

Spaces: An Introduction to Real Analysis

Tom L. Lindstrøm

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF OSLO, BOX 1053 BLINDERN, NO-0316 OSLO, NORWAY

E-mail address: `lindstro@math.uio.no`

2010 *Mathematics Subject Classification.* Primary 26-01, 28-01, 42-01, 46-01,
54E35, 26E15

Key words and phrases. Real analysis, metric spaces, function spaces, normed
spaces, measure theory, Fourier series

Contents

Preface	i
Introduction – Mainly to the Students	1
Chapter 1. Preliminaries: Proofs, Sets, and Functions	5
1.1. Proofs	5
1.2. Sets and boolean operations	8
1.3. Families of sets	11
1.4. Functions	13
1.5. Relations and partitions	17
1.6. Countability	20
Notes and references to Chapter 1	22
Chapter 2. The Foundation of Calculus	23
2.1. Epsilon-delta and all that	24
2.2. Completeness	29
2.3. Four important theorems	37
Notes and references to Chapter 2	42
Chapter 3. Metric Spaces	43
3.1. Definitions and examples	43
3.2. Convergence and continuity	48
3.3. Open and closed sets	52
3.4. Complete spaces	59
3.5. Compact sets	63
3.6. An alternative description of compactness	69
3.7. The completion of a metric space	71

Notes and references to Chapter 3	76
Chapter 4. Spaces of Continuous Functions	79
4.1. Modes of continuity	79
4.2. Modes of convergence	81
4.3. Integrating and differentiating sequences	86
4.4. Applications to power series	92
4.5. Spaces of bounded functions	98
4.6. Spaces of bounded, continuous functions	101
4.7. Applications to differential equations	103
4.8. Compact sets of continuous functions	107
4.9. Differential equations revisited	112
4.10. Polynomials are dense in the continuous function	116
4.11. The Stone-Weierstrass Theorem	123
Notes and references to Chapter 4	131
Bibliography	133
Index	135

Preface

While most calculus books are so similar that they seem to have been tested in the same wind tunnel, there is a lot more variety between books on real analysis, both with respect to content and level. If we start with levels, it is easy to distinguish at least three. The most elementary one is made up of books whose main purpose is to redo single-variable calculus in a more rigorous way – classical examples are Frank Morgan’s *Real Analysis*, Colin Clark’s *The Theoretical Side of Calculus*, and Stephen Abbott’s *Understanding Analysis*. On the intermediate level we have undergraduate texts like Walter Rudin’s *Principles of Mathematical Analysis*, Tom Körner’s *A Companion to Analysis*, and Kenneth R. Davidson and Allan P. Donsig’s *Real Analysis and Applications*, just to mention a few. In these texts, metric or normed spaces usually play a central part. On the third level we find graduate texts like H. L. Royden’s classic *Real Analysis* (now in a new edition by Patrick Fitzpatrick), Gerald B. Folland’s *Real Analysis: Modern Techniques and Their Applications*, and John N. McDonald and Neil A. Weiss: *A Course in Real Analysis*; books where measure theory is usually the point of departure. Above these again we have research level texts on different aspects of real analysis.

The present book is intended to be on the second level – it is written for students with a good background in (advanced) calculus and linear algebra but not more (although many students would undoubtedly benefit from a course on proofs and mathematical thinking). Books on this level are to a varying degree forward-looking or backward-looking, where backward-looking means reflecting on material in previous courses from a more advanced point of view, and forward-looking means providing the tools necessary for the next courses. The distinction is neatly summed up in the subtitle of Körner’s book: *A Second First or a First Second Course in Analysis*. While Körner aims to balance the two aspects, this book is unabashedly forward-looking – it is definitely intended as a first second course in analysis. For that reason I have dropped some of the staple ingredients of courses on this level

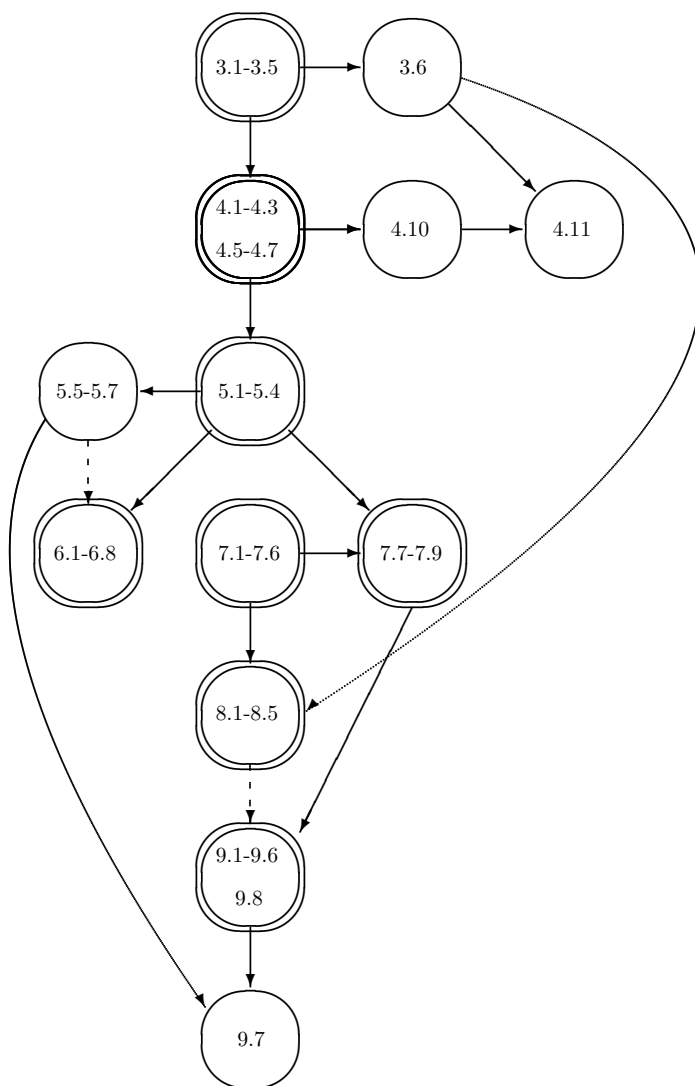
in favor of more advanced material; for example, I don't redo Riemann integration but go directly to Lebesgue integrals, and I do differentiation in normed spaces rather than refining differentiation in euclidean spaces. Although the exposition is still aimed at students on the second level, these choices bring in material that are usually taught on the third level, and I have tried to compensate by putting a lot of emphasis on examples and motivation, and by writing out arguments in greater detail than what is usually done in books on the third level. I have also included an introductory chapter on the foundation of calculus for students who have not had much previous exposure to the theoretical side of the subject.

The central concepts of the book are completeness, compactness, convergence, and continuity, and the students get to see them from many different perspectives – first in the context of metric spaces, then in normed spaces, and finally in measure theory and Fourier analysis. As the book is forward-looking, my primary aim has been to provide the students with the platform they need to understand applications, read more advanced texts, and follow more specialized courses. Although the book is definitely not a full-fledged course in functional analysis or measure theory, it does provide the students with many of the tools they need in more advanced courses, such as Banach's Fixed Point Theorem, the Arzelà-Ascoli Theorem, the Stone-Weierstrass Theorem, Baire's Category Theorem, the Open Mapping Theorem, the Inverse and Implicit Function Theorems, Lebesgue's Dominated Convergence Theorem, the Riesz-Fischer Theorem on the completeness of L^p , Carathéodory's Extension Theorem, Fubini's Theorem, the L^2 -convergence of Fourier series, Fejér's Theorem, and Dini's Test for pointwise convergence.

The main danger with a forward-looking course of this kind, is that it becomes all method and no content – that the only message to the students is: "Believe me, you will need this when you grow up!" This is definitely a danger also with the present text, but I have tried to include a selection of examples and applications (mainly to differential equations and Fourier analysis) that I hope will convince the students that all the theory is worthwhile.

Various versions of the text have been used for a fourth semester course at the University of Oslo, but I should warn you that I have never been able to cover all the material in the same semester – some years the main emphasis has been on measure theory (chapters ?? and ??) and other years on normed spaces (chapters ?? and ??). The chart below shows the major dependencies between the main chapters 3-??, but before we turn to it, it may be wise to say a few words about the introductory chapters 1 and 2. Chapter 1 is a short introduction to sets, functions, and relations from an abstract point of view. As most of our students don't have this background, I usually cover it during the first week of classes. The second chapter is meant as a service to students who lack a conceptual grasp of calculus, either because they have taken a more computational oriented calculus sequence, or because they haven't really understood the theoretical parts of their courses. I have never lectured on this chapter as our students are supposed to have the background needed to go directly to Chapter 3 on metric spaces, but my feeling is that many have found it useful for review and consolidation. One small point: I always have to pick up the material on \liminf and \limsup in Section 2.2 as it is not covered by our calculus courses.

Let us now turn to the chart showing the main logical dependencies between the various parts of the book. It is not as complicated as it may seem at first glance. The doubly ringed parts form the theoretical backbone of the book. This doesn't mean that the other parts are uninteresting (in fact, you will find deep and important theorems such as the Stone-Weierstrass Theorem and the Baire Category Theorem in these parts), but they are less important for the continuity.



The two dotted arrows indicate less important dependencies – chapter ?? only depends on Sections ??-?? through the Bounded Inverse Theorem in Section ??, and Chapter ?? only depends on Chapter ?? through Theorem ?? which states that the continuous functions are dense in $L^p([a, b], \mu)$. In my opinion, both these

results can be postulated. Note that some sections such as 3.7, 4.4, 4.8-4.9, and ??-?? don't appear in the chart at all. This just means that no later sections depend directly on them, and that I don't consider them part of the core of the book.

At the end of each chapter there is a brief section with a historical summary and suggestions for further reading. I have on purpose made the reading lists short as I feel that long lists are more intimidating than helpful at this level. You will probably find many of your favorite books missing (so are some of mine!), but I had to pick the ones I like and find appropriate for the level.

Acknowledgements. The main acknowledgments should probably go to all the authors I have read and all the lecturers I have listened to, but I have a feeling that the more important their influence is, the less I am aware of it – some of the core material “is just there”, and I have no recollection of learning it for the first time. During the writing of the book, I have looked up innumerable texts, some on real analysis and some on more specialized topics, but I hope I have managed to make all the material my own. An author often has to choose between different approaches, and in most cases I have chosen what to me seems intuitive and natural rather than sleek and elegant.

There are probably some things an author should not admit, but let me do it anyway. I never had any plans for a book on real analysis until the textbook for the course I was teaching in the Spring of 2011 failed to show up. I started writing notes in the hope that the books would be there in a week or two, but when the semester was over, the books still hadn't arrived, and I had almost 200 pages of class notes. Over the intervening years, I and others have taught from ever expanding versions of the notes, some years with an emphasis on measure theory, other years with an emphasis on functional analysis and differentiability.

I would like to thank everybody who has made constructive suggestions or pointed out mistakes and weaknesses during this period, in particular Snorre H. Christiansen, Geir Ellingsrud, Klara Hveberg, Erik Løv, Nils Henrik Risebro, Nikolai Bjørnestøl Hansen, Bernt Ivar Nødland, Simon Foldvik, Marius Jonsson (who also helped with the figure of vibrating strings in Chapter 9), Daniel Aubert, Lisa Eriksen, and Imran Ali. I would also like to extend my thanks to anonymous but very constructive referees who have helped improve the text in a number of ways, and to my enthusiastic editor Ina Mette and the helpful staff of the AMS.

If you find a misprint or a more serious mistake, please send a note to
`lindstro@math.uio.no`.

Oslo, April 19th, 2017

Tom Lindstrøm

Introduction – Mainly to the Students

This is a book on real analysis, and real analysis is a continuation of calculus. Many of the words you find in this book, such as “continuity”, “convergence”, “derivative”, and “integral”, are familiar from calculus, but they will now appear in new contexts and with new interpretations.

The biggest change from calculus to real analysis is a shift in emphasis from calculations to arguments. If you thumb through the book, you will find surprisingly few long calculations and fancy graphs, and instead a lot of technical words and unfamiliar symbols. This is what advanced mathematics books usually look like – calculations never lose their importance, but they become less dominant. However, this doesn’t mean that the book reads like a novel; as always, you have to read mathematics with pencil and paper at hand.

Your calculus courses have probably been divided into single variable and multi-variable calculus. Single variable calculus is about functions of one variable x while multivariable calculus is about functions of several variables x_1, x_2, \dots, x_n . Another way of looking at it, is to say that multivariable calculus is still about functions of one variable, but that this variable is now a vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$.

Real analysis covers single and multivariable calculus in one sweep and at the same time opens the door to even more general situations – functions of infinitely many variables! This is not as daunting as it may sound: Just as functions of several variables can be thought of as functions of a single, vector-valued variable, functions of infinitely many variables can often be thought of as functions of a single, functioned-valued variable (intuitively, a function is infinite dimensional as its graph consists of infinitely many points). Hence you should be prepared to deal with functions of the form $F(y)$ where y is a function.

As real analysis deals with functions of one, several, and infinitely many variables at the same time, it is necessarily rather abstract. It turns out that in order

to study such notions as convergence and continuity, we don't really need to specify what kinds of objects we are dealing with (numbers, vectors, functions, etc.) – all we need to know is that there is a reasonable way to measure the distance between them. This leads to the theory of *metric spaces* that will be the foundation for most of what we shall be doing in this book. If we also want to differentiate functions, we need the notion of a *normed space*, i.e. a metric space that is also a vector space (recall linear algebra). For integration, we shall invent another kind of space, called a *measure space*, which is tailored to measuring the size of sets. These spaces will again give rise to new kinds of normed spaces.

What I have just written probably doesn't make too much sense to you at this stage. What is a space, after all? Well, in mathematics a space is just a set (i.e. a collection of objects) with some additional structure that allows us to operate with the objects. In linear algebra, you have met vector spaces which are just collections of objects that can be added and multiplied by numbers in the same way that ordinary vectors can. The metric spaces that we shall study in this book, are just collections of objects equipped with a function that measures the distance between them in a reasonable manner. In the same way, the measure spaces we shall study toward the end of the book, consists of a set and a function that measures the size of (some of the) subsets of that set.

Spaces are abstractly defined by rules (often called axioms); anything that satisfies the rules is a space of the appropriate kind, and anything that does *not* satisfy the rules, is not a space of this kind. These abstract definitions give real analysis a different flavor from calculus – in calculus it often suffices to have an intuitive understanding of a concept; in real analysis you need to read the definitions carefully as they are all you have to go by. As the theory develops, we get more information about the spaces we study. This information is usually formulated as propositions or theorems, and you need to read these propositions and theorems carefully to see when they apply and what they mean.

Students often complain that there are too few examples in books on advanced mathematics. That is true in one sense and false in another. It's true in the sense that if you count the labelled examples in this book, there are far fewer of them than you are used to from calculus. However, there are lots of examples under a different label – and that is the label “proof”. Advanced mathematics is about arguments and proofs, and every proof is an example you can learn from. The aim of your mathematics education is to make you able of producing your own mathematical arguments, and the only practical way to learn how to make proofs is to read and understand proofs. Also, I should add, knowing mathematics is much more about knowing ways to argue than about knowing theorems and propositions.

So how does one read proofs? There are probably many ways, but the important thing is to try to understand the idea behind the proof, and how that idea can be turned into a logically valid argument. A trick that helped me as a student, was to read the proof one day, understand it as well as I could, and then return to it a day or two later to see if I could do it on my own without looking at the book. As I don't have a photographic memory, this technique forced me to concentrate on the ideas of the proof. If I had understood the main idea (which can usually be summed up in a sentence or a drawing once you have understood it), I could usually

reconstruct the rest of the proof without any problem. If I had not understood the main idea, I would be hopelessly lost.

Let us take a closer look at the contents of the book. The first two chapters contain preliminary material that is not really about real analysis as such. The first chapter gives a quick introduction to proofs, sets, and functions. If you have taken a course in mathematical reasoning or the foundations of mathematics, there is probably little new here, otherwise you should read it carefully. The second chapter reviews the theoretical foundation of calculus. How much you have to read here, depends on the calculus sequence you have taken. If it was fairly theoretical, this chapter may just be review; if it was mainly oriented toward calculations, it's probably a good idea to work carefully through most of this chapter. I'm sure your instructor will advise you on what to do.

The real contents of the book start with Chapter 3 on metric spaces. This is the theoretical foundation for the rest of the book, and it is important that you understand the basic ideas and become familiar with the concepts. Pay close attention to the arguments – they will reappear with small variations throughout the text. Chapter 4 is a continuation of Chapter 3 and focuses on spaces where the elements are continuous functions. This chapter is less abstract than Chapter 3 as it deals with objects that you are already familiar with (continuous functions, sequences, power series, differential equations), but some of the arguments are perhaps tougher as we have more structure to work with and try to tackle problems that are closer to “real life”.

In Chapter ?? we turn to normed spaces which are an amalgamation of metric spaces and the vector spaces you know from linear algebra. The big difference between this chapter and linear algebra is that we are now primarily interested in infinite dimensional spaces. The last two sections are quite theoretical, otherwise this is a rather friendly chapter. In Chapter ?? we use tools from Chapter ?? to study derivatives of functions between normed spaces in a way that generalizes many of the concepts you know from calculus (the chain rule, directional derivatives, partial derivatives, higher order derivatives, Taylor's formula). We also prove two important theorems on inverse and implicit functions that you may not have seen before.

Chapter ?? deals with integration and is a new start in two ways – both because most of the chapter is independent of the previous chapters, and also because it presents an approach to integration that is totally different from what you have seen in calculus. This new approach is based on the notion of measure, which is a very general way of assigning size to sets. Toward the end of the chapter, you will see how these measures lead to a new class of normed spaces with attractive properties. Chapter ?? is a continuation of Chapter ??. Here you will learn how to construct measures and see some important applications.

The final chapter is on Fourier analysis. It shows you an aspect of real analysis that has to some degree been neglected in the previous chapters – the power of concrete calculations. It also brings together techniques from most of the other

chapters in the book, and illustrates in a striking manner a phenomenon that appears again and again throughout the text: The convergence of a sequence or series of functions is a tricky business!

At the end of each chapter there is a short section with notes and references. Here you will find a brief historical summary and some suggestions for further reading. If you want to be a serious student of mathematics, I really recommend that you take a look at its history. Mathematics – and particularly the abstracts parts – is so much easier to appreciate when you know where it comes from. In fact, learning mathematics without knowing something of its history, is a bit like watching a horror movie with the sound turned off: You see that people get scared and take their precautions, but you don't understand why. This is particularly true of real analysis where much of the theory developed out of a need to deal with (what at the time felt like) counter-intuitive examples.

I hope you will enjoy the book. I know it's quite tough and requires hard work, but I have done my best to explain things as clearly as I can. Good Luck!

Preliminaries: Proofs, Sets, and Functions

Chapters with the word "preliminaries" in the title are never much fun, but they are useful – they provide the readers with the background information they need to enjoy the rest of the text. This chapter is no exception, but I have tried to keep it short and to the point; everything you find here will be needed at some stage, and most of the material will show up throughout the book.

Real analysis is a continuation of calculus, but it is more abstract and therefore in need of a larger vocabulary and more precisely defined concepts. You have undoubtedly dealt with proofs, sets, and functions in your previous mathematics courses, but probably in a rather casual fashion. Now they become the centerpiece of the theory, and there is no way to understand what is going on if you don't have a good grasp of them: The subject matter is so abstract that you can no longer rely on drawings and intuition; you simply have to be able to understand the concepts and to read, make and write proofs. Fortunately, this is not as difficult as it may sound if you have never tried to take proofs and formal definitions seriously before.

1.1. Proofs

There is nothing mysterious about mathematical proofs; they are just chains of logically irrefutable arguments that bring you from things you already know to whatever you want to prove. Still there are a few tricks of the trade that are useful to know about.

Many mathematical statements are of the form "If A, then B". This simply means that whenever statement A holds, statement B also holds, but not necessarily vice versa. A typical example is: "If $n \in \mathbb{N}$ is divisible by 14, then n is divisible by 7". This is a true statement since any natural number that is divisible by 14, is also divisible by 7. The opposite statement is not true as there are numbers that are divisible by 7, but not by 14 (e.g. 7 and 21).

Instead of “If A, then B”, we often say that “A implies B” and write $A \implies B$. As already observed, $A \implies B$ and $B \implies A$ mean two different things. If they are both true, A and B hold in exactly the same cases, and we say that A and B are *equivalent*. In words, we say “A if and only if B”, and in symbols, we write $A \iff B$. A typical example is:

“A triangle is equilateral if and only if all three angles are 60° ”

When we want to prove that $A \iff B$, it is often convenient to prove that $A \implies B$ and $B \implies A$ separately. Another method is to show that $A \implies B$ and that $\text{not-}A \implies \text{not-}B$ (why is this sufficient?).

If you think a little, you will realize that “ $A \implies B$ ” and “ $\text{not-}B \implies \text{not-}A$ ” mean exactly the same thing – they both say that whenever A happens, so does B. This means that instead of proving “ $A \implies B$ ”, we might just as well prove “ $\text{not-}B \implies \text{not-}A$ ”. This is called a *contrapositive proof*, and is convenient when the hypothesis $\text{not-}B$ gives us more to work with than the hypothesis A. Here is a typical example.

Proposition 1.1.1. *If n^2 is an even number, so is n .*

Proof. We prove the contrapositive statement: “If n is odd, so is n^2 ”: If n is odd, it can be written as $n = 2k + 1$ for a nonnegative integer k . But then

$$n^2 = (2k + 1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1$$

which is clearly odd. □

It should be clear why a contrapositive proof is best in this case: The hypothesis “ n is odd” is much easier to work with than the original hypothesis “ n^2 is even”.

A related method of proof is *proof by contradiction* or *reductio ad absurdum*. In these proofs, we assume the *opposite* of what we want to show, and prove that it leads to a contradiction. Hence our assumption must be false, and the original claim is established. Here is a well-known example.

