty$ with $\frac{1}{p} + \frac{1}{q} = 1$. Let $g \in L^q(X, \mu)$, and define $T : L^p(X, \mu) \to \mathbf{C}$ by

$$Tf = \int_X fg\, d\mu.$$

Prove that T is a bounded operator, and determine its operator norm.

Exercise 6.3.3. Let V and W be normed vector spaces. We say an operator $T : V \to W$ is **isometric** if $\|Tx\| = \|x\|$ for all $x \in V$. Prove that any isometric operator is injective.

Exercise 6.3.4. Let V and W be Banach spaces, let $V_0 \subseteq V$ be a dense subspace, and suppose $T : V_0 \to W$ is bounded. Show that there is a unique bounded operator $\overline{T} : V \to W$ extending T (i.e., $\overline{T}x = Tx$ for all $x \in V_0$), and that $\|\overline{T}\| = \|T\|$.

Exercise 6.3.5. Let $\mathcal{X}$ be a Banach space, and let $I : \mathcal{X} \to \mathcal{X}$ denote the identity operator. Suppose $T \in B(\mathcal{X})$ satisfies $\|I - T\| < 1$. Prove that the series

$$\sum_{n=0}^{\infty} (I - T)^n$$

converges to an operator $S \in B(\mathcal{X})$ (here we interpret $(I - T)^0$ as I), and that $ST = TS = I$. Conclude that T is invertible.

6.4 Three Fundamental Theorems

In this section we will set about proving three of the most important theorems in functional analysis: the Open Mapping Theorem, the Closed Graph Theorem, and the Principle of Uniform Boundedness. All of them have interesting applications to the study of bounded operators on Banach spaces.

We begin our investigation with the Open Mapping Theorem. First, recall that if X and Y are metric spaces, a map $f : X \to Y$ is called an **open map** if $f(U)$ is open for every open set $U \subseteq X$.

Theorem 6.4.1 (Open Mapping Theorem). *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces and that $T \in B(\mathcal{X}, \mathcal{Y})$ is surjective. Then T is an open map.*

The proof of the Open Mapping Theorem relies on the Baire Category Theorem. We will follow the approach of [Ped89] and prove a small technical lemma first.

Lemma 6.4.2. *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces and $T \in B(\mathcal{X}, \mathcal{Y})$. Given $r > 0$, suppose $T(B(0;1))$ contains a dense subset of the open ball $B(0;r) \subseteq Y$. Then for any $0 < \delta < r$,*

$$B(0;\delta) \subseteq T(B(0;1)).$$

Proof. For each $0 < \varepsilon < 1$, we will prove that

$$B(0;(1-\varepsilon)r) \subseteq T(B(0;1)).$$

Let $A = T(B(0;1))$. Given $y \in B(0;r)$, there is a $y_1 \in A$ such that

$$\|y - y_1\| < \varepsilon r,$$

since we have assumed that A contains a dense subset of $B(0;r)$. Then

$$\left\|\tfrac{1}{\varepsilon}y - \tfrac{1}{\varepsilon}y_1\right\| < r.$$

Thus $\frac{1}{\varepsilon}y - \frac{1}{\varepsilon}y_1 \in B(0;r)$, so there exists $y_2' \in A$ such that

$$\left\|\tfrac{1}{\varepsilon}y - \tfrac{1}{\varepsilon}y_1 - y_2'\right\| < \varepsilon r,$$

and from the homogeneity of the norm we obtain

$$\left\|y - y_1 - \varepsilon y_2'\right\| < \varepsilon^2 r.$$

If we set $y_2 = \varepsilon y_2' \in \varepsilon A$, we see that $\|y - y_1 - y_2\| < \varepsilon^2 r$. Proceeding inductively, we can construct a sequence $(y_k)_{k=1}^{\infty}$ such that $y_k \in \varepsilon^{k-1}A$ for all k and

$$\left\|y - \sum_{k=1}^{n} y_k\right\| < \varepsilon^n r$$

for all n. Since $y_k \in \varepsilon^{k-1}A$, we claim that there exists $x_k \in \mathfrak{X}$ such that $Tx_k = y_k$ and $\|x_k\| < \varepsilon^{k-1}$. To see this, we observe that $\frac{1}{\varepsilon^{k-1}}y_k \in A$, so there exists $x_k' \in B(0;1)$ such that $Tx_k' = \frac{1}{\varepsilon^{k-1}}y_k$. Consequently, $y_k = \varepsilon^{k-1}Tx_k' = T(\varepsilon^{k-1}x_k')$. Now set $x_k = \varepsilon^{k-1}x_k'$, and observe that $\|x_k\| = \varepsilon^{k-1}\|x_k'\| < \varepsilon^{k-1}$. It follows that the series $\sum_{k=1}^{\infty} x_k$ is absolutely convergent, hence it converges to some point $x \in \mathfrak{X}$. Then

$$\begin{aligned}
\|y - Tx\| &= \lim_{n\to\infty}\left\|y - \sum_{k=1}^{n} Tx_k\right\| \\
&= \lim_{n\to\infty}\left\|y - \sum_{k=1}^{n} y_k\right\| \\
&\le \limsup_{n\to\infty} \varepsilon^n r \\
&= 0.
\end{aligned}$$

Thus $Tx = y$. Furthermore,

$$\|x\| \le \sum_{k=1}^{\infty}\|x_k\| \le \sum_{k=1}^{\infty}\varepsilon^{k-1} = \frac{1}{1-\varepsilon},$$

so $x \in \frac{1}{1-\varepsilon}B(0;1)$. Hence

$$y \in T\left(\tfrac{1}{1-\varepsilon}B(0;1)\right) = \tfrac{1}{1-\varepsilon}T(B(0;1)) = \tfrac{1}{1-\varepsilon}A.$$

Therefore, $B(0;r) \subseteq \frac{1}{1-\varepsilon}A$, or equivalently,

$$(1-\varepsilon)B(0;r) = B(0;(1-\varepsilon)r)) \subseteq A. \qquad \square$$

Proof of the Open Mapping Theorem. We need to show that $T(U)$ is open in $\mathcal{Y}$ whenever U is open in $\mathcal{X}$. If $x \in U$, then there exists $r > 0$ such that $B(x;r) \subseteq U$, and it suffices to see that $T(B(x;r))$ is a neighborhood of Tx. But $B(x;r) = x + B(0;r)$, so we just need to show that $T(B(0;r))$ is a neighborhood of 0 in $\mathcal{Y}$.

Since T is surjective, we can write

$$\mathcal{Y} = T(\mathcal{X}) = \bigcup_{n=1}^{\infty} T(B(0;n)).$$

By the Baire Category Theorem, it cannot be the case that each of the sets $T(B(0;n))$ is nowhere dense. Hence there exist $n \in \mathbf{N}$, $y \in \mathcal{Y}$, and $\varepsilon > 0$ such that

$$B(y;\varepsilon) \subseteq \overline{T(B(0;n))}.$$

Thus $T(B(0;n))$ is dense in $B(y;\varepsilon)$.[4] Consequently, $T(B(0;1))$ is dense in $B(y;\frac{\varepsilon}{n})$. If $y_1 \in B(0;\frac{\varepsilon}{n})$, then $2y_1 = y + y_1 - (y - y_1)$, so

$$B(y;\tfrac{\varepsilon}{n}) - B(y;\tfrac{\varepsilon}{n}) \subseteq 2B(0;\tfrac{\varepsilon}{n}).$$

Now we claim that $T(B(0;1))$ is dense in $B(0;\frac{\varepsilon}{n})$. Since $2y_1 \in B(y;\frac{\varepsilon}{n}) - B(y;\frac{\varepsilon}{n})$, we can write $2y_1 = a - b$ for some $a, b \in B(y;\frac{\varepsilon}{n})$. Furthermore, since $T(B(0;1))$ is dense in $B(y;\frac{\varepsilon}{n})$, there exist $x_1, x_2 \in B(0;1)$ such that

$$\|Tx_1 - a\| < \varepsilon$$

and

$$\|Tx_2 - b\| < \varepsilon.$$

Thus

$$\|Tx_1 - Tx_2 - 2y_1\| = \|Tx_1 - Tx_2 - (a-b)\| < \varepsilon + \varepsilon = 2\varepsilon,$$

and by the linearity of T and homogeneity, we see that

$$\|Tx_1 - Tx_2 - 2y_1\| = \|T(x_1 - x_2) - 2y_1\| = 2\big\|T\left(\tfrac{1}{2}(x_1 - x_2)\right) - y_1\big\|,$$

so

$$\big\|T\left(\tfrac{1}{2}(x_1 - x_2)\right) - y_1\big\| < \varepsilon.$$

Note that $\big\|\frac{1}{2}(x_1 - x_2)\big\| \le \frac{1}{2}(\|x_1\| + \|x_2\|) < 1$, so $T(B(0;1))$ is dense in $B(0;\frac{\varepsilon}{n})$. The hypotheses of the previous lemma are thus satisfied, so if $0 < \delta < \frac{\varepsilon}{n}$, then

$$B(0;\delta) \subseteq T(B(0;1)).$$

Therefore, for any $r > 0$, $T(B(0;r)) = rT(B(0;r)) \supseteq B(0;r\delta)$. Hence $T(B(0;r))$ is a neighborhood of 0, and we are done. $\square$

[4] More appropriately, $T(B(0;n)) \cap B(y;\varepsilon)$ is dense in $B(y;\varepsilon)$.

The Open Mapping Theorem has several important consequences. First, recall from linear algebra that if $T \in B(\mathcal{X}, \mathcal{Y})$ and T is bijective, then the inverse map $T^{-1} : \mathcal{Y} \to \mathcal{X}$ is also linear. In fact, it must also be bounded.

Corollary 6.4.3. *Suppose that $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces and that $T \in B(\mathcal{X}, \mathcal{Y})$ is a bijection. Then $T^{-1} \in B(\mathcal{Y}, \mathcal{X})$.*

Proof. By the Open Mapping Theorem, T is an open map. Therefore, if $U \subseteq \mathcal{X}$ is open, then $(T^{-1})^{-1}(U) = T(U)$ is open in $\mathcal{Y}$. Hence T^{-1} is continuous. But continuous linear maps are bounded, so T^{-1} is bounded. □

Corollary 6.4.4. *Suppose V is a vector space equipped with two different norms $\|\cdot\|_1$ and $\|\cdot\|_2$, each of which make V into a Banach space. If there exists $\alpha > 0$ such that*

$$\|x\|_1 \leq \alpha \|x\|_2$$

for all $x \in V$, then there exists $\beta > 0$ such that

$$\|x\|_2 \leq \beta \|x\|_1$$

for all $x \in V$.

Proof. We simply apply Corollary 6.4.3 to the identity map $\mathrm{id} : V \to V$. If there exists $\alpha \geq 0$ such that $\|x\|_1 \leq \alpha \|x\|_2$ for all $x \in V$, then id is bounded as a map from $(V, \|\cdot\|_1)$ to $(V, \|\cdot\|_2)$. Hence the inverse map (which is also the identity map) is bounded, so there is a $\beta \geq 0$ such that $\|x\|_2 \leq \beta \|x\|_1$ for all $x \in V$. □

The second of our three major theorems, the Closed Graph Theorem, also follows from the Open Mapping Theorem. However, we need a little discussion first. Suppose $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces. Then the Cartesian product $\mathcal{X} \times \mathcal{Y}$ is a vector space in a natural way, and it is easy to see that

$$\|(x, y)\| = \max\{\|x\|, \|y\|\}$$

defines a norm on $\mathcal{X} \times \mathcal{Y}$. In fact, it is straightforward to check (Exercise 6.4.1) that $\mathcal{X} \times \mathcal{Y}$ is complete with respect to this norm, hence it is a Banach space.

Theorem 6.4.5 (Closed Graph Theorem)**.** *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces and that $T : \mathcal{X} \to \mathcal{Y}$ is linear. If the **graph** of T,*

$$\Gamma(T) = \big\{(x, T(x)) \in \mathcal{X} \times \mathcal{Y} : x \in \mathcal{X}\big\}$$

is closed in the Banach space $\mathcal{X} \times \mathcal{Y}$, then T is bounded.

Proof. Assume $\Gamma(T)$ is closed in $X \times Y$. It is easy to check that $\Gamma(T)$ is a linear subspace of $\mathcal{X} \times \mathcal{Y}$, so it is a Banach space in its own right. Define $P_1 : \Gamma(T) \to \mathcal{X}$ by $P_1(x, Tx) = x$ and $P_2 : \Gamma(T) \to \mathcal{Y}$ by $P_2(x, Tx) = Tx$. Both maps are linear and bounded, since

$$\|x\|, \|Tx\| \leq \max\{\|x\|, \|Tx\|\} = \|(x, Tx)\|.$$

Furthermore, P_1 is clearly a bijection, so P_1^{-1} is bounded by Corollary 6.4.3. But we can express T as $T = P_2 \circ P_1^{-1}$, so T is bounded. □

In practice, one often uses the following variation of the Closed Graph Theorem to verify boundedness.

Corollary 6.4.6. *Let $\mathcal{X}$ and $\mathcal{Y}$ be Banach spaces, and suppose $T : \mathcal{X} \to \mathcal{Y}$ is linear. Then T is bounded if and only if whenever $(x_n)_{n=1}^{\infty}$ is a sequence in $\mathcal{X}$ such that $x_n \to x$ and $Tx_n \to y$ for some $y \in \mathcal{Y}$, we have $y = Tx$.*

Proof. The forward direction is obvious—if T is bounded, then it is continuous, so $x_n \to x$ implies $Tx_n \to Tx$. Conversely, if $x_n \to x$ and $Tx_n \to y$, then $(x_n, Tx_n) \to (x, y)$ in the product space $\mathcal{X} \times \mathcal{Y}$. But (x_n, Tx_n) belongs to the graph of T for all n, and it follows from the Closed Graph Theorem that $(x, y) \in \Gamma(T)$. Hence $y = Tx$. □

Remark 6.4.7. The last result might seem kind of peculiar. An operator T is bounded if and only if it is continuous, which holds if and only if it takes convergent sequences to convergent sequences. However, observe that if we are to show $x_n \to x$ implies $Tx_n \to Tx$, we must prove first that the sequence $(Tx_n)_{n=1}^{\infty}$ converges. The last result allows us to *assume* $(Tx_n)_{n=1}^{\infty}$ converges, and we are simply left to verify that its limit is Tx.

Our third theorem is the Principle of Uniform Boundedness, which is sometimes called the Banach-Steinhaus theorem. Intuitively, it says that a family of operators that is pointwise bounded is in fact uniformly bounded.

Theorem 6.4.8 (Principle of Uniform Boundedness)**.** *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces and that $\{T_\alpha\}_{\alpha \in I}$ is a family of operators in $B(\mathcal{X}, \mathcal{Y})$. Suppose that for each $x \in \mathcal{X}$, the set $\{T_\alpha x\}_{\alpha \in I}$ is bounded. Then $\{\|T_\alpha\|\}_{\alpha \in I}$ is bounded.*

Proof. For each natural number n, define

$$A_n = \big\{x \in \mathcal{X} \;\big|\; \|T_\alpha x\| \leq n \text{ for all } \alpha \in I\big\}.$$

Notice that we can write

$$A_n = \bigcap_{\alpha \in I} \{x \in \mathcal{X} \mid \|T_\alpha x\| \leq n\},$$

so each A_n is closed. Furthermore, we have $\mathcal{X} = \bigcup_{n=1}^{\infty} A_n$, so the Baire Category Theorem implies that one of the A_n must have nonempty interior. That is, there exists $n \in \mathbf{N}$, $y \in \mathcal{X}$, and $\varepsilon > 0$ such that $B(y;\varepsilon) \subseteq A_n$.

Now suppose $x \in \mathcal{X}$ satisfies $\|x\| < \varepsilon$. Then $x + y \in B(y;\varepsilon)$, so for all $\alpha \in I$ we have

$$\|T_\alpha x\| - \|T_\alpha y\| \le \|T_\alpha(x+y)\| \le n$$

by the reverse triangle inequality. Thus

$$\|T_\alpha x\| \le n + \|T_\alpha y\| \le 2n$$

for all α. Given any nonzero $x \in \mathcal{X}$, set $z = \frac{\varepsilon}{2\|x\|}x$. Then $\|z\| = \frac{\varepsilon}{2} < \varepsilon$, so $\|T_\alpha z\| \le 2n$ for all $\alpha \in I$. It follows from homogeneity that

$$\frac{\varepsilon}{2\|x\|}\|T_\alpha x\| \le 2n,$$

or

$$\|T_\alpha x\| \le \frac{4n}{\varepsilon}\|x\|$$

for all $\alpha \in I$. Therefore, $\|T_\alpha\| \le \frac{4n}{\varepsilon}$ for all α, so the set $\{\|T_\alpha\|\}_{\alpha \in I}$ is bounded. □

One interesting application of the Principle of Uniform Boundedness is the following fact regarding pointwise convergent sequences of operators.

Corollary 6.4.9. *Let $\mathcal{X}$ and $\mathcal{Y}$ be Banach spaces, and suppose $(T_n)_{n=1}^{\infty}$ is a sequence in $B(\mathcal{X},\mathcal{Y})$ such that $(T_n x)_{n=1}^{\infty}$ converges for each $x \in \mathcal{X}$. If we define an operator $T \colon \mathcal{X} \to \mathcal{Y}$ by*

$$Tx = \lim_{n\to\infty} T_n x,$$

then $T \in B(\mathcal{X},\mathcal{Y})$.