Proposition 1.1.2. *$\sqrt{2}$ is an irrational number.*

Proof. We assume for contradiction that $\sqrt{2}$ is rational. This means that

$$\sqrt{2} = \frac{m}{n}$$

for natural numbers m and n . By canceling as much as possible, we may assume that m and n have no common factors.

If we square the equality above and multiply by n^2 on both sides, we get

$$2n^2 = m^2$$

This means that m^2 is even, and by the previous proposition, so is m . Hence $m = 2k$ for some natural number k , and if we substitute this into the last formula above and cancel a factor 2, we see that

$$n^2 = 2k^2$$

This means that n^2 is even, and by the previous proposition n is even. Thus we have proved that both m and n are even, which is impossible as we assumed that they have no common factors. The assumption that $\sqrt{2}$ is rational hence leads to a contradiction, and $\sqrt{2}$ must therefore be irrational. \square

Let me end this section by reminding you of a technique you have certainly seen before, *proof by induction*. We use this technique when we want to prove that a certain statement $P(n)$ holds for all natural numbers $n = 1, 2, 3, \dots$. A typical statement one may want to prove in this way, is

$$P(n) : 1 + 2 + 3 + \dots + n = \frac{n(n+1)}{2}$$

The basic observation behind the technique is:

1.1.3. Induction Principle: Assume that for each natural number $n = 1, 2, 3, \dots$ we have a statement $P(n)$ such that the following two conditions are satisfied:

- (i) $P(1)$ is true
- (ii) If $P(k)$ is true for a natural number k , then $P(k+1)$ is also true.

Then $P(n)$ holds for all natural numbers n .

Let us see how we can use the principle to prove that

$$P(n) : 1 + 2 + 3 + \dots + n = \frac{n(n+1)}{2}$$

holds for all natural numbers n .

First we check that the statement holds for $n = 1$: In this case the formula says

$$1 = \frac{1 \cdot (1+1)}{2}$$

which is obviously true. Assume now that $P(k)$ holds for some natural number k , i.e.

$$1 + 2 + 3 + \dots + k = \frac{k(k+1)}{2}$$

We then have

$$1 + 2 + 3 + \dots + k + (k+1) = \frac{k(k+1)}{2} + (k+1) = \frac{(k+1)(k+2)}{2}$$

which means that $P(k+1)$ is true. By the Induction Principle, $P(n)$ holds for all natural numbers n .

Remark: If you are still uncertain about what constitutes a proof, the best advice is to read proofs carefully and with understanding – you have to grasp *why* they force the conclusion. And then you have to start making your own proofs. The exercises in this book will give you plenty of opportunities!

Exercises for Section 1.1

1. Assume that the product of two integers x and y is even. Show that at least one of the numbers is even.
2. Assume that the sum of two integers x and y is even. Show that x and y are either both even or both odd.
3. Show that if n is a natural number such that n^2 is divisible by 3, then n is divisible by 3. Use this to show that $\sqrt{3}$ is irrational.
4. In this problem, we shall prove some basic properties of rational numbers. Recall that a real number r is *rational* if $r = \frac{a}{b}$ where a, b are integers and $b \neq 0$. A real number that is not rational, is called *irrational*.
 - a) Show that if r, s are rational numbers, so are $r + s$, $r - s$, rs , and (provided $s \neq 0$) $\frac{r}{s}$.
 - b) Assume that r is a rational number and a is an irrational number. Show that $r + a$ and $r - a$ are irrational. Show also that if $r \neq 0$, then ra , $\frac{r}{a}$, and $\frac{a}{r}$ are irrational.
 - c) Show by example that if a, b are irrational numbers, then $a + b$ and ab can be rational or irrational depending on a and b .

1.2. Sets and boolean operations

In the systematic development of mathematics, *set* is usually taken as the fundamental notion from which all other concepts are developed. We shall not be so ambitious, but just think naively of a set as a collection of mathematical objects. A set may be finite, such as the set

$$\{1, 2, 3, 4, 5, 6, 7, 8, 9\}$$

of all natural numbers less than 10, or infinite as the set $(0, 1)$ of all real numbers between 0 and 1.

We shall write $x \in A$ to say that x is an *element* of the set A , and $x \notin A$ to say that x is *not* an element of A . Two sets are *equal* if they have exactly the same elements, and we say that A is *subset* of B (and write $A \subseteq B$) if all elements of A are elements of B , but not necessarily vice versa. Note that there is no requirement that A is *strictly* included in B , and hence it is correct to write $A \subseteq B$ when $A = B$ (in fact, a standard technique for showing that $A = B$ is first to show that $A \subseteq B$ and then that $B \subseteq A$). By \emptyset we shall mean the *empty set*, i.e. the set with no elements (you may feel that a set with no elements is a contradiction in terms, but mathematical life would be much less convenient without the empty set).

Many common sets have a standard name and notation such as

$\mathbb{N} = \{1, 2, 3, \dots\}$, the set of natural numbers

$\mathbb{Z} = \{\dots - 3, -2, -1, 0, 1, 2, 3, \dots\}$, the set of all integers

\mathbb{Q} , the set of all rational numbers

\mathbb{R} , the set of all real numbers

\mathbb{C} , the set of all complex numbers

\mathbb{R}^n , the set of all real n -tuples

To specify other sets, we shall often use expressions of the kind

$$A = \{a \mid P(a)\}$$

which means the set of all objects satisfying condition P . Often it is more convenient to write

$$A = \{a \in B \mid P(a)\}$$

which means the set of all elements in B satisfying the condition P . Examples of this notation are

$$[-1, 1] = \{x \in \mathbb{R} \mid -1 \leq x \leq 1\}$$

and

$$A = \{2n - 1 \mid n \in \mathbb{N}\}$$

where A is the set of all odd numbers. To increase readability, I shall occasionally replace the vertical bar \mid by a colon $:$ and write $A = \{a : P(a)\}$ and $A = \{a \in B : P(a)\}$ instead of $A = \{a \mid P(a)\}$ and $A = \{a \in B \mid P(a)\}$, e.g. in expressions like $\{\|\alpha \mathbf{x}\| : |\alpha| < 1\}$ where there are lots of vertical bars already.

If A_1, A_2, \dots, A_n are sets, their *union* and *intersection* are given by

$$A_1 \cup A_2 \cup \dots \cup A_n = \{a \mid a \text{ belongs to at least one of the sets } A_1, A_2, \dots, A_n\}$$

and

$$A_1 \cap A_2 \cap \dots \cap A_n = \{a \mid a \text{ belongs to all the sets } A_1, A_2, \dots, A_n\},$$

respectively. Unions and intersections are often called *boolean operations* after the English logician George Boole (1815-1864). Two sets are called *disjoint* if they do not have elements in common, i.e. if $A \cap B = \emptyset$.

When we calculate with numbers, the *distributive law* tells us how to move common factors in and out of parentheses:

$$b(a_1 + a_2 + \dots + a_n) = ba_1 + ba_2 + \dots + ba_n$$

Unions and intersections are distributive both ways, i.e. we have:

Proposition 1.2.1 (Distributive laws). *For all sets B, A_1, A_2, \dots, A_n*

$$(1.2.1) \quad B \cap (A_1 \cup A_2 \cup \dots \cup A_n) = (B \cap A_1) \cup (B \cap A_2) \cup \dots \cup (B \cap A_n)$$

and

$$(1.2.2) \quad B \cup (A_1 \cap A_2 \cap \dots \cap A_n) = (B \cup A_1) \cap (B \cup A_2) \cap \dots \cap (B \cup A_n)$$

Proof. I'll prove the first formula and leave the second as an exercise. The proof is in two steps: first we prove that the set on the left is a subset of the one on the right, and then we prove that the set on the right is a subset of the one on the left.

Assume first that x is an element of the set on the left, i.e. $x \in B \cap (A_1 \cup A_2 \cup \dots \cup A_n)$. Then x must be in B and at least one of the sets A_i . But then $x \in B \cap A_i$, and hence $x \in (B \cap A_1) \cup (B \cap A_2) \cup \dots \cup (B \cap A_n)$. This proves that

$$B \cap (A_1 \cup A_2 \cup \dots \cup A_n) \subseteq (B \cap A_1) \cup (B \cap A_2) \cup \dots \cup (B \cap A_n)$$

To prove the opposite inclusion, assume that $x \in (B \cap A_1) \cup (B \cap A_2) \cup \dots \cup (B \cap A_n)$. Then $x \in B \cap A_i$ for at least one i , and hence $x \in B$ and $x \in A_i$. But if $x \in A_i$ for some i , then $x \in A_1 \cup A_2 \cup \dots \cup A_n$, and hence $x \in B \cap (A_1 \cup A_2 \cup \dots \cup A_n)$. This proves that

$$B \cap (A_1 \cup A_2 \cup \dots \cup A_n) \supseteq (B \cap A_1) \cup (B \cap A_2) \cup \dots \cup (B \cap A_n)$$

As we now have inclusion in both directions, formula (1.2.1) follows. \square

Remark: It is possible to prove formula (1.2.1) in one sweep by noticing that all steps in the argument are equivalences and not only implications, but most people are more prone to making mistakes when they work with chains of equivalences than with chains of implications.

There are also other algebraic rules for unions and intersections, but most of them are so obvious that we do not need to state them here (an exception is De Morgan's laws which we shall return to in a moment).

The *set theoretic difference* $A \setminus B$ (also written $A - B$) is defined by

$$A \setminus B = \{a \mid a \in A, a \notin B\}$$

In many situations we are only interested in subsets of a given set U (often referred to as the *universe*). The *complement* A^c of a set A with respect to U is defined by

$$A^c = U \setminus A = \{a \in U \mid a \notin A\}$$

We can now formulate *De Morgan's laws*:

Proposition 1.2.2 (De Morgan's laws). *Assume that A_1, A_2, \dots, A_n are subsets of a universe U . Then*

$$(1.2.3) \quad (A_1 \cup A_2 \cup \dots \cup A_n)^c = A_1^c \cap A_2^c \cap \dots \cap A_n^c$$

and

$$(1.2.4) \quad (A_1 \cap A_2 \cap \dots \cap A_n)^c = A_1^c \cup A_2^c \cup \dots \cup A_n^c$$

(These rules are easy to remember if you observe that you can distribute the c outside the parentheses on the individual sets provided you turn all \cup 's into \cap 's and all \cap 's into \cup 's).

Proof. Again I'll prove the first part and leave the second as an exercise. The strategy is as indicated above; we first show that any element of the set on the left must also be an element of the set on the right, and then vice versa.

Assume that $x \in (A_1 \cup A_2 \cup \dots \cup A_n)^c$. Then $x \notin A_1 \cup A_2 \cup \dots \cup A_n$, and hence for all i , $x \notin A_i$. This means that for all i , $x \in A_i^c$, and hence $x \in A_1^c \cap A_2^c \cap \dots \cap A_n^c$.

Assume next that $x \in A_1^c \cap A_2^c \cap \dots \cap A_n^c$. This means that $x \in A_i^c$ for all i , in other words: for all i , $x \notin A_i$. Thus $x \notin A_1 \cup A_2 \cup \dots \cup A_n$ which means that $x \in (A_1 \cup A_2 \cup \dots \cup A_n)^c$. \square

We end this section with a brief look at cartesian products. If we have two sets, A and B , the *cartesian product* $A \times B$ consists of all ordered pairs (a, b) where $a \in A$ and $b \in B$. If we have more sets A_1, A_2, \dots, A_n , the cartesian product $A_1 \times A_2 \times \dots \times A_n$ consists of all n -tuples (a_1, a_2, \dots, a_n) where $a_1 \in A_1, a_2 \in A_2$ etc.

$A_2, \dots, a_n \in A_n$. If all the sets are the same (i.e. $A_i = A$ for all i), we usually write A^n instead of $A \times A \times \dots \times A$. Hence \mathbb{R}^n is the set of all n -tuples of real numbers, just as you are used to, and \mathbb{C}^n is the set of all n -tuples of complex numbers.

Exercises for Section 1.2

1. Show that $[0, 2] \cup [1, 3] = [0, 3]$ and that $[0, 2] \cap [1, 3] = [1, 2]$
2. Let $U = \mathbb{R}$ be the universe. Explain that $(-\infty, 0)^c = [0, \infty)$
3. Show that $A \setminus B = A \cap B^c$.
4. The *symmetric difference* $A \triangle B$ of two sets A, B consists of the elements that belong to *exactly one* of the sets A, B . Show that

$$A \triangle B = (A \setminus B) \cup (B \setminus A)$$

5. Prove formula (1.2.2).
6. Prove formula (1.2.4).
7. Prove that if U is the universe, then $A_1 \cup A_2 \cup \dots \cup A_n = U$ if and only if $A_1^c \cap A_2^c \cap \dots \cap A_n^c = \emptyset$.
8. In this exercise, all sets are subsets of a universe U . Use the distributive laws and De Morgan's laws to show that:
 - a) $(A^c \cup B)^c = A \setminus B$.
 - b) $A \cap (B^c \cap A)^c = A \cap B$.
 - c) $A^c \cap (B \cup C) = (B \setminus A) \cup (C \setminus A)$
9. Prove that $(A \cup B) \times C = (A \times C) \cup (B \times C)$ and $(A \cap B) \times C = (A \times C) \cap (B \times C)$.

1.3. Families of sets

A collection of sets is usually called a *family*. An example is the family

$$\mathcal{A} = \{[a, b] \mid a, b \in \mathbb{R}\}$$

of all closed and bounded intervals on the real line. Families may seem abstract, but you have to get used to them as they appear in all parts of higher mathematics. We can extend the notions of union and intersection to families in the following way: If \mathcal{A} is a family of sets, we define

$$\bigcup_{A \in \mathcal{A}} A = \{a \mid a \text{ belongs to at least one set } A \in \mathcal{A}\}$$

and

$$\bigcap_{A \in \mathcal{A}} A = \{a \mid a \text{ belongs to all sets } A \in \mathcal{A}\}$$

The distributive laws extend to this case in the obvious way, i.e.,

$$B \cap \left(\bigcup_{A \in \mathcal{A}} A \right) = \bigcup_{A \in \mathcal{A}} (B \cap A) \quad \text{and} \quad B \cup \left(\bigcap_{A \in \mathcal{A}} A \right) = \bigcap_{A \in \mathcal{A}} (B \cup A)$$

and so do the laws of De Morgan:

$$\left(\bigcup_{A \in \mathcal{A}} A \right)^c = \bigcap_{A \in \mathcal{A}} A^c \quad \text{and} \quad \left(\bigcap_{A \in \mathcal{A}} A \right)^c = \bigcup_{A \in \mathcal{A}} A^c$$

Families are often given as *indexed sets*. This means we have a basic set I , and that the family consists of one set A_i for each element i in I . We then write the family as

$$\mathcal{A} = \{A_i \mid i \in I\} \quad \text{or} \quad \mathcal{A} = \{A_i\}_{i \in I},$$

and use notation such as

$$\bigcup_{i \in I} A_i \quad \text{and} \quad \bigcap_{i \in I} A_i$$

or alternatively

$$\bigcup \{A_i : i \in I\} \quad \text{and} \quad \bigcap \{A_i : i \in I\}$$

for unions and intersections

A rather typical example of an indexed set is $\mathcal{A} = \{B_r \mid r \in [0, \infty)\}$ where $B_r = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = r^2\}$. This is the family of all circles in the plane with center at the origin.

Exercises for Section 1.3

1. Show that $\bigcup_{n \in \mathbb{N}} [-n, n] = \mathbb{R}$
2. Show that $\bigcap_{n \in \mathbb{N}} (-\frac{1}{n}, \frac{1}{n}) = \{0\}$.
3. Show that $\bigcup_{n \in \mathbb{N}} [\frac{1}{n}, 1] = (0, 1]$
4. Show that $\bigcap_{n \in \mathbb{N}} (0, \frac{1}{n}] = \emptyset$
5. Prove the distributive laws for families. i.e.,

$$B \cap \left(\bigcup_{A \in \mathcal{A}} A \right) = \bigcup_{A \in \mathcal{A}} (B \cap A) \quad \text{and} \quad B \cup \left(\bigcap_{A \in \mathcal{A}} A \right) = \bigcap_{A \in \mathcal{A}} (B \cup A)$$

6. Prove De Morgan's laws for families:

$$\left(\bigcup_{A \in \mathcal{A}} A \right)^c = \bigcap_{A \in \mathcal{A}} A^c \quad \text{and} \quad \left(\bigcap_{A \in \mathcal{A}} A \right)^c = \bigcup_{A \in \mathcal{A}} A^c$$

7. Later in the book we shall often study families of sets with given properties, and it may be worthwhile to take a look at an example here. If X is a nonempty set and \mathcal{A} is a family of subsets of X , we call \mathcal{A} an *algebra of sets* if the following three properties are satisfied:

- (i) $\emptyset \in \mathcal{A}$.
- (ii) If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$ (all complements are with respect to the universe X ; hence $A^c = X \setminus A$).
- (iii) If $A, B \in \mathcal{A}$, then $A \cup B \in \mathcal{A}$.

In the rest of the problem, we assume that \mathcal{A} is an algebra of sets on X .

- a) Show that $X \in \mathcal{A}$.
- b) Show that if $A_1, A_2, \dots, A_n \in \mathcal{A}$ for an $n \in \mathbb{N}$, then

$$A_1 \cup A_2 \cup \dots \cup A_n \in \mathcal{A}$$

(Hint: Use induction.)

- c) Show that if $A_1, A_2, \dots, A_n \in \mathcal{A}$ for an $n \in \mathbb{N}$, then

$$A_1 \cap A_2 \cap \dots \cap A_n \in \mathcal{A}$$

(Hint: Use b), property (ii), and one of De Morgan's laws.)

1.4. Functions

Functions can be defined in terms of sets, but for our purposes it suffices to think of a function $f: X \rightarrow Y$ from a set X to a set Y as an *assignment* which to each element $x \in X$ assigns an element $y = f(x)$ in Y .¹ A function is also called a *map* or a *mapping*. Formally, functions and maps are exactly the same thing, but people tend to use the word “map” when they are thinking geometrically, and the word “function” when they are thinking more in terms of formulas and calculations. If we have a formula or an expression $H(x)$, it is sometimes convenient to write $x \mapsto H(x)$ for the function it defines.




Figure 1.4.1. A function f from X to Y

When we are dealing with functions between general sets, there is usually no sense in trying to picture them as graphs in a coordinate system. Instead, we shall picture them as shown in Figure 1.4.1 where the function f maps the point x in X to the point $y = f(x)$ in Y .




Figure 1.4.2. Composition of functions

If we have three sets X, Y, Z and functions $f: X \rightarrow Y$ and $g: Y \rightarrow Z$, we can define a *composite function* $h: X \rightarrow Z$ by $h(x) = g(f(x))$ (see Figure 1.4.2). This composite function is often denoted by $g \circ f$, and hence $g \circ f(x) = g(f(x))$. You may recall composite functions from the chain rule in calculus.

If A is subset of X , the set $f(A) \subseteq Y$ defined by

$$f(A) = \{f(a) \mid a \in A\}$$

¹Set-theoretically, a function from X to Y is a subset f of $X \times Y$ such that for each $x \in X$, there is exactly one $y \in Y$ such that $(x, y) \in f$. For $x \in X$, we then define $f(x)$ to be the unique element $y \in Y$ such that $(x, y) \in f$, and we are back to our usual notation.

is called the *image of A under f* . Figure 1.4.3 shows how f maps $A \subseteq X$ into $f(A) \subseteq Y$.




Figure 1.4.3. The image $f(A)$ of $A \subseteq X$

If B is subset of Y , the set $f^{-1}(B) \subseteq X$ defined by

$$f^{-1}(B) = \{x \in X \mid f(x) \in B\}$$

is called the *inverse image of B under f* . Figure 1.4.4 shows $f^{-1}(B)$ as the set of all elements in X that are being mapped into $B \subseteq Y$ by f .




Figure 1.4.4. The inverse image $f^{-1}(B)$ of $B \subseteq Y$

In analysis, images and inverse images of sets play important parts, and it is useful to know how these operations relate to the boolean operations of union and intersection. Let us begin with the good news.

Proposition 1.4.1. *Let \mathcal{B} be a family of subset of Y . Then for all functions $f: X \rightarrow Y$ we have*

$$f^{-1}\left(\bigcup_{B \in \mathcal{B}} B\right) = \bigcup_{B \in \mathcal{B}} f^{-1}(B) \quad \text{and} \quad f^{-1}\left(\bigcap_{B \in \mathcal{B}} B\right) = \bigcap_{B \in \mathcal{B}} f^{-1}(B)$$

We say that inverse images commute with arbitrary unions and intersections.

Proof. I prove the first part; the second part is proved similarly. Assume first that $x \in f^{-1}(\bigcup_{B \in \mathcal{B}} B)$. This means that $f(x) \in \bigcup_{B \in \mathcal{B}} B$, and consequently there must

be at least one $B' \in \mathcal{B}$ such that $f(x) \in B'$. But then $x \in f^{-1}(B')$, and hence $x \in \bigcup_{B \in \mathcal{B}} f^{-1}(B)$. This proves that $f^{-1}(\bigcup_{B \in \mathcal{B}} B) \subseteq \bigcup_{B \in \mathcal{B}} f^{-1}(B)$.

To prove the opposite inclusion, assume that $x \in \bigcup_{B \in \mathcal{B}} f^{-1}(B)$. There must be at least one $B' \in \mathcal{B}$ such that $x \in f^{-1}(B')$, and hence $f(x) \in B'$. This implies that $f(x) \in \bigcup_{B \in \mathcal{B}} B$, and hence $x \in f^{-1}(\bigcup_{B \in \mathcal{B}} B)$. \square

For forward images the situation is more complicated:

Proposition 1.4.2. *Let \mathcal{A} be a family of subset of X . Then for all functions $f: X \rightarrow Y$ we have*

$$f\left(\bigcup_{A \in \mathcal{A}} A\right) = \bigcup_{A \in \mathcal{A}} f(A) \quad \text{and} \quad f\left(\bigcap_{A \in \mathcal{A}} A\right) \subseteq \bigcap_{A \in \mathcal{A}} f(A)$$

In general, we do not have equality in the latter case. Hence forward images commute with unions, but not always with intersections.