Proof. It is easy to check that T is linear, so we just need to verify that it is bounded. Well, notice first that since each sequence $(T_n x)_{n=1}^{\infty}$ converges in $\mathcal{Y}$, the sequence $(\|T_n x\|)_{n=1}^{\infty}$ is bounded. Thus the Principle of Uniform Boundedness guarantees that there is a constant $K \ge 0$ such that $\|T_n\| \le K$ for all n. Therefore, we have

$$\|Tx\| = \lim_{n\to\infty} \|T_n x\| \le \limsup_{n\to\infty} \|T_n\|\|x\| \le K\|x\|$$

for all $x \in \mathcal{X}$. Hence T is bounded. □

Exercises for Section 6.4

Exercise 6.4.1. Let $\mathcal{X}$ and $\mathcal{Y}$ be Banach spaces. Show that $\mathcal{X} \times \mathcal{Y}$ is a Banach space with respect to the norm

$$\|(x, y)\| = \max\{\|x\|, \|y\|\}.$$

Exercise 6.4.2. Suppose $\{\mathcal{X}_j\}_{j=J}$ is a family of Banach spaces, and let

$$\overline{\prod_{j\in J}}\mathcal{X}_j = \left\{ x \in \prod_{j\in J} \mathcal{X}_j : \sup_{j\in J} \|x(j)\| < \infty \right\}.$$

Prove that $\overline{\prod}_{j\in J}\mathcal{X}_j$ is a Banach space with respect to the norm

$$\|x\|_\infty = \sup_{j\in J} \|x(j)\|.$$

This space is called the **direct product** of the family $\{\mathcal{X}_j\}_{j\in J}$ (in the category of Banach spaces).

Exercise 6.4.3. Let $\mathcal{X}$ and $\mathcal{Y}$ be Banach spaces and suppose T is a **densely-defined** operator from $\mathcal{X}$ to $\mathcal{Y}$, meaning we have a dense subspace $D(T)$ of $\mathcal{X}$ and a linear operator $T : D(T) \to \mathcal{Y}$. We say that T is a **closed operator** if its graph is closed in $\mathcal{X} \times \mathcal{Y}$. If T is closed, prove that T is bounded if and only if $D(T) = X$.

Exercise 6.4.4. Let $\mathcal{X}$ be a Banach space, and suppose V is a vector space equipped with two norms $\|\cdot\|_\alpha$ and $\|\cdot\|_\beta$ such that:

1. there is a constant $K > 0$ such that $\|x\|_\beta \geq K\|x\|_\alpha$ for all $x \in V$;
2. V is a Banach space under $\|\cdot\|_\beta$.

Suppose $T : \mathcal{X} \to V$ is a linear operator that is bounded with respect to the norm $\|\cdot\|_\alpha$ on V. Show that T is also bounded with respect to $\|\cdot\|_\beta$.

Exercise 6.4.5. Let $\mathcal{X}$ be a Banach space (with norm $\|\cdot\|_\alpha$) and $V \subseteq \mathcal{X}$ a subspace. Suppose there is a second norm $\|\cdot\|_\beta$ on V that makes V into a Banach space, and that there is a $K > 0$ such that $\|x\|_\beta \geq K\|x\|_\alpha$ for all $x \in V$. (For a concrete example, one could take $\mathcal{X} = L^2([0, 1], \mu)$ and $V = C([0, 1])$, with $\|\cdot\|_\alpha = \|\cdot\|_2$ and $\|\cdot\|_\beta = \|\cdot\|_\infty$.) Suppose $T : \mathcal{X} \to \mathcal{X}$ is a bounded operator and that $T(V) \subseteq V$. Prove that $T|_V : V \to V$ is bounded, where the domain and codomain are both equipped with the norm $\|\cdot\|_\beta$. (**Hint:** Use Exercise 6.4.4.)

6.5 Dual Spaces

We are now going to investigate another important area in the study of Banach spaces, namely a particular cadre of operators known as *linear functionals*. We have actually encountered linear functionals before, though we have not formally defined them yet.

Definition 6.5.1. Let V be a vector space over a field F. A **linear functional** is a linear map $\varphi : V \to F$.

Linear functionals are an important tool in linear algebra and functional analysis for many reasons. For one, they give us way of assigning scalar values to vectors in an algebraic way, thus granting us a way of "measuring" vectors (beyond the usual notion of a norm). A classic example of this idea arises in quantum mechanics, where physicists represent *observables* (physical quantities that can be measured, such as position or momentum) as elements of a Banach space (or more precisely, operators on a particular Banach space). *Measuring* an observable corresponds to applying a linear functional to it. Such functionals are usually called *quantum states*.

Example 6.5.2. Fix $x \in \mathbf{R}$, and define $\varphi : C_0(\mathbf{R}) \to \mathbf{C}$ to be the map given by evaluation at x_0:

$$\varphi(f) = f(x_0).$$

It is easily checked that φ is a linear functional.

Example 6.5.3. Let $(X, \mathcal{M}, \mu)$ be a measure space, and define $\varphi : L^1(X, \mu) \to \mathbf{C}$ by

$$\varphi(f) = \int_X f \, d\mu.$$

Due to the linear properties of the integral, it is clear that φ is a linear functional.

Example 6.5.4. Fix $n \geq 1$, and define $\varphi : C([-\pi, \pi], \mathbf{R}) \to \mathbf{R}$ by

$$\varphi(f) = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(nx) \, dx.$$

Then φ is a linear functional, which one might recognize as the nth Fourier sine coefficient for f.

Example 6.5.5. Fix $n \in \mathbf{N}$, and let $V = \mathrm{M_n}(\mathbf{R})$ denote the $\mathbf{R}$-vector space of $n \times n$-matrices with real coefficients. Then the **trace**, which is defined for a matrix $A = [a_{ij}]$ by

$$\mathrm{tr}(A) = \sum_{i=1}^{n} a_{ii},$$

is a linear functional.

Much of what we know about operators can be used to immediately say things about linear functionals. Let V be a normed vector space over $F = \mathbf{R}$ or $F = \mathbf{C}$. We will say that a linear functional $\varphi : V \to F$ is **bounded** if $\varphi \in B(V, F)$, and the **norm** of a linear functional is simply its operator norm:

$$\|\varphi\| = \sup_{\|x\| \leq 1} |\varphi(x)|.$$

It is not the case that every linear functional on a normed vector space is bounded; we will produce an example soon.

Just as with operators, it is a fruitful endeavor to study all of the bounded operators on a normed vector space simultaneously.

Definition 6.5.6. Let V be a normed vector space (over $\mathbf{C}$). The collection of bounded linear functionals $V^* = B(V, \mathbf{C})$ is called the **dual space** of V.

Notice that since $\mathbf{C}$ is complete, we have the following immediate corollary to Theorem 6.3.16.

Theorem 6.5.7. *For any normed vector space V over $\mathbf{C}$, the dual space V^* is a Banach space.*

An important problem is to characterize the dual space of a given Banach space. That is, if $\mathcal{X}$ is a Banach space, can we recognize $\mathcal{X}^*$ as some other familiar Banach space? In general, this question should be quite difficult to answer. However, there are some specific situations where we can say what the dual of $\mathcal{X}$ looks like.

Naturally enough, we will consider the case of finite-dimensional spaces first. Recall that if V is finite-dimensional, then *any* linear functional $\varphi : V \to \mathbf{C}$ is bounded. In other words, the *algebraic* dual of V (the set of all linear functionals on V) equals the *continuous* dual (which we are just calling the dual here). In fact, we can say much more.

Theorem 6.5.8. *Let V be a finite-dimensional normed vector space over $\mathbf{C}$. Then V and V^* are isomorphic as $\mathbf{C}$-vector spaces.*

Proof. Let $n = \dim(V)$, and choose a basis $\{x_1, x_2, \dots, x_n\}$ for V. For $1 \le i \le n$, we define a linear functional $\varphi_i : V \to \mathbf{C}$ by

$$\varphi_i(x_j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j, \end{cases}$$

and then extending linearly to all of V. We claim that the set $\{\varphi_1, \varphi_2, \dots, \varphi_n\}$ defines a basis for V^*. First, suppose we have scalars $\alpha_1, \alpha_2, \dots, \alpha_n$ such that

$$\alpha_1\varphi_1 + \alpha_2\varphi_2 + \cdots + \alpha_n\varphi_n = 0.$$

Then for all $1 \le i \le n$ we have

$$\alpha_i = \alpha_i\varphi_i(x_i) = (\alpha_1\varphi_1 + \alpha_2\varphi_2 + \cdots + \alpha_n\varphi_n)(x_i) = 0,$$

so the set $\{\varphi_1, \varphi_2, \ldots, \varphi_n\}$ is linearly independent. We still need to check that this set spans V^*. Let $\varphi \in V^*$, and observe that if $x = \sum_{i=1}^n \alpha_i x_i \in V$, then

$$\varphi(x) = \sum_{i=1}^n \alpha_i \varphi(x_i).$$

On the other hand,

$$\varphi_j(x) = \sum_{i=1}^n \alpha_i \varphi_j(x_i) = \alpha_j$$

for all $1 \leq j \leq n$, so we have

$$\varphi(x) = \sum_{i=1}^n \varphi(x_i)\varphi_i(x).$$

Hence $\varphi = \sum_{i=1}^n \varphi(x_i)\varphi_i$, so $\varphi \in \operatorname{span}\{\varphi_1, \varphi_2, \ldots, \varphi_n\}$. Therefore, the φ_i form a basis for V^*.

We have thus far shown that $\dim(V^*) = \dim(V)$. It follows from results in linear algebra that $V^* \cong V$. In fact, we could define an explicit isomorphism of V onto V^* by sending $x_i \mapsto \varphi_i$ for each i and extending linearly to V. $\square$

Remark 6.5.9. The basis that we constructed for V^* in the last proof is called the **dual basis** relative to the basis $\{x_1, x_2, \ldots, x_n\}$ for V. Its elements are often written as

$$\{x^1, x^2, \ldots, x^n\}.$$

It is also worth noting that the last proof is quite messy, and there is not a canonical isomorphism between V and V^*. The isomorphism we arrived at depended on the choice of basis for V. For this reason, this result would be termed *unnatural* by an algebraist.

There is not generally an isomorphism between an infinite-dimensional vector space V and its continuous dual. However, there are at least two classes of infinite-dimensional Banach spaces for which we can characterize the dual spaces. The proof in both cases requires a bit of work, so we will postpone the details and simply give a rough outline of how things will work.

Example 6.5.10. Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $1 < p, q < \infty$ are conjugate exponents. Then $L^p(X, \mu)^*$ is isometrically isomorphic to $L^q(X, \mu)$. It is not hard to see that each function $g \in L^q(X, \mu)$ defines a linear functional φ_g on $L^p(X, \mu)$ via

$$\varphi_g(f) = \int_X fg\, d\mu.$$

A straightforward computation shows that φ_g is bounded, and it can also be shown that $\|\varphi_g\| = \|g\|_q$. We will verify all of these details later, and we will also see that the map $g \mapsto \varphi_g$ defines an isometric isomorphism of $L^q(X, \mu)$ onto $L^p(X, \mu)^*$.

The situation where $p = 1$ and $q = \infty$ is much more complicated. The map $g \mapsto \varphi_g$ defines an isometric embedding of $L^1(X, \mu)$ into $L^\infty(X, \mu)^*$, but it is rarely surjective. On the other hand, $L^1(X, \mu)^*$ can be identified with $L^\infty(X, \mu)$, provided μ is σ-finite.

Example 6.5.11. Let X be a compact metric space. Recall that any complex Borel measure μ on X defines a linear functional on $C(X)$ via

$$\varphi(f) = \int_X f \, d\mu.$$

Furthermore, it can be checked that

$$|\varphi(f)| \leq \|f\|_\infty \cdot |\mu|(X)$$

for all $f \in C(X)$. Moreover, $\|\varphi\| = |\mu|(X) = \|\mu\|$. In fact, every bounded linear functional on $C(X)$ arises in this manner. If we place extra requirements on our Borel measures (namely, that they are *Radon* measures), then we can guarantee a one-to-one correspondence between such measures and bounded linear functionals on $C(X)$.

We finish this section with one last preliminary fact about dual spaces. If V is a normed vector space, then so is V^*, and we could then talk about bounded linear functionals *on the dual space*. This leads us to consider the **double dual** or **second dual** of V, which is precisely

$$V^{**} = (V^*)^*.$$

It turns out that every element of V defines a linear functional on V^* as follows: given $x \in V$ and $\varphi \in V^*$, define $\iota(x) \in V^{**}$ by

$$\iota(x)(\varphi) = \varphi(x).$$

Then $\iota(x)$ clearly maps from V^* to $\mathbf{C}$, and it is easy to check that it is linear:

$$\iota(x)(\varphi + \psi) = (\varphi + \psi)(x) = \varphi(x) + \psi(x) = \iota(x)(\varphi) + \iota(x)(\psi)$$

and

$$\iota(x)(\alpha\varphi) = (\alpha\varphi)(x) = \alpha\varphi(x) = \alpha\iota(x)(\varphi)$$

for all $\varphi, \psi \in V^*$ and all $\alpha \in \mathbf{C}$. Furthermore, $\iota(x)$ is bounded for each $x \in V$:

$$|\iota(x)(\varphi)| = |\varphi(x)| \leq \|\varphi\|\|x\|$$

for all $\varphi \in V^*$. This computation also shows that $\|\iota(x)\| \leq \|x\|$. Therefore, $\iota(x)$ indeed defines an element of V^{**}, and we have a map $\iota : V \to V^{**}$. In fact, it is easy to check that ι is itself linear: if $x, y \in V$, then

$$\iota(x + y)(\varphi) = \varphi(x + y) = \varphi(x) + \varphi(y) = \iota(x)(\varphi) + \iota(y)(\varphi)$$

for all $\varphi \in V^*$, so $\iota(x+y) = \iota(x) + \iota(y)$. Similarly,

$$\iota(\alpha x)(\varphi) = \varphi(\alpha x) = \alpha\varphi(x) = \alpha\iota(x)(\varphi)$$

for all $\alpha \in \mathbf{C}$, $x \in V$, and $\varphi \in V^*$. Finally, we have already seen that

$$\|\iota(x)\| \leq \|x\|$$

for all $x \in V$, so ι is bounded with $\|\iota\| \leq 1$. In fact, ι defines an isometric injection of V into V^{**}. To prove it, we will need some results from the next section. Therefore, we will revisit this discussion there.

Exercises for Section 6.5

Exercise 6.5.1. Prove that a linear functional $\varphi : V \to \mathbf{C}$ is bounded if and only if its **kernel**,

$$\ker\varphi = \{x \in V \mid \varphi(x) = 0\},$$

is closed in V.

Exercise 6.5.2. Let V be a vector space over $\mathbf{C}$, and suppose $\varphi : V \to \mathbf{C}$ is a linear functional. Prove that

$$\|x\| = |\varphi(x)|$$

defines a seminorm on V.

Exercise 6.5.3. Let V and W be normed vector spaces, and let $T \in B(V, W)$. Define a map $T^* : W^* \to V^*$ (called the **adjoint** of T) by

$$T^*(\varphi) = \varphi \circ T.$$

Prove that $T^* \in B(W^*, V^*)$, and that $\|T^*\| = \|T\|$.

6.6 The Crown Jewel

This section is devoted to a single theorem regarding dual spaces—the Hahn-Banach theorem. This theorem is known as one of the most fundamental results in functional analysis, and it has several important consequences for Banach spaces and their duals. Due to its influence in analysis, it is has been called the “crown jewel of functional analysis”.

Before we even attempt to state the Hahn-Banach theorem, let us remark that its proof relies heavily on the Axiom of Choice. In particular, it is usually proven using the variant of the Axiom of Choice known as Zorn’s Lemma. Therefore, we make a slight digression to discuss the necessary background on this tool.

Definition 6.6.1. Let P be a set. A **partial order** on P is a relation $\preccurlyeq$ satisfying the following conditions:

1. (Reflexivity) $a \preccurlyeq a$ for all $a \in P$;
2. (Antisymmetry) $a \preccurlyeq b$ and $b \preccurlyeq a$ implies $a = b$ for all $a, b \in P$;
3. (Transitivity) $a \preccurlyeq b$ and $b \preccurlyeq c$ implies $a \preccurlyeq c$ for all $a, b, c \in P$.

The pair $(P, \preccurlyeq)$ is called a **partially ordered set**, or a **poset**.

Partial orders are meant to generalize our usual notions of ordering, such as the natural orderings on $\mathbf{N}$ and $\mathbf{R}$. However, it is not required that every pair of elements is comparable—there may be elements $a, b \in A$ such that none of the statements $a \preccurlyeq b$, $b \preccurlyeq a$, or $a = b$ are true. A partial order with the additional property that any two elements are comparable is called a **linear** or **total order**.

Example 6.6.2. A fundamental example of a partial order is defined in terms of subset inclusion. Given a set X, we can define a partial order on the power set $\mathcal{P}(X)$ by $A \preccurlyeq B$ if and only if $A \subseteq B$. It is easy to check that the conditions for a partial order are satisfied, and that $\preccurlyeq$ is not a linear order. (For example, if $A \neq \emptyset$, then A and A^c are not comparable.)

Example 6.6.3. The relation $\preccurlyeq$ on $\mathbf{N}$ defined by $a \preccurlyeq b$ if and only if $a \mid b$ (i.e., a divides b) is a partial order.

Before we can properly state Zorn's lemma, we need a few more definitions regarding posets.

Definition 6.6.4. Let $(P, \preccurlyeq)$ be a poset.

- An element $a \in P$ is **maximal** if $a \preccurlyeq b$ implies $a = b$.
- A linearly ordered subset of P is called a **chain**.
- If $A \subseteq P$, an **upper bound** for A is an element $b \in P$ satisfying $a \preccurlyeq b$ for all $a \in A$.

It is worth noting that if $(P, \preccurlyeq)$ is not linearly ordered, then maximal elements need not be unique. However, any two distinct maximal elements cannot be comparable. Also, a poset need not contain a maximal element in the first place. Zorn's Lemma provides us with conditions that guarantee the existence of a maximal element.

Theorem 6.6.5 (Zorn's Lemma)**.** *Let $(P, \preccurlyeq)$ be a poset. If every chain in P has an upper bound, then P contains a maximal element.*

It is perhaps a bit disingenuous to label Zorn's Lemma as a "theorem", since it is equivalent to an *axiom* from set theory, namely the Axiom of Choice. We omit the proof of this equivalence, though we will give one example of an argument that uses Zorn's Lemma before returning to the Hahn-Banach theorem.

Theorem 6.6.6. *Every vector space has a basis.*

Proof. Let V be a vector space over a field F, and let $\mathcal{B}$ denote the set of all linearly independent subsets of V. Then $\mathcal{B}$ is a poset with respect to subset inclusion. Let $\mathcal{C} \subseteq \mathcal{B}$ be a chain. In order to invoke Zorn's lemma, we need to show that $\mathcal{C}$ has an upper bound in $\mathcal{B}$. Write $\mathcal{C} = \{A_i\}_{i \in I}$, and define

$$A = \bigcup_{i \in I} A_i.$$

We claim that A is an upper bound for $\mathcal{C}$. Clearly $A_i \subseteq A$ for all i, but it is not clear that A actually lies in $\mathcal{B}$. To this end, suppose $x_1, x_2, \ldots, x_n \in A$, and that there exist scalars $\alpha_1, \alpha_2, \ldots, \alpha_n \in F$ such that

$$\alpha_1 x_1 + \alpha_2 x_2 + \ldots + \alpha_n x_n = 0.$$

Each vector x_j belongs to some set $A_{i_j} \in \mathcal{C}$, and since $\mathcal{C}$ is a chain, the collection $\{A_{i_j}\}_{j=1}^n$ has a largest element. In other words, there is an index $i_0 \in I$ such that $x_j \in A_{i_0}$ for all j. Since A_{i_0} is a linearly independent set, it follows that $\alpha_j = 0$ for all j. Hence A is linearly independent.

Since every chain in $\mathcal{B}$ has an upper bound in $\mathcal{B}$, Zorn's lemma guarantees that $\mathcal{B}$ contains a maximal element B. In other words, B is a maximal linearly independent subset of V. The maximality implies that B spans V: if there was a vector $x \notin \operatorname{span} B$, then $B \cup \{x\}$ would be linearly independent and strictly larger than B, contradicting the maximality of B among linearly independent sets. Therefore, B is a basis for V. □

Notice that the proof of Theorem 6.6.6 is not constructive—we can guarantee the existence of a basis, but it might not be possible to write one down. (This is usually the case with proofs that rely upon the Axiom of Choice or its equivalents.) Consequently, we can define other equally intangible objects in terms of this mysterious basis.

Theorem 6.6.7. *Let V be an infinite-dimensional normed vector space over* $\mathbf{C}$*. There exists an unbounded linear functional on V.*

Proof. By Theorem 6.6.6, we know that V has a basis B, which is necessarily infinite. Moreover, by normalizing the elements of B we may assume that $\|x\| = 1$ for all $x \in B$. Let $\{x_n\}_{n=1}^{\infty}$ be a countably infinite subset of B. For each $n \geq 1$, define $\varphi(x_n) = n$, and define φ to be zero on all other elements of B. Extend φ by linearity to obtain a linear functional $\varphi\colon V \to \mathbf{C}$. Then φ is not bounded, since

$$\sup_{\|x\| \leq 1} |\varphi(x)| = \sup_{n \geq 1} |\varphi(x_n)| = \infty. \qquad \square$$

Remark 6.6.8. As we observed in the last proof, any algebraic basis for an infinite-dimensional space is clearly infinite. We will see later that if $\mathcal{X}$ is an infinite-dimensional Banach space, then any basis for $\mathcal{X}$ must actually be *uncountable.*

Remark 6.6.9. The existence of a basis for any vector space produces other strange consequences as well. Many of these arise from the observation that $\mathbf{R}$, when viewed as a vector space over $\mathbf{Q}$, has a basis over $\mathbf{Q}$. Some examples are the following.

- There exist additive functions $f\colon \mathbf{R} \to \mathbf{R}$ (i.e., $f(x+y) = f(x) + f(y)$ for all $x, y \in \mathbf{R}$) that are not continuous. In fact, there are uncountably many of these functions, and any such function fails to be Lebesgue measurable and its graph is dense in $\mathbf{R}^2$. (See Section 5.1 of [Her06].)
- There exists a function $f\colon \mathbf{R} \to \mathbf{R}$ that maps any open interval (a, b) onto $\mathbf{R}$. [Bel, Proposition 1]
- There is an isomorphism between the additive groups $\mathbf{R}$ and $\mathbf{R}^2$. [Bel, Proposition 8]

Let us now return to our discussion about linear functionals. With Zorn's lemma in hand, we can now head toward a proof of the Hahn-Banach theorem. This theorem will say that if $\mathcal{X}$ is a Banach space and $\mathcal{X}_0$ is a closed subspace, then we can always extend linear functionals from $\mathcal{X}_0$ to $\mathcal{X}$. First we need a definition that slightly generalizes the idea of a linear functional for real vector spaces

Definition 6.6.10. Suppose V is a vector space over $\mathbf{R}$. We say that a function $m\colon V \to \mathbf{R}$ is a **sublinear functional** if

1. $m(x+y) \leq m(x) + m(y)$ for all $x, y \in V$;
2. $m(\alpha x) = \alpha m(x)$ whenever $x \in V$ and $\alpha \geq 0$.

Example 6.6.11. If $\|\cdot\|$ is a seminorm on V, then the triangle inequality guarantees that $m(x) = \|x\|$ defines a sublinear functional on V. In fact, a seminorm is the chief example of a sublinear functional.

Lemma 6.6.12. *Suppose V is a vector space over $\mathbf{R}$, W is a subspace of V, and $m : V \to \mathbf{R}$ is a sublinear functional. Let $\varphi : W \to \mathbf{R}$ be a linear functional on W such that*

$$\varphi(y) \leq m(y)$$

for all $y \in W$. Then there exists a linear functional Φ on V such that

(a) $\Phi(x) \leq m(x)$ *for all $x \in V$; and*

(b) $\Phi(y) = \varphi(y)$ *for all $y \in W$.*

Proof. Let $x \in V \backslash W$. We can extend φ to a linear functional $\tilde{\varphi}$ on the subspace $W + \mathbf{R}x$ by

$$\tilde{\varphi}(y + tx) = \varphi(y) + t\alpha,$$

where $\alpha \in \mathbf{R}$. The problem is to choose α such that

$$\varphi(y) + t\alpha \leq m(y + tx)$$

for all $y \in W$ and $t \in \mathbf{R}$. We claim first that it suffices to choose α such that

$$\varphi(y) + \alpha \leq m(y + x)$$

for all $y \in W$ and

$$\varphi(y) - \alpha \leq m(y - x)$$

for all $y \in W$. To see this, we observe that if $t \geq 0$, the first condition implies that

$$\begin{aligned}\varphi(y) + t\alpha &= t\left(\varphi\left(\tfrac{1}{t}y\right) + \alpha\right)\\ &\leq tm\left(\tfrac{1}{t}y + x\right)\\ &= m\left(y + tx\right)\end{aligned}$$

for all $y \in W$. Similarly, if $t \leq 0$, the second condition gives

$$\begin{aligned}\varphi(y) + t\alpha &= \varphi(y) - (-t)\alpha\\ &= -t\left(\varphi(-\tfrac{1}{t}y) - \alpha\right)\\ &\leq -tm(-\tfrac{1}{t}y - x)\\ &= m(y + tx).\end{aligned}$$

Therefore, the two conditions together imply that $\varphi(y) + t\alpha \leq m(y + tx)$ for all $y \in W$ and $t \in \mathbf{R}$. Thus we simply need to show that it is possible to find an α such that

$$\varphi(z) - m(z - x) \leq \alpha \leq -\varphi(y) + m(y + x)$$

for all $y, z \in W$. But subtracting the left hand side from the right gives

$$\begin{aligned} -\varphi(y) + m(y+x) - \varphi(z) + m(z-x) &= m(y+x) + m(z-x) - \varphi(y+z) \\ &\geq m(y+z) - \varphi(y+z) \\ &\geq 0. \end{aligned}$$

Therefore, if we set

$$b = \sup_{y \in W} \{-\varphi(y) + m(y+x)\}$$

and

$$a = \sup_{z \in W} \{\varphi(z) - m(z-x)\},$$

then we have $a \leq b$. Hence any $\alpha \in [a, b]$ will do.

Now let A be the collection of pairs (E, ψ) where $E \supseteq W$ is a subspace of V and $\psi|_W = \varphi$, and moreover that $\psi(x) \leq m(x)$ for all $x \in E$. We define a partial order on A by

$$(E_1, \psi_1) \preccurlyeq (E_2, \psi_2)$$

if $E_1 \subseteq E_2$ and $\psi_2|_{E_1} = \psi_1$. It is straightforward to check that $(A, \preccurlyeq)$ is a poset. Now suppose $\{(E_j, \psi_j)\}_{j \in J}$ is a chain in A. Then set

$$\mathcal{E}_0 = \bigcup_{j \in J} E_j$$

and define $\psi_0 : \mathcal{E}_0 \to \mathbf{R}$ by

$$\psi_0(x) = \psi_j(x)$$

if $x \in E_j$. Then $(\mathcal{E}_0, \psi_0) \in A$ and $(E_j, \psi_j) \preccurlyeq (\mathcal{E}_0, \psi_0)$ for all $j \in J$. This shows that every chain in A has an upper bound in A, so Zorn's Lemma implies that there is a maximal element $(\mathcal{E}, \Phi) \in A$. If $\mathcal{E} \neq V$, then the first part of the proof shows that we can choose $x \in V \backslash \mathcal{E}$ and extend Φ to a functional $\tilde{\Phi}$ on $\mathcal{E} + x\mathbf{R}$, which contradicts the maximality of $(\mathcal{E}, \Phi)$. Therefore, $\mathcal{E} = V$ and we are done. $\square$

Theorem 6.6.13 (Hahn-Banach)**.** *Suppose $\|\cdot\|$ is a seminorm on a vector space V over $\mathbf{R}$ or $\mathbf{C}$. Let W be a subspace of V, let φ be a linear functional on W, and suppose there is a constant $K > 0$ such that*

$$|\varphi(y)| \leq K\|y\|$$

for all $y \in W$. Then there exists a linear functional $\tilde{\varphi}$ on V such that

1. *$\tilde{\varphi}(y) = \varphi(y)$ for all $y \in W$; and*
2. *$|\tilde{\varphi}(x)| \leq K\|x\|$ for all $x \in V$.*

Proof. If we are working over $\mathbf{R}$, then Lemma 6.6.12 applies with $m(x) = K\|x\|$. Thus there exists a functional $\varphi_0 : V \to \mathbf{R}$ extending φ such that $\varphi_0(x) \le K\|x\|$ for all $x \in V$. But

$$-\varphi_0(x) = \varphi_0(-x) \le K\|-x\| = K\|x\|$$

as well, so $|\varphi_0(x)| \le K\|x\|$ for all $x \in X$.

Now assume that V is a vector space over $\mathbf{C}$. By restricting scalar multiplication, we can treat V as a real vector space. Then consider the linear functional

$$\varphi_0(y) = \operatorname{Re}\varphi(y).$$

Then

$$|\varphi_0(y)| = |\operatorname{Re}\varphi(y)| \le |\varphi(y)| \le K\|y\|$$

for all $y \in W$. Using the first part of our proof, we can get an $\mathbf{R}$-linear map $\psi : V \to \mathbf{R}$ such that

$$\psi(y) = \varphi_0(y)$$

for all $y \in W$ and

$$|\psi(x)| \le K\|x\|$$

for all $x \in V$. Now we define $\tilde{\varphi} : V \to \mathbf{C}$ by

$$\tilde{\varphi}(x) = \psi(x) - i\psi(ix).$$

Note that

$$\tilde{\varphi}(ix) = \psi(ix) - i\psi(-x) = \psi(ix) + i\psi(x) = i(-i\psi(ix) + \psi(x)) = i\tilde{\varphi}(x),$$

so $\tilde{\varphi}$ is in fact $\mathbf{C}$-linear. Now if $y \in W$, we have

$$\begin{aligned}
\tilde{\varphi}(y) &= \psi(y) - i\psi(iy) \\
&= \operatorname{Re}\varphi(y) - i\operatorname{Re}\varphi(iy) \\
&= \operatorname{Re}\varphi(y) - i\operatorname{Re}(i\varphi(y)) \\
&= \operatorname{Re}\varphi(y) + i\operatorname{Im}\varphi(y) \\
&= \varphi(y).
\end{aligned}$$

Now fix $x \in V$, and choose $\alpha \in \mathbf{C}$ such that $|\alpha| = 1$ and

$$|\tilde{\varphi}(x)| = \alpha\tilde{\varphi}(x) = \tilde{\varphi}(\alpha x).$$

Then

$$|\tilde{\varphi}(x)| = \tilde{\varphi}(\alpha x) = \psi(\alpha x) \le K\|\alpha x\| = K\|x\|,$$

and we are done. $\square$

Remark 6.6.14. If $\|\cdot\|$ is a full-fledged norm on V, then the condition $|\varphi(y)| \leq K\|y\|$ for all $y \in W$ simply says that φ is a bounded linear functional on W. It then turns out that the norm of the extension $\tilde{\varphi}$ agrees with that of φ. Since $|\varphi(y)| \leq \|\varphi\|\|y\|$ for all $y \in W$, the Hahn-Banach theorem guarantees that $|\tilde{\varphi}(x)| \leq \|\varphi\|\|x\|$ for all $x \in V$. Thus $\|\tilde{\varphi}\| \leq \|\varphi\|$. On the other hand, it is clear from the definition of the operator norm that $\|\varphi\| \leq \|\tilde{\varphi}\|$, the relevant supremum being taken over a larger set of vectors when computing $\|\tilde{\varphi}\|$. Hence $\|\tilde{\varphi}\| = \|\varphi\|$.

The Hahn-Banach theorem has several important consequences. Perhaps foremost among them is the conclusion that any normed vector space has "enough" linear functionals. There are a couple of ways to interpret this statement; the first says that the linear functionals can be used to detect the norm of an element, while the second says that there are enough linear functionals to separate points.

Corollary 6.6.15. *Let V be a normed vector space (over $\mathbf{C}$), and let $x \in V$ with $x \neq 0$. Then there exists $\varphi \in V^*$ such that $\|\varphi\| = 1$ and $\varphi(x) = \|x\|$.*

Proof. Let $W = \mathbf{C}x$. Define φ_0 on W by

$$\varphi_0(\alpha x) = \alpha\|x\|.$$

Since

$$|\varphi_0(\alpha x)| = |\alpha||\varphi_0(x)| = |\alpha|\|x\| = \|\alpha x\|$$

for all $\alpha \in \mathbf{C}$, we see that $\|\varphi_0\| = 1$. Now the Hahn-Banach theorem guarantees that there is an extension of φ_0 to a linear functional $\varphi \in V^*$ such that $\|\varphi\| = 1$. $\square$

Corollary 6.6.16. *If V is a normed vector space, then V^* separates points of V. That is, if $x \neq y$, then there exists $\varphi \in V^*$ such that $\varphi(x) \neq \varphi(y)$.*

Proof. If $x \neq y$, then $z = x - y \neq 0$, and there exists $\varphi \in V^*$ such that $\varphi(z) = \|z\| \neq 0$ by the previous corollary. It then follows that $\varphi(x) \neq \varphi(y)$. $\square$

The next corollary requires a little context first. Let V be a normed vector space, and suppose W is a closed subspace of V. Recall that we can form the **quotient** vector space

$$V/W = \{x + W \mid x \in V\},$$

whose elements are the cosets of W in V. We can also equip V/W with a norm via

$$\|x + W\| = \inf\{\|x - y\| \mid y \in W\}.$$

The assumption that W is closed ensures positive definiteness; if W is not closed, we are only guaranteed a seminorm. (See Exercise 6.6.1.) As a consequence of the

Hahn-Banach theorem, the quotient norm can be described using linear functionals on V.

Corollary 6.6.17. *Let W be a closed subspace of a normed vector space V. If $x \in V \backslash W$, there exists $\varphi \in V^*$ such that $\|\varphi\| = 1$, $\varphi|_W = 0$, and*

$$\varphi(x) = \|x + W\| = \inf_{y \in W} \|x - y\|.$$

Proof. We can apply Corollary 6.6.15 to the quotient V/W to obtain $\tilde{\varphi} \in (V/W)^*$ such that $\|\tilde{\varphi}\| = 1$ and $\tilde{\varphi}(x + W) = \|x + W\|$. Now let $q : V \to V/W$ denote the canonical quotient map, and define $\varphi \in V^*$ by $\varphi = \tilde{\varphi} \circ q$. Clearly φ vanishes on W and $\|\varphi\| \leq \|\tilde{\varphi}\|\|q\| = 1$. It remains to show that $\|\varphi\|$ is indeed 1. Well, given $\varepsilon > 0$ there exists $y \in W$ such that

$$\|q(x)\| > (1 - \varepsilon)\|x - y\|.$$

Now observe that

$$|\varphi(x - y)| = |\varphi(x)| = |\tilde{\varphi}(q(x))| = \|q(x)\| > (1 - \varepsilon)\|x - y\|$$

since φ vanishes on W. Since ε was arbitrary, it follows that $\|\varphi\| = 1$. □

We now come to some leftover business regarding the second dual. In the previous section we constructed a linear map ι from a normed vector space V into its second dual V^{**}, which was bounded with $\|\iota\| \leq 1$. We can now show that ι is actually injective.

Corollary 6.6.18. *Let V be a normed vector space. The map $\iota : V \to V^{**}$ defined by*

$$\iota(x)(\varphi) = \varphi(x)$$

is an isometric injection, hence $\iota(V)$ is a subspace of V^{**}.

Proof. We already know that $\|\iota(x)\| \leq \|x\|$ for all $x \in V$. Given $x \in X$, Corollary 6.6.15 implies that there exists $\varphi \in V^*$ such that

$$\varphi(x) = \|x\|.$$

Hence $\iota(x)(\varphi) = \|x\|$, so $\|\iota(x)\| = \|x\|$. This shows that ι is isometric, which also guarantees that it is injective by Exercise 6.3.3. □

We have thus shown that every normed vector space embeds into its second dual in a natural way. For infinite dimensional spaces, it is rare that this map is an isomorphism.

Definition 6.6.19. A Banach space $\mathcal{X}$ is called **reflexive** if the natural embedding $\iota : \mathcal{X} \to \mathcal{X}^{**}$ is surjective.

Though reflexive Banach spaces are relatively rare, we will see a whole family of reflexive spaces quite soon. In particular, the L^p spaces (except $p = 1$ and $p = \infty$) turn out to be reflexive.

While this discussion of the second dual might seem odd, it actually has an interesting application. We obtain a quick proof that any normed vector space can be completed to obtain a Banach space.