Proof. To prove the statement about unions, we first observe that since $A \subseteq \bigcup_{A \in \mathcal{A}} A$ for all $A \in \mathcal{A}$, we have $f(A) \subseteq f(\bigcup_{A \in \mathcal{A}} A)$ for all such A . Since this inclusion holds for all A , we must also have $\bigcup_{A \in \mathcal{A}} f(A) \subseteq f(\bigcup_{A \in \mathcal{A}} A)$. To prove the opposite inclusion, assume that $y \in f(\bigcup_{A \in \mathcal{A}} A)$. This means that there exists an $x \in \bigcup_{A \in \mathcal{A}} A$ such that $f(x) = y$. This x has to belong to at least one $A' \in \mathcal{A}$, and hence $y \in f(A') \subseteq \bigcup_{A \in \mathcal{A}} f(A)$.

To prove the inclusion for intersections, just observe that since $\bigcap_{A \in \mathcal{A}} A \subseteq A$ for all $A \in \mathcal{A}$, we must have $f(\bigcap_{A \in \mathcal{A}} A) \subseteq f(A)$ for all such A . Since this inclusion holds for all A , it follows that $f(\bigcap_{A \in \mathcal{A}} A) \subseteq \bigcap_{A \in \mathcal{A}} f(A)$. The example below shows that the opposite inclusion does not always hold. \square

Example 1: Let $X = \{x_1, x_2\}$ and $Y = \{y\}$. Define $f: X \rightarrow Y$ by $f(x_1) = f(x_2) = y$, and let $A_1 = \{x_1\}, A_2 = \{x_2\}$. Then $A_1 \cap A_2 = \emptyset$ and consequently $f(A_1 \cap A_2) = \emptyset$. On the other hand $f(A_1) = f(A_2) = \{y\}$, and hence $f(A_1) \cap f(A_2) = \{y\}$. This means that $f(A_1 \cap A_2) \neq f(A_1) \cap f(A_2)$. \clubsuit

The problem in this example stems from the fact that y belongs to both $f(A_1)$ and $f(A_2)$, but only as the image of two *different* elements $x_1 \in A_1$ and $x_2 \in A_2$; there is no *common* element $x \in A_1 \cap A_2$ which is mapped to y . To see how it's sometimes possible to avoid this problem, define a function $f: X \rightarrow Y$ to be *injective* if $f(x_1) \neq f(x_2)$ whenever $x_1 \neq x_2$.

Corollary 1.4.3. *Let \mathcal{A} be a family of subset of X . Then for all injective functions $f: X \rightarrow Y$ we have*

$$f\left(\bigcap_{A \in \mathcal{A}} A\right) = \bigcap_{A \in \mathcal{A}} f(A)$$

Proof. To prove the missing inclusion $f(\bigcap_{A \in \mathcal{A}} A) \supseteq \bigcap_{A \in \mathcal{A}} f(A)$, assume that $y \in \bigcap_{A \in \mathcal{A}} f(A)$. For each $A \in \mathcal{A}$ there must be an element $x_A \in A$ such that $f(x_A) = y$. Since f is injective, all these $x_A \in A$ must be the same element x , and hence $x \in A$ for all $A \in \mathcal{A}$. This means that $x \in \bigcap_{A \in \mathcal{A}} A$, and since $y = f(x)$, we have proved that $y \in f(\bigcap_{A \in \mathcal{A}} A)$. \square

Taking complements is another operation that commutes with inverse images, but not (in general) with forward images.

Proposition 1.4.4. *Assume that $f: X \rightarrow Y$ is a function and that $B \subseteq Y$. Then $f^{-1}(B^c) = (f^{-1}(B))^c$. (Here, of course, $B^c = Y \setminus B$ is the complement with respect to the universe Y , while $(f^{-1}(B))^c = X \setminus f^{-1}(B)$ is the complement with respect to the universe X).*

Proof. An element $x \in X$ belongs to $f^{-1}(B^c)$ if and only if $f(x) \in B^c$. On the other hand, it belongs to $(f^{-1}(B))^c$ if and only if $f(x) \notin B$, i.e. if and only if $f(x) \in B^c$. \square

We also observe that being disjoint is a property that is conserved under inverse images; if $A \cap B = \emptyset$, then $f^{-1}(A) \cap f^{-1}(B) = \emptyset$. Again the corresponding property for forward images fails in general.

We end this section by taking a look at three important properties a function can have. We have already defined a function $f: X \rightarrow Y$ to be *injective* (or *one-to-one*) if $f(x_1) \neq f(x_2)$ whenever $x_1 \neq x_2$. It is called *surjective* (or *onto*) if for all $y \in Y$, there is an $x \in X$ such that $f(x) = y$, and it is called *bijective* (or a *one-to-one correspondence*) if it is both injective and surjective. Injective, surjective, and bijective functions are also referred to as *injections*, *surjections*, and *bijections*, respectively.

If $f: X \rightarrow Y$ is bijective, there is for each $y \in Y$ exactly one $x \in X$ such that $f(x) = y$. Hence we can define a function $g: Y \rightarrow X$ by

$$g(y) = x \quad \text{if and only if} \quad f(x) = y$$

This function g is called the *inverse function* of f and is often denoted by f^{-1} . Note that the inverse function g is necessarily a bijection, and that $g^{-1} = f$.

Remark: Note that the *inverse function* f^{-1} is only defined when the function f is bijective, but that the *inverse images* $f^{-1}(B)$ that we studied earlier in this section, are defined for all functions f .

The following observation is often useful.

Proposition 1.4.5. *If $f: X \rightarrow Y$ and $g: Y \rightarrow Z$ are bijective, so is their composition $g \circ f$, and $(g \circ f)^{-1} = (f^{-1}) \circ (g^{-1})$.*

Proof. Left to the reader (see Exercise 8 below). \square

Exercises for Section 1.4

1. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be the function $f(x) = x^2$. Find $f([-1, 2])$ and $f^{-1}([-1, 2])$.
2. Let $g: \mathbb{R}^2 \rightarrow \mathbb{R}$ be the function $g(x, y) = x^2 + y^2$. Find $g([-1, 1] \times [-1, 1])$ and $g^{-1}([0, 4])$.
3. Show that the function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x^2$ is neither injective nor surjective. What if we change the definition to $f(x) = x^3$?

4. Show that a strictly increasing function $f: \mathbb{R} \rightarrow \mathbb{R}$ is injective. Does it have to be surjective?
5. Prove the second part of Proposition 1.4.1.
6. Find a function $f: X \rightarrow Y$ and a set $A \subseteq X$ such that we have neither $f(A^c) \subseteq f(A)^c$ nor $f(A)^c \subseteq f(A^c)$.
7. Let X, Y be two nonempty sets and consider a function $f: X \rightarrow Y$.
 - a) Show that if $B \subseteq Y$, then $f(f^{-1}(B)) = B$.
 - b) Show that if $A \subseteq X$, then $f^{-1}(f(A)) \supseteq A$. Find an example where $f^{-1}(f(A)) \neq A$.
8. In this problem f, g are functions $f: X \rightarrow Y$ and $g: Y \rightarrow Z$.
 - a) Show that if f and g are injective, so is $g \circ f$.
 - b) Show that if f and g are surjective, so is $g \circ f$.
 - c) Explain that if f and g are bijective, so is $g \circ f$, and show that $(g \circ f)^{-1} = (f^{-1}) \circ (g^{-1})$.
9. Given a set Z , we let $\text{id}_Z: Z \rightarrow Z$ be the *identity map* $\text{id}_Z(z) = z$ for all $z \in Z$.
 - a) Show that if $f: X \rightarrow Y$ is bijective with inverse function $g: Y \rightarrow X$, then $g \circ f = \text{id}_X$ and $f \circ g = \text{id}_Y$.
 - b) Assume that $f: X \rightarrow Y$ and $g: Y \rightarrow X$ are two functions such that $g \circ f = \text{id}_X$ and $f \circ g = \text{id}_Y$. Show that f and g are bijective, and that $g = f^{-1}$.
10. As pointed out in the remark above, we are using the symbol f^{-1} in two slightly different ways. It may refer to the inverse of a bijective function $f: X \rightarrow Y$, but it may also be used to denote inverse images $f^{-1}(B)$ of sets under arbitrary functions $f: X \rightarrow Y$. The only instances where this might have caused real confusion, is when $f: X \rightarrow Y$ is a bijection and we write $C = f^{-1}(B)$ for a subset B of Y . This can then be interpreted as: a) C is the inverse image of B under f and b) C is the (direct) image of B under f^{-1} . Show that these two interpretations of C coincide.

1.5. Relations and partitions

In mathematics there are lots of relations between objects; numbers may be smaller or larger than each other, lines may be parallel, vectors may be orthogonal, matrices may be similar and so on. Sometimes it is convenient to have an abstract definition of what we mean by a relation.

Definition 1.5.1. *By a relation on a set X , we mean a subset R of the cartesian product $X \times X$. We usually write xRy instead of $(x, y) \in R$ to denote that x and y are related. The symbols \sim and \equiv are often used to denote relations, and we then write $x \sim y$ and $x \equiv y$.*

At first glance this definition may seem strange as very few people think of relations as subsets of $X \times X$, but a little thought will convince you that it gives us a convenient starting point, especially if I add that in practice relations are rarely arbitrary subsets of $X \times X$, but have much more structure than the definition indicates.

Example 1. Equality $=$ and “less than” $<$ are relations on \mathbb{R} . To see that they fit into the formal definition above, note that they can be defined as

$$R = \{(x, y) \in \mathbb{R}^2 \mid x = y\}$$

for equality and

$$S = \{(x, y) \in \mathbb{R}^2 \mid x < y\}$$

for “less than”.



We shall take a look at an important class of relations, the *equivalence relations*. Equivalence relations are used to partition sets into subsets, and from a pedagogical point of view, it is probably better to start with the related notion of a partition.

Informally, a partition is what we get if we divide a set into non-overlapping pieces. More precisely, if X is a set, a *partition* \mathcal{P} of X is a family of nonempty subset of X such that each element in x belongs to exactly one set $P \in \mathcal{P}$. The elements P of \mathcal{P} are called *partition classes* of \mathcal{P} .

Given a partition of X , we may introduce a relation \sim on X by

$$x \sim y \iff x \text{ and } y \text{ belong to the same set } P \in \mathcal{P}$$

It is easy to check that \sim has the following three properties:

- (i) $x \sim x$ for all $x \in X$.
- (ii) If $x \sim y$, then $y \sim x$.
- (iii) If $x \sim y$ and $y \sim z$, then $x \sim z$.

We say that \sim is the relation *induced by* the partition \mathcal{P} .

Let us now turn the tables around and start with a relation on X satisfying conditions (i)-(iii):

Definition 1.5.2. An equivalence relation on X is a relation \sim satisfying the following conditions:

- (i) Reflexivity: $x \sim x$ for all $x \in X$,
- (ii) Symmetry: If $x \sim y$, then $y \sim x$.
- (iii) Transitivity: If $x \sim y$ and $y \sim z$, then $x \sim z$.

Given an equivalence relation \sim on X , we may for each $x \in X$ define the *equivalence class* (also called the *partition class*) $[x]$ of x by:

$$[x] = \{y \in X \mid x \sim y\}$$

The following result tells us that there is a one-to-one correspondence between partitions and equivalence relations – just as all partitions induce an equivalence relation, all equivalence relations define a partition.

Proposition 1.5.3. If \sim is an equivalence relation on X , the collection of equivalence classes

$$\mathcal{P} = \{[x] : x \in X\}$$

is a partition of X .

Proof. We must prove that each x in X belongs to exactly one equivalence class. We first observe that since $x \sim x$ by (i), $x \in [x]$ and hence belongs to at least one equivalence class. To finish the proof, we have to show that if $x \in [y]$ for some other element $y \in X$, then $[x] = [y]$.

We first prove that $[y] \subseteq [x]$. To this end assume that $z \in [y]$. By definition, this means that $y \sim z$. On the other hand, the assumption that $x \in [y]$ means that $y \sim x$, which by (ii) implies that $x \sim y$. We thus have $x \sim y$ and $y \sim z$, which by (iii) means that $x \sim z$. Thus $z \in [x]$, and hence we have proved that $[y] \subseteq [x]$.

The opposite inclusion $[x] \subseteq [y]$ is proved similarly: Assume that $z \in [x]$. By definition, this means that $x \sim z$. On the other hand, the assumption that $x \in [y]$ means that $y \sim x$. We thus have $y \sim x$ and $x \sim z$, which by (iii) implies that $y \sim z$. Thus $z \in [y]$, and we have proved that $[x] \subseteq [y]$. \square

The main reason why this theorem is useful is that it is often more natural to describe situations through equivalence relations than through partitions. The following example assumes that you remember a little linear algebra:

Example 2: Let V be a vector space and U a subspace. Define a relation on V by

$$x \sim y \iff y - x \in U$$

Let us show that \sim is an equivalence relation by checking the three conditions (i)-(iii) in the definition:

- (i) *Reflexive:* Since $x - x = 0 \in U$, we see that $x \sim x$ for all $x \in V$.
- (ii) *Symmetric:* Assume that $x \sim y$. This means that $y - x \in U$, and consequently $x - y = (-1)(y - x) \in U$ as subspaces are closed under multiplication by scalars. Hence $y \sim x$.
- (iii) *Transitive:* If $x \sim y$ and $y \sim z$, then $y - x \in U$ and $z - y \in U$. Since subspaces are closed under addition, this means that $z - x = (z - y) + (y - x) \in U$, and hence $x \sim z$.

As we have now proved that \sim is an equivalence relation, the equivalence classes of \sim form a partition of V . The equivalence class of an element x is

$$[x] = \{x + u \mid u \in U\}$$

(check that this really is the case!)



If \sim is an equivalence relation on X , we let X/\sim denote the set of all equivalence classes of \sim . Such *quotient constructions* are common in all parts of mathematics, and you will see a few examples later in the book.

Exercises to Section 1.5

1. Let \mathcal{P} be a partition of a set A , and define a relation \sim on A by

$$x \sim y \iff x \text{ and } y \text{ belong to the same set } P \in \mathcal{P}$$

Check that \sim really is an equivalence relation.

2. Assume that \mathcal{P} is the partition defined by an equivalence relation \sim . Show that \sim is the equivalence relation induced by \mathcal{P} .
3. Let \mathcal{L} be the collection of all lines in the plane. Define a relation on \mathcal{L} by saying that two lines are equivalent if and only if they are parallel or equal. Show that this is an equivalence relation on \mathcal{L} .

4. Define a relation on \mathbb{C} by

$$z \sim y \iff |z| = |y|$$

Show that \sim is an equivalence relation. What do the equivalence classes look like?

5. Define a relation \sim on \mathbb{R}^3 by

$$(x, y, z) \sim (x', y', z') \iff 3x - y + 2z = 3x' - y' + 2z'$$

Show that \sim is an equivalence relation and describe the equivalence classes of \sim .

6. Let m be a natural number. Define a relation \equiv on \mathbb{Z} by

$$x \equiv y \iff x - y \text{ is divisible by } m$$

Show that \equiv is an equivalence relation on \mathbb{Z} . How many equivalence classes are there, and what do they look like?

7. Let \mathcal{M} be the set of all $n \times n$ matrices. Define a relation \sim on \mathcal{M} by

$$A \sim B \iff \text{there exists an invertible matrix } P \text{ such that } A = P^{-1}BP$$

Show that \sim is an equivalence relation.

1.6. Countability

A set A is called *countable* if it is possible to make a list $a_1, a_2, \dots, a_n, \dots$ which contains all elements of A . A set that is not countable is called *uncountable*. The infinite countable sets are the smallest infinite sets, and we shall later in this section see that the set \mathbb{R} of real numbers is too large to be countable.

Finite sets $A = \{a_1, a_2, \dots, a_m\}$ are obviously countable² as they can be listed

$$a_1, a_2, \dots, a_m, a_m, a_m, \dots$$

(you may list the same elements many times). The set \mathbb{N} of all natural numbers is also countable as it is automatically listed by

$$1, 2, 3, \dots$$

It is a little less obvious that the set \mathbb{Z} of all integers is countable, but we may use the list

$$0, 1, -1, 2, -2, 3, -3, \dots$$

It is also easy to see that a subset of a countable set must be countable, and that the image $f(A)$ of a countable set is countable (if $\{a_n\}$ is a listing of A , then $\{f(a_n)\}$ is a listing of $f(A)$).

The next result is perhaps more surprising:

Proposition 1.6.1. *If the sets A, B are countable, so is the cartesian product $A \times B$.*

Proof. Since A and B are countable, there are lists $\{a_n\}$, $\{b_n\}$ containing all the elements of A and B , respectively. But then

$$\{(a_1, b_1), (a_2, b_1), (a_1, b_2), (a_3, b_1), (a_2, b_2), (a_1, b_3), (a_4, b_1), (a_3, b_2), \dots, \}$$

is a list containing all elements of $A \times B$ (observe how the list is made; first we list the (only) element (a_1, b_1) where the indices sum to 2, then we list the elements

²Some books exclude the finite sets from the countable and treat them as a separate category, but that would be impractical for our purposes.

$(a_2, b_1), (a_1, b_2)$ where the indices sum to 3, then the elements $(a_3, b_1), (a_2, b_2), (a_1, b_3)$ where the indices sum to 4 etc.) \square

Remark: If A_1, A_2, \dots, A_n is a finite collection of countable sets, then the cartesian product $A_1 \times A_2 \times \dots \times A_n$ is countable. This can be proved directly by using the “index trick” in the proof above, or by induction using that $A_1 \times \dots \times A_k \times A_{k+1}$ is essentially the same set as $(A_1 \times \dots \times A_k) \times A_{k+1}$.

The “index trick” can also be used to prove the next result:

Proposition 1.6.2. *If the sets $A_1, A_2, \dots, A_n, \dots$ are countable, so is their union $\bigcup_{n \in \mathbb{N}} A_n$. Hence a countable union of countable sets is itself countable.*

Proof. Let $A_i = \{a_{i1}, a_{i2}, \dots, a_{in}, \dots\}$ be a listing of the i -th set. Then

$$\{a_{11}, a_{21}, a_{12}, a_{31}, a_{22}, a_{13}, a_{41}, a_{32}, \dots\}$$

is a listing of $\bigcup_{i \in \mathbb{N}} A_i$. \square

Proposition 1.6.1 can also be used to prove that the rational numbers are countable:

Proposition 1.6.3. *The set \mathbb{Q} of all rational numbers is countable.*

Proof. According to Proposition 1.6.1, the set $\mathbb{Z} \times \mathbb{N}$ is countable and can be listed $(a_1, b_1), (a_2, b_2), (a_3, b_3), \dots$. But then $\frac{a_1}{b_1}, \frac{a_2}{b_2}, \frac{a_3}{b_3}, \dots$ is a list of all the elements in \mathbb{Q} (due to cancellations, all rational numbers will appear infinitely many times in this list, but that doesn’t matter). \square

Finally, we prove an important result in the opposite direction:

Theorem 1.6.4. *The set \mathbb{R} of all real numbers is uncountable.*

Proof. (Cantor’s diagonal argument) Assume for contradiction that \mathbb{R} is countable and can be listed r_1, r_2, r_3, \dots . Let us write down the decimal expansions of the numbers on the list:

$$\begin{aligned} r_1 &= w_1.a_{11}a_{12}a_{13}a_{14}\dots \\ r_2 &= w_2.a_{21}a_{22}a_{23}a_{24}\dots \\ r_3 &= w_3.a_{31}a_{32}a_{33}a_{34}\dots \\ r_4 &= w_4.a_{41}a_{42}a_{43}a_{44}\dots \\ &\vdots \quad \quad \quad \vdots \end{aligned}$$

(w_i is the integer part of r_i , and $a_{i1}, a_{i2}, a_{i3}, \dots$ are the decimals). To get our contradiction, we introduce a new decimal number $c = 0.c_1c_2c_3c_4\dots$ where the decimals are defined by:

$$c_i = \begin{cases} 1 & \text{if } a_{ii} \neq 1 \\ 2 & \text{if } a_{ii} = 1 \end{cases}$$

This number has to be different from the i -th number r_i on the list as the decimal expansions disagree in the i -th place (as c has only 1 and 2 as decimals, there are

no problems with non-uniqueness of decimal expansions). This is a contradiction as we assumed that *all* real numbers were on the list. \square

Exercises to Section 1.6

1. Show that a subset of a countable set is countable.
2. Show that if A_1, A_2, \dots, A_n are countable, then $A_1 \times A_2 \times \dots \times A_n$ is countable.
3. Show that the set of all finite sequences (q_1, q_2, \dots, q_k) , $k \in \mathbb{N}$, of rational numbers is countable.
4. Show that if A is an *infinite*, countable set, then there is a list a_1, a_2, a_3, \dots which only contains elements in A and where each element in A appears only once. Show that if A and B are two infinite, countable sets, there is a bijection (i.e. an injective and surjective function) $f: A \rightarrow B$.
5. Show that the set of all subsets of \mathbb{N} is uncountable (*Hint*: Try to modify the proof of Theorem 1.6.4.)

Notes and references to Chapter 1

I have tried to make this introductory chapter as brief and concise as possible, but if you think it is too brief, there are many books that treat the material at greater length and with more examples. You may want to try Lakins' book [25] or Hammack's [17] (the latter can be downloaded free of charge).³

Set theory was created by the German mathematician Georg Cantor (1845-1918) in the second half of the 19th century and has since become the most popular foundation for mathematics. Halmos' classic book [16] is still a very readable introduction.

³Numbers in square brackets refer to the bibliography at the end of the book

The Foundation of Calculus

In this chapter we shall take a look at some of the fundamental ideas of calculus that we shall build on throughout the book. How much new material you will find here, depends on your calculus courses. If you have followed a fairly theoretical calculus sequence or taken a course in advanced calculus, almost everything may be familiar, but if your calculus courses were only geared towards calculations and applications, you should work through this chapter before you approach the more abstract theory in Chapter 3.

What we shall study in this chapter is a mixture of theory and technique. We begin by looking at the ϵ - δ -technique for making definitions and proving theorems. You may have found this an incomprehensible nuisance in your calculus courses, but when you get to real analysis, it becomes an indispensable tool that you have to master – the subject matter is now so abstract that you can no longer base your work on geometrical figures and intuition alone. We shall see how the ϵ - δ -technique can be used to treat such fundamental notions as convergence and continuity.

The next topic we shall look at is completeness of \mathbb{R} and \mathbb{R}^n . Although it is often undercommunicated in calculus courses, this is the property that makes calculus work – it guarantees that there are enough real numbers to support our belief in a one-to-one correspondence between real numbers and points on a line. There are two ways to introduce the completeness of \mathbb{R} – by least upper bounds and Cauchy sequences – and we shall look at them both. Least upper bounds will be an important tool throughout the book, and Cauchy sequences will show us how completeness can be extended to more general structures.

In the last section we shall take a look at four important theorems from calculus: the Intermediate Value Theorem, the Bolzano-Weierstrass Theorem, the Extreme Value Theorem, and the Mean Value Theorem. All these theorems are based on the completeness of the real numbers, and they introduce themes that will be important later in the book.

2.1. Epsilon-delta and all that

One often hears that the fundamental concept of calculus is that of a *limit*, but the notion of limit is based on an even more fundamental concept, that of the *distance* between points. When something approaches a limit, the distance between this object and the limit point decreases to zero. To understand limits, we first of all have to understand the notion of distance.

Norms and distances

As you know, the distance between two points $\mathbf{x} = (x_1, x_2, \dots, x_m)$ and $\mathbf{y} = (y_1, y_2, \dots, y_m)$ in \mathbb{R}^m is

$$\|\mathbf{x} - \mathbf{y}\| = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \dots + (x_m - y_m)^2}$$

If we have two numbers x, y on the real line, this expression reduces to

$$|x - y|$$

Note that the order of the points doesn't matter: $\|\mathbf{x} - \mathbf{y}\| = \|\mathbf{y} - \mathbf{x}\|$ and $|x - y| = |y - x|$. This simply means that the distance from \mathbf{x} to \mathbf{y} is the same as the distance from \mathbf{y} to \mathbf{x} .

If you don't like absolute values and norms, these definitions may have made you slightly uncomfortable, but don't despair – there isn't really that much you need to know about absolute values and norms to begin with.

The first thing I would like to emphasize is:

*Whenever you see expressions of the form $\|\mathbf{x} - \mathbf{y}\|$,
think of the distance between \mathbf{x} and \mathbf{y} .*

Don't think of norms or individual points; think of the distance between the points! The same goes for expressions of the form $|x - y|$ where $x, y \in \mathbb{R}$: Don't think of numbers and absolute values; think of the distance between two points on the real line!

The next thing you need to know, is the *triangle inequality* which says that if $\mathbf{x}, \mathbf{y} \in \mathbb{R}^m$, then

$$\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$$

If we put $\mathbf{x} = \mathbf{u} - \mathbf{w}$ and $\mathbf{y} = \mathbf{w} - \mathbf{v}$, this inequality becomes

$$\|\mathbf{u} - \mathbf{v}\| \leq \|\mathbf{u} - \mathbf{w}\| + \|\mathbf{w} - \mathbf{v}\|$$

Try to understand this inequality geometrically. It says that if you are given three points $\mathbf{u}, \mathbf{v}, \mathbf{w}$ in \mathbb{R}^m , the distance $\|\mathbf{u} - \mathbf{v}\|$ of going directly from \mathbf{u} to \mathbf{v} is always less than or equal to the combined distance $\|\mathbf{u} - \mathbf{w}\| + \|\mathbf{w} - \mathbf{v}\|$ of first going from \mathbf{u} to \mathbf{w} and then continuing from \mathbf{w} to \mathbf{v} .

The triangle inequality is important because it allows us to control the size of the sum $\mathbf{x} + \mathbf{y}$ if we know the size of the individual parts \mathbf{x} and \mathbf{y} .

Remark: It turns out that the notion of distance is so central that we can build a theory of convergence and continuity on it alone. This is what we are going to do in the next chapter where we introduce the concept of a metric space. Roughly

speaking, a metric space is a set with a measure of distance that satisfies the triangle inequality.

Convergence of sequences

As a first example of how the notion of distance can be used to define limits, we'll take a look at convergence of sequences. How do we express that a sequence $\{x_n\}$ of real numbers converges to a number a ? The intuitive idea is that we can get x_n as close to a as we want by going sufficiently far out in the sequence; i.e., we can get the distance $|x_n - a|$ as small as we want by choosing n sufficiently large. This means that if our wish is to get the distance $|x_n - a|$ smaller than some chosen number $\epsilon > 0$, there is a number $N \in \mathbb{N}$ (indicating what it means to be "sufficiently large") such that if $n \geq N$, then $|x_n - a| < \epsilon$. Let us state this as a formal definition.

Definition 2.1.1. *A sequence $\{x_n\}$ of real numbers converges to $a \in \mathbb{R}$ if for every $\epsilon > 0$ (no matter how small), there is an $N \in \mathbb{N}$ such that $|x_n - a| < \epsilon$ for all $n \geq N$. We write $\lim_{n \rightarrow \infty} x_n = a$.*

The definition says that for every $\epsilon > 0$, there should be $N \in \mathbb{N}$ satisfying a certain requirement. This N will usually depend on ϵ – the smaller ϵ gets, the larger we have to choose N . Some books emphasize this relationship by writing $N(\epsilon)$ for N . This may be a good pedagogical idea in the beginning, but as it soon becomes a burden, I shall not follow it in this book.

If we think of $|x_n - a|$ as the distance between x_n and a , it's fairly obvious how to extend the definition to sequences $\{\mathbf{x}_n\}$ of points in \mathbb{R}^m .

Definition 2.1.2. *A sequence $\{\mathbf{x}_n\}$ of points in \mathbb{R}^m converges to $\mathbf{a} \in \mathbb{R}^m$ if for every $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $\|\mathbf{x}_n - \mathbf{a}\| < \epsilon$ for all $n \geq N$. Again we write $\lim_{n \rightarrow \infty} \mathbf{x}_n = \mathbf{a}$.*

Note that if we want to show that $\{\mathbf{x}_n\}$ does not converge to $\mathbf{a} \in \mathbb{R}^m$, we have to find an $\epsilon > 0$ such that no matter how large we choose $N \in \mathbb{N}$, there is always an $n \geq N$ such that $\|\mathbf{x}_n - \mathbf{a}\| \geq \epsilon$.

Remark: Some people like to think of the definition above as a game between two players, I and II. Player I wants to show that the sequence $\{\mathbf{x}_n\}$ does *not* converge to \mathbf{a} , while Player II wants to show that it does. The game is very simple: Player I chooses a number $\epsilon > 0$, and player II responds with a number $N \in \mathbb{N}$. Player II wins if $\|\mathbf{x}_n - \mathbf{a}\| < \epsilon$ for all $n \geq N$, otherwise player I wins.

If the sequence $\{\mathbf{x}_n\}$ converges to \mathbf{a} , player II has a winning strategy in this game: No matter which $\epsilon > 0$ player I chooses, player II has a response N that wins the game. If the sequence does not converge to \mathbf{a} , it's player I that has a winning strategy – she can play an $\epsilon > 0$ that player II cannot parry.

Let us take a look at a simple example of how the triangle inequality can be used to prove results about limits.

Proposition 2.1.3. *Assume that $\{\mathbf{x}_n\}$ and $\{\mathbf{y}_n\}$ are two sequences in \mathbb{R}^m converging to \mathbf{a} and \mathbf{b} , respectively. Then the sequence $\{\mathbf{x}_n + \mathbf{y}_n\}$ converges to $\mathbf{a} + \mathbf{b}$.*

Proof. We must show that given an $\epsilon > 0$, we can always find an $N \in \mathbb{N}$ such that $\|(\mathbf{x}_n + \mathbf{y}_n) - (\mathbf{a} + \mathbf{b})\| < \epsilon$ for all $n \geq N$. We start by collecting the terms that “belong together”, and then use the triangle inequality:

$$\|(\mathbf{x}_n + \mathbf{y}_n) - (\mathbf{a} + \mathbf{b})\| = \|(\mathbf{x}_n - \mathbf{a}) + (\mathbf{y}_n - \mathbf{b})\| \leq \|\mathbf{x}_n - \mathbf{a}\| + \|\mathbf{y}_n - \mathbf{b}\|$$

As \mathbf{x}_n converges to \mathbf{a} , we know that there is an $N_1 \in \mathbb{N}$ such that $\|\mathbf{x}_n - \mathbf{a}\| < \frac{\epsilon}{2}$ for all $n \geq N_1$ (if you don’t understand this, see the remark below). As \mathbf{y}_n converges to \mathbf{b} , we can in the same way find an $N_2 \in \mathbb{N}$ such that $\|\mathbf{y}_n - \mathbf{b}\| < \frac{\epsilon}{2}$ for all $n \geq N_2$. If we put $N = \max\{N_1, N_2\}$, we see that when $n \geq N$, then

$$\|(\mathbf{x}_n + \mathbf{y}_n) - (\mathbf{a} + \mathbf{b})\| \leq \|\mathbf{x}_n - \mathbf{a}\| + \|\mathbf{y}_n - \mathbf{b}\| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

This is what we set out to show, and the proposition is proved. \square

Remark: Many get confused when $\frac{\epsilon}{2}$ shows up in the proof above and takes over the rôle of ϵ : We are finding an N_1 such that $\|\mathbf{x}_n - \mathbf{a}\| < \frac{\epsilon}{2}$ for all $n \geq N_1$. But there is nothing irregular in this; since $\mathbf{x}_n \rightarrow \mathbf{a}$, we can tackle any “epsilon-challenge”, including half of the original epsilon.

The proof above illustrates an important aspect of the ϵ - N -definition of convergence, namely that it provides us with a *recipe* for proving that a sequence converges: Given an (arbitrary) $\epsilon > 0$, we simply have to produce an $N \in \mathbb{N}$ that satisfies the condition. This practical side of the definition is often overlooked by students, but as the theory unfolds, you will see it used over and over again.

Continuity

Let us now see how we can use the notion of distance to define continuity. Intuitively, one often says that a function $f: \mathbb{R} \rightarrow \mathbb{R}$ is continuous at a point a if $f(x)$ approaches $f(a)$ as x approaches a , but this is not a precise definition (at least not until one has agreed on what it means for $f(x)$ to “approach” $f(a)$). A better alternative is to say that f is continuous at a if we can get $f(x)$ as close to $f(a)$ as we want by choosing x sufficiently close to a . This means that if we want $f(x)$ to be so close to $f(a)$ that the distance $|f(x) - f(a)|$ is less than some number $\epsilon > 0$, it should be possible to find a $\delta > 0$ such that if the distance $|x - a|$ from x to a is less than δ , then $|f(x) - f(a)|$ is indeed less than ϵ . This is the formal definition of continuity:

Definition 2.1.4. A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is continuous at a point $a \in \mathbb{R}$ if for every $\epsilon > 0$ (no matter how small) there is a $\delta > 0$ such that if $|x - a| < \delta$, then $|f(x) - f(a)| < \epsilon$.

Again we may think of a game between two players: player I who wants to show that the function is discontinuous at a , and player II who wants to show that it is continuous at a . The game is simple: Player I first picks a number $\epsilon > 0$, and player II responds with a $\delta > 0$. Player I wins if there is an x such that $|x - a| < \delta$ and $|f(x) - f(a)| \geq \epsilon$, and player II wins if $|f(x) - f(a)| < \epsilon$ whenever $|x - a| < \delta$. If the function is continuous at a , player II has a winning strategy – she can always parry an ϵ with a judicious choice of δ . If the function is discontinuous at a , player

I has a winning strategy – he can choose an $\epsilon > 0$ that no choice of $\delta > 0$ will parry.

Let us now consider a situation where player I wins, i.e. where the function f is *not* continuous.

Example 1: Let

$$f(x) = \begin{cases} 1 & \text{if } x \leq 0 \\ 2 & \text{if } x > 0 \end{cases}$$

Intuitively this function has a discontinuity at 0 as it makes a jump there, but how is this caught by the ϵ - δ -definition? We see that $f(0) = 1$, but that there are points arbitrarily near 0 where the function value is 2. If we now (acting as player I) choose an $\epsilon < 1$, player II cannot parry: No matter how small she chooses $\delta > 0$, there will be points x , $0 < x < \delta$ where $f(x) = 2$, and consequently $|f(x) - f(0)| = |2 - 1| = 1 > \epsilon$. Hence f is discontinuous at 0. ♣

We shall now take a look at a more complex example of the ϵ - δ -technique where we combine convergence and continuity. Note that the result gives a precise interpretation of our intuitive idea that f is continuous at a if and only if $f(x)$ approaches $f(a)$ whenever x approaches a .

Proposition 2.1.5. *The function $f: \mathbb{R} \rightarrow \mathbb{R}$ is continuous at a if and only if $\lim_{n \rightarrow \infty} f(x_n) = f(a)$ for all sequences $\{x_n\}$ that converge to a .*

Proof. Assume first that f is continuous at a , and that $\lim_{n \rightarrow \infty} x_n = a$. We must show that $f(x_n)$ converges to $f(a)$, i.e., that for a given $\epsilon > 0$, there is always an $N \in \mathbb{N}$ such that $|f(x_n) - f(a)| < \epsilon$ when $n \geq N$. Since f is continuous at a , there is a $\delta > 0$ such that $|f(x) - f(a)| < \epsilon$ whenever $|x - a| < \delta$. As x_n converges to a , there is an $N \in \mathbb{N}$ such that $|x_n - a| < \delta$ when $n \geq N$ (observe that δ now plays the part that usually belongs to ϵ , but that's unproblematic). We now see that if $n \geq N$, then $|x_n - a| < \delta$, and hence $|f(x_n) - f(a)| < \epsilon$, which proves that $\{f(x_n)\}$ converges to $f(a)$.

It remains to show that if f is *not* continuous at a , then there is at least one sequence $\{x_n\}$ that converges to a without $\{f(x_n)\}$ converging to $f(a)$. Since f is discontinuous at a , there is an $\epsilon > 0$ such that no matter how small we choose $\delta > 0$, there is a point x such that $|x - a| < \delta$, but $|f(x) - f(a)| \geq \epsilon$. If we choose $\delta = \frac{1}{n}$, there is thus a point x_n such that $|x_n - a| < \frac{1}{n}$, but $|f(x_n) - f(a)| \geq \epsilon$. The sequence $\{x_n\}$ converges to a , but $\{f(x_n)\}$ *does not* converge to $f(a)$ (since $f(x_n)$ always has distance at least ϵ to $f(a)$). \square

The proof above shows how we can combine different forms of dependence. Note in particular how old quantities reappear in new rôles – suddenly δ is playing the part that usually belongs to ϵ . This is unproblematic as what symbol we are using to denote a quantity, is irrelevant; what we usually call ϵ , could just as well have been called a , b – or δ . The reason why we are always trying to use the same symbol for quantities playing fixed rôles, is that it simplifies our mental processes

– we don't have to waste effort on remembering what the symbols stand for.

Let us also take a look at continuity in \mathbb{R}^n . With our “distance philosophy”, this is just a question of reinterpreting the definition in one dimension:

Definition 2.1.6. A function $\mathbf{F}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuous at the point \mathbf{a} if for every $\epsilon > 0$, there is a $\delta > 0$ such that $\|\mathbf{F}(\mathbf{x}) - \mathbf{F}(\mathbf{a})\| < \epsilon$ whenever $\|\mathbf{x} - \mathbf{a}\| < \delta$.

You can test your understanding by proving the following higher dimensional version of Proposition 2.1.5:

Proposition 2.1.7. The function $\mathbf{F}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuous at \mathbf{a} if and only if $\lim_{k \rightarrow \infty} \mathbf{F}(\mathbf{x}_k) = \mathbf{F}(\mathbf{a})$ for all sequences $\{\mathbf{x}_k\}$ that converge to \mathbf{a} .

For simplicity, I have so far only defined continuity for functions defined on all of \mathbb{R} or all of \mathbb{R}^n , but later in the chapter we shall meet functions that are only defined on subsets, and we need to know what it means for them to be continuous. All we have to do, is to relativize the definition above:

Definition 2.1.8. Assume that A is a subset of \mathbb{R}^n and that \mathbf{a} is an element of A . A function $\mathbf{F}: A \rightarrow \mathbb{R}^m$ is continuous at the point \mathbf{a} if for every $\epsilon > 0$, there is a $\delta > 0$ such that $\|\mathbf{F}(\mathbf{x}) - \mathbf{F}(\mathbf{a})\| < \epsilon$ whenever $\|\mathbf{x} - \mathbf{a}\| < \delta$ and $\mathbf{x} \in A$.

All the results above continue to hold as long as we restrict our attention to points in A .

Estimates

There are several reasons why students find ϵ - δ -arguments difficult. One reason is that they find the basic definitions hard to grasp, but I hope the explanations above have helped you overcome these difficulties, at least to a certain extent. Another reason is that ϵ - δ -arguments are often technically complicated and involve a lot of estimation, something most student find difficult. I'll try to give you some help with this part by working carefully through an example.

Before we begin, I would like to emphasize that when we are doing an ϵ - δ -argument, we are looking for *some* $\delta > 0$ that does the job, and there is usually no sense in looking for the *best* (i.e. the largest) δ . This means that we can often simplify the calculations by using estimates instead of exact values, e.g., by saying things like “this factor can never be larger than 10, and hence it suffices to choose δ equal to $\frac{\epsilon}{10}$.”

Let us take a look at the example:

Proposition 2.1.9. Assume that $g: \mathbb{R} \rightarrow \mathbb{R}$ is continuous at the point a , and that $g(a) \neq 0$. Then the function $h(x) = \frac{1}{g(x)}$ is continuous at a .

Proof. Given an $\epsilon > 0$, we must show that there is a $\delta > 0$ such that $|\frac{1}{g(x)} - \frac{1}{g(a)}| < \epsilon$ when $|x - a| < \delta$.

Let us first write the expression on a more convenient form. Combining the fractions, we get

$$\left| \frac{1}{g(x)} - \frac{1}{g(a)} \right| = \frac{|g(a) - g(x)|}{|g(x)||g(a)|}$$

Since $g(x) \rightarrow g(a)$, we can get the numerator as small as we wish by choosing x sufficiently close to a . The problem is that if the denominator is small, the fraction can still be large (remember that small denominators produce large fractions – we have to think upside down here!) One of the factors in the denominator, $|g(a)|$, is easily controlled as it is constant. What about the other factor $|g(x)|$? Since $g(x) \rightarrow g(a) \neq 0$, this factor can't be too small when x is close to a ; there must, e.g., be a $\delta_1 > 0$ such that $|g(x)| > \frac{|g(a)|}{2}$ when $|x - a| < \delta_1$ (think through what is happening here – it is actually a separate little ϵ - δ -argument). For all x such that $|x - a| < \delta_1$, we thus have

$$\left| \frac{1}{g(x)} - \frac{1}{g(a)} \right| = \frac{|g(a) - g(x)|}{|g(x)||g(a)|} < \frac{|g(a) - g(x)|}{\frac{|g(a)|}{2}|g(a)|} = \frac{2}{|g(a)|^2} |g(a) - g(x)|$$

How can we get this expression less than ϵ ? We obviously need to get $|g(a) - g(x)| < \frac{|g(a)|^2}{2}\epsilon$, and since g is continuous at a , we know there is a $\delta_2 > 0$ such that $|g(a) - g(x)| < \frac{|g(a)|^2}{2}\epsilon$ whenever $|x - a| < \delta_2$. If we choose $\delta = \min\{\delta_1, \delta_2\}$, we get

$$\left| \frac{1}{g(x)} - \frac{1}{g(a)} \right| \leq \frac{2}{|g(a)|^2} |g(a) - g(x)| < \frac{2}{|g(a)|^2} \frac{|g(a)|^2}{2} \epsilon = \epsilon$$

and the proof is complete. \square

Exercises for Section 2.1

1. Show that if the sequence $\{x_n\}$ converges to a , then the sequence $\{Mx_n\}$ (where M is a constant) converges to Ma . Use the definition of convergence and explain carefully how you find N when ϵ is given.
2. Assume that $\{x_n\}$, $\{y_n\}$, and $\{z_n\}$ are three sequences of real numbers such that $x_n \leq y_n \leq z_n$ for all $n \in \mathbb{N}$. Use the definition of convergence to show that if $\{x_n\}$ and $\{z_n\}$ converge to the same number a , then $\{y_n\}$ also converges to a (this is sometimes called the *squeeze law*).
3. Use the definition of continuity to show that if $f: \mathbb{R} \rightarrow \mathbb{R}$ is continuous at a point a , then the function $g(x) = Mf(x)$, where M is a constant, is also continuous at a .
4. Use the definition of continuity to show that if $f, g: \mathbb{R} \rightarrow \mathbb{R}$ are continuous at a point a , then so is $f + g$.
5.
 - a) Use the definition of continuity to show that if $f, g: \mathbb{R} \rightarrow \mathbb{R}$ are continuous at the point a , then so is fg . (*Hint:* Write $|f(x)g(x) - f(a)g(a)| = |(f(x)g(x) - f(a)g(x)) + (f(a)g(x) - f(a)g(a))|$ and use the triangle inequality.)
 - b) Combine the result in a) with Proposition 2.1.9 to show that if f and g are continuous at a and $g(a) \neq 0$, then $\frac{f}{g}$ is continuous at a .
6. Use the definition of continuity to show that if $f(x) = \frac{1}{\sqrt{x}}$ is continuous at all points $a > 0$.
7. Use the triangle inequality to prove that $\|\mathbf{a}\| - \|\mathbf{b}\| \leq \|\mathbf{a} - \mathbf{b}\|$ for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$.