Theorem 6.6.20. *Let V be a normed vector space. There exists a Banach space $\overline{V}$, called the* ***completion*** *of V, such that V is isometrically isomorphic to a dense subspace of $\overline{V}$. Moreover, $\overline{V}$ is unique up to isometric isomorphism.*

Proof. Recall that V^{**} is a Banach space, and that we have an isometric embedding $\iota : V \to V^{**}$ by the previous corollary. Thus the closure $\overline{\iota(V)}$ is a Banach space, which contains an isometric copy of V as a dense subspace. Therefore, we can take this space to be $\overline{V}$.

Now suppose that $\overline{V_1}$ and $\overline{V_2}$ are two completions of V. Then we have isometric embeddings
$$T_j : V \to \overline{V_j}$$
onto dense subspaces for $j = 1$ and 2. Thus $T_0 = T_2 \circ T_1^{-1}$ is an isometric map of $T_1(V) \subseteq \overline{V}_1$ onto $T_2(V) \subseteq \overline{V}_2$. Since $T_1(V)$ is dense in $\overline{V_1}$, T_0 extends to a linear map $T : \overline{V}_1 \to \overline{V}_2$. Since T_0 is isometric, it is not hard to verify that T is isometric. It is also easy to check that $T(\overline{V}_1)$ is complete in $\overline{V}_2$. Thus $T(\overline{V}_1)$ is closed and contains $T_2(V)$, which is dense in $\overline{V_2}$. Therefore, T is surjective, hence an isometric isomorphism as required. □

Exercises for Section 6.6

Exercise 6.6.1. Let V be a normed vector space and $W \subseteq V$ a subspace. Prove that
$$\|x + W\| = \inf \{\|x - y\| \mid y \in W\}.$$
defines a seminorm on V/W, which is a norm if W is closed. Moreover, show that if V is a Banach space and W is closed, then V/W is a Banach space with respect to the quotient norm.

Exercise 6.6.2. Suppose $\mathcal{X}$ and $\mathcal{Y}$ are Banach spaces and that $T : \mathcal{X} \to \mathcal{Y}$ and $S : \mathcal{Y}^* \to \mathcal{X}^*$ are linear operators satisfying
$$\varphi(Tx) = S(\varphi(x))$$

for all $x \in \mathcal{X}$ and all $\varphi \in \mathcal{Y}^*$. Prove that T and S are bounded and $S = T^*$. (**Hint:** Use the Closed Graph Theorem and the fact that $\mathcal{Y}^*$ separates points.)

6.7 Duality for L^p Spaces

We now return to the problem of identifying the dual spaces of certain familiar Banach spaces. We begin by studying the L^p spaces associated to a given measure space.

Theorem 6.7.1. *Let $(X, \mathcal{M}, \mu)$ be a measure space, and suppose $1 \le p, q \le \infty$, where p and q are conjugate exponents. Given $g \in L^q(X, \mu)$, the function*

$$\varphi_g(f) = \int_X fg\, d\mu$$

for $f \in L^p(X, \mu)$ defines a bounded linear functional on $L^p(X, \mu)$. Furthermore, we have $\|\varphi_g\| = \|g\|_q$ when $1 \le q < \infty$, and the same is true for $q = \infty$ if μ is semifinite.

Proof. The result clearly holds if $g = 0$ μ-a.e., so we may assume that $g \neq 0$. Suppose first that $1 \le q < \infty$, and let $f \in L^p(X, \mu)$. Recall that Hölder's inequality implies that $fg \in L^1(X, \mu)$, and that $\|fg\|_1 \le \|f\|_p \|g\|_q$. Thus φ_g is well-defined, and

$$|\varphi_g(f)| \le \int_X |fg|\, d\mu \le \|f\|_p \|g\|_q,$$

so φ_g is bounded with $\|\varphi_g\| \le \|g\|_q$. To establish equality, we define $\omega \colon X \to \mathbf{C}$ by

$$\omega(x) = \begin{cases} \overline{\frac{g(x)}{|g(x)|}} & \text{if } g(x) \neq 0 \\ 0 & \text{if } g(x) = 0 \end{cases}$$

consider the function

$$f = \frac{|g|^{q-1}\omega}{\|g\|_q^{q-1}}.$$

Notice that

$$\int_X |f|^p\, d\mu = \frac{1}{\|g\|_q^{(q-1)p}} \int_X |g|^{(q-1)p}\, d\mu = \frac{1}{\|g\|_q^q} \int_X |g|^q\, d\mu = 1,$$

so $f \in L^p(X, \mu)$ and $\|f\|_p = 1$. Furthermore, it is easily seen that

$$fg = \frac{|g|^q}{\|g\|_q^{q-1}},$$

so

$$\int_X fg\,d\mu = \frac{1}{\|g\|_q^{q-1}} \int_X |g|^q\,d\mu = \|g\|_q.$$

It follows that $\|\varphi_g\| = \|g\|_q$.

Now assume μ is semifinite, and let $q = \infty$. By the definition of the essential supremum, given $\varepsilon > 0$ the set

$$\{x \in X \mid |g(x)| \geq \|g\|_\infty - \varepsilon\}$$

has nonzero measure. Since μ is semifinite, it contains a subset E satisfying $0 < \mu(E) < \infty$. Let ω be as above, and define

$$f = \frac{\omega \cdot \chi_E}{\mu(E)}.$$

Then

$$\int_X |f|\,d\mu = \frac{1}{\mu(E)} \int_X \chi_E\,d\mu = 1,$$

so $f \in L^1(X,\mu)$ with $\|f\|_1 = 1$. Furthermore,

$$\varphi_g(f) = \int_X fg\,d\mu = \frac{1}{\mu(E)} \int_X |g|\chi_E\,d\mu \geq \frac{1}{\mu(E)} \cdot (\|g\|_\infty - \varepsilon)\mu(E) = \|g\|_\infty - \varepsilon,$$

and it follows that $\|\varphi_g\| \geq \|g\|_\infty - \varepsilon$. Since $\varepsilon > 0$ was arbitrary, we can conclude that $\|\varphi_g\| = \|g\|_\infty$. □

The previous theorem shows that if p and q are conjugate exponents, then there is a natural map $\Phi : L^q(X,\mu) \to L^p(X,\mu)^*$, given by $\Phi(g) = \varphi_g$. Moreover, it is not hard to check that Φ is linear, and we know that Φ is isometric when $1 \leq q < \infty$ (and also when $q = \infty$ if μ is assumed to be semifinite). Hence we have an isometric embedding of $L^q(X,\mu)$ into $L^p(X,\mu)^*$.

In fact, the next theorem shows that Φ turns out to be surjective most of the time. As a result, we will have an identification of $L^p(X,\mu)^*$ with $L^q(X,\mu)$ whenever $1 < p, q < \infty$ are conjugate exponents. The situation where $p = 1$ and $q = \infty$ turns out to be a little more delicate. Though this result holds more generally, we will prove it under the assumption that μ is σ-finite, since we can give a much more elegant proof (modeled on that of Rudin) in this case.

Theorem 6.7.2. *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space and suppose $1 < p, q < \infty$, where p and q are conjugate exponents. Then every bounded linear functional on $L^p(X,\mu)$ is of the form*

$$\varphi_g(f) = \int_X fg\,d\mu$$

for some $g \in L^q(X,\mu)$. Consequently, we have an isometric isomorphism $\Phi : L^q(X,\mu) \to L^p(X,\mu)^*$ given by $\Phi(g) = \varphi(g)$.

Proof. Assume first that $\mu(X) < \infty$, and let $\varphi \in L^p(X,\mu)^*$. We need to produce a function $g \in L^q(X,\mu)$ such that $\varphi_g = \varphi$. To this end, we begin by defining a complex measure ν on $\mathcal{M}$ as follows: for each $E \in \mathcal{M}$, set

$$\nu(E) = \varphi(\chi_E).$$

Clearly $\nu(\emptyset) = 0$. Suppose $\{E_j\}_{j=1}^{\infty}$ is a countable collection of disjoint sets in $\mathcal{M}$, and for each $k \geq 1$ put $A_k = \bigcup_{j=1}^{k} E_j$. Observe that

$$\chi_{A_k} = \sum_{j=1}^{k} \chi_{E_j}$$

for each k, so

$$\nu(A_k) = \varphi(\chi_{A_k}) = \sum_{j=1}^{k} \varphi(\chi_{E_j}) = \sum_{j=1}^{k} \nu(E_j).$$

In other words, ν is finitely additive. Now if we let $E = \bigcup_{j=1}^{\infty} E_j$, we have

$$\|\chi_E - \chi_{A_k}\|_p = \left(\int_X |\chi_E - \chi_{A_k}|^p \, d\mu \right)^{1/p} = \mu(E \backslash A_k)^{1/p}$$

for all $k \geq 1$. By continuity of measure, $\mu(E\backslash A_k) \to 0$ as $k \to \infty$. Hence $\chi_{A_k} \to \chi_E$ in L^p, so $\varphi(\chi_{A_k}) \to \varphi(\chi_E)$ since φ is continuous. Therefore,

$$\nu(E) = \lim_{k\to\infty} \nu(A_k) = \lim_{k\to\infty} \sum_{j=1}^{k} \nu(E_j) = \sum_{j=1}^{\infty} \nu(E_j),$$

so ν is countably additive. Thus ν defines a complex measure on $\mathcal{M}$.

Suppose $E \in \mathcal{M}$ with $\mu(E) = 0$. Then $\chi_E = 0$ μ-a.e., so χ_E belongs to the equivalence class of the zero function in $L^p(X,\mu)$. Hence

$$\nu(E) = \varphi(\chi_E) = 0,$$

and it follows that $\nu \ll \mu$. The Radon-Nikodym theorem then guarantees that there is a function $g \in L^1(X,\mu)$ satisfying $d\nu = g\,d\mu$. As a result, we have

$$\varphi(\chi_E) = \nu(E) = \int_X \chi_E \cdot g \, d\mu$$

for all $E \in \mathcal{M}$. The same sort of equation holds for all measurable simple functions due to the linearity of the integral. Moreover, we claim that it holds for all $f \in L^\infty(X,\mu)$. Since the measurable simple functions are dense in $L^\infty(X,\mu)$, we can approximate any element $f \in L^\infty(X,\mu)$ uniformly with a sequence of simple functions $(f_j)_{j=1}^\infty$. Then

$$\|f_j - f\|_p^p = \int_X |f_j - f|^p \, d\mu \le \|f_j - f\|_\infty \cdot \mu(X) \to 0$$

as $j \to \infty$, so $f_j \to f$ in L^p. Since φ is continuous, we have

$$\varphi(f) = \lim_{j\to\infty} \varphi(f_j) = \lim_{j\to\infty} \int_X f_j g \, d\mu = \int_X fg \, d\mu,$$

and our claim holds.

Next, we aim to show that g actually belongs to $L^q(X,\mu)$. First, for each $n \in \mathbf{N}$ we define

$$E_n = \big\{x \in X \;\big|\; |g(x)| \ge n\big\},$$

and let ω be as in the proof of Theorem 6.7.1. Then $\omega g = |g|$, and we set

$$f_n = \chi_{E_n} |g|^{q-1} \omega.$$

for each $n \in \mathbf{N}$. Then f_n is bounded, hence $f_n \in L^\infty(X,\mu)$, and

$$\int_X \chi_{E_n} |g|^q \, d\mu = \int_X f_n g \, d\mu = \varphi(f_n) \le \|\varphi\| \|f_n\|_p. \tag{6.1}$$

Now observe that

$$|f_n|^p = \chi_{E_n} |g|^{(q-1)p} = \chi_{E_n} |g|^q,$$

so

$$\|f_n\|_p = \left(\int_X \chi_{E_n} |g|^q \, d\mu\right)^{1/p}.$$

Thus the inequality in (6.1) becomes

$$\int_X \chi_{E_n} |g|^q \, d\mu \le \|\varphi\| \left(\int_X \chi_{E_n} |g|^q \, d\mu\right)^{1/p},$$

or dividing through by the right hand side and raising to the power q,

$$\int_X \chi_{E_n} |g|^q \, d\mu \le \|\varphi\|^q.$$

This inequality holds for all n, so an application of the Monotone Convergence Theorem yields

$$\int_X |g|^q \, d\mu \le \|\varphi\|^q.$$

It follows that $g \in L^q(X,\mu)$ (and in fact, $\|g\|_q \le \|\varphi\|$). Since $L^\infty(X,\mu)$ is dense in $L^p(X,\mu)$ (as the measurable simple functions are dense) and $\varphi(f) = \varphi_g(f)$ for all $f \in L^\infty(X,\mu)$, it follows from continuity that $\varphi = \varphi_g$.

Now suppose $\mu(X) = \infty$. Since ν is σ-finite, there is a finite measure ν on $\mathcal{M}$ such that ν and μ are mutually absolutely continuous. Suppose $\varphi \in L^p(X,\mu)^*$ and $f \in L^p(X,\nu)$. Then $|f|^p \cdot \frac{d\nu}{d\mu} \in L^1(X,\mu)$, so $f \cdot (\frac{d\nu}{d\mu})^{1/p} \in L^p(X,\mu)$. Consequently,

$$\psi(f) = \varphi\big(f \cdot (\tfrac{d\nu}{d\mu})^{1/p}\big)$$

defines a linear functional on $L^p(X,\nu)$, which is easily checked to be bounded. As a result, we can find a function $\tilde{g} \in L^q(X,\nu)$ such that

$$\psi(f) = \int_X f\tilde{g}\,d\nu = \int_X f\tilde{g} \cdot \frac{d\nu}{d\mu}\,d\mu$$

for all $f \in L^p(X,\nu)$. Set $g = \tilde{g} \cdot (\frac{d\nu}{d\mu})^{1/q}$. Then $g \in L^q(X,\mu)$, and for all $f \in L^p(X,\mu)$ we have

$$\int_X fg\,d\mu = \int_X fg \cdot \tfrac{d\mu}{d\nu}\,d\nu = \int_X f\big(\tfrac{d\mu}{d\nu}\big)^{1/p} \cdot g\big(\tfrac{d\mu}{d\nu}\big)^{1/q}\,d\nu = \int_X f\big(\tfrac{d\mu}{d\nu}\big)^{1/p} \cdot \tilde{g}\,d\nu.$$

Now the right hand side is precisely

$$\psi\big(f \cdot (\tfrac{d\mu}{d\nu})^{1/p}\big) = \varphi(f),$$

so we have

$$\int_X fg\,d\mu = \varphi(f)$$

for all $f \in L^p(X,\mu)$. Thus $\varphi = \varphi_g$, and we are done. $\square$

As mentioned above, the question of duality for $L^1(X,\mu)$ and $L^\infty(X,\mu)$ is much more delicate. A rehashing of the arguments above does show that $L^1(X,\mu)^*$ can be identified with $L^\infty(X,\mu)$.

Theorem 6.7.3. *Let $(X, \mathcal{M}, \mu)$ be a σ-finite measure space. Every bounded linear functional on $L^1(X,\mu)$ is of the form*

$$\varphi_g(f) = \int_X fg\,d\mu$$

for some $g \in L^\infty(X,\mu)$. Consequently, there is an isometric isomorphism $\Phi : L^\infty(X,\mu) \to L^1(X,\mu)^*$ given by $\Phi(g) = \varphi(g)$.

The σ-finiteness assumption in the last theorem is crucial; it is left to the reader to find an example where $L^1(X,\mu)^*$ is not isomorphic to $L^\infty(X,\mu)$.

Remark 6.7.4. There is always an isometric embedding of $L^1(X,\mu)$ into $L^\infty(X,\mu)^*$, but the dual of $L^\infty(X,\mu)$ is usually much larger than $L^1(X,\mu)$. In fact, it is only in fairly trivial cases that we have $L^\infty(X,\mu)^* \cong L^1(X,\mu)$. For example, if X is a finite set and μ is the counting measure on X, then $L^\infty(X,\mu)$, $L^\infty(X,\mu)^*$, and $L^1(X,\mu)$ are all finite-dimensional vector spaces of the same dimension, so they are isomorphic.

Recall that a Banach space $\mathcal{X}$ is *reflexive* if the natural embedding $\iota : \mathcal{X} \to \mathcal{X}^{**}$ is surjective. As we look at Theorem 6.7.2, we can see that the L^p spaces are reflexive when $1 < p < \infty$, since

$$L^p(X,\mu)^{**} \cong L^q(X,\mu)^* \cong L^p(X,\mu).$$

However, we should be very careful here—this isomorphism needs to be given by the canonical map $\iota : L^p(X,\mu) \to L^p(X,\mu)^{**}$. Fortunately, the details are not hard to check.

Theorem 6.7.5. *Let $(X,\mathcal{M},\mu)$ be a measure space. For all $1 < p < \infty$, the Banach space $L^p(X,\mu)$ is reflexive.*

Proof. Suppose $1 < p < \infty$, and let q be the conjugate exponent of p. Let $\Phi : L^q(X,\mu) \to L^p(X,\mu)^*$ be the isometric isomorphism afforded by Theorem 6.7.2, and let $F \in L^p(X,\mu)^{**}$. Then $F \circ \Phi$ defines a bounded linear functional on $L^q(X,\mu)$, which simply takes the form

$$F \circ \Phi(g) = F(\varphi_g)$$

for all $g \in L^q(X,\mu)$. Consequently there is a function $f \in L^p(X,\mu)$ such that

$$F(\varphi_g) = F \circ \Phi(g) = \int_X fg\,d\mu = \varphi_g(f)$$

for all $g \in L^q(X,\mu)$. However, recall from the definition of $\iota : L^p(X,\mu) \to L^p(X,\mu)^{**}$ that

$$\iota(f)(\varphi_g) = \varphi_g(f),$$

so we have $F(\varphi_g) = \iota(f)(\varphi_g)$ for all $g \in L^q(X,\mu)$. It follows that $F = \iota(f)$, so ι is surjective. □

Since it is rarely the case that $L^\infty(X,\mu)^* \cong L^1(X,\mu)$, it is almost never the case that $L^1(X,\mu)$ is reflexive. Even though we have not identified the dual of $L^\infty(X,\mu)$ in general, our conclusion about $L^1(X,\mu)$ is enough to guarantee that $L^\infty(X,\mu)$ is usually not reflexive either. In fact, a Banach space $\mathcal{X}$ is reflexive if and only if $\mathcal{X}^*$ is reflexive.