2.2. Completeness

Completeness is probably the most important concept in this book. It will be introduced in full generality in the next chapter, but in this section we shall take a brief look at what it's like in \mathbb{R} and \mathbb{R}^n .

The Completeness Principle

Assume that A is a nonempty subset of \mathbb{R} . We say that A is *bounded above* if there is a number $b \in \mathbb{R}$ such that $b \geq a$ for all $a \in A$, and we say that A is *bounded below* if there is a number $c \in \mathbb{R}$ such that $c \leq a$ for all $a \in A$. We call b and c an *upper* and a *lower bound* of A , respectively.

If b is an upper bound for A , all larger numbers will also be upper bounds. How far can we push it in the opposite direction? Is there a *least upper bound*, i.e. an upper bound d such that $d < b$ for all other upper bounds b ? The Completeness Principle says that there is:

The Completeness Principle: *Every nonempty subset A of \mathbb{R} that is bounded above, has a least upper bound.*

The least upper bound of A is also called the *supremum* of A and is denoted by

$$\sup A$$

We shall sometimes use this notation even when A is not bounded above, and we then put

$$\sup A = \infty$$

This doesn't mean that we count ∞ as a number; it is just a short way of expressing that A stretches all the way to infinity.

We also have a completeness property for lower bounds, but we don't have to state that as a separate principle as it follows from the Completeness Principle above (see Exercise 2 for help with the proof).

Proposition 2.2.1 (The Completeness Principle for Lower Bounds). *Every nonempty subset A of \mathbb{R} that is bounded below, has a greatest lower bound.*

The greatest lower bound of A is also called the *infimum* of A and is denoted by

$$\inf A$$

We shall sometimes use this notation even when A is not bounded below, and we then put

$$\inf A = -\infty$$

Here is a simple example showing some of the possibilities:

Example 1: We shall describe $\sup A$ and $\inf A$ for the following sets.

- (i) $A = [0, 1]$: We have $\sup A = 1$ and $\inf A = 0$. Note that in this case both $\sup A$ and $\inf A$ are elements of A .
- (ii) $A = (0, 1]$: We have $\sup A = 1$ and $\inf A = 0$ as above, but in this case $\sup A \in A$ while $\inf A \notin A$.
- (iii) $A = \mathbb{N}$: We have $\sup A = \infty$ and $\inf A = 1$. In this case $\sup A \notin A$ ($\sup A$ isn't even a real number) while $\inf A \in A$. ♣

The first obstacle in understanding the Completeness Principle is that it seems so obvious – doesn't it just tell us the trivial fact that a bounded set has to stop somewhere? Well, it actually tells us a little bit more; it says that there is a real number that marks where the set ends. To see the difference, let us take a look at an example.

Example 2: The set

$$A = \{x \in \mathbb{R} \mid x^2 < 2\}$$

has $\sqrt{2}$ as its least upper bound. Although this number is not an element of A , it marks in a natural way where the set ends. Consider instead the set

$$B = \{x \in \mathbb{Q} \mid x^2 < 2\}$$

If we are working in \mathbb{R} , $\sqrt{2}$ is still the least upper bound. However, if we insist on working with only the rational numbers \mathbb{Q} , the set B will not have a least upper bound (in \mathbb{Q}) – the only candidate is $\sqrt{2}$ which isn't a rational number. The point is that there isn't a number in \mathbb{Q} that marks where B ends – only a gap that is filled by $\sqrt{2}$ when we extend \mathbb{Q} to \mathbb{R} . This means that \mathbb{Q} doesn't satisfy the Completeness Principle. ♣

Now that we have understood why the Completeness Principle isn't obvious, we may wonder why it is true. This depends on our approach to real numbers. In some books, the real numbers are constructed from the rational numbers, and the Completeness Principle is then a consequence of the construction that has to be proved. In other books, the real numbers are described by a list of axioms (a list of properties we want the system to have), and the Completeness Principle is then one of these axioms. A more everyday approach is to think of the real numbers as the set of all decimal numbers, and the argument in the following example then gives us a good feeling for why the Completeness Principle is true.

Example 3: Let A be a nonempty set of real numbers that has an upper bound b , say $b = 134.27$. We now take a look at the integer parts of the numbers in A . Clearly none of the integer parts can be larger than 134, and probably they don't even go that high. Let's say 87 is the largest integer part we find. We next look at all the elements in A with integer part 87 and ask what is the largest first decimal among these numbers. It cannot be more than 9, and is probably smaller, say 4. We then look at all numbers in A that starts with 87.4 and ask for the biggest second decimal. If it is 2, we next look at all numbers in A that starts with 87.42 and ask for the largest third decimal. Continuing in this way, we produce an infinite decimal expansion 87.42... which gives us the least upper bound of A .

Although I have chosen to work with specific numbers in this example, it is clear that the procedure will work for all bounded sets. ♣

Which of the approaches to the Completeness Principle you prefer, doesn't matter for the rest of the book – we shall just take it to be an established property of the

real numbers. To understand the importance of this property, one has to look at its consequences in different areas of calculus, and we start with sequences.

Monotone sequences, \limsup , and \liminf

A sequence $\{a_n\}$ of real numbers is *increasing* if $a_{n+1} \geq a_n$ for all n , and its *decreasing* if $a_{n+1} \leq a_n$ for all n . We say that a sequence is *monotone* if it's either increasing or decreasing. We also say that $\{a_n\}$ is *bounded* if there is a number $M \in \mathbb{R}$ such that $|a_n| \leq M$ for all n .

Our first result on sequences looks like a triviality, but is actually a very powerful tool.

Theorem 2.2.2. *Every monotone, bounded sequence in \mathbb{R} converges to a number in \mathbb{R} .*

Proof. We consider increasing sequences; the decreasing ones can be dealt with in the same manner. Since the sequence $\{a_n\}$ is bounded, the set

$$A = \{a_1, a_2, a_3, \dots, a_n, \dots\}$$

consisting of all the elements in the sequence, is also bounded and hence has a least upper bound $a = \sup A$. To show that the sequence converges to a , we must show that for each $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $|a - a_n| < \epsilon$ whenever $n \geq N$.

This is not so hard. As a is the *least* upper bound of A , $a - \epsilon$ can not be an upper bound, and hence there must be an a_N such that $a_N > a - \epsilon$. Since the sequence is increasing, this means that $a - \epsilon < a_n \leq a$ for all $n \geq N$, and hence $|a - a_n| < \epsilon$ for such n . \square

Note that the theorem does not hold if we replace \mathbb{R} by \mathbb{Q} : The sequence

$$1, \quad 1.4, \quad 1.41, \quad 1.414, \quad 1.4142, \quad \dots,$$

consisting of longer and longer decimal approximations to $\sqrt{2}$, is a bounded, increasing sequence of rational numbers, but it does not converge to a number in \mathbb{Q} (it converges to $\sqrt{2}$ which is not in \mathbb{Q}).

The theorem above doesn't mean that all sequences converge – unbounded sequences may go to ∞ or $-\infty$, and oscillating sequences may refuse to settle down anywhere. Even when a sequence does not converge, it is possible to say something about its asymptotic behavior (that is the behavior as $n \rightarrow \infty$) by looking at its *upper* and *lower limits*, also known as *limit superior*, \limsup , and *limit inferior*, \liminf . These notions are usually not treated in calculus courses, but as we shall need them now and then later in the book, I'll take this opportunity to introduce them.

Given a sequence $\{a_k\}$ of real numbers, we define two new sequences $\{M_n\}$ and $\{m_n\}$ by

$$M_n = \sup\{a_k \mid k \geq n\}$$

and

$$m_n = \inf\{a_k \mid k \geq n\}$$

We allow that $M_n = \infty$ and that $m_n = -\infty$ as may well occur. The upper sequence $\{M_n\}$ measures how large the original sequence $\{a_k\}$ can become “after” n , and the lower sequence $\{m_n\}$ measures in the same way how small $\{a_k\}$ can become.

Observe that the sequence $\{M_n\}$ is decreasing (as we are taking suprema over smaller and smaller sets), and that $\{m_n\}$ is increasing (as we are taking infima over increasingly smaller sets). Since the sequences are monotone, the limits

$$\lim_{n \rightarrow \infty} M_n \quad \text{and} \quad \lim_{n \rightarrow \infty} m_n$$

exist (we allow them to be ∞ or $-\infty$). We now define the *limit superior* of the original sequence $\{a_n\}$ to be

$$\limsup_{n \rightarrow \infty} a_n = \lim_{n \rightarrow \infty} M_n$$

and the *limit inferior* to be

$$\liminf_{n \rightarrow \infty} a_n = \lim_{n \rightarrow \infty} m_n$$

The intuitive idea is that as n goes to infinity, the sequence $\{a_n\}$ may oscillate and not converge to a limit, but the oscillations will be asymptotically bounded by $\limsup a_n$ above and $\liminf a_n$ below. Figure 2.2.1 shows the graph of a sequence $\{x_n\}$ where the top points converge to an upper limit M and the bottom points to a lower limit m .




Figure 2.2.1. Upper and lower limits

The following relationship should be no surprise:

Proposition 2.2.3. *Let $\{a_n\}$ be a sequence of real numbers. Then*

$$\lim_{n \rightarrow \infty} a_n = b$$

if and only if

$$\limsup_{n \rightarrow \infty} a_n = \liminf_{n \rightarrow \infty} a_n = b$$

(we allow b to be a real number or $\pm\infty$.)

Proof. Assume first that $\limsup_{n \rightarrow \infty} a_n = \liminf_{n \rightarrow \infty} a_n = b$. Since $m_n \leq a_n \leq M_n$, and

$$\begin{aligned}\lim_{n \rightarrow \infty} m_n &= \liminf_{n \rightarrow \infty} a_n = b, \\ \lim_{n \rightarrow \infty} M_n &= \limsup_{n \rightarrow \infty} a_n = b,\end{aligned}$$

we clearly have $\lim_{n \rightarrow \infty} a_n = b$ by “squeezing” (if you are unfamiliar with squeezing, see Exercise 2 in the previous section).

We now assume that $\lim_{n \rightarrow \infty} a_n = b$ where $b \in \mathbb{R}$ (the cases $b = \pm\infty$ are left to the reader). Given an $\epsilon > 0$, there exists an $N \in \mathbb{N}$ such that $|a_n - b| < \epsilon$ for all $n \geq N$. In other words

$$b - \epsilon < a_n < b + \epsilon$$

for all $n \geq N$. But then

$$b - \epsilon \leq m_n < b + \epsilon$$

and

$$b - \epsilon < M_n \leq b + \epsilon$$

for all $n \geq N$. Since this holds for every $\epsilon > 0$, we have $\limsup_{n \rightarrow \infty} a_n = \liminf_{n \rightarrow \infty} a_n = b$. \square

Cauchy sequences

We now want to extend the notion of completeness from \mathbb{R} to \mathbb{R}^m . As there is no natural way to order the points in \mathbb{R}^m when $m > 1$, it is not convenient to use upper and lower bounds to describe the completeness of \mathbb{R}^m . Instead we shall use the notion of Cauchy sequences which also has the advantage of generalizing nicely to the more abstract structures we shall study later in the book. Let us begin with the definition.

Definition 2.2.4. A sequence $\{\mathbf{x}_n\}$ in \mathbb{R}^m is called a Cauchy sequence if for every $\epsilon > 0$ there is an $N \in \mathbb{N}$ such that $\|\mathbf{x}_n - \mathbf{x}_k\| < \epsilon$ when $n, k \geq N$.

Intuitively, a Cauchy sequence is a sequence where the terms are squeezed tighter and tighter the further out in the sequence we get.

The completeness of \mathbb{R}^m will be formulated as a theorem:

Theorem 2.2.5 (Completeness of \mathbb{R}^m). A sequence $\{\mathbf{x}_n\}$ in \mathbb{R}^m converges if and only if it is a Cauchy sequence.

At first glance it is not easy to see the relationship between this theorem and the Completeness Principle for \mathbb{R} , but there is at least a certain similarity on the conceptual level – in a space “without holes”, the terms in a Cauchy sequence ought to be squeezed toward a limit point.

We shall use the Completeness Principle to prove the theorem above, first for \mathbb{R} and then for \mathbb{R}^m . Note that the theorem doesn’t hold in \mathbb{Q} (or in \mathbb{Q}^m for $m > 1$); the sequence

$$1, \quad 1.4, \quad 1.41, \quad 1.414, \quad 1.4142, \quad \dots,$$

of approximations to $\sqrt{2}$ is a Cauchy sequence in \mathbb{Q} that doesn't converge to a number in \mathbb{Q} .

We begin by proving the easy implication.

Proposition 2.2.6. *All convergent sequences in \mathbb{R}^m are Cauchy sequences.*

Proof. Assume that $\{\mathbf{a}_n\}$ converges to \mathbf{a} . Given an $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $\|\mathbf{a}_n - \mathbf{a}\| < \frac{\epsilon}{2}$ for all $n \geq N$. If $n, k \geq N$, we then have

$$\|\mathbf{a}_n - \mathbf{a}_k\| = \|(\mathbf{a}_n - \mathbf{a}) + (\mathbf{a} - \mathbf{a}_k)\| \leq \|\mathbf{a}_n - \mathbf{a}\| + \|\mathbf{a} - \mathbf{a}_k\| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon,$$

and hence $\{\mathbf{a}_n\}$ is a Cauchy sequence. \square

Note that the proof above doesn't rely on the Completeness Principle; it works equally well in \mathbb{Q}^m . The same holds for the next result which we only state for sequences in \mathbb{R} , although it holds for sequences in \mathbb{R}^m (and \mathbb{Q}^m).

Lemma 2.2.7. *Every Cauchy sequence in \mathbb{R} is bounded.*

Proof. We can use the definition of a Cauchy sequence with any ϵ , say $\epsilon = 1$. According to the definition, there is an $N \in \mathbb{N}$ such that $|a_n - a_k| < 1$ whenever $n, k \geq N$. In particular, we have $|a_n - a_N| < 1$ for all $n > N$. This means that

$$K = \max\{a_1, a_2, \dots, a_{N-1}, a_N + 1\}$$

is an upper bound for the sequence and that

$$k = \min\{a_1, a_2, \dots, a_{N-1}, a_N - 1\}$$

is a lower bound. \square

We can now complete the first part of our program. The proof relies on the Completeness Principle through Theorem 2.2.2 and Proposition 2.2.3.

Proposition 2.2.8. *All Cauchy sequences in \mathbb{R} converge.*

Proof. Let $\{a_n\}$ be a Cauchy sequence. Since $\{a_n\}$ is bounded, the upper and lower limits

$$M = \limsup_{n \rightarrow \infty} a_n \quad \text{and} \quad m = \liminf_{n \rightarrow \infty} a_n$$

are finite, and according to Proposition 2.2.3, it suffices to show that $M = m$.

Given an $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $|a_n - a_k| < \epsilon$ whenever $n, k \geq N$. In particular, we have $|a_n - a_N| < \epsilon$ for all $n \geq N$, and hence $m_k \geq a_N - \epsilon$ and $M_k \leq a_N + \epsilon$ for $k \geq N$. Consequently $M_k - m_k \leq 2\epsilon$ for all $k \geq N$. This means that $M - m \leq 2\epsilon$, and since ϵ is an arbitrary, positive number, this is only possible if $M = m$. \square

We are now ready to prove the main theorem:

Proof of Theorem 2.2.5. As we have already proved that all convergent sequences are Cauchy sequences, it only remains to prove that any Cauchy sequence $\{\mathbf{a}_n\}$ converges. If we write out the components of \mathbf{a}_n as

$$\mathbf{a}_n = (a_n^{(1)}, a_n^{(2)}, \dots, a_n^{(m)})$$

the component sequences $\{a_n^{(k)}\}$ are Cauchy sequences in \mathbb{R} and hence convergent according to the previous result. But if the components converge, so does the original sequence $\{\mathbf{a}_n\}$ (see Exercise 10). \square

The argument above shows how we can use the Completeness Principle to prove that all Cauchy sequences converge. It's possible to turn the argument around – to start by assuming that all Cauchy sequences in \mathbb{R} converge and deduce the Completeness Principle (strictly speaking, we then also have to use something called the Archimedean property of the real numbers, but that's something you would probably take for granted anyway). The Completeness Principle and Theorem 2.2.5 can therefore be seen as describing the same notion from two different angles; they capture the phenomenon of completeness in alternative ways. They both have their advantages and disadvantages: The Completeness Principle is simpler and easier to grasp, but convergence of Cauchy sequences is easier to generalize to other structures. In the next chapter we shall generalize it to the setting of metric spaces.

It is probably not clear at this point why completeness is such an important property, but in the next section we shall prove four natural and important theorems that all rely on completeness.

Exercises for section 2.2

1. Explain that $\sup [0, 1) = 1$ and $\sup [0, 1] = 1$. Note that 1 is an element in the latter set, but not in the former.
2. Prove Proposition 2.2.1. (*Hint:* Define $B = \{-a : a \in A\}$ and let $b = \sup B$. Show that $-b$ is the greatest lower bound of A).
3. Prove Theorem 2.2.2 for decreasing sequences.
4. Let $a_n = (-1)^n$. Find $\limsup_{n \rightarrow \infty} a_n$ and $\liminf_{n \rightarrow \infty} a_n$.
5. Let $a_n = \cos \frac{n\pi}{2}$. Find $\limsup_{n \rightarrow \infty} a_n$ and $\liminf_{n \rightarrow \infty} a_n$.
6. Complete the proof of Proposition 2.2.3 for the cases $b = \infty$ and $b = -\infty$.
7. Show that

$$\limsup_{n \rightarrow \infty} (a_n + b_n) \leq \limsup_{n \rightarrow \infty} a_n + \limsup_{n \rightarrow \infty} b_n$$

and

$$\liminf_{n \rightarrow \infty} (a_n + b_n) \geq \liminf_{n \rightarrow \infty} a_n + \liminf_{n \rightarrow \infty} b_n$$

and find examples which show that we do not in general have equality. State and prove a similar result for the product $\{a_n b_n\}$ of two *positive* sequences.

8. Assume that the sequence $\{a_n\}$ is nonnegative and converges to a , and that $b = \limsup_{n \rightarrow \infty} b_n$ is finite and positive. Show that $\limsup_{n \rightarrow \infty} a_n b_n = ab$ (the result holds without the condition that b is positive, but the proof becomes messy). What happens if the sequence $\{a_n\}$ is negative?
9. We shall see how we can define \limsup and \liminf for functions $f: \mathbb{R} \rightarrow \mathbb{R}$. Let $a \in \mathbb{R}$, and define (note that we exclude $x = a$ in these definitions)

$$M_\epsilon = \sup\{f(x) \mid x \in (a - \epsilon, a + \epsilon), x \neq a\}$$

$$m_\epsilon = \inf\{f(x) \mid x \in (a - \epsilon, a + \epsilon), x \neq a\}$$

for $\epsilon > 0$ (we allow $M_\epsilon = \infty$ and $m_\epsilon = -\infty$).

- a) Show that M_ϵ decreases and m_ϵ increases as $\epsilon \rightarrow 0$.

- b) Show that $\lim_{\epsilon \rightarrow 0^+} M_\epsilon$ and $\lim_{\epsilon \rightarrow 0^+} m_\epsilon$ exist (we allow $\pm\infty$ as values).
 We now define $\limsup_{x \rightarrow a} f(x) = \lim_{\epsilon \rightarrow 0^+} M_\epsilon$ and $\liminf_{x \rightarrow a} f(x) = \lim_{\epsilon \rightarrow 0^+} m_\epsilon$.
 c) Show that $\lim_{x \rightarrow a} f(x) = b$ if and only if

$$\limsup_{x \rightarrow a} f(x) = \liminf_{x \rightarrow a} f(x) = b$$

- d) Find $\liminf_{x \rightarrow 0} \sin \frac{1}{x}$ and $\limsup_{x \rightarrow 0} \sin \frac{1}{x}$
 10. Assume that $\{\mathbf{a}_n\}$ is a sequence in \mathbb{R}^m , and write the terms on component form

$$\mathbf{a}_n = (a_n^{(1)}, a_n^{(2)}, \dots, a_n^{(m)})$$

Show that $\{\mathbf{a}_n\}$ converges if and only if all of the component sequences $\{a_n^{(k)}\}$, $k = 1, 2, \dots, m$ converge.

2.3. Four important theorems

We shall end this chapter by taking a look at some famous and important theorems of single- and multivariable calculus: The Intermediate Value Theorem, the Bolzano-Weierstrass Theorem, the Extreme Value Theorem, and the Mean Value Theorem. These results are both a foundation and an inspiration for much of what is going to happen later in the book. Some of them you have probably seen before, others you may not.

The Intermediate Value Theorem

This theorem says that a continuous function $f: \mathbb{R} \rightarrow \mathbb{R}$ cannot change sign without intersecting the x -axis.

Theorem 2.3.1 (The Intermediate Value Theorem). *Assume that $f: [a, b] \rightarrow \mathbb{R}$ is continuous and that $f(a)$ and $f(b)$ have opposite sign. Then there is a point $c \in (a, b)$ such that $f(c) = 0$.*

Proof. We shall consider the case where $f(a) < 0 < f(b)$; the other case can be treated similarly. Let

$$A = \{x \in [a, b] : f(x) < 0\}$$

and put $c = \sup A$. We shall show that $f(c) = 0$. Observe first that since f is continuous and $f(b)$ is strictly positive, our point c has to be strictly less than b . This means that the elements of the sequence $x_n = c + \frac{1}{n}$ lie in the interval $[a, b]$ for all sufficiently large n . Hence $f(x_n) > 0$ for all such n . By Proposition 2.1.5, $f(c) = \lim_{n \rightarrow \infty} f(x_n)$, and as $f(x_n) > 0$, we must have $f(c) \geq 0$.