Exercises for Section 6.7

Exercise 6.7.1. This problem will give an example to show that $L^1(X,\mu)$ and $L^\infty(X,\mu)$ are not generally reflexive spaces.

View $\mathbf{N}$ as both a locally compact metric space (under the discrete metric) and a measure space (with counting measure μ). Recall that $L^1(\mathbf{N},\mu) = \ell^1$ and $L^\infty(\mathbf{N},\mu) = \ell^\infty$. Also, if we define

$$\mathfrak{c}_0 = \left\{ (x_n)_{n=1}^\infty \;\middle|\; \lim_{n\to\infty} x_n = 0 \right\},$$

then it is easy to check that $\mathfrak{c}_0 = C_0(\mathbf{N})$. Note that $\mathfrak{c}_0 \subseteq \ell^\infty$.

(a) For each $x \in \ell^1$, define $\varphi_x : \mathfrak{c}_0 \to \mathbf{C}$ by

$$\varphi_x(y) = \sum_{n=1}^\infty x_n y_n.$$

Prove that φ_x is a bounded linear functional on $\mathfrak{c}_0$.

(b) Define $\iota : \ell^1 \to \mathfrak{c}_0^*$ by $\iota(x) = \varphi_x$. Prove that ι defines an isometric isomorphism of ℓ^1 with $\mathfrak{c}_0^*$.

(c) Define

$$\mathfrak{c} = \left\{ (x_n)_{n=1}^\infty \;\middle|\; \lim_{n\to\infty} x_n \text{ exists} \right\}.$$

Note that $\mathfrak{c} \subseteq \ell^\infty$. Show that every element $x \in \ell^1$ gives rise to an element of $\mathfrak{c}^*$ via the formula

$$\varphi_x(y) = \sum_{n=1}^\infty x_n y_n$$

for $y \in \mathfrak{c}$. Use this fact to show that there is an isometric embedding of ℓ^1 into $\mathfrak{c}^*$.

(d) Show that the bounded linear functional on $\mathfrak{c}$ defined by

$$\varphi(x) = \lim_{n\to\infty} x_n$$

is not induced by an element of ℓ^1. Conclude that ℓ^1 is a *proper* subspace of $\mathfrak{c}^*$.

(e) Show that every element of ℓ^∞ induces a bounded linear functional on $\mathfrak{c}$. (**Hint:** Use the fact that $(\ell^1)^* = \ell^\infty$ together with the Hahn-Banach theorem.) Prove that we then have an embedding of ℓ^∞ into $\mathfrak{c}^{**}$.

(f) Argue that $\mathfrak{c}$ is a proper subspace of $\mathfrak{c}^{**}$, so $\mathfrak{c}$ is not reflexive. Give a similar proof to show that $\mathfrak{c}_0$ is not reflexive.

Chapter 7

Hilbert Spaces

There is a particular class of Banach spaces for which some of our earlier investigations look particularly nice. These spaces—called *Hilbert spaces*—have additional structure that allows one to discuss geometric notions (such as angles and orthogonality) that generalize those of Euclidean space. The additional property that makes Hilbert spaces special is the presence of an *inner product*.

7.1 Inner Products

Recall that in the standard Euclidean space (i.e., $\mathbf{R}^n$ or $\mathbf{C}^n$ equipped with its Euclidean norm) we have a way of "pairing" two vectors together to obtain a scalar via the dot product: if $x = (x_1, x_2, \ldots, x_n)$ and $y = (y_1, y_2, \ldots, y_n)$ are vectors in $\mathbf{R}^n$ or $\mathbf{C}^n$, then

$$x \cdot y = x_1y_1 + x_2y_2 + \cdots + x_ny_n.$$

Recall that the dot product allows one to measure angles between vectors, and two vectors are said to be *orthogonal* (i.e., perpendicular) if $x \cdot y = 0$. Thus a pairing such as the dot product allows gives one a more refined notion of geometry within a vector space.

We will be interested in generalizations of the dot product known as *inner products*. However, we will start off a little more generally, since such a discussion will prove fruitful later on. Throughout this discussion we will be working with vector spaces over a field F, where it is understood that $F = \mathbf{R}$ or $F = \mathbf{C}$ unless otherwise noted.

Definition 7.1.1. Let X be a vector space over a field F. A **sesquilinear form** on X is a function $(\cdot \mid \cdot) : X \times X \to F$ satisfying the following conditions.

1. The map $x \mapsto (x \mid y)$ is linear for all $y \in X$, meaning that

$$(\alpha x + z \mid y) = \alpha(x \mid y) + (z \mid y)$$

for all $x, z \in X$ and $\alpha \in F$.

2. If $F = \mathbf{R}$, the map $y \mapsto (x \mid y)$ is linear for all $x \in X$.
3. If $F = \mathbf{C}$, the map $y \mapsto (x \mid y)$ is *conjugate linear* for all $x \in X$, meaning that
$$(x \mid \alpha y + z) = \overline{\alpha}(x \mid y) + (x \mid z)$$
for all $x, y, z \in X$ and $\alpha \in \mathbf{C}$.

Remark 7.1.2. Note that if $F = \mathbf{R}$, then we are simply saying that a sesquilinear form is linear in both variables, hence it is in fact a **bilinear** form.

Before moving on, let us introduce a bit of useful terminology. If $(\cdot \mid \cdot)$ is a sesquilinear form on a complex vector space X that satisfies
$$(y \mid x) = \overline{(x \mid y)}$$
for all $x, y \in X$, then we say that $(\cdot \mid \cdot)$ is **Hermitian**. For a vector space over $\mathbf{R}$, this condition simply reduces to
$$(y \mid x) = (x \mid y),$$
in which case we say the bilinear form $(\cdot \mid \cdot)$ is **symmetric**.

Proposition 7.1.3. *Let X be a vector space over $\mathbf{C}$. A sesquilinear form $(\cdot \mid \cdot)$ on X is Hermitian if and only if $(x \mid x) \in \mathbf{R}$ for all $x \in X$.*

Proof. First assume that $(\cdot \mid \cdot)$ is Hermitian. Then for all $x \in X$ we have
$$\overline{(x \mid x)} = (x \mid x),$$
which implies that $(x \mid x) \in \mathbf{R}$. Now assume $(x \mid x)$ is real for all $x \in X$, and let $x, y \in X$. Then
$$(x + y \mid x + y) = (x \mid x) + (y \mid y) + (x \mid y) + (y \mid x),$$
and since $(x + y \mid x + y)$, $(x \mid x)$, and $(y \mid y)$ are all real, we can deduce that $(x \mid y) + (y \mid x) \in \mathbf{R}$. Thus $\operatorname{Im}(x \mid y) = -\operatorname{Im}(y \mid x)$. On the other hand,
$$(x + iy \mid x + iy) = (x \mid x) + (y \mid y) - i(x \mid y) + i(y \mid x),$$
and again it must be the case that $-i(x \mid y) + i(y \mid x) \in \mathbf{R}$. Hence we have $\operatorname{Re}(x \mid y) = \operatorname{Re}(y \mid x)$, and we can conclude that $\overline{(y \mid x)} = (x \mid y)$. Therefore, the form is Hermitian. $\square$

Computations like those in the last proof actually allow us to say even more about Hermitian forms. There is a natural way of decomposing any Hermitian form on a vector space, called *polarization.*

Proposition 7.1.4 (Polarization Identity)**.** *Suppose X is a vector space over F and $(\cdot \mid \cdot)$ is a Hermitian sesquilinear form on X.*

1. *If $F = \mathbf{C}$, then for all $x, y \in X$ we have*

$$(x \mid y) = \frac{1}{4}\sum_{k=0}^{3} i^k\bigl(x + i^k y \bigm| x + i^k y\bigr).$$

2. *If $F = \mathbf{R}$, then for all $x, y \in X$ we have*

$$(x \mid y) = \frac{1}{4}\bigl((x + y \mid x + y) - (x - y \mid x - y)\bigr).$$

Proof. In both cases the polarization identity really just follows from simple algebraic manipulations. Suppose first that $X = \mathbf{C}$. Note that for all $x, y \in X$ we have

$$\begin{aligned}(x + y \mid x + y) &= (x \mid x) + (y \mid y) + (x \mid y) + (y \mid x)\\ &= (x \mid x) + (y \mid y) + (x \mid y) + \overline{(x \mid y)}\\ &= (x \mid x) + (y \mid y) + 2\operatorname{Re}(x \mid y),\end{aligned}$$

while

$$\begin{aligned}\bigl(x + i^2 y \bigm| x + i^2 y\bigr) &= (x - y \mid x - y)\\ &= (x \mid x) + (y \mid y) - (x \mid y) - (y \mid x)\\ &= (x \mid x) + (y \mid y) - (x \mid y) - \overline{(x \mid y)}\\ &= (x \mid x) + (y \mid y) - 2\operatorname{Re}(x \mid y).\end{aligned}$$

Hence

$$\begin{aligned}(x + y \mid x + y) + i^2\bigl(x + i^2 y \bigm| x + i^2 y\bigr) &= (x + y \mid x + y) - \bigl(x + i^2 y \bigm| x + i^2 y\bigr)\\ &= 4\operatorname{Re}(x \mid y).\end{aligned}$$

Similarly,

$$\begin{aligned}(x + iy \mid x + iy) &= (x \mid x) + (y \mid y) - i(x \mid y) + i(y \mid x)\\ &= (x \mid x) + (y \mid y) - i(x \mid y) + i\overline{(x \mid y)}\end{aligned}$$

$$= (x \mid x) + (y \mid y) - i \cdot 2i \operatorname{Im}(x \mid y),$$

and

$$\begin{aligned}(x + i^3 y | x + i^3 y) &= (x - iy \mid x - iy) \\ &= (x \mid x) + (y \mid y) + i(x \mid y) - i(y \mid x) \\ &= (x \mid x) + (y \mid y) + i(x \mid y) - i\overline{(x \mid y)} \\ &= (x \mid x) + (y \mid y) - i \cdot 2i \operatorname{Im}(x \mid y),\end{aligned}$$

so

$$\begin{aligned}i(x + iy \mid x + iy) + i^3(x + i^3 y \mid x + i^3 y) &= i(x + iy \mid x + iy) - i(x + i^3 y \mid x + i^3 y) \\ &= 4i \operatorname{Im}(x \mid y).\end{aligned}$$

It then follows that

$$\sum_{k=0}^{3} i^k \big(x + i^k y \mid x + i^k y\big) = 4 \operatorname{Re}(x \mid y) + 4i \operatorname{Im}(x \mid y) = 4(x \mid y).$$

Now suppose $F = \mathbf{R}$. Since a Hermitian form on X is really just a symmetric bilinear form, our earlier computations show that

$$(x + y \mid x + y) = (x \mid x) + (y \mid y) + 2(x \mid y)$$

and

$$(x - y \mid x - y) = (x \mid x) + (y \mid y) - 2(x \mid y).$$

Thus

$$(x + y \mid x + y) - (x - y \mid x - y) = 4(x \mid y),$$

and the result follows. □

An inner product is a sesquilinear form satisfying two additional properties, which allow us to define a norm using the form.

Definition 7.1.5. A sesquilinear form is said to be **positive** if $(x \mid x) \geq 0$ for all $x \in X$.

Since $(x \mid x) \geq 0$ implies $(x \mid x) \in \mathbf{R}$, if X is a $\mathbf{C}$-vector space then any positive sesquilinear form on X is automatically Hermitian by Proposition 7.1.3. The same is not true for sesquilinear forms on vector spaces over $\mathbf{R}$. For example, define $(\cdot \mid \cdot) : \mathbf{R}^2 \to \mathbf{R}$ by

$$(x \mid y) = 2x_1 y_1 + 2x_2 y_1 + 2x_2 y_2.$$

It is easy to check that this form is bilinear, and for all $x \in \mathbf{R}^2$ we have

$$(x \mid x) = 2x_1^2 + 2x_1x_2 + 2x_2^2 = x_1^2 + (x_1 + x_2)^2 + x_2^2 \geq 0.$$

However, $(\cdot \mid \cdot)$ is not symmetric—if we take $x = (1, 0)$ and $y = (1, 1)$, then

$$(x \mid y) = 2 + 0 + 0 = 2,$$

while

$$(y \mid x) = 2 + 2 + 0 = 4.$$

Note that this example shows that Proposition 7.1.3 also fails for $\mathbf{R}$-vector spaces.

Definition 7.1.6. Let X be a vector space over a field F. A positive, Hermitian sesquilinear form $(\cdot \mid \cdot)$ on X is called a **pre-inner product**. We say $(\cdot \mid \cdot)$ is an **inner product** if in addition $(x \mid x) = 0$ implies $x = 0$ for all $x \in X$.

Example 7.1.7. The usual dot product on $\mathbf{R}^n$ is an example of an inner product. The analog of the dot product for $\mathbf{C}^n$, given by

$$x \cdot y = x_1\overline{y}_1 + x_2\overline{y}_2 + \cdots + x_n\overline{y}_n$$

also defines an inner product.

Example 7.1.8. There is an easy way to generalize the dot products on $\mathbf{R}^n$ and $\mathbf{C}^n$ to obtain other inner products. Let $A \in M_n(\mathbf{R})$, and define $(\cdot \mid \cdot) : \mathbf{R}^n \times \mathbf{R}^n \to \mathbf{R}$ by

$$(x \mid y) = x^T Ay = x \cdot (Ay).$$

Since matrix multiplication distributes over vector addition and the dot product is bilinear, it is easy to see that $(\cdot \mid \cdot)$ is a bilinear form on $\mathbf{R}^n$. It can be shown that this form is symmetric if and only if A is a symmetric matrix, and if A is a positive-definite matrix (i.e., it is symmetric and all of its eigenvalues are positive), then $(\cdot \mid \cdot)$ is an inner product.

In the complex case, we can similarly define a form on $\mathbf{C}^n$ by taking a matrix $A \in M_n(\mathbf{C})$ and setting

$$(x \mid y) = x^T A\overline{y}$$

for all $x, y \in \mathbf{C}^n$. This form is Hermitian precisely when A is a Hermitian matrix, meaning that $A = A^*$, where A^* denotes the conjugate transpose. Any such matrix has real eigenvalues, and $(\cdot \mid \cdot)$ will be an inner product precisely when those eigenvalues are all positive.

Example 7.1.9. Let $(X, \mathcal{M}, \mu)$ be a measure space, and define $(\cdot \mid \cdot)$ on $L^2(X, \mu)$ by

$$(f \mid g) = \int_X f\overline{g}\, d\mu.$$

Recall that the Cauchy-Schwarz inequality guarantees that $f\overline{g} \in L^1(X, \mu)$, so this form is well-defined. It is also easily checked to be sesquilinear, and the definition of the integral guarantees that

$$(g \mid f) = \int_X g\overline{f}\, d\mu = \int_X \overline{f\overline{g}}\, d\mu = \overline{\int_X f\overline{g}\, d\mu} = \overline{(f \mid g)},$$

so we have a Hermitian form. Finally, observe that for all $f \in L^2(X, \mu)$,

$$(f \mid f) = \int_X f\overline{f}\, d\mu = \int_X |f|^2\, d\mu \geq 0,$$

with equality if and only if $f = 0$ almost everywhere. Hence $(\cdot \mid \cdot)$ is an inner product.

Recall that the dot product on $\mathbf{R}^n$ can be used to recover the Euclidean norm. If $x \in \mathbf{R}^n$, then

$$x \cdot x = x_1^2 + x_2^2 + \cdots + x_n^2 = \|x\|_2^2.$$

It turns out that any inner product defines a norm in a similar way: we simply set

$$\|x\| = (x \mid x)^{1/2}$$

for all $x \in X$. However, verifying that we actually obtain a norm (specifically, checking that the triangle inequality holds) requires us to first introduce a new version of an old friend.

Proposition 7.1.10 (Cauchy-Schwarz Inequality)**.** *Let X be a vector space over a field F. Suppose $(\cdot \mid \cdot)$ is a pre-inner product on X and define $\|x\| = (x \mid x)^{1/2}$ for all $x \in X$. Then for all $x, y \in X$,*

$$|(x \mid y)| \leq \|x\|\|y\|. \tag{7.1}$$

Proof. Fix $x, y \in X$ and let $\alpha \in F$. Then we have

$$\begin{aligned} 0 \leq (\alpha x + y \mid \alpha x + y) &= |\alpha|^2\|x\|^2 + \alpha(x \mid y) + \overline{\alpha}(y \mid x) + \|y\|^2 \\ &= |\alpha|^2\|x\|^2 + 2\operatorname{Re}(\alpha(x \mid y)) + \|y\|^2. \end{aligned}$$

Now let $\tau \in F$ with $|\tau| = 1$ and $\tau(x \mid y) = |(x \mid y)|$. Now let $\lambda \in \mathbf{R}$, and set $\alpha = \tau\lambda$ in the inequality above. Then we see that

$$\lambda^2\|x\|^2 + 2\lambda|(x \mid y)| + \|y\|^2 \geq 0$$

for all $\lambda \in \mathbf{R}$. The left hand side is a quadratic polynomial in λ which is always nonnegative, hence its discriminant is at most zero. More precisely,

$$4|(x \mid y)|^2 - 4\|x\|^2\|y\|^2 \leq 0,$$

and the result follows. $\square$

Corollary 7.1.11. *Suppose $(\cdot \mid \cdot)$ is a pre-inner product on X. Then*

$$\|x\| = (x \mid x)^{1/2}$$

defines a seminorm on X. If $(\cdot \mid \cdot)$ is an inner product, then $\|\cdot\|$ defines a norm.

Proof. It is obvious that $\|\cdot\|$ is nonnegative and homogeneous, so we simply need to verify that the triangle inequality holds. Given $x, y \in X$, observe that

$$\|x+y\|^2 = (x+y \mid x+y) = \|x\|^2 + 2\operatorname{Re}(x \mid y) + \|y\|^2 \leq \|x\|^2 + 2|(x \mid y)| + \|y\|^2.$$

The Cauchy-Schwarz inequality then gives

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

and the triangle inequality follows. Hence $\|\cdot\|$ is a seminorm. It is then easy to see that $\|\cdot\|$ will be positive definite (hence a norm) precisely when $(\cdot \mid \cdot)$ is an inner product. $\square$

There is actually an interesting way of characterizing norms that come from inner products. Any norm induced by an inner product must satisfy a geometric condition known as the *parallelogram law.*

Proposition 7.1.12. *Suppose X is an inner product space over a field F, and let $\|\cdot\|$ be the norm induced by the inner product. Then this norm satisfies the* ***parallelogram law****:*

$$\|x+y\|^2 + \|x-y\|^2 = 2\|x\|^2 + 2\|y\|^2$$

for all $x, y \in X$.

Proof. We observe that

$$\|x+y\|^2 = (x+y \mid x+y) = (x \mid x) + 2\operatorname{Re}(x \mid y) + (y \mid y)$$

and

$$\|x-y\|^2 = (x-y \mid x-y) = (x \mid x) - 2\operatorname{Re}(x \mid y) + (y \mid y),$$

so adding these two equations gives

$$\|x+y\|^2 + \|x-y\|^2 = 2(x \mid x) + 2(y \mid y) = 2\|x\|^2 + 2\|y\|^2. \qquad \square$$

Conversely, if X is a normed vector space such that the parallelogram law holds, then there exists an inner product on X that induces the given norm. In fact, the inner product is defined using a form of the polarization identity—one defines

$$(x \mid y) = \frac{1}{4}\sum_{k=0}^{3} i^k \left\|x + i^k y\right\|^2,$$

and the parallelogram law allows one to show that $(\cdot \mid \cdot)$ is an inner product. This result is known as the Jordan-von Neumann theorem, and its proof is quite involved. See [Wil] for a detailed treatment.

Definition 7.1.13. A **Hilbert space** is an inner product space $\mathcal{H}$ that is complete with respect to the norm induced from its inner product.

Example 7.1.14. The usual dot products on $\mathbf{R}^n$ and $\mathbf{C}^n$ induce the Euclidean norm, which makes these spaces into Hilbert spaces.

Example 7.1.15. Let $(X, \mathcal{M}, \mu)$ be a measure space. Since the inner product $(f \mid g) = \int_X f\overline{g}\,d\mu$ on $L^2(X, \mu)$ clearly induces the usual L^2-norm, it follows from the Riesz-Fischer theorem that $L^2(X, \mu)$ is a Hilbert space. In particular, ℓ^2 is a Hilbert space with respect to the inner product

$$(x \mid y) = \sum_{j=1}^{\infty} x_j \overline{y}_j.$$

Remark 7.1.16. Since the Banach space completion of an inner product space is a Hilbert space, we will often use the term *pre-Hilbert space* to refer to inner product spaces that are not complete.

7.2 The Riesz Representation Theorem

Recall from linear algebra that the dot product allows one to determine when two vectors in $\mathbf{R}^n$ are orthogonal. Naturally enough, we have a similar notion of orthogonality in any inner product space.

Definition 7.2.1. Let X be an inner product space. Two elements $x, y \in X$ are said to be **orthogonal**, denoted by $x \perp y$, if $(x \mid y) = 0$. We say that two sets $U \subseteq X$ and $V \subseteq X$ are **orthogonal**, written $U \perp V$, if $(x \mid y) = 0$ for all $x \in U$ and $y \in V$.

Example 7.2.2. Consider the Hilbert space $L^2([0, 2\pi], \mu)$. The functions $f(x) = \cos x$ and $g(x) = \sin x$ are orthogonal, since

$$(f \mid g) = \int_0^{2\pi} \cos x \sin x \, dx = \tfrac{1}{2} \sin^2 x \Big|_0^{2\pi} = 0.$$

If we define

$$U = \left\{ f \in L^2([0, 2\pi], \mu) \;\middle|\; f(x) = 0 \text{ for } \mu - \text{a.e.} 0 \leq x \leq \pi \right\}$$

and

$$V = \left\{ f \in L^2([0, 2\pi], \mu) \;\middle|\; f(x) = 0 \text{ for } \mu - \text{a.e.} \pi \leq x \leq 2\pi \right\},$$

then it is easy to see that $U \perp V$, since $fg = 0$ μ-a.e. whenever $f \in U$ and $g \in V$.

Remark 7.2.3. Observe that if $x, y \in X$ and $x \perp y$, then

$$(x + y \mid x + y) = (x \mid x) + (y \mid y).$$

More succinctly, we have

$$\|x + y\|^2 = \|x\|^2 + \|y\|^2,$$

which is sometimes called the **Pythagorean identity**.

Suppose $\mathcal{H}$ is a Hilbert space, and let $X \subseteq \mathcal{H}$. It is often useful to consider the set of all vectors that are orthogonal to X:

$$X^\perp = \left\{ h \in \mathcal{H} \;\middle|\; (h \mid x) = 0 \text{ for all } x \in X \right\}.$$

In the case that X is a subspace of $\mathcal{H}$, we will call $X^\perp$ the **orthogonal complement** of X. It turns out that $X^\perp$ is always a closed subspace of $\mathcal{H}$, regardless of whether X is even a subspace.

Theorem 7.2.4. *Let $\mathcal{H}$ be a Hilbert space. For any subset $X \subseteq \mathcal{H}$, $X^\perp$ is a closed subspace of $\mathcal{H}$.*

Proof. Let $X \subseteq \mathcal{H}$, and let $h_1, h_2 \in X^\perp$. Then for all $x \in X$,

$$(h_1 + h_2 \mid x) = (h_1 \mid x) + (h_2 \mid x) = 0,$$

so $h_1 + h_2 \in X^\perp$. Similarly, if $\alpha \in F$ and $h \in X^\perp$, then

$$(\alpha h \mid x) = \alpha (h \mid x) = 0$$

for all $x \in X$. Hence $X^\perp$ is a subspace of $\mathcal{H}$.

Suppose now that $(h_j)_{j=1}^\infty$ is a sequence in $X^\perp$, and that $h_j \to h$ for some $h \in \mathcal{H}$. The Cauchy-Schwarz inequality implies that the map $y \mapsto (x \mid y)$ is continuous, so

$$(x \mid h) = \lim_{j \to \infty} (x \mid h_j) = 0.$$

Thus $h \in X^\perp$, and it follows that $X^\perp$ is closed. $\square$

If W is a closed subspace of $\mathcal{H}$, then we can say even more about $W^\perp$. There is always a natural decomposition of $\mathcal{H}$ involving W and $W^\perp$. If $\mathcal{H}$ and $\mathcal{K}$ are Hilbert spaces, we define the **direct sum** of $\mathcal{H}$ and $\mathcal{K}$, denoted $\mathcal{H} \oplus \mathcal{K}$, by equipping the vector space $\mathcal{H} \times \mathcal{K}$ with the inner product

$$((h_1, k_1) \mid (h_2, k_2)) = (h_1 \mid h_2) + (k_1 \mid k_2). \tag{7.2}$$

Consequently, the norm on $\mathcal{H} \oplus \mathcal{K}$ takes the form

$$\|(h,k)\| = \sqrt{\|h\|^2 + \|k\|^2}.$$

It is left to the reader to verify that (7.2) defines an inner product and that $\mathcal{H} \oplus \mathcal{K}$ is a Hilbert space. Observe that $\mathcal{H}$ and $\mathcal{K}$ can be viewed as subspaces of $\mathcal{H} \oplus \mathcal{K}$ via the embeddings $h \mapsto (h, 0)$ and $k \mapsto (0, k)$, and moreover that $\mathcal{H} \perp \mathcal{K}$ inside $\mathcal{H} \oplus \mathcal{K}$.

Given a Hilbert space $\mathcal{H}$, we can also form the *internal* direct sum of two orthogonal subspaces W and V in a similar fashion. However, it is more customary in this situation to write the elements of $W \oplus V$ as $h + k$, where $h \in W$ and $k \in V$. The fact that $W \perp V$ guarantees that every element of $W \oplus V$ can be written uniquely in this way.

Our next goal is to show that if W is a closed subspace of a Hilbert space $\mathcal{H}$, then we can decompose $\mathcal{H}$ as the direct sum $W \oplus W^\perp$. To do so, we first need a small lemma.

Lemma 7.2.5. *Let $\mathcal{H}$ be a Hilbert space, and let W be a closed subspace of $\mathcal{H}$. Then for all $h \in \mathcal{H}$, there is a unique element $x \in W$ satisfying*

$$\|x - h\| = \inf_{z \in W} \|z - h\|.$$

Proof. Let $h \in \mathcal{H}$ and set $\alpha = \inf_{z \in W} \|z - h\|$. Let $(x_n)_{n=1}^\infty$ be a sequence in W such that $\|x_n - h\| \to \alpha$. The parallelogram law implies that for all $n, m \in \mathbf{N}$ we have

$$\begin{aligned} 2\|x_n - h\|^2 + 2\|x_m - h\|^2 &= \|x_n + x_m - 2h\|^2 + \|x_n - x_m\|^2 \\ &= 4\left\|\tfrac{1}{2}(x_n + x_m) - h\right\|^2 + \|x_n - x_m\|^2 \\ &\geq 4\alpha^2 + \|x_n - x_m\|^2, \end{aligned}$$

so

$$2\|x_n - h\|^2 + 2\|x_m - h\|^2 - 4\alpha^2 \geq \|x_n - x_m\|^2.$$

Since the left hand side goes to zero as $n, m \to \infty$, it follows that $(x_n)_{n=1}^\infty$ is Cauchy. Thus $x_n \to x$ for some $x \in W$, and $\|x - h\| = \alpha$. If $y \in W$ is any other element satisfying $\|y - h\| = \alpha$, then the same argument as above using the parallelogram law shows that

$$\|x - y\|^2 \leq 2\|x - h\|^2 + 2\|y - h\|^2 - 4\alpha^2 = 0.$$

Thus $x = y$, so x is unique. $\square$

Theorem 7.2.6. *Suppose W is a closed subspace of a Hilbert space $\mathcal{H}$. Any element $h \in \mathcal{H}$ can be written uniquely as a sum $h = h_1 + h_2$, with $h_1 \in W$ and $h_2 \in W^\perp$. Consequently,*

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

and $(W^\perp)^\perp = W$.

Proof. Fix $h \in \mathcal{H}$ and let h_1 be the closest point in W to h, as in Lemma 7.2.5. Put $h_2 = h - h_1$. If $z \in W$ and $\varepsilon > 0$, then

$$\begin{aligned}
\|h_2\|^2 &= \|h - h_1\|^2 \\
&\le \|h - (h_1 + \varepsilon z)\|^2 \\
&= \|h_2 - \varepsilon z\|^2 \\
&= \|h_2\|^2 - 2\varepsilon \operatorname{Re}(h_2 \mid z) + \varepsilon^2 \|z\|^2.
\end{aligned}$$

Thus $2\operatorname{Re}(h_2 \mid z) \le \varepsilon \|z\|^2$ for all $\varepsilon > 0$, so $2\operatorname{Re}(h_2 \mid z) \le 0$ for all $z \in W$. Now choose $\alpha \in \mathbf{C}$ such that $|\alpha| = 1$ and $\alpha(h_2 \mid z) \ge 0$. Since W is a subspace, $\overline{\alpha} z \in W$ and

$$\alpha(h_2 \mid z) = (h_2 \mid \overline{\alpha} z) = \operatorname{Re}(h_2 \mid \overline{\alpha} z) = 0.$$

It follows that $(h_2 \mid z) = 0$ for all $z \in W$, hence $h_2 \in W^\perp$.

Now if $h, k \in \mathcal{H}$ and we write $h = h_1 + h_2$ and $k = k_1 + k_2$ with $h_1, k_1 \in W$ and $h_2, k_2 \in W^\perp$, then the Pythagorean identity implies that

$$(h \mid k) = (h_1 \mid k_1) + (h_2 \mid k_2).$$

In particular, we have $\|h\|^2 = \|h_1\|^2 + \|h_2\|^2$. It follows that $\mathcal{H} = W \oplus W^\perp$.

It remains to see that the decomposition of each $h \in \mathcal{H}$ is unique. If $h = z_1 + z_2$ is another such decomposition, then

$$(z_1 - h_1) + (z_2 - h_2) = 0.$$

By the Pythagorean identity, $\|z_1 - h_1\| = \|z_2 - h_2\| = 0$, so $h_1 = z_1$ and $h_2 = z_2$.

Now suppose $h \in (W^\perp)^\perp$. We can write $h = h_1 + h_2$ as above with $h_1 \in W$ and $h_2 \in W^\perp$. But

$$\|h_2\|^2 = (h_2 \mid h_1 + h_2) = (h_2 \mid h) = 0.$$

Thus $h_2 = 0$, so $h \in W$. □

Corollary 7.2.7. *Let X be any subset of a Hilbert space $\mathcal{H}$. Then $(X^\perp)^\perp$ is the smallest closed subspace containing X. Thus if W is a subspace of $\mathcal{H}$,*

$$\overline{W} = (W^\perp)^\perp.$$

Proof. If W is any closed subspace containing X, then $X \subseteq W$ implies that $W^\perp \subseteq X^\perp$, hence $(X^\perp)^\perp \subseteq (W^\perp)^\perp = W$. □

As a byproduct of our investigation of orthogonal complements, we can now prove our first major result on Hilbert spaces. This result will show us that any Hilbert space is naturally isomorphic to its own dual space.

Theorem 7.2.8 (Riesz Representation Theorem)**.** *Let $\mathcal{H}$ be a Hilbert space and $\varphi \in \mathcal{H}^*$. Then there is a unique $y \in \mathcal{H}$ such that $\varphi(x) = (x \mid y)$ for all $x \in \mathcal{H}$. In particular, the map $\Phi : \mathcal{H} \to \mathcal{H}^*$ given by*

$$\Phi(y)(x) = (x \mid y)$$

defines a conjugate linear isomorphism of $\mathcal{H}$ with $\mathcal{H}^*$.

Proof. It is easy to check that the map Φ is conjugate linear. Also,

$$|\Phi(y)(x)| = |(x \mid y)| \leq \|x\|\|y\|$$

for all $x, y \in \mathcal{H}$ by the Cauchy-Schwarz inequality, which implies that $\|\Phi(y)\| \leq \|y\|$ for all $y \in \mathcal{H}$. But $\Phi(y)(y) = (y \mid y) = \|y\|^2$, so $\|\Phi(y)\| = \|y\|$.

Now let $\varphi \in \mathcal{H}^* \backslash \{0\}$. Then $\ker \varphi$ is a proper closed subspace of $\mathcal{H}$. Thus there exists $y \in (\ker \varphi)^\perp$ such that $\varphi(y) = 1$. Then for any $x \in \mathcal{H}$, $x - \varphi(x)y \in \ker \varphi$. Hence

$$(x \mid y) = (x - \varphi(x)y + \varphi(x)y \mid y) = \varphi(x)\|y\|^2.$$

It then follows that

$$\varphi(x) = (x \mid \|y\|^{-2}y) = \Phi(\|y\|^{-2}y)(x),$$

for all $x \in \mathcal{H}$. Hence $\Phi(\|y\|^{-2}y) = \varphi$, so Φ is surjective. □

7.3 Orthonormal Bases

As another consequence of the rigidity supplied by an inner product, we are about to see that vectors in a Hilbert space $\mathcal{H}$ can always be expressed in terms of an appropriate basis for $\mathcal{H}$. This statement might seem obvious—we know every vector space has a basis, so of course we can expand vectors as finite linear combinations of the elements of our chosen basis. However, if $\mathcal{H}$ is infinite-dimensional, any such algebraic basis for $\mathcal{H}$ must be *uncountable* as a consequence of the Baire category theorem. Furthermore, the basis might be impossible to describe. If we instead allow for the possibility of expressing vectors as *infinite* linear combinations (i.e., series) of basis vectors, then it becomes much easier to describe and work with such bases in a Hilbert space.

Definition 7.3.1. Let $\mathcal{H}$ be a Hilbert space. A set $\{e_j\}_{j\in J}$ of vectors in $\mathcal{H}$ is called **orthonormal** if

- $\|e_j\| = 1$ for all $j \in J$, and
- $(e_i \mid e_j) = 0$ if $i \neq j$.

If $\operatorname{span}\{e_j\}_{j\in J}$ is dense in $\mathcal{H}$, we call $\{e_j\}_{j\in J}$ an **orthonormal basis** for $\mathcal{H}$.

Remark 7.3.2. To reiterate our earlier discussion, an orthonormal basis is not a true basis for $\mathcal{H}$. For this reason, it is common to refer to an algebraic basis as a *Hamel basis*, while a "basis" that has dense span in a Banach space is called a *Schauder basis*.

We will see soon that if a Hilbert space $\mathcal{H}$ is separable (i.e., it has a countable dense subset), then it possesses a *countable* orthonormal basis. However, there are Hilbert spaces that are not separable, and we must therefore be able to talk about infinite sums of vectors over *uncountable* index sets. Let J be a set, and let Λ denote the set of all finite subsets of J. If $f : J \to X$ is a function from J into a normed vector space X, then we say the series

$$\sum_{j\in J} f(j)$$

converges to x if given any $\varepsilon > 0$, there exists $\lambda_0 \in \Lambda$ such that

$$\left\| x - \sum_{j\in\lambda} f(j) \right\| < \varepsilon$$

for all $\lambda \in \Lambda$ with $\lambda \supseteq \lambda_0$.[1] In this case it is also customary to write

$$x = \lim_{\lambda\in\Lambda} \sum_{j\in\lambda} f(j).$$

Example 7.3.3. Suppose $f : J \to [0, \infty) \subseteq \mathbf{R}$. Then we say that

$$\sum_{j\in J} f(j) = \infty$$

if given $M > 0$, there exists λ_0 such that $\lambda \supseteq \lambda_0$ implies that

$$\sum_{j\in\lambda} f(j) \geq M.$$

[1] This really says that if we make Λ into a directed set via inclusion, then $(\sum_{j\in\lambda} f(j))_{\lambda\in\Lambda}$ defines a *net* which converges to x in X.