On the other hand, by definition of c there must for each $n \in \mathbb{N}$ be an element $z_n \in A$ such that $c - \frac{1}{n} < z_n \leq c$. Hence $f(z_n) < 0$ and $z_n \rightarrow c$. Using proposition 2.1.5 again, we get $f(c) = \lim_{n \rightarrow \infty} f(z_n)$, and since $f(z_n) < 0$, this means that $f(c) \leq 0$. But then we have both $f(c) \geq 0$ and $f(c) \leq 0$, which means that $f(c) = 0$. \square

The Intermediate Value Theorem may seem geometrically obvious, but the next example indicates that it isn't.

Example 1: Define a function $f: \mathbb{Q} \rightarrow \mathbb{Q}$ by $f(x) = x^2 - 2$. Then $f(0) = -2 < 0$ and $f(2) = 2 > 0$, but still there isn't a rational number c between 0 and 2 such

that $f(c) = 0$. Hence the Intermediate Value Theorem fails when \mathbb{R} is replaced by \mathbb{Q} . ♣

What is happening here? The function graph sneaks through the x -axis at $\sqrt{2}$ where the rational line has a gap. The Intermediate Theorem tells us that this isn't possible when we are using the real numbers. If you look at the proof, you will see that the reason is that the Completeness Principle allows us to locate a point c where the function value is 0.

The Bolzano-Weierstrass Theorem

To state and prove this theorem, we need the notion of a *subsequence*. If we are given a sequence $\{\mathbf{x}_n\}$ in \mathbb{R}^m , we get a subsequence $\{\mathbf{y}_k\}$ by picking infinitely many (but usually not all) of the terms in $\{\mathbf{x}_n\}$ and then combining them to a new sequence $\{\mathbf{y}_k\}$. More precisely, if

$$n_1 < n_2 < \dots < n_k < \dots$$

are the indices of the terms we pick, then our subsequence is $\{\mathbf{y}_k\} = \{\mathbf{x}_{n_k}\}$.

Recall that a sequence $\{\mathbf{x}_n\}$ in \mathbb{R}^m is *bounded* if there is a number $K \in \mathbb{R}$ such that $\|\mathbf{x}_n\| \leq K$ for all n . The Bolzano-Weierstrass Theorem says that all bounded sequences in \mathbb{R}^m have a convergent subsequence. This is a preview of the notion of compactness that will play an important part later in the book.

Let us first prove the Bolzano-Weierstrass Theorem for \mathbb{R} .

Proposition 2.3.2. *Every bounded sequence in \mathbb{R} has a convergent subsequence.*

Proof. Since the sequence is bounded, there is a finite interval $I_0 = [a_0, b_0]$ that contains all the terms x_n . If we divide this interval into two equally long subintervals $[a_0, \frac{a_0+b_0}{2}]$, $[\frac{a_0+b_0}{2}, b_0]$, at least one of them must contain infinitely many terms from the sequence. Call this interval I_1 (if both subintervals contain infinitely many terms, just choose one of them). We now divide I_1 into two equally long subintervals in the same way, and observe that at least one of them contains infinitely many terms of the sequence. Call this interval I_2 . Continuing in this way, we get an infinite succession of intervals $\{I_n\}$, all containing infinitely many terms of the sequence. Each interval is a subinterval of the previous one, and the lengths of the intervals tend to 0.

We are now ready to define the subsequence. Let y_1 be the first element of the original sequence $\{x_n\}$ that lies in I_1 . Next, let y_2 be the first element after y_1 that lies in I_2 , then let y_3 be the first element after y_2 that lies in I_3 etc. Since all intervals contain infinitely many terms of the sequence, such a choice is always possible, and we obtain a subsequence $\{y_k\}$ of the original sequence. As the y_k 's lie nested in shorter and shorter intervals, $\{y_k\}$ is a Cauchy sequence and hence converges. \square

We are now ready for the main theorem.

Theorem 2.3.3 (The Bolzano-Weierstrass Theorem). *Every bounded sequence in \mathbb{R}^m has a convergent subsequence.*

Proof. Let $\{\mathbf{x}_n\}$ be our sequence, and write it on component form

$$\mathbf{x}_n = (x_n^{(1)}, x_n^{(2)}, \dots, x_n^{(m)})$$

According to the proposition above, there is a subsequence $\{\mathbf{x}_{n_k}\}$ where the first components $\{x_{n_k}^{(1)}\}$ converge. If we use the proposition again, we get a subsequence of $\{\mathbf{x}_{n_k}\}$ where the second components converge (the first components will continue to converge to the same limit as before). Continuing in this way, we end up with a subsequence where all components converge, and then the subsequence itself converges. \square

In the proof of the next result, we shall see a typical example of how the Bolzano-Weierstrass Theorem is used.

The Extreme Value Theorem

Finding maximal and minimal values of functions is important in many parts of mathematics. Before one sets out to find them, it's often smart to check that they exist, and then the Extreme Value Theorem is a useful tool. The theorem has a version that works in \mathbb{R}^m , but as I don't want to introduce extra concepts just for this theorem, I'll stick to the one-dimensional version.

Theorem 2.3.4 (The Extreme Value Theorem). *Assume that $[a, b]$ is a closed, bounded interval, and that $f: [a, b] \rightarrow \mathbb{R}$ is a continuous function. Then f has maximum and minimum points, i.e. there are points $c, d \in [a, b]$ such that*

$$f(d) \leq f(x) \leq f(c)$$

for all $x \in [a, b]$.

Proof. We show that f has a maximum point; the argument for a minimum point is similar.

Let

$$M = \sup\{f(x) \mid x \in [a, b]\}$$

(as we don't know yet that f is bounded, we have to consider the possibility that $M = \infty$). Choose a sequence $\{x_n\}$ in $[a, b]$ such that $f(x_n) \rightarrow M$ (such a sequence exists regardless of whether M is finite or not). Since $[a, b]$ is bounded, $\{x_n\}$ has a convergent subsequence $\{y_k\}$ by the Bolzano-Weierstrass Theorem, and since $[a, b]$ is closed, the limit $c = \lim_{k \rightarrow \infty} y_k$ belongs to $[a, b]$. By construction $f(y_k) \rightarrow M$, but on the other hand, $f(y_k) \rightarrow f(c)$ according to Proposition 2.1.5. Hence $f(c) = M$, and as $M = \sup\{f(x) \mid x \in [a, b]\}$, we have found a maximum point c for f on $[a, b]$. \square

The Mean Value Theorem

The last theorem we are going to look at, differs from the others in that it involves differentiable (and not only continuous) functions. Recall that the derivative of a function f at a point a is defined by

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

The function f is *differentiable* at a if the limit on the right exists (otherwise the function doesn't have a derivative at a).

We need a few lemmas. The first should come as no surprise.

Lemma 2.3.5. *Assume that $f: [a, b] \rightarrow \mathbb{R}$ has a maximum or minimum at an inner point $c \in (a, b)$ where the function is differentiable. Then $f'(c) = 0$.*

Proof. We need to show that we can neither have $f'(c) > 0$ nor $f'(c) < 0$. I'll treat the former case and leave the latter (and similar one) to you. So assume for contradiction that $f'(c) > 0$. Since

$$f'(c) = \lim_{x \rightarrow c} \frac{f(x) - f(c)}{x - c},$$

we must have $\frac{f(x) - f(c)}{x - c} > 0$ for all x sufficiently close to c . If $x > c$, this means that $f(x) > f(c)$, and if $x < c$, it means that $f(x) < f(c)$. Hence c is neither a maximum nor a minimum for f , contradiction. \square

For the proof of the next lemma, we bring in the Extreme Value Theorem.

Lemma 2.3.6 (Rolle's Theorem). *Assume that $f: [a, b] \rightarrow \mathbb{R}$ is continuous in all of $[a, b]$ and differentiable at all inner points $x \in (a, b)$. Assume further that $f(a) = f(b)$. Then there is a point $c \in (a, b)$ where $f'(c) = 0$,*

Proof. According to the Extreme Value Theorem, the function has minimum and maximum points, and since $f(a) = f(b)$, at least one of these must be at an inner point c . According to the previous lemma, $f'(c) = 0$. \square

We are now ready to prove the theorem. It says that for a differentiable function f there is in each interval $[a, b]$ a point c where the instantaneous growth of the function equals its average growth over the interval, i.e.

$$f'(c) = \frac{f(b) - f(a)}{b - a}.$$

Figure 2.3.1 shows what this means geometrically: The slope of the secant through the points $(a, f(a))$ and $(b, f(b))$ on the graph equals the slope of the tangent at some point $c \in (a, b)$.

Theorem 2.3.7 (The Mean Value Theorem). *Assume that $f: [a, b] \rightarrow \mathbb{R}$ is continuous in all of $[a, b]$ and differentiable at all inner points $x \in (a, b)$. Then there is a point $c \in (a, b)$ such that*

$$f'(c) = \frac{f(b) - f(a)}{b - a}$$

Proof. Let g be the function

$$g(x) = f(x) - \frac{f(b) - f(a)}{b - a}(x - a).$$

It is easy to check that $g(a)$ and $g(b)$ are both equal to $f(a)$, and according to Rolle's Theorem there is a point $c \in (a, b)$ where $g'(c) = 0$. As

$$g'(x) = f'(x) - \frac{f(b) - f(a)}{b - a},$$




Figure 2.3.1. The Mean Value Theorem

this means that

$$f'(c) = \frac{f(b) - f(a)}{b - a} \quad \square$$

The Mean Value Theorem is an extremely useful tool in single variable calculus, and in Chapter 6 we shall meet a version of it that also works for functions taking values in higher (including infinite!) dimensional spaces.

Exercises for section 2.3

In exercises 1-4 you are asked to show that the results above would not have held if we had insisted on only working with rational numbers. As the Completeness Principle is the only property that really separates \mathbb{R} from \mathbb{Q} , they underline the importance of this principle. In these exercises, we shall be using the notation

$$[a, b]_{\mathbb{Q}} = \{x \in \mathbb{Q} \mid a \leq x \leq b\}$$

for the set of *rational* numbers between a and b .

1. Show that the function $f: \mathbb{Q} \rightarrow \mathbb{Q}$ defined by $f(x) = \frac{1}{x^2-2}$ is continuous at all $x \in \mathbb{Q}$, but that it is unbounded on $[0, 2]_{\mathbb{Q}}$. Compare to the Extremal Value Theorem.
2. Show that the function $f: \mathbb{Q} \rightarrow \mathbb{Q}$ defined by $f(x) = x^3 - 6x$ is continuous at all $x \in \mathbb{Q}$, but that it does not have a maximum in $[0, 2]_{\mathbb{Q}}$. Compare to the Extremal Value Theorem.
3. Show that the function $f: \mathbb{Q} \rightarrow \mathbb{Q}$ defined by $f(x) = x^3 - 9x$ satisfies $f(0) = f(3) = 0$, but that there are no points in $[0, 3]_{\mathbb{Q}}$ where the derivative is 0. Compare to the Mean Value Theorem.
4. Find a bounded sequence in \mathbb{Q} which does not have a subsequence converging to a point in \mathbb{Q} . Compare to the Bolzano-Weierstrass Theorem.
5. Carry out the proof of the Intermediate Value Theorem in the case where $f(a) > 0 > f(b)$.
6. Explain why the sequence $\{y_k\}$ in the proof of Proposition 2.3.2 is a Cauchy sequence.
7. Explain why there has to be a sequence $\{x_n\}$ as in the proof of the Extremal Value Theorem. Treat the cases $M = \infty$ and $M \neq \infty$ separately.
8. Carry out the proof of Lemma 2.3.5 when $f'(c) < 0$.

9. Assume that f and f' are continuous on the interval $[a, b]$. Show that there is a constant M such that $|f(x) - f(y)| \leq M|x - y|$ for all $x, y \in [a, b]$.
10. In this exercise, we shall prove the following result:

The Heine-Borel Theorem: Assume that \mathcal{I} is a family of open intervals such that $[0, 1] \subseteq \bigcup_{I \in \mathcal{I}} I$. Then there is a finite collection of intervals $I_1, I_2, \dots, I_n \in \mathcal{I}$ such that

$$[0, 1] \subseteq I_1 \cup I_2 \cup \dots \cup I_n$$

To prove this result, let A be the set of points x in $[0, 1]$ with the property that there is a finite collection of intervals $I_1, I_2, \dots, I_n \in \mathcal{I}$ such that $[0, x] \subseteq I_1 \cup I_2 \cup \dots \cup I_n$. Let $c = \sup A$.

- a) Show that $c > 0$.
- b) Show that c cannot be an element in $(0, 1)$ and conclude that $c = 1$.
- c) Prove the theorem.
- d) Explain that the theorem continues to hold if you replace $[0, 1]$ by an arbitrary closed and bounded interval $[a, b]$. Does it hold for open intervals?

Notes and references to Chapter 2

Calculus was developed by Isaac Newton (1642-1727) and Gottfried Wilhelm Leibniz (1646-1716) as an extremely powerful computational tool. In the hands of mathematicians like Jakob Bernoulli (1654-1705), Johann Bernoulli (1667-1748), Leonhard Euler (1707-1783), Joseph Louis Lagrange (1736-1813), and Pierre-Simon Laplace (1749-1827) it revolutionized mathematics, physics, and astronomy in the 18th century. It took a long time, however, to understand the logical foundations for calculus, partly because one didn't have a good grasp of the real number system.

The key figures in giving calculus a rigorous foundation were Augustin Louis Cauchy (1789-1857), Bernhard Bolzano (1781-1848), and Karl Theodor Wilhelm Weierstrass (1815-1897). You will recognize their names from some of the main concepts and results in this chapter (Cauchy sequences and the Bolzano-Weierstrass Theorem). A major reason their program finally succeeded was a better understanding of the real number system developed by Richard Dedekind (1831-1916) and others. Gray's book [14] gives an excellent introduction to the development of mathematical analysis in the 19th century.

The calculus books by Apostol [2] and Spivak [34] have clear but rather long expositions of the theoretical foundations of the theory. Morgan's little book [28] is shorter, but also more condensed. If you really want to understand the key concept of completeness from many different perspectives, Körner's real analysis text [21] has a very thorough discussion, and the book by Hubbard and Hubbard [19] will show you many different ways completeness appears in applied and numerical mathematics. Bressoud [7] has written an introduction to real analysis from a historical perspective, and his text is an excellent companion to this book as it will show you the challenges that led to the modern theory of analysis.

Metric Spaces

Many of the arguments you have seen in several variable calculus are almost identical to the corresponding arguments in single variable calculus, especially arguments concerning convergence and continuity. The reason is that the notions of convergence and continuity can be formulated in terms of distance, and that the notion of distance between numbers that you need in single variable theory, is very similar to the notion of distance between points or vectors that you need in the theory of functions of several variables. In more advanced mathematics, we need to find the distance between more complicated objects than numbers and vectors, e.g. between sequences, sets and functions. These new notions of distance leads to new notions of convergence and continuity, and these again lead to new arguments surprisingly similar to those you have already seen in single and several variable calculus.

After a while it becomes quite boring to perform almost the same arguments over and over again in new settings, and one begins to wonder if there is general theory that covers all these examples – is it possible to develop a general theory of distance where we can prove the results we need once and for all? The answer is yes, and the theory is called the theory of metric spaces.

A metric space is just a set X equipped with a function d of two variables which measures the distance between points: $d(x, y)$ is the distance between two points x and y in X . It turns out that if we put mild and natural conditions on the function d , we can develop a general notion of distance that covers distances between numbers, vectors, sequences, functions, sets and much more. Within this theory we can formulate and prove results about convergence and continuity once and for all. The purpose of this chapter is to develop the basic theory of metric spaces. In later chapters we shall meet some of the applications of the theory.

3.1. Definitions and examples

As already mentioned, a metric space is just a set X equipped with a function $d: X \times X \rightarrow \mathbb{R}$ that measures the distance $d(x, y)$ between points $x, y \in X$. For the

theory to work, we need the function d to have properties similar to the distance functions we are familiar with. So what properties do we expect from a measure of distance?

First of all, the distance $d(x, y)$ should be a nonnegative number, and it should only be equal to zero if $x = y$. Second, the distance $d(x, y)$ from x to y should equal the distance $d(y, x)$ from y to x . Note that this is not always a reasonable assumption – if we, e.g., measure the distance from x to y by the time it takes to walk from x to y , $d(x, y)$ and $d(y, x)$ may be different – but we shall restrict ourselves to situations where the condition is satisfied. The third condition we shall need, says that the distance obtained by going directly from x to y , should always be less than or equal to the distance we get when we go via a third point z , i.e.

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

It turns out that these conditions are the only ones we need, and we sum them up in a formal definition.

Definition 3.1.1. A metric space (X, d) consists of a non-empty set X and a function $d: X \times X \rightarrow [0, \infty)$ such that:

- (i) (Positivity) For all $x, y \in X$, we have $d(x, y) \geq 0$ with equality if and only if $x = y$.
- (ii) (Symmetry) For all $x, y \in X$, we have $d(x, y) = d(y, x)$.
- (iii) (Triangle inequality) For all $x, y, z \in X$, we have

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

A function d satisfying conditions (i)-(iii) is called a metric on X .

Comment: When it is clear – or irrelevant – which metric d we have in mind, we shall often refer to “the metric space X ” rather than “the metric space (X, d) ”.

Let us take a look at some examples of metric spaces.

Example 1: If we let $d(x, y) = |x - y|$, then (\mathbb{R}, d) is a metric space. The first two conditions are obviously satisfied, and the third follows from the ordinary triangle inequality for real numbers:

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

Example 2: If we let

$$d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\| = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$$

then (\mathbb{R}^n, d) is a metric space. The first two conditions are obviously satisfied, and the third follows from the triangle inequality for vectors the same way as above :

$$d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\| = \|(\mathbf{x} - \mathbf{z}) + (\mathbf{z} - \mathbf{y})\| \leq \|\mathbf{x} - \mathbf{z}\| + \|\mathbf{z} - \mathbf{y}\| = d(\mathbf{x}, \mathbf{z}) + d(\mathbf{z}, \mathbf{y})$$

Example 3: Assume that we want to move from one point $\mathbf{x} = (x_1, x_2)$ in the plane to another $\mathbf{y} = (y_1, y_2)$, but that we are only allowed to move horizontally and

vertically. If we first move horizontally from (x_1, x_2) to (y_1, x_2) and then vertically from (y_1, x_2) to (y_1, y_2) , the total distance is

$$d(\mathbf{x}, \mathbf{y}) = |y_1 - x_1| + |y_2 - x_2|$$

This gives us a metric on \mathbb{R}^2 which is different from the usual metric in Example 2. It is often referred to as the *Manhattan metric* or the *taxi cab metric*.

Also in this case the first two conditions of a metric space are obviously satisfied. To prove the triangle inequality, observe that for any third point $\mathbf{z} = (z_1, z_2)$, we have

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

where we have used the ordinary triangle inequality for real numbers to get from the second to the third line. ♣

Example 4: We shall now take a look at an example of a different kind. Assume that we want to send messages in a language with N symbols (letters, numbers, punctuation marks, space, etc.) We assume that all messages have the same length K (if they are too short or too long, we either fill them out or break them into pieces). We let X be the set of all messages, i.e. all sequences of symbols from the language of length K . If $\mathbf{x} = (x_1, x_2, \dots, x_K)$ and $\mathbf{y} = (y_1, y_2, \dots, y_K)$ are two messages, we define

$$d(\mathbf{x}, \mathbf{y}) = \text{the number of indices } n \text{ such that } x_n \neq y_n$$

It is not hard to check that d is a metric. It is usually referred to as the *Hamming-metric*, and is much used in communication theory where it serves as a measure of how much a message gets distorted during transmission. ♣

Example 5: There are many ways to measure the distance between functions, and in this example we shall look at some. Let X be the set of all continuous functions $f: [a, b] \rightarrow \mathbb{R}$. Then

$$d_1(f, g) = \sup\{|f(x) - g(x)| : x \in [a, b]\}$$

is a metric on X . This metric determines the distance between two functions by measuring it at the x -value where the graphs are most apart, and hence the distance between the functions may be large even if they in average are quite close. The metric

$$d_2(f, g) = \int_a^b |f(x) - g(x)| dx$$

instead sums up the distance between $f(x)$ and $g(x)$ at all points. A third popular metric is

$$d_3(f, g) = \left(\int_a^b |f(x) - g(x)|^2 dx \right)^{\frac{1}{2}}$$

This metric is a generalization of the usual (*eucclidean*) metric in \mathbb{R}^n :

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2} = \left(\sum_{i=1}^n (x_i - y_i)^2 \right)^{\frac{1}{2}}$$

(think of the integral as a generalized sum). That we have more than one metric on X , doesn't mean that one of them is "right" and the others "wrong", but that they are useful for different purposes. ♣

Example 6: The metrics in this example may seem rather strange. Although they don't appear much in applications, they are still quite useful as they are totally different from the other metrics we have seen. If you want to check whether a phenomenon from \mathbb{R}^n generalizes to all metric spaces, it's often a good idea first to see what happens in these spaces.

Let X be any non-empty set, and define:

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

It is not hard to check that d is a metric on X , usually referred to as the *discrete* metric. ♣

Example 7: There are many ways to make new metric spaces from old. The simplest is the subspace metric: If (X, d) is a metric space and A is a non-empty subset of X , we can make a metric d_A on A by putting $d_A(x, y) = d(x, y)$ for all $x, y \in A$ – we simply restrict the metric to A . It is trivial to check that d_A is a metric on A . In practice, we rarely bother to change the name of the metric and refer to d_A simply as d , but remember in the back of our head that d is now restricted to A . ♣

There are many more types of metric spaces than we have seen so far, but the hope is that the examples above will give you a certain impression of the variety of the concept. In the next section we shall see how we can define convergence and continuity for sequences and functions in metric spaces. When we prove theorems about these concepts, they automatically hold in all metric spaces, saving us the labor of having to prove them over and over again each time we introduce new spaces.