With this convention, we claim that

$$\sum_{j\in J} f(j) = \sup_{\lambda\in\Lambda} \sum_{j\in\lambda} f(j).$$

Let $\alpha = \sup_{\lambda\in\Lambda} \sum_{j\in\lambda} f(j)$. First we note that if $\sum_{j\in J} f(j) = \infty$, then $\alpha = \infty$. On the other hand, if $\sum_{j\in J} f(j)$ is finite then we clearly have $\lim_{\lambda\in\Lambda} \sum_{j\in\lambda} f(j) \le \alpha$ and for any $\varepsilon > 0$, there exists $\lambda_0 \in \Lambda$ such that

$$\sum_{j\in\lambda_0} f(j) > \alpha - \varepsilon.$$

But if $\lambda \supseteq \lambda_0$, then we see that

$$\sum_{j\in\lambda_0} f(j) \le \sum_{j\in\lambda} f(j) \le \alpha,$$

so

$$\left| \sum_{j\in\lambda} f(j) - \alpha \right| < \varepsilon.$$

Thus $\sum_{j\in J} f(j)$ converges to α.

If we take $J = \mathbf{N}$ in the previous example, then $f : \mathbf{N} \to [0,\infty)$ is simply a sequence of nonnegative real numbers, and our calculations show that

$$\sum_{j\in\mathbf{N}} f(j) = \sum_{j=1}^{\infty} f(j) = \lim_{n\to\infty} \sum_{j=1}^{n} f(j).$$

However, these sums need not agree in general if $f : \mathbf{N} \to \mathbf{R}$, though they will if we assume that $\sum_{j=1}^{\infty} f(j)$ converges absolutely. To establish this fact, it is helpful to introduce a little measure theory.

Remark 7.3.4. Let J be a set, and let ν denote the counting measure on $\mathcal{P}(J)$. Then if $1 \le p < \infty$, it is natural to write $\ell^p(J)$ in place of $L^p(J,\nu)$. It can then be shown that

$$\|f\|_p = \left(\sum_{j\in J} |f(j)|^p \right)^{1/p},$$

and if $f \in \ell^1(J)$, then

$$\int_J f\, d\nu = \sum_{j\in J} f(j).$$

To see the first assertion, observe that

$$\|f\|_p^p = \int_J |f|^p\, d\nu = \sup_{0\le\varphi\le|f|^p} \int_J \varphi\, d\nu,$$

where the supremum is taken over all simple functions $0 \leq \varphi \leq |f|^p$. Let Λ denote the collection of all finite subsets of J, and for each $\lambda \in \Lambda$, define $f_\lambda = f \cdot \chi_\lambda$. Then each f_λ is a simple function, hence

$$\sup_{0\leq\varphi\leq|f|^p} \int_J \varphi \, d\nu \geq \sup_{\lambda\in\Lambda} \int_J |f_\lambda|^p \, d\nu = \sup_{\lambda\in\Lambda} \sum_{j\in\lambda} |f_\lambda(j)|^p = \sup_{\lambda\in\Lambda} \sum_{j\in\lambda} |f(j)|^p.$$

It follows that $\|f\|_p^p \geq \sum_{j\in J} |f(j)|^p$. On the other hand, if φ is a simple function with $0 \leq \varphi \leq |f|^p$, then there exists $\lambda \in \Lambda$ such that $\varphi \leq |f_\lambda|^p \leq |f|^p$. (More specifically, $\varphi \in \ell^1(J)$, which implies that $\operatorname{supp}(\varphi)$ is finite. Hence we could take $\lambda = \operatorname{supp}(\varphi)$.) Therefore,

$$\sup_{0\leq\varphi\leq|f|^p} \int_J \varphi \, d\nu \leq \sup_{\lambda\in\Lambda} \int_J |f|^p \, d\nu,$$

and we thus have $\|f\|_p^p \leq \sum_{j\in J} |f(j)|^p$. It then follows that

$$\|f\|_p = \left(\sum_{j\in J} |f(j)|^p \, d\nu(j) \right)^{1/p}.$$

The second assertion now follows by considering positive and negative parts.

Corollary 7.3.5. *If $(a_n)_{n=1}^\infty \in \ell^1(\mathbf{N})$, then*

$$\|a\|_1 = \sum_{n\in\mathbf{N}} |a_n| = \sum_{n=1}^\infty |a_n|.$$

Moreover, the series $\sum_{n\in\mathbf{N}} a_n$ and $\sum_{n=1}^\infty a_n$ both converge and are equal.

Finally, we make one almost obvious observation regarding series with uncountably many terms.

Proposition 7.3.6. *Suppose J is a set and $f : J \to [0,\infty)$. If $\sum_{j\in J} f(j) < \infty$, then $f(j) = 0$ for all but countably many j.*

Proof. Let $J_n = \{j \in J : |f(j)| \geq 1/n\}$. Since the series converges, J_n is necessarily finite for all n. But then

$$\{j \in J : f(j) \neq 0\} = \bigcup_{n=1}^\infty J_n$$

is a countable union of finite sets, hence countable. $\square$

We have said more than enough about uncountably infinite sums. Let us see what we can deduce about orthonormal bases.

Theorem 7.3.7 (Bessel's Inequality)**.** *Let $\mathcal{H}$ be a Hilbert space, and let $\{e_j\}_{j\in J}$ be an orthonormal set in $\mathcal{H}$. Then for any $x \in \mathcal{H}$,*

$$\|x\|^2 \geq \sum_{j\in J} |(x \mid e_j)|^2.$$

Proof. Suppose $\lambda \subseteq J$ is finite. Then we have

$$\begin{aligned} 0 &\leq \left\| x - \sum_{j\in\lambda} (x|e_j)e_j \right\|^2 \\ &= \|x\|^2 - 2\,\mathrm{Re}\big(x \mid \sum_{j\in\lambda}(x \mid e_j)e_j\big) + \left\|\sum_{j\in\lambda}(x \mid e_j)e_j\right\|^2 \\ &= \|x\|^2 - 2\,\mathrm{Re}\sum_{j\in\lambda}(x \mid (x \mid e_j)e_j) + \left\|\sum_{j\in\lambda}(x \mid e_j)e_j\right\|^2 \\ &= \|x\|^2 - 2\,\mathrm{Re}\sum_{j\in\lambda}|(x \mid e_j)|^2 + \sum_{j\in\lambda}|(x \mid e_j)|^2 \\ &= \|x\|^2 - \sum_{j\in\lambda}|(x|e_j)|^2 \end{aligned}$$

By the Pythagorean identity. It follows that

$$\|x\|^2 \geq \sup_{\lambda} \sum_{j\in\lambda} |(x \mid e_j)|^2$$

where the supremum is taken over all finite $\lambda \subseteq J$, whence the result. $\square$

Corollary 7.3.8. *Suppose $\mathcal{H}$ is a Hilbert space and $\{e_j\}_{j\in J}$ is an orthonormal set in $\mathcal{H}$. Then for any $x \in \mathcal{H}$, the sum*

$$\sum_{j\in J}(x \mid e_j)e_j$$

converges in $\mathcal{H}$.

Proof. By Bessel's inequality, $\sum_{j\in J}|(x \mid e_j)|^2 < \infty$, hence $(x \mid e_j) = 0$ for all but countably many j. Let $j_1, j_2, \ldots$ be an enumeration of those j such that $(x \mid e_j) \neq 0$.

Now by the Pythagorean identity,

$$\left\| \sum_{k=n}^{m} (x \mid e_{j_k}) \right\|^2 = \sum_{k=n}^{m} |(x \mid e_{j_k})|^2,$$

and the right hand side tends to 0 as $n, m \to \infty$ since $\sum_{k=1}^{\infty} |(x \mid e_{j_k})|^2$ converges. Since $\mathcal{H}$ is complete, there exists $y \in \mathcal{H}$ such that

$$y = \sum_{k=1}^{\infty} (x \mid e_{j_k}) e_{j_k}.$$

Let $s_N = \sum_{k=1}^{N} (x \mid e_{j_k})$, and let Λ denote the collection of all finite subsets of J. We see that $s_N \to y$ in $\mathcal{H}$, so given $\varepsilon > 0$, there exists N_0 such that $N, M \geq N_0$ implies that $\|s_N - y\| < \varepsilon$ and $\|s_N - s_M\| < \varepsilon$. Now let $\lambda_0 = \{1, 2, \ldots, N_0\}$. Then if $\lambda \supseteq \lambda_0$, we have

$$\left\| \sum_{j \in \lambda} (x \mid e_j) e_j - y \right\| < \varepsilon.$$

Thus

$$y = \sum_{j \in J} (x \mid e_j) e_j,$$

and we are done. □

Proposition 7.3.9. *Let $\mathcal{H}$ be a Hilbert space, and suppose $\{e_j\}_{j \in J}$ is an orthonormal set in $\mathcal{H}$. Then the following are equivalent.*

1. *The set $\{e_j\}_{j \in J}$ forms an orthonormal basis for $\mathcal{H}$.*
2. *If $(x \mid e_j) = 0$ for all $j \in J$, then $x = 0$.*
3. *(Parseval's Identity) For all $x \in \mathcal{H}$,*

$$\|x\|^2 = \sum_{j \in J} |(x \mid e_j)|^2.$$

4. *For all $x \in \mathcal{H}$, $x = \sum_{j \in J} (x \mid e_j) e_j$.*

Proof. Let $S = \operatorname{span}\{e_j\}_{j \in J}$. We first claim that S is dense in $\mathcal{H}$ if and only if $S^{\perp} = \{0\}$. To see this, first suppose that S is dense in $\mathcal{H}$. Let $y \in \mathcal{H}$ and choose ε such that $0 < \varepsilon < \|y\|$. Then there exists $x \in S$ such that $\|x - y\| < \varepsilon$. But then

$$\varepsilon^2 > \|x - y\|^2 = \|x\|^2 - 2\operatorname{Re}(x \mid y) + \|y\|^2,$$

and since $\|y\|^2 > \varepsilon^2$, we must have $\mathrm{Re}(x \mid y) \neq 0$. Thus $y \notin S^\perp$. On the other hand, if $y \in S^\perp$ and $0 < \varepsilon < \|y\|$, then for all $x \in S$ we have

$$\|x - y\|^2 = \|x\|^2 + \|y\|^2 > \varepsilon^2$$

by the Pythagorean identity. Hence $\|x - y\| > \varepsilon$ for all $x \in S$, so S is not dense in $\mathcal{H}$. It is now easy to see from our claim that (1) and (2) are equivalent.

Now we will show that (2) implies (4). By Corollary 7.3.8, the series $y = \sum_{j\in J}(x \mid e_j)e_j$ converges in $\mathcal{H}$. The map $z \mapsto (z \mid e_k)$ defines a bounded linear functional on $\mathcal{H}$, hence

$$(y \mid e_k) = \sum_{j\in J}(x \mid e_j)(e_j \mid e_k) = (x \mid e_k).$$

Thus $(x - y \mid e_j) = 0$ for all $j \in J$, so $x = y$.

Next we show that (4) implies (3). Well, observe that if $\lambda \subseteq J$ is finite, then we have

$$\left\| x - \sum_{j\in\lambda}(x \mid e_j)e_j \right\|^2 = \|x\|^2 - \sum_{j\in\lambda} |(x \mid e_j)|^2.$$

By assumption, the left hand side tends to zero as λ increases, so

$$\|x\|^2 = \sum_{j\in J} |(x \mid e_j)|^2$$

and Parseval's identity holds. Finally, we note that (3) clearly implies (2), completing the proof. $\square$

Just as one can always extend a linearly independent set to an algebraic basis for a vector space, we can extend orthonormal sets in Hilbert spaces to obtain orthonormal bases.

Proposition 7.3.10. *Every orthonormal set in a Hilbert space $\mathcal{H}$ can be extended to an orthonormal basis.*

Proof. Let $\{e_j\}_{j\in J_0}$ be an orthonormal set in $\mathcal{H}$, and let Λ denote the collection of all orthonormal sets in $\mathcal{H}$ that contain $\{e_j\}_{j\in J_0}$. Partially order Λ by containment. Let $\mathcal{C} \subseteq \Lambda$ be a chain in Λ, and put

$$E_0 = \bigcup_{E\in\mathcal{C}} E.$$

Clearly E_0 contains $\{e_j\}_{j\in J_0}$ and $E \preccurlyeq E_0$ for all $E \in \mathcal{C}$. Furthermore, if $e, d \in E_0$, then $e, d \in E$ for some $E \in \mathcal{C}$, so $(e \mid d) = 1$ if $e = d$ and $(e \mid d) = 0$ if $e \neq d$.

Therefore, E_0 is an orthonormal set, hence $E_0 \in \Lambda$ and E_0 is an upper bound for $\mathcal{C}$. Zorn's Lemma now applies, so Λ contains a maximal element. Let $\{e_j\}_{j\in J}$ denote this element, and let

$$W = \overline{\operatorname{span}\{e_j\}_{j\in J}}.$$

Suppose $W \neq \mathcal{H}$, and let $e \in W^\perp$ be a unit vector. Then $\{e_j\}_{j\in J} \cup \{e\}$ is an orthonormal set which properly contains $\{e_j\}_{j\in J}$. This contradicts the maximality of $\{e_j\}_{j\in J}$, so we must have $W = \mathcal{H}$. It follows that $\{e_j\}_{j\in J}$ forms an orthonormal basis for $\mathcal{H}$. $\square$

Definition 7.3.11. Let $\mathcal{H}$ and $\mathcal{K}$ be Hilbert spaces over a field F. A linear surjection $U : \mathcal{H} \to \mathcal{K}$ is called a **unitary transformation** if

$$(Ux \mid Uy) = (x \mid y)$$

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

Note that a unitary transformation is actually $U : \mathcal{H} \to \mathcal{K}$ an isometric isomorphism of $\mathcal{H}$ onto $\mathcal{K}$, since it preserve inner products: if $x \in \mathcal{H}$, then

$$\|Ux\|^2 = (Ux \mid Ux) = (x \mid x) = \|x\|^2.$$

Furthermore, it is not hard to see that the inverse $U^{-1} : \mathcal{K} \to \mathcal{H}$ is also unitary. Thus unitaries provide the appropriate notion of isomorphism between Hilbert spaces.

Lemma 7.3.12. *Suppose $\mathcal{H}$ and $\mathcal{K}$ are Hilbert spaces and that $U : \mathcal{H} \to \mathcal{K}$ is a surjective isometry. Then U is unitary.*

Proof. The result really just follows from the polarization identity. If $F = \mathbf{C}$, then

$$\begin{aligned}
(Ux \mid Uy) &= \frac{1}{4}\sum_{k=0}^{3} i^k (Ux + i^k Uy \mid Ux + i^k Uy) \\
&= \frac{1}{4}\sum_{k=0}^{3} i^k (U(x + i^k y) \mid U(x + i^k y)) \\
&= \frac{1}{4}\sum_{k=0}^{3} i^k \big\|U(x + i^k y)\big\|^2 \\
&= \frac{1}{4}\sum_{k=0}^{3} i^k \big\|x + i^k y\big\|^2 \\
&= (x \mid y).
\end{aligned}$$

Similarly, if $F = \mathbf{R}$, we have

$$\begin{aligned}(Ux \mid Uy) &= \tfrac{1}{4}\left(\|Ux + Uy\|^2 - \|x - y\|^2\right)\\ &= \tfrac{1}{4}\left(\|U(x + y)\|^2 - \|U(x - y)\|^2\right)\\ &= \tfrac{1}{4}\left(\|x + y\|^2 - \|x - y\|^2\right)\\ &= (x \mid y),\end{aligned}$$

so U is unitary. $\square$

Recall from linear algebra that one can define a linear transformation between vector spaces by simply specifying where a particular basis should be sent. A similar approach works with orthonormal bases for a Hilbert space, provided one takes care with boundedness issues. When defining unitary operators, there are no such problems.

Proposition 7.3.13. *Suppose $\mathcal{H}$ and $\mathcal{K}$ are Hilbert spaces over a field F with orthonormal bases $\{e_j\}_{j\in J}$ and $\{f_i\}_{i\in I}$, respectively. If $\operatorname{card} J = \operatorname{card} I$, then there is a unitary transformation $U : \mathcal{H} \to \mathcal{K}$.*

Proof. If $\operatorname{card} I = \operatorname{card} J$, then there exists a bijection $\gamma : I \to J$. Let $W = \operatorname{span}\{e_j\}_{j\in J}$ and $V = \operatorname{span}\{f_i\}_{i\in I}$. Then since $\{e_j\}_{j\in J}$ is linearly independent, we can define $U_0 : W \to V$ by

$$U_0\left(\sum_{j\in\lambda} \alpha_j e_j\right) = \sum_{j\in\lambda} \alpha_j f_{\gamma(j)}.$$

It is clear that U_0 is linear by definition, and Parseval's identity guarantees it is isometric. Furthermore, since γ is surjective, U_0 maps W onto V. Since W is dense in $\mathcal{H}$, U_0 extends to an isometry $U : \mathcal{H} \to \mathcal{K}$. It is immediate that U is surjective since V is dense in $\mathcal{K}$. But then Lemma 7.3.12 implies that U is unitary. $\square$

Recall that the statement $\operatorname{card} I = \operatorname{card} J$ simply means that there is a bijection $f : I \to J$. Of course when I and J are finite, we are simply saying that the two sets have the same number of elements. For finite sets, it is also natural to talk about the cardinality of one set being less than or equal to that of another set. This notion also extends to infinite sets—we say that $\operatorname{card} I \leq \operatorname{card} J$ if there is an injection $f : I \to J$. It turns out that if $\operatorname{card} I \leq \operatorname{card} J$ and $\operatorname{card} J \leq \operatorname{card} I$, then $\operatorname{card} I = \operatorname{card} J$. This result is known as the **Schröder-Bernstein theorem**.

Proposition 7.3.14. *Suppose* $\{e_j\}_{j\in J}$ *and* $\{f_i\}_{i\in I}$ *are orthonormal bases for a Hilbert space* $\mathcal{H}$. *Then* $\operatorname{card} I = \operatorname{card} J$.

Proof. For each $j \in J$, let

$$A_j = \{i \in I : (e_j \mid f_i) \neq 0\}.$$

Note that A_j is countable as a consequence of Bessel's inequality. Since $\operatorname{span}\{e_j\}_{j\in J}$ is dense in $\mathcal{H}$, for each $i \in I$ there is a $j \in J$ such that $(e_j \mid f_i) \neq 0$. Thus

$$I = \bigcup_{j\in J} A_j,$$

and since $\operatorname{card} \bigcup_{j\in J} A \leq \operatorname{card} J$, it follows that $\operatorname{card} I \leq \operatorname{card} J$. This argument is symmetric in I and J, so $\operatorname{card} J \leq \operatorname{card} I$. The Schröder-Bernstein theorem then implies that $\operatorname{card} I = \operatorname{card} J$. □

Since the cardinality of any orthonormal basis for a given Hilbert space is fixed, we have a well-defined notion of the "size" of a Hilbert space.

Definition 7.3.15. The **dimension** of a Hilbert space $\mathcal{H}$ is the cardinality of any orthonormal basis for $\mathcal{H}$.

Note that our definition of dimension is not the same as the usual algebraic definition, since an orthonormal basis is not really a basis for $\mathcal{H}$. In particular, $\mathcal{H}$ could have a countably infinite orthonormal basis (so $\dim \mathcal{H} = \aleph_0$), but its algebraic dimension would have to be at least $\aleph_1$.

Proposition 7.3.16. *A Hilbert space* $\mathcal{H}$ *is separable if and only if* $\dim \mathcal{H} \leq \aleph_0$.

Proof. Suppose first that $\dim \mathcal{H} \leq \aleph_0$. Then $\mathcal{H}$ has an orthonormal basis $\{e_j\}_{j\in J}$ with J countable, and $\operatorname{span}\{e_j\}_{j\in J}$ is dense in $\mathcal{H}$. Consequently, the set

$$\mathcal{H}_0 = \left\{ \sum_{j\in J} \alpha_j e_j \;\middle|\; \alpha_j \in \mathbf{Q} + \mathbf{Q}i \right\},$$

i.e., the "rational span" of $\{e_j\}_{j\in J}$ is easily checked to be dense in $\mathcal{H}$. Thus $\mathcal{H}$ is separable.

Now suppose $\mathcal{H}$ is separable, and let $D = \{d_n\}_{n=1}^{\infty}$ be a countable dense subset of $\mathcal{H}$. Let $\{e_j\}_{j\in J}$ be an orthonormal basis for $\mathcal{H}$, and define

$$A_n = \bigl\{ j \in J \bigm| (e_j \mid d_n) \neq 0 \bigr\}$$

for each $n \in \mathbf{N}$. As a consequence of Parseval's identity, each A_n is countable. Since D is dense in $\mathcal{H}$, it must be the case that

$$J = \bigcup_{n=1}^{\infty} A_n,$$

so J is countable. □

Example 7.3.17. Consider the Hilbert space ℓ^2. For each $n \in \mathbf{N}$, define $\delta_n = \chi_{\{n\}}$. Then we have

$$\delta_n \delta_m = \begin{cases} \delta_n & \text{if } n = m \\ 0 & \text{if } n \neq m, \end{cases}$$

so

$$(\delta_n \mid \delta_m) = \sum_{j=1}^{\infty} \delta_n(j)\delta_m(j) = \begin{cases} 1 & \text{if } n = m \\ 0 & \text{if } n \neq m. \end{cases}$$

It is also easy to see that if $x = (x_n)_{n=1}^{\infty}$ belongs to ℓ^2, then

$$x = \sum_{n=1}^{\infty} x_n \delta_n.$$

Hence $\{\delta_n\}_{n=1}^{\infty}$ is an orthonormal basis for ℓ^2, and ℓ^2 is a separable Hilbert space. In fact, since Proposition 7.3.13 says that two Hilbert spaces are isomorphic (via a unitary) if and only if they have the same dimension, *any* separable Hilbert space is isomorphic to ℓ^2.

7.4 Operators on Hilbert Space

The final statement in Example 7.3.17 makes Hilbert spaces seem quite boring. In particular, we know from Proposition 7.3.13 that any two Hilbert spaces of the same dimension are isomorphic. Not to worry though—analysts are interested in the study of *operators* on Hilbert spaces.

Since a Hilbert space is a special kind of Banach space, we already have a notion of boundedness for Hilbert space operators. We also know quite a bit about bounded operators from our previous study, but the presence of an inner product will allow us to say even more. We first need to make a connection between bounded operators and sesquilinear forms.

Definition 7.4.1. Let $\mathcal{H}$ be a Hilbert space over a field F. A sesquilinear form $B : \mathcal{H} \times \mathcal{H} \to F$ is said to be **bounded** if there exists $\alpha \geq 0$ such that

$$|B(x, y)| \leq \alpha \|x\| \|y\|$$

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

Clearly the Cauchy-Schwarz inequality tells us that any pre-inner product is a bounded sesquilinear form, where we can take $\alpha = 1$. In general, we can define the **norm** of a bounded sesquilinear form by

$$\|B\| = \sup_{\substack{\|x\|\leq 1\\ \|y\|\leq 1}} |B(x,y)|.$$

Then just as for bounded operators, we have

$$|B(x,y)| \leq \|B\|\|x\|\|y\|$$

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

Proposition 7.4.2. *Let $\mathcal{H}$ be a Hilbert space. For each $T \in B(\mathcal{H})$, define $B_T : \mathcal{H} \times \mathcal{H} \to F$ by*

$$B_T(x,y) = (x \mid Ty).$$

Then B_T is a bounded sesquilinear form, and $\|B_T\| = \|T\|$. Moreover, the map $T \mapsto B_T$ gives a one-to-one correspondence between the elements of $B(\mathcal{H})$ and bounded sesquilinear forms on $\mathcal{H}$.

Proof. Let $T \in B(\mathcal{H})$. Then B_T is clearly sesquilinear, and the Cauchy-Schwarz inequality gives

$$|B_T(x,y)| = |(x \mid Ty)| \leq \|x\|\|Ty\| \leq \|T\|\|x\|\|y\|,$$

for all $x, y \in \mathcal{H}$. Thus B_T is bounded with $\|B_T\| \leq \|T\|$. On the other hand, for all $x \in \mathcal{H}$ we have

$$\|Tx\|^2 = (Tx \mid Tx) = B_T(Tx, x) \leq \|B_T\|\|Tx\|\|x\| \leq \|B_T\|\|T\|\|x\|^2.$$

Thus

$$\|T\|^2 = \sup_{\|x\|\leq 1} \|Tx\|^2 \leq \|B_T\|\|T\|,$$

so $\|T\| \leq \|B_T\|$, and it follows that $\|B_T\| = \|T\|$.

Now assume that B is a bounded sesquilinear form on $\mathcal{H}$. Then for each $y \in \mathcal{H}$, the map $x \mapsto B(x,y)$ defines a bounded linear functional on $\mathcal{H}$. Thus the Riesz Representation Theorem guarantees that there is a unique $Ty \in \mathcal{H}$ such that $B(x,y) = (x \mid Ty)$. It is easy to see that the map $y \to Ty$ is linear, for if $y_1, y_2 \in \mathcal{H}$, then

$$B(x, y_1 + y_2) = (x \mid T(y_1 + y_2)),$$

while we also have

$$B(x, y_1 + y_2) = B(x, y_1) + B(x, y_2) = (x \mid Ty_1) + (x \mid Ty_2) = (x \mid Ty_1 + Ty_2).$$

Since $T(y_1 + y_2)$ is unique by the Riesz Representation Theorem, it follows that $T(y_1 + y_2) = Ty_1 + Ty_2$. Now observe that for any $y \in \mathcal{H}$,

$$\|Ty\|^2 = |B(Ty, y)| \leq \|B\|\|Ty\|\|y\|,$$

so $\|Ty\| \leq \|B\|\|y\|$. Therefore, $T \in B(\mathcal{H})$ and $B = B_T$, so we are done. $\square$

An important consequence of the last result is the existence of an *adjoint* for any bounded linear operator.

Theorem 7.4.3. *Let $\mathcal{H}$ be a Hilbert space, and let $T \in B(\mathcal{H})$. There exists a unique $T^* \in B(\mathcal{H})$, called the* ***adjoint*** *of T, such that*

$$(Tx \mid y) = (x \mid T^*y)$$

for all $x, y \in \mathcal{H}$. Furthermore, $\|T^*\| = \|T\|$.

Proof. It is easy to see that $B(x, y) = (Tx \mid y)$ defines a bounded sesquilinear form on $\mathcal{H}$. Thus the previous proposition guarantees that there is a unique operator $T^* \in B(\mathcal{H})$ satisfying

$$(Tx \mid y) = B(x, y) = (x \mid T^*y)$$

for all $x, y \in \mathcal{H}$. Clearly $\|T^*T\| \leq \|T^*\|\|T\|$, and

$$\|Tx\|^2 = (Tx \mid Tx) = (T^*Tx \mid x) \leq \|T^*T\|\|x\|^2.$$

Taking suprema over all x with $\|x\| \leq 1$, we get $\|T\|^2 \leq \|T^*T\|$. Hence $\|T\|^2 \leq \|T^*\|\|T\|$, so $\|T\| \leq \|T^*\|$. By symmetry, $\|T^*\| \leq \|T^{**}\| = \|T\|$, so $\|T^*\| = \|T\|$. $\square$

Example 7.4.4. Let $\mathcal{H} = \mathbf{C}^n$ and let $\beta = \{e_j\}_{j=1}^n$ be its standard orthonormal basis. If $T \in B(\mathbf{C}^n)$, then recall that the matrix representation $[T]_\beta$ of T with respect to the basis $\{e_j\}_{j=1}^n$ is the matrix (a_{ij}), where

$$a_{ij} = (Te_j \mid e_i).$$

Then $[T^*]_\beta = (b_{ij})$, where

$$b_{ij} = (T^*e_j \mid e_i) = (e_j \mid Te_i) = \overline{(Te_i \mid e_j)} = \overline{a_{ji}}.$$

Then the matrix $[T^*]_\beta$ of T^* is just the conjugate transpose of $[T]_\beta$. Thus the adjoint of an operator generalizes the conjugate transpose (and hence the transpose) of a matrix.

Proposition 7.4.5. *Let $\mathcal{H}$ be a Hilbert space over F. Then for all $T, S \in B(\mathcal{H})$, we have:*

1. $T^{**} = T$;
2. $(TS)^* = S^*T^*$;
3. *if* $\alpha \in F$, *then* $(T + \alpha S)^* = T^* + \overline{\alpha}S^*$;
4. $\|T^*T\| = \|T\|^2$.

Proof. Let $T \in B(\mathcal{H})$, and observe that

$$(x \mid T^{**}x) = (T^*x \mid y) = \overline{(y \mid T^*x)} = \overline{(Ty \mid x)} = (x \mid Ty)$$

for all $x, y \in \mathcal{H}$. Thus we must have $T^{**} = T$, proving (1). To establish (2), note that

$$(x \mid (TS)^*y) = (TSx \mid y) = (Sx \mid T^*y) = (x \mid S^*T^*y)$$

for all $x, y \in \mathcal{H}$, so $(TS)^* = S^*T^*$. Similarly, for (3) we have

$$\begin{aligned}(x \mid (T + \alpha S)^*y) &= ((T + \alpha S)x \mid y) \\ &= (Tx + \alpha Sx \mid y) \\ &= (Tx \mid y) + \alpha(Sx \mid y) \\ &= (x \mid T^*y + \overline{\alpha}S^*y)\end{aligned}$$

for all $x, y \in \mathcal{H}$ and $\alpha \in F$. Thus $(T + \alpha S)^* = T^* + \overline{\alpha}S^*$. Finally, observe that

$$\|T\|^2 \leq \|T^*T\| \leq \|T^*\|\|T\| = \|T\|^2,$$

so $\|T^*T\| = \|T\|^2$, and (4) holds. □

The equation in part 4 of the previous proposition is known as the C^**-identity*. The reason for this name is that such an identity is the defining characteristic of a C^**-algebra*.

Definition 7.4.6. A C^***-algebra** is a Banach algebra A over $\mathbf{C}$ with the following additional properties.

1. There is a conjugate linear map $^* : A \to A$ satisfying

$$(a^*)^* = a$$

and

$$(ab)^* = b^*a^*$$

for all $a, b \in A$.[2]

2. The C^***-identity** holds, i.e.,
$$\|a^*a\| = \|a\|^2$$
for all $a \in A$.

Of we can now rephrase Proposition 7.4.5 as saying that $B(\mathcal{H})$ is a C^*-algebra for any Hilbert space $\mathcal{H}$. We have seen other examples of C^*-algebras as well—if X is a locally compact metric space, then $C_0(X)$ and $C_b(X)$ are both C^*-algebras under the involution
$$f^*(x) = \overline{f(x)}.$$
In fact, they are both examples of *commutative* C^*-algebras.

The C^*-identity alone introduces a certain rigidity into the study of C^*-algebras that often produces seemingly magical results. One simple example involves homomorphisms between C^*-algebras.

Definition 7.4.7. Suppose A and B are C^*-algebras. A map $\varphi : A \to B$ is called a $*$**-homomorphism** if

1. φ is linear;
2. $\varphi(ab) = \varphi(a)\varphi(b)$ for all $a, b \in A$;
3. $\varphi(a^*) = \varphi(a)^*$ for all $a \in A$.

In other words, a $*$-homomorphism is simply the correct notion of "homomorphism" in the category of C^*-algebras—it is a map that preserves all the relevant operations on a C^*-algebra.

Theorem 7.4.8. *Let A and B be C^*-algebras, and suppose $\varphi : A \to B$ is a $*$-homomorphism. Then φ is bounded and $\|\varphi\| \leq 1$.*

It turns out that our two earlier examples of C^*-algebras are prototypical. First we remark that every C^*-algebra can be realized as an algebra of operators on some Hilbert space.

Theorem 7.4.9 (Gelfand-Naimark)**.** *If A is a C^*-algebra, then for some Hilbert space $\mathcal{H}$ there exists an injective $*$-homomorphism $\varphi : A \to B(\mathcal{H})$. Consequently, A is isomorphic to a self-adjoint subalgebra of $B(\mathcal{H})$.*

[2] The condition that $a^{**} = a$ is usually abbreviated by saying that $^* : A \to A$ is an **involution**.

For commutative C^*-algebras, we have a more refined characterization in terms of continuous function spaces.

Theorem 7.4.10 (Abstract Spectral Theorem)**.** *Suppose A is a commutative C^*-algebra. Then there is a locally compact Hausdorff space X such that A is $*$-isomorphic to $C_0(X)$.*

The Abstract Spectral Theorem actually has an interesting application in the realm of noncommutative C^*-algebras as well. Let A be a C^*-algebra, and assume for simplicity that A is *unital*, meaning that it possesses a multiplicative identity 1_A. For an element $a \in A$, we define the **spectrum** of a to be

$$\sigma(a) = \big\{\lambda \in \mathbf{C} \bigm| a - \lambda 1_A \text{ is not invertible}\big\}.$$

This should look an awful lot like the definition of an eigenvalue. Indeed, if $A \subseteq B(\mathcal{H})$ for some Hilbert space $\mathcal{H}$, then any eigenvalue of an operator $T \in A$ will belong to the spectrum $\sigma(T)$. However, if $\mathcal{H}$ is infinite-dimensional then the spectrum may contain complex numbers that are not eigenvalues of T. (In fact, T might not have any eigenvalues at all!)

Definition 7.4.11. Let A be a C^*-algebra. An element $a \in A$ is called **normal** if $a^*a = aa^*$.

Suppose $a \in A$ is normal. Then the C^*-algebra *generated* by a and the identity, which we write as $C^*(a, 1_A)$, is commutative. Consequently, we can apply the Abstract Spectral Theorem and identify $C^*(a, 1_A)$ with $C_0(X)$ for some locally compact Hausdorff space X. In this case, X actually turns out to be compact, so we really have $C^*(a, 1_A) \cong C(X)$. In fact, we can take X to be $\sigma(a)$!

Theorem 7.4.12 (Functional calculus)**.** *Suppose A is a unital C^*-algebra and $a \in A$ is normal. Then there is an isometric $*$-isomorphism*

$$\Phi : C(\sigma(a)) \to C^*(a, 1_A)$$

such that $\Phi(\mathrm{id}) = a$.

Bibliography

[AL17] Kassie Archer and Scott M. LaLonde. Allowed patterns of symmetric tent maps via commuter functions. *SIAM J. Discrete Math*, 31(1):317–334, 2017. arXiv:1606.01317 [math.DS].

[Bel] Jim Belk. Non-measurable sets. Math 461 course notes, Bard College. `http://math.bard.edu/belk/math461/`.

[DMRV06] O. Dovgoshey, O. Martio, V. Ryazanov, and M. Vuorinen. The Cantor function. *Expo. Math.*, 24:1–37, 2006.

[Fol99] Gerald B. Folland. *Real Analysis: Modern Techniques and Their Applications*. Wiley Interscience, second edition, 1999.

[Her06] Horst Herrlich. *Axiom of Choice*, volume 1876 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin Heidelberg, 2006.

[HS91] Norman B. Haaser and Joseph A. Sullivan. *Real Analysis*. Dover Books on Mathematics. Dover Publications, 1991.

[KF70] A. N. Kolmogorov and S. V. Fomin. *Introductory Real Analysis*. Dover Books on Mathematics. Dover Publications, 1970.

[Ped89] Gert K. Pedersen. *Analysis Now*, volume 118 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1989.

[RF10] H. L. Royden and P. M. Fitzpatrick. *Real Analysis*. Pearson, 4th edition, 2010.

[Rud76] Walter Rudin. *Principles of Mathematical Analysis*. International Series in Pure & Applied Mathematics. McGraw-Hill, 3rd edition, 1976.

[Wad03] William R. Wade. *An Introduction to Analysis*. Prentice Hall, third edition, 2003.

[Wil] Dana P. Williams. The Jordan-von Neumann theorem. `https://math.dartmouth.edu/archive/m113w10/public_html/jordan-vneumann-thm.pdf`.

[Wil07] Dana P. Williams. *Crossed Products of C^*-algebras.* Number 134 in Math. Surveys Monogr. American Mathematical Society, Providence, RI, 2007.

[WZ77] Richard L. Wheeden and Antoni Zygmund. *Measure and Integral: An Introduction to Real Analysis.* CRC Press, 1st edition, 1977.