An important question is when two metric spaces (X, d_X) and (Y, d_Y) are the same. The easy answer is to say that we need the sets X, Y and the functions d_X, d_Y to be equal. This is certainly correct if one interprets "being the same" in the strictest sense, but it is often more appropriate to use a looser definition – in mathematics we are usually not interested in what the elements of a set are, but only in the relationship between them (you may, e.g., want to ask yourself what the natural number 3 "is").

An *isometry* between two metric spaces is a bijection which preserves what is important for metric spaces: the distance between points. More precisely:

Definition 3.1.2. Assume that (X, d_X) and (Y, d_Y) are metric spaces. An isometry between (X, d_X) to (Y, d_Y) is a bijection $i: X \rightarrow Y$ such that $d_X(x, y) = d_Y(i(x), i(y))$ for all $x, y \in X$. We say that (X, d_X) and (Y, d_Y) are isometric if there exists an isometry from (X, d_X) to (Y, d_Y) .

In many situations it is convenient to think of two metric spaces as “the same” if they are isometric. Note that if i is an isometry from (X, d_X) to (Y, d_Y) , then the inverse i^{-1} is an isometry from (Y, d_Y) to (X, d_X) , and hence being isometric is a symmetric relation.

A map which preserves distance, but does not necessarily hit all of Y , is called an *embedding*:

Definition 3.1.3. Assume that (X, d_X) and (Y, d_Y) are metric spaces. An embedding of (X, d_X) into (Y, d_Y) is an injection $i: X \rightarrow Y$ such that $d_X(x, y) = d_Y(i(x), i(y))$ for all $x, y \in X$.

Note that an embedding i can be regarded as an isometry between X and its image $i(X)$.

We end this section with an important consequence of the triangle inequality.

Proposition 3.1.4 (Inverse Triangle Inequality). For all elements x, y, z in a metric space (X, d) , we have

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

Proof. Since the absolute value $|d(x, y) - d(x, z)|$ is the largest of the two numbers $d(x, y) - d(x, z)$ and $d(x, z) - d(x, y)$, it suffices to show that they are both less than or equal to $d(y, z)$. By the triangle inequality

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

and hence $d(x, y) - d(x, z) \leq d(z, y) = d(y, z)$. To get the other inequality, we use the triangle inequality again,

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

and hence $d(x, z) - d(x, y) \leq d(y, z)$. □

Exercises for Section 3.1

1. Show that (X, d) in Example 4 is a metric space.
2. Show that (X, d_1) in Example 5 is a metric space.
3. Show that (X, d_2) in Example 5 is a metric space.
4. Show that (X, d) in Example 6 is a metric space.
5. A sequence $\{x_n\}_{n \in \mathbb{N}}$ of real numbers is called *bounded* if there is a number $M \in \mathbb{R}$ such that $|x_n| \leq M$ for all $n \in \mathbb{N}$. Let X be the set of all bounded sequences. Show that

$$d(\{x_n\}, \{y_n\}) = \sup\{|x_n - y_n| : n \in \mathbb{N}\}$$

is a metric on X .

6. If V is a vector space over \mathbb{R} or \mathbb{C} , a function $\|\cdot\|: V \rightarrow \mathbb{R}$ is called a *norm* if the following conditions are satisfied:

- (i) For all $x \in V$, $\|x\| \geq 0$ with equality if and only if $x = 0$.
 - (ii) $\|\alpha x\| = |\alpha|\|x\|$ for all $\alpha \in \mathbb{R}$ and all $x \in V$.
 - (iii) $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in V$.
- Show that if $\|\cdot\|$ is a norm, then $d(x, y) = \|x - y\|$ defines a metric on V .

7. Show that if x_1, x_2, \dots, x_n are points in a metric space (X, d) , then

$$d(x_1, x_n) \leq d(x_1, x_2) + d(x_2, x_3) + \cdots + d(x_{n-1}, x_n)$$

8. Assume that d_1 and d_2 are two metrics on X . Show that

$$d(x, y) = d_1(x, y) + d_2(x, y)$$

is a metric on X .

9. Assume that (X, d_X) and (Y, d_Y) are two metric spaces. Define a function

$$d: (X \times Y) \times (X \times Y) \rightarrow \mathbb{R}$$

by

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

Show that d is a metric on $X \times Y$.

10. Let X be a non-empty set, and let $\rho: X \times X \rightarrow \mathbb{R}$ be a function satisfying:

- (i) $\rho(x, y) \geq 0$ with equality if and only if $x = y$.
- (ii) $\rho(x, y) \leq \rho(x, z) + \rho(z, y)$ for all $x, y, z \in X$.

Define $d: X \times X \rightarrow \mathbb{R}$ by

$$d(x, y) = \max\{\rho(x, y), \rho(y, x)\}$$

Show that d is a metric on X .

11. Let $a \in \mathbb{R}$. Show that the function $f(x) = x + a$ is an isometry from \mathbb{R} to \mathbb{R} .
12. Recall that an $n \times n$ matrix U is *orthogonal* if $U^{-1} = U^T$. Show that if U is orthogonal and $\mathbf{b} \in \mathbb{R}^n$, then the mapping $i: \mathbb{R}^n \rightarrow \mathbb{R}^n$ given by $i(\mathbf{x}) = U\mathbf{x} + \mathbf{b}$ is an isometry.

3.2. Convergence and continuity

We shall begin our study of metric spaces by defining convergence. A *sequence* $\{x_n\}$ in a metric space X is just an ordered collection $\{x_1, x_2, x_3, \dots, x_n, \dots\}$ of elements in X enumerated by the natural numbers. We want to define what it means for a sequence $\{x_n\}$ to converge to a point a in X .

As we want a definition that works in all metric spaces, and the only thing all metric spaces have in common is a metric, our definition must necessarily be in terms of metrics. Fortunately, we can just mimic the definition we use in \mathbb{R}^m (if this is unfamiliar, take a look at Section 2.1):

Definition 3.2.1. Let (X, d) be a metric space. A sequence $\{x_n\}$ in X converges to a point $a \in X$ if there for every $\epsilon > 0$ (no matter how small) exists an $N \in \mathbb{N}$ such that $d(x_n, a) < \epsilon$ for all $n \geq N$. We write $\lim_{n \rightarrow \infty} x_n = a$ or $x_n \rightarrow a$.

Note that this definition is in accordance with the way we usually think of convergence: We can get x_n as close to a as we wish (i.e. closer than ϵ) by choosing n sufficiently large (i.e. larger than N).

Here is an alternative way to formulate the definition that is often useful.

Lemma 3.2.2. *A sequence $\{x_n\}$ in a metric space (X, d) converges to a if and only if $\lim_{n \rightarrow \infty} d(x_n, a) = 0$.*

Proof. The distances $\{d(x_n, a)\}$ form a sequence of nonnegative numbers. This sequence converges to 0 if and only if there for every $\epsilon > 0$ exists an $N \in \mathbb{N}$ such that $d(x_n, a) < \epsilon$ when $n \geq N$. But this is exactly what the definition above says. \square

May a sequence converge to more than one point? We know that it cannot in \mathbb{R}^n , but some of these new metric spaces are so strange that we can not be certain without a proof.

Proposition 3.2.3. *A sequence in a metric space can not converge to more than one point.*

Proof. Assume that $\lim_{n \rightarrow \infty} x_n = a$ and $\lim_{n \rightarrow \infty} x_n = b$. We must show that this is only possible if $a = b$. According to the triangle inequality

$$d(a, b) \leq d(a, x_n) + d(x_n, b)$$

Taking limits, we get

$$d(a, b) \leq \lim_{n \rightarrow \infty} d(a, x_n) + \lim_{n \rightarrow \infty} d(x_n, b) = 0 + 0 = 0$$

Consequently, $d(a, b) = 0$, and according to point (i) (positivity) in the definition of metric spaces, $a = b$. \square

Observe how we used the conditions in Definition 3.1.1 in the proof above. So far they are all we know about metric spaces. As the theory develops, we shall get more and more tools to work with.

We can also phrase the notion of convergence in more geometric terms. If a is an element of a metric space X , and r is a positive number, the (open) *ball centered at a with radius r* is the set

$$B(a; r) = \{x \in X \mid d(x, a) < r\}$$

As the terminology suggests, we think of $B(a; r)$ as a ball around a with radius r . Note that $x \in B(a; r)$ means exactly the same as $d(x, a) < r$.

The definition of convergence can now be rephrased by saying that $\{x_n\}$ converges to a if the elements of the sequence $\{x_n\}$ eventually end up inside any ball $B(a; \epsilon)$ around a .

Our next task is to define continuity of functions from one metric space to another. We follow the same strategy as above – mimic the definition we have from \mathbb{R}^m (if this is unfamiliar, take a look at Section 2.1).

Definition 3.2.4. *Assume that (X, d_X) , (Y, d_Y) are two metric spaces. A function $f: X \rightarrow Y$ is continuous at a point $a \in X$ if for every $\epsilon > 0$ there is a $\delta > 0$ such that $d_Y(f(x), f(a)) < \epsilon$ whenever $d_X(x, a) < \delta$.*

As already indicated, this definition says exactly the same as the usual definitions of continuity for functions of one or several variables: We can get the distance between $f(x)$ and $f(a)$ smaller than ϵ by choosing x such that the distance between

x and a is smaller than δ . The only difference is that we are now using the metrics d_X and d_Y to measure the distances.

A more geometric formulation of the definition is to say that for any open ball $B(f(a); \epsilon)$ around $f(a)$, there is an open ball $B(a; \delta)$ around a such that $f(B(a; \delta)) \subseteq B(f(a); \epsilon)$ (see Figure 3.2.1 where the dotted curve indicates the boundary of the image $f(B(a; \delta))$).




Figure 3.2.1. Continuity at the point a

There is a close connection between continuity and convergence that reflects our intuitive feeling that f is continuous at a point a if $f(x)$ approaches $f(a)$ whenever x approaches a .

Proposition 3.2.5. *The following are equivalent for a function $f: X \rightarrow Y$ between metric spaces:*

- (i) f is continuous at a point $a \in X$.
- (ii) For all sequences $\{x_n\}$ converging to a , the sequence $\{f(x_n)\}$ converges to $f(a)$.

Proof. Assume first that f is continuous at a and that $\{x_n\}$ is a sequence converging to a . We must show that for any $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $d_Y(f(x_n), f(a)) < \epsilon$ when $n \geq N$. Since f is continuous at a , there is a $\delta > 0$ such that $d_Y(f(x), f(a)) < \epsilon$ whenever $d_X(x, a) < \delta$. Since x_n converges to a , there is an $N \in \mathbb{N}$ such that $d_X(x_n, a) < \delta$ when $n \geq N$. But then $d_Y(f(x_n), f(a)) < \epsilon$ for all $n \geq N$.

Assume now that f is *not* continuous at a . We shall show that there is a sequence $\{x_n\}$ converging to a such that $\{f(x_n)\}$ does *not* converge to $f(a)$. That f is not continuous at a , means that there is an $\epsilon > 0$ such that no matter how small we choose $\delta > 0$, there is an x such that $d_X(x, a) < \delta$, but $d_Y(f(x), f(a)) \geq \epsilon$. In particular, we can for each $n \in \mathbb{N}$ find an x_n such that $d_X(x_n, a) < \frac{1}{n}$ but $d_Y(f(x_n), f(a)) \geq \epsilon$. Then $\{x_n\}$ converges to a , but $\{f(x_n)\}$ does not converge to $f(a)$. \square

As an example of how this result can be applied, we use it to prove that the composition of two continuous functions is continuous.

Proposition 3.2.6. *Let (X, d_X) , (Y, d_Y) , (Z, d_Z) be three metric spaces. Assume that $f: X \rightarrow Y$ and $g: Y \rightarrow Z$ are two functions, and let $h: X \rightarrow Z$ be the*

composition $h(x) = g(f(x))$. If f is continuous at the point $a \in X$ and g is continuous at the point $b = f(a)$, then h is continuous at a .

Proof. Assume that $\{x_n\}$ converges to a . Since f is continuous at a , the sequence $\{f(x_n)\}$ converges to $f(a)$, and since g is continuous at $b = f(a)$, the sequence $\{g(f(x_n))\}$ converges to $g(f(a))$, i.e. $\{h(x_n)\}$ converges to $h(a)$. By the proposition above, h is continuous at a . \square

As in calculus, a function is called continuous if it is continuous at all points:

Definition 3.2.7. A function $f: X \rightarrow Y$ between two metrics spaces is called continuous if it is continuous at all points x in X .

Occasionally, we need to study functions that are only defined on a subset A of our metric space X . We define continuity of such functions by restricting the conditions to elements in A :

Definition 3.2.8. Assume that (X, d_X) , (Y, d_Y) are two metric spaces and that A is a subset of X . A function $f: A \rightarrow Y$ is continuous at a point $a \in A$ if for every $\epsilon > 0$ there is a $\delta > 0$ such that $d_Y(f(x), f(a)) < \epsilon$ whenever $x \in A$ and $d_X(x, a) < \delta$. We say that f is continuous if it is continuous at all $a \in A$.

There is another way of formulating this definition that is sometimes useful: We can think of f as a function from the metric space (A, d_A) (recall Example 7 in Section 3.1) to (Y, d_Y) and use the original definition of continuity in 3.2.4. By just writing it out, it is easy to see that this definition says exactly the same as the one above. The advantage of the second definition is that it makes it easier to transfer results from the full to the restricted setting, e.g., it is now easy to see that Proposition 3.2.5 can be generalized to:

Proposition 3.2.9. Assume that (X, d_X) and (Y, d_Y) are metric spaces and that $A \subseteq X$. Then the following are equivalent for a function $f: A \rightarrow Y$:

- (i) f is continuous at a point $a \in A$.
- (ii) For all sequences $\{x_n\}$ in A converging to a , the sequence $\{f(x_n)\}$ converges to $f(a)$.

Exercises to Section 3.2

1. Assume that (X, d) is a discrete metric space (recall Example 6 in Section 3.1). Show that the sequence $\{x_n\}$ converges to a if and only if there is an $N \in \mathbb{N}$ such that $x_n = a$ for all $n \geq N$.
2. Prove Proposition 3.2.6 without using Proposition 3.2.5, i.e. use only the definition of continuity.
3. Prove Proposition 3.2.9.
4. Assume that (X, d) is a metric space, and let \mathbb{R} have the usual metric $d_{\mathbb{R}}(x, y) = |x - y|$. Assume that $f, g: X \rightarrow \mathbb{R}$ are continuous functions.
 - a) Show that cf is continuous for all constants $c \in \mathbb{R}$.
 - b) Show that $f + g$ is continuous.
 - c) Show that fg is continuous.

5. Let (X, d) be a metric space and choose a point $a \in X$. Show that the function $f: X \rightarrow \mathbb{R}$ given by $f(x) = d(x, a)$ is continuous (we are using the usual metric $d_{\mathbb{R}}(x, y) = |x - y|$ on \mathbb{R}).
6. Let (X, d_X) and (Y, d_Y) be two metric spaces. A function $f: X \rightarrow Y$ is said to be a *Lipschitz function* if there is a constant $K \in \mathbb{R}$ such that $d_Y(f(u), f(v)) \leq K d_X(u, v)$ for all $u, v \in X$. Show that all Lipschitz functions are continuous.
7. Let $d_{\mathbb{R}}$ be the usual metric on \mathbb{R} and let d_{disc} be the discrete metric on \mathbb{R} . Let $id: \mathbb{R} \rightarrow \mathbb{R}$ be the identity function $id(x) = x$. Show that

$$id: (\mathbb{R}, d_{\text{disc}}) \rightarrow (\mathbb{R}, d_{\mathbb{R}})$$

is continuous, but that

$$id: (\mathbb{R}, d_{\mathbb{R}}) \rightarrow (\mathbb{R}, d_{\text{disc}})$$

is not continuous. Note that this shows that the inverse of a bijective, continuous function is not necessarily continuous.

8. In this problem you might want to use the Inverse Triangle Inequality 3.1.4.
 - a) Assume that $\{x_n\}$ is a sequence in a metric space X converging to x . Show that $d(x_n, y) \rightarrow d(x, y)$ for all $y \in X$.
 - b) Assume that $\{x_n\}$ and $\{y_n\}$ are sequences in X converging to x and y , respectively. Show that $d(x_n, y_n) \rightarrow d(x, y)$.
9. Assume that d_1 and d_2 are two metrics on the same space X . We say that d_1 and d_2 are *equivalent* if there are constants K and M such that $d_1(x, y) \leq K d_2(x, y)$ and $d_2(x, y) \leq M d_1(x, y)$ for all $x, y \in X$.
 - a) Assume that d_1 and d_2 are equivalent metrics on X . Show that if $\{x_n\}$ converges to a in one of the metrics, it also converges to a in the other metric.
 - b) Assume that d_1 and d_2 are equivalent metrics on X , and that (Y, d) is a metric space. Show that if $f: X \rightarrow Y$ is continuous when we use the d_1 -metric on X , it is also continuous when we use the d_2 -metric.
 - c) We are in the same setting as in part b), but this time we have a function $g: Y \rightarrow X$. Show that if g is continuous when we use the d_1 -metric on X , it is also continuous when we use the d_2 -metric.
 - d) Assume that d_1, d_2 and d_3 are three metrics on X . Show that if d_1 and d_2 are equivalent, and d_2 and d_3 are equivalent, then d_1 and d_3 are equivalent.
 - e) Show that

$$d_1(\mathbf{x}, \mathbf{y}) = |x_1 - y_1| + |x_2 - y_2| + \dots + |x_n - y_n|$$

$$d_2(\mathbf{x}, \mathbf{y}) = \max\{|x_1 - y_1|, |x_2 - y_2|, \dots, |x_n - y_n|\}$$

$$d_3(\mathbf{x}, \mathbf{y}) = \sqrt{|x_1 - y_1|^2 + |x_2 - y_2|^2 + \dots + |x_n - y_n|^2}$$

are equivalent metrics on \mathbb{R}^n .

3.3. Open and closed sets

In this and the following sections, we shall study some of the most important classes of subsets of metric spaces. We begin by recalling and extending the definition of balls in a metric space:

Definition 3.3.1. Let a be a point in a metric space (X, d) , and assume that r is a positive, real number. The (open) ball centered at a with radius r is the set

$$B(a; r) = \{x \in X : d(x, a) < r\}$$

The closed ball centered at a with radius r is the set

$$\overline{B}(a; r) = \{x \in X : d(x, a) \leq r\}$$

In many ways, balls in metric spaces behave just the way we are used to, but geometrically they may look quite different from ordinary balls. A ball in the Manhattan metric (Example 3 in Section 3.1) looks like an ace of diamonds, while a ball in the discrete metric (Example 6 in Section 3.1) consists either of only one point or the entire space X .¹

Given a point x in X and a subset A of X , there are intuitively three possibilities for the relative positions of the point and the set (see Figure 3.3.1 for an illustration):




Figure 3.3.1. Interior point a , exterior point b , and boundary point c .

- (i) There is a ball $B(x; r)$ around x which is contained in A . In this case x is called an *interior* point of A (the point a in Figure 3.3.1).
- (ii) There is a ball $B(x; r)$ around x which is contained in the complement A^c (all complements are with respect to X , hence $A^c = X \setminus A$). In this case x is called an *exterior* point of A (the point b in Figure 3.3.1).
- (iii) All balls $B(x; r)$ around x contain points in A as well as points in the complement A^c . In this case x is a *boundary point* of A (the point c in Figure 3.3.1).

Note that an interior point *always* belongs to A , while an exterior point *never* belongs to A . A boundary point will some times belong to A , and some times to A^c .

We can now define the important concepts of open and closed sets:

Definition 3.3.2. A subset A of a metric space is open if it does not contain any of its boundary points, and it is closed if it contains all its boundary points.

Most sets contain some, but not all of their boundary points, and are hence neither open nor closed. Figure 3.3.2 illustrates the difference between closed sets, open sets, and sets that are neither closed nor open (whole lines indicate parts of the boundary that belong to the set, dashed lines indicate parts of the boundary that do not belong to the set).

The empty set \emptyset and the entire space X are both open and closed as they do not have any boundary points. Here is an obvious, but useful reformulation of the definition of an open set.

¹As an undergraduate I once came across an announcement on the Mathematics Department's message board saying that somebody would give a talk on "Square balls in Banach spaces". I can still remember the mixture of excitement and disbelief.




Figure 3.3.2. Closed, open, and neither closed nor open set.

Proposition 3.3.3. *A subset A of a metric space X is open if and only if it only consists of interior points, i.e. for all $a \in A$, there is a ball $B(a; r)$ around a which is contained in A .*

Observe that a set A and its complement A^c have exactly the same boundary points. This leads to the following useful result.

Proposition 3.3.4. *A subset A of a metric space X is open if and only if its complement A^c is closed.*

Proof. If A is open, it does not contain any of the (common) boundary points. Hence they all belong to A^c , and A^c must be closed.

Conversely, if A^c is closed, it contains all boundary points, and hence A can not have any. This means that A is open. \square

We can turn an arbitrary set A into an open or closed set by subtracting or adding boundary points. More precisely, we define the *interior* A° of A by

$$A^\circ = \{x \mid x \text{ is an interior point of } A\},$$

and the *closure* \overline{A} of A by

$$\overline{A} = \{x \mid x \in A \text{ or } x \text{ is a boundary point of } A\}$$

Proposition 3.3.5. *A° is an open set and \overline{A} is a closed set.*

Proof. I leave this to the reader with the warning that it is not quite as obvious as it may seem. \square

The following observation may also seem obvious for semantic reasons, but needs to be proved:

Lemma 3.3.6. *All open balls $B(a; r)$ are open sets, while all closed balls $\overline{B}(a; r)$ are closed sets.*

Proof. We prove the statement about open balls and leave the other as an exercise. Assume that $x \in B(a; r)$; we must show that there is a ball $B(x; \epsilon)$ around x which is contained in $B(a; r)$. If we choose $\epsilon = r - d(x, a)$, we see that if $y \in B(x; \epsilon)$ then by the triangle inequality

$$d(y, a) \leq d(y, x) + d(x, a) < \epsilon + d(x, a) = (r - d(x, a)) + d(x, a) = r$$

Thus $d(y, a) < r$, and hence $B(x; \epsilon) \subseteq B(a; r)$. \square

Remark: If you wonder why I chose $\epsilon = r - d(x, a)$ for the radius of the smaller ball in the proof above, draw the situation in \mathbb{R}^2 . Such drawings are often – but not always – helpful for finding the right idea, but they always have to be checked by calculations as other metric spaces may have geometries that are quite different from \mathbb{R}^2 and \mathbb{R}^3 .

The next result shows that closed sets are indeed closed as far as sequences are concerned – at least in the sense that a sequence cannot escape from a closed set:

Proposition 3.3.7. *Assume that F is a subset of a metric space X . The following are equivalent:*

- (i) F is closed.
- (ii) For every convergent sequence $\{x_n\}$ of elements in F , the limit $a = \lim_{n \rightarrow \infty} x_n$ also belongs to F .

Proof. Assume that F is closed and that a does not belong to F . We must show that a sequence from F cannot converge to a . Since F is closed and contains all its boundary points, a has to be an exterior point, and hence there is a ball $B(a; \epsilon)$ around a which only contains points from the complement of F . But then a sequence from F can never get inside $B(a, \epsilon)$, and hence cannot converge to a .

Assume now that F is *not* closed. We shall construct a sequence from F that converges to a point outside F . Since F is not closed, there is a boundary point a that does not belong to F . For each $n \in \mathbb{N}$, we can find a point x_n from F in $B(a; \frac{1}{n})$. Then $\{x_n\}$ is a sequence from F that converges to a point a that is not in F . \square

Characterizations of continuity

We shall use the rest of this section to describe continuous functions in terms of open and closed sets. These descriptions will be useful on several occasions later in the book, but their main importance is that they in later courses will show you how to extend the notion of continuity to even more abstract spaces – so-called *topological spaces*.

Let us begin with some useful terminology. An open set containing x is called a *neighborhood* of x .² The first result is rather silly, but also quite useful.

Lemma 3.3.8. *Let U be a subset of the metric space X , and assume that each $x_0 \in U$ has a neighborhood $U_{x_0} \subseteq U$. Then U is open.*

Proof. We must show that any $x_0 \in U$ is an interior point. Since U_{x_0} is open, there is an $r > 0$ such that $B(x_0, r) \subseteq U_{x_0}$. But then $B(x_0, r) \subseteq U$, which shows that x_0 is an interior point of U . \square

We can use neighborhoods to describe continuity at a point:

Proposition 3.3.9. *Let $f: X \rightarrow Y$ be a function between metric spaces, and let x_0 be a point in X . Then the following are equivalent:*

²In some books, a *neighborhood* of x is not necessarily open, but does contain a ball centered at x . What we have defined, is then referred to as an *open neighborhood*

- (i) f is continuous at x_0 .
- (ii) For all neighborhoods V of $f(x_0)$, there is a neighborhood U of x_0 such that $f(U) \subseteq V$.

Proof. (i) \implies (ii): Assume that f is continuous at x_0 . If V is a neighborhood of $f(x_0)$, there is a ball $B_Y(f(x_0), \epsilon)$ centered at $f(x_0)$ and contained in V . Since f is continuous at x_0 , there is a $\delta > 0$ such that $d_Y(f(x), f(x_0)) < \epsilon$ whenever $d_X(x, x_0) < \delta$. But this means that $f(B_X(x_0, \delta)) \subseteq B_Y(f(x_0), \epsilon) \subseteq V$. Hence (ii) is satisfied if we choose $U = B(x_0, \delta)$.

(ii) \implies (i) We must show that for any given $\epsilon > 0$, there is a $\delta > 0$ such that $d_Y(f(x), f(x_0)) < \epsilon$ whenever $d_X(x, x_0) < \delta$. Since $V = B_Y(f(x_0), \epsilon)$ is a neighborhood of $f(x_0)$, there must be a neighborhood U of x_0 such that $f(U) \subseteq V$. Since U is open, there is a ball $B(x_0, \delta)$ centered at x_0 and contained in U . Assume that $d_X(x, x_0) < \delta$. Then $x \in B_X(x_0, \delta) \subseteq U$, and hence $f(x) \in V = B_Y(f(x_0), \epsilon)$, which means that $d_Y(f(x), f(x_0)) < \epsilon$. Hence we have found a $\delta > 0$ such that $d_Y(f(x), f(x_0)) < \epsilon$ whenever $d_X(x, x_0) < \delta$, and thus f is continuous at x_0 . \square

We can also use open sets to characterize global continuity of functions:

Proposition 3.3.10. *The following are equivalent for a function $f: X \rightarrow Y$ between two metric spaces:*

- (i) f is continuous.
- (ii) Whenever V is an open subset of Y , the inverse image $f^{-1}(V)$ is an open set in X .

Proof. (i) \implies (ii): Assume that f is continuous and that $V \subseteq Y$ is open. We shall prove that $f^{-1}(V)$ is open. For any $x_0 \in f^{-1}(V)$, $f(x_0) \in V$, and we know from the previous theorem that there is a neighborhood U_{x_0} of x_0 such that $f(U_{x_0}) \subseteq V$. But then $U_{x_0} \subseteq f^{-1}(V)$, and by Lemma 3.3.8, $f^{-1}(V)$ is open.

(ii) \implies (i) Assume that the inverse images of open sets are open. To prove that f is continuous at an arbitrary point x_0 , Proposition 3.3.9 tells us that it suffices to show that for any neighborhood V of $f(x_0)$, there is a neighborhood U of x_0 such that $f(U) \subseteq V$. But this is easy: Since the inverse image of an open set is open, we can simply choose $U = f^{-1}(V)$. \square

The description above is useful in many situations. Using that inverse images commute with complements (recall Proposition 1.4.4), and that closed sets are the complements of open sets, we can translate it into a statement about closed sets:

Proposition 3.3.11. *The following are equivalent for a function $f: X \rightarrow Y$ between two metric spaces:*

- (i) f is continuous.
- (ii) Whenever F is a closed subset of Y , the inverse image $f^{-1}(F)$ is a closed set in X .

Proof. (i) \implies (ii): Assume that f is continuous and that $F \subseteq Y$ is closed. Then F^c is open, and by the previous proposition, $f^{-1}(F^c)$ is open. Since inverse images

commute with complements, $f^{-1}(F^c) = (f^{-1}(F))^c$. This means that $f^{-1}(F)$ has an open complement and hence is closed.

(ii) \implies (i) Assume that the inverse images of closed sets are closed. According to the previous proposition, it suffices to show that the inverse image of any open set $V \subseteq Y$ is open. But if V is open, the complement V^c is closed, and hence by assumption $f^{-1}(V^c)$ is closed. Since inverse images commute with complements, $f^{-1}(V^c) = (f^{-1}(V))^c$. This means that the complement of $f^{-1}(V)$ is closed, and hence $f^{-1}(V)$ is open. \square

Mathematicians usually sum up the last two theorems by saying that open and closed sets are preserved under inverse, continuous images. Beware that they are *not* preserved under continuous, *direct* images; even if f is continuous, the image $f(U)$ of an open set U need not be open, and the image $f(F)$ of a closed F need not be closed:

Example 1: Let $f, g: \mathbb{R} \rightarrow \mathbb{R}$ be the continuous functions defined by

$$f(x) = x^2 \quad \text{and} \quad g(x) = \arctan x$$

The set \mathbb{R} is both open and closed, but $f(\mathbb{R})$ equals $[0, \infty)$ which is not open, and $g(\mathbb{R})$ equals $(-\frac{\pi}{2}, \frac{\pi}{2})$ which is not closed. Hence the continuous image of an open set need not be open, and the continuous image of a closed set need not be closed. \clubsuit

As already mentioned, the results above are important when we try to extend the notion of continuity to topological spaces. In such spaces we don't have a notion of distance, only a concept of open sets. We shall not look at topological spaces in this book, but I would still like to point out that from a topological point of view, the crucial properties of open and closed sets are the following:

Proposition 3.3.12. *Let (X, d) be a metric space.*

- a) *If \mathcal{G} is a (finite or infinite) collection of open sets, then the union $\bigcup_{G \in \mathcal{G}} G$ is open.*
- b) *If G_1, G_2, \dots, G_n is a finite collection of open sets, then the intersection $G_1 \cap G_2 \cap \dots \cap G_n$ is open.*

Proof. Left to the reader (see Exercise 11, where you are also asked to show that the intersection of infinitely many open sets is not necessarily open). \square

Proposition 3.3.13. *Let (X, d) be a metric space.*

- a) *If \mathcal{F} is a (finite or infinite) collection of closed sets, then the intersection $\bigcap_{F \in \mathcal{F}} F$ is closed.*
- b) *If F_1, F_2, \dots, F_n is a finite collection of closed sets, then the union $F_1 \cup F_2 \cup \dots \cup F_n$ is closed.*

Proof. Left to the reader (see Exercise 12, where you are also asked to show that the union of infinitely many closed sets is not necessarily closed). \square

Exercises to Section 3.3

1. Assume that (X, d) is a discrete metric space.
 - a) Show that an open ball in X is either a set with only one element (a *singleton*) or all of X .
 - b) Show that all subsets of X are both open and closed.
 - c) Assume that (Y, d_Y) is another metric space. Show that all functions $f: X \rightarrow Y$ are continuous.
2. Give a geometric description of the ball $B(a; r)$ in the Manhattan metric (see Example 3 in Section 3.1). Make a drawing of a typical ball. Show that the Manhattan metric and the usual metric in \mathbb{R}^2 have exactly the same open sets.
3. Assume that F is a non-empty, closed and bounded subset of \mathbb{R} (with the usual metric $d(x, y) = |y - x|$). Show that $\sup F \in F$ and $\inf F \in F$. Give an example of a bounded, but not closed set F such that $\sup F \in F$ and $\inf F \in F$.
4.
 - a) Prove Proposition 3.3.5.
 - b) Show that $x \in \bar{A}$ if and only if there is a sequence $\{a_n\}$ of points in A converging to x .
 - c) Show that $\overline{A \cup B} = \bar{A} \cup \bar{B}$. Give an example of $\overline{A \cap B} \neq \bar{A} \cap \bar{B}$.
5. Prove the second part of Lemma 3.3.6, i.e. prove that a closed ball $\bar{B}(a; r)$ is always a closed set.
6. Assume that $f: X \rightarrow Y$ and $g: Y \rightarrow Z$ are continuous functions. Use Proposition 3.3.10 to show that the composition $g \circ f: X \rightarrow Z$ is continuous.
7. Assume that A is a subset of a metric space (X, d) . Show that the interior points of A are the exterior points of A^c , and that the exterior points of A are the interior points of A^c . Check that the boundary points of A are the boundary points of A^c .
8. Let (X, d) be a metric space, and let A be a subset of X . We shall consider A with the subset metric d_A .
 - a) Assume that $G \subseteq A$ is open in (X, d) . Show that G is open in (A, d_A) .
 - b) Find an example which shows that although $G \subseteq A$ is open in (A, d_A) it need not be open in (X, d_X) .
 - c) Show that if A is an open set in (X, d_X) , then a subset G of A is open in (A, d_A) if and only if it is open in (X, d_X) .
9. Let (X, d) be a metric space, and let A be a subset of X . We shall consider A with the subset metric d_A .
 - a) Assume that $F \subseteq A$ is closed in (X, d) . Show that F is closed in (A, d_A) .
 - b) Find an example which shows that although $F \subseteq A$ is closed in (A, d_A) it need not be closed in (X, d_X) .
 - c) Show that if A is a closed set in (X, d_X) , then a subset F of A is closed in (A, d_A) if and only if it is closed in (X, d_X) .
10. Let (X, d) be a metric space and give \mathbb{R} the usual metric. Assume that $f: X \rightarrow \mathbb{R}$ is continuous.
 - a) Show that the set

$$\{x \in X \mid f(x) < a\}$$
 is open for all $a \in \mathbb{R}$.
 - a) Show that the set

$$\{x \in X \mid f(x) \leq a\}$$
 is closed for all $a \in \mathbb{R}$.
11. Prove Proposition 3.3.12. Find an example of an infinite collection of open sets G_1, G_2, \dots whose intersection is *not* open.

12. Prove Proposition 3.3.13. Find an example of an infinite collection of closed sets F_1, F_2, \dots whose union is *not* closed.
- 13 A metric space (X, d) is said to be *disconnected* if there are two nonempty open set O_1, O_2 such that

$$X = O_1 \cup O_2 \quad \text{and} \quad O_1 \cap O_2 = \emptyset$$

A metric space that is *not* disconnected is said to be *connected*.

- a) Let $X = (-1, 1) \setminus \{0\}$ and let d be the usual metric $d(x, y) = |x - y|$ on X . Show that (X, d) is disconnected.
- b) Let $X = \mathbb{Q}$ and let again d be the usual metric $d(x, y) = |x - y|$ on X . Show that (X, d) is disconnected.
- c) Assume that (X, d) is a connected metric space and that $f: X \rightarrow Y$ is continuous and surjective. Show that Y is connected.

A metric space (X, d) is called *path-connected* if for every pair x, y of points in X , there is a continuous function $r: [0, 1] \rightarrow X$ such that $r(0) = x$ and $r(1) = y$ (such a function is called a *path* from x to y).

- d) Let d be the usual metric on \mathbb{R}^n :

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

Show that (\mathbb{R}^n, d) is path-connected.

- e) Show that every path-connected metric space is connected. (*Hint:* Argue contrapositively: Assume that (X, d) is not connected. Let O_1, O_2 be two nonempty, open sets such that $X = O_1 \cup O_2$ and $O_1 \cap O_2 = \emptyset$, and pick points $x \in O_1, y \in O_2$. Show that there doesn't exist a path from x to y).

Just for your information, there are connected spaces that are not path-connected. A famous example is "the topologist's sine curve" where X consists of all points on the graph $y = \sin \frac{1}{x}$, $x \neq 0$, plus the point $(0, 0)$, and the metric is the one inherited from \mathbb{R}^2 .

3.4. Complete spaces

The main reason why calculus in \mathbb{R} and \mathbb{R}^n is so successful, is that that these spaces are complete. In order to follow up the success, we shall now generalize the notion of completeness to metric spaces. If you are not familiar with the completeness of \mathbb{R} and \mathbb{R}^n , you should take a look at Section 2.2 before you continue.

There are two standard ways to describe the completeness of \mathbb{R} : by least upper bounds and by Cauchy sequences. As metric spaces are usually not ordered, the least upper bound description is impossible to generalize, and we need to use Cauchy sequences.

Definition 3.4.1. A sequence $\{x_n\}$ in a metric space (X, d) is a Cauchy sequence if for each $\epsilon > 0$ there is an $N \in \mathbb{N}$ such that $d(x_n, x_m) < \epsilon$ whenever $n, m \geq N$.

We begin by a simple observation:

Proposition 3.4.2. Every convergent sequence is a Cauchy sequence.

Proof. If a is the limit of the sequence, there is for any $\epsilon > 0$ a number $N \in \mathbb{N}$ such that $d(x_n, a) < \frac{\epsilon}{2}$ whenever $n \geq N$. If $n, m \geq N$, the triangle inequality tells us that

$$d(x_n, x_m) \leq d(x_n, a) + d(a, x_m) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

and consequently $\{x_n\}$ is a Cauchy sequence. □

The converse of the proposition above does not hold in all metric spaces, and we make the following crucial definition:

Definition 3.4.3. *A metric space is called complete if all Cauchy sequences converge.*

Example 1: We know from Section 2.2 that \mathbb{R}^n is complete, but that \mathbb{Q} is not. ♣

Example 2: The metric is important when we are dealing with completeness. In Example 5 in Section 3.1, we looked at three different metrics d_1, d_2, d_3 on the space X of all continuous $f: [a, b] \rightarrow \mathbb{R}$, but only d_1 is complete (we shall prove this in the next chapter). However, by introducing a stronger notion of integral (the Lebesgue integral, see Chapter ??) we can extend d_2 and d_3 to complete metrics by making them act on richer spaces of functions. In Section 3.7, we shall study an abstract method for making incomplete spaces complete by adding new points. ♣

The complete spaces are in many ways the “nice” metric spaces, and we shall spend much time studying their properties. The main reason why completeness is so important, is that in mathematics we often solve problems and construct objects by approximation – we find better and better approximate solutions and then obtain the true solution as the limit of these approximations. Completeness is usually necessary to guarantee that the limit exists.

Before we take a closer look at an example of such approximations, we pick up the following technical proposition that is often useful. Remember that if A is a subset of X , then d_A is the subspace metric obtained by restricting d to A (see Example 7 in Section 3.1).

Proposition 3.4.4. *Assume that (X, d) is a complete metric space. If A is a subset of X , (A, d_A) is complete if and only if A is closed.*

Proof. Assume first that A is closed. If $\{a_n\}$ is a Cauchy sequence in A , $\{a_n\}$ is also a Cauchy sequence in X , and since X is complete, $\{a_n\}$ converges to a point $a \in X$. Since A is closed, Proposition 3.3.7 tells us that $a \in A$. But then $\{a_n\}$ converges to a in (A, d_A) , and hence (A, d_A) is complete.

If A is not closed, there is a boundary point a that does not belong to A . Each ball $B(a, \frac{1}{n})$ must contain an element a_n from A . In X , the sequence $\{a_n\}$ converges to a , and must be a Cauchy sequence. However, since $a \notin A$, the sequence $\{a_n\}$ does not converge to a point in A . Hence we have found a Cauchy sequence in (A, d_A) that does not converge to a point in A , and hence (A, d_A) is incomplete. \square

Let us now take a look at a situation where we use completeness to find a solution by repeated approximation. Assume that X is a metric space and that $f: X \rightarrow X$ is function mapping X to itself. Given a point x_0 in X , we can construct a sequence $x_0, x_1, x_2, \dots, x_n, \dots$ by putting $x_1 = f(x_0), x_2 = f(x_1)$ and so on. We say that we *iterate* f with initial condition x_0 . It is often helpful to think of the sequence as a system evolving in time: x_0 is the state of the system at time 0, x_1 is the state of the system at time 1 etc. A *fixed point* for f is an element $a \in X$ such that $f(a) = a$. If we think in terms of an evolving system, a fixed point is an

equilibrium state – a state that doesn't change with time. As many systems have a tendency to converge to equilibrium, it is interesting to ask when our iterated sequence x_0, x_1, x_2, \dots converges to a fixed point a .

To formulate a theorem, we need a few more definitions. A function $f : X \rightarrow X$ is called a *contraction* if there is a positive number $s < 1$ such that

$$d(f(x), f(y)) \leq s d(x, y) \quad \text{for all } x, y \in X$$

We call s a *contraction factor* for f (note that the same s should work for all $x, y \in X$). All contractions are continuous (prove this!), and by induction it is easy to see that

$$d(f^{\circ n}(x), f^{\circ n}(y)) \leq s^n d(x, y)$$

where $f^{\circ n}(x) = f(f(\dots f(x) \dots))$ is the result of iterating f exactly n times.

Theorem 3.4.5 (Banach's Fixed Point Theorem). *Assume that (X, d) is a complete metric space and that $f : X \rightarrow X$ is a contraction. Then f has a unique fixed point a , and no matter which starting point $x_0 \in X$ we choose, the sequence*

$$x_0, x_1 = f(x_0), x_2 = f^{\circ 2}(x_0), \dots, x_n = f^{\circ n}(x_0), \dots$$

converges to a .

Proof. Let us first show that f cannot have more than one fixed point. If a and b are two fixed points, and s is a contraction factor for f , we have

$$d(a, b) = d(f(a), f(b)) \leq s d(a, b)$$

Since $0 < s < 1$, this is only possible if $d(a, b) = 0$, i.e. if $a = b$.

To show that f has a fixed point, choose a starting point x_0 in X and consider the sequence

$$x_0, x_1 = f(x_0), x_2 = f^{\circ 2}(x_0), \dots, x_n = f^{\circ n}(x_0), \dots$$

Assume, for the moment, that we can prove that this is a Cauchy sequence. Since (X, d) is complete, the sequence must converge to a point a . To prove that a is a fixed point, observe that we have $x_{n+1} = f(x_n)$ for all n , and taking the limit as $n \rightarrow \infty$, we get $a = f(a)$. Hence a is a fixed point of f , and the theorem must hold. Thus it suffices to prove our assumption that $\{x_n\}$ is a Cauchy sequence.

Choose two elements x_n and x_{n+k} of the sequence. By repeated use of the triangle inequality (see Exercise 3.1.7 if you need help), we get

$$\begin{aligned} d(x_n, x_{n+k}) &\leq d(x_n, x_{n+1}) + d(x_{n+1}, x_{n+2}) + \dots + d(x_{n+k-1}, x_{n+k}) \\ &= d(f^{\circ n}(x_0), f^{\circ n}(x_1)) + d(f^{\circ(n+1)}(x_0), f^{\circ(n+1)}(x_1)) + \dots \\ &\quad \dots + d(f^{\circ(n+k-1)}(x_0), f^{\circ(n+k-1)}(x_1)) \\ &\leq s^n d(x_0, x_1) + s^{n+1} d(x_0, x_1) + \dots + s^{n+k-1} d(x_0, x_1) \\ &= \frac{s^n(1 - s^k)}{1 - s} d(x_0, x_1) \leq \frac{s^n}{1 - s} d(x_0, x_1) \end{aligned}$$

where we have summed a geometric series to get to the last line. Since $s < 1$, we can get the last expression as small as we want by choosing n large enough. Given an

$\epsilon > 0$, we can in particular find an N such that $\frac{s^N}{1-s} d(x_0, x_1) < \epsilon$. For $n, m = n + k$ larger than or equal to N , we thus have

$$d(x_n, x_m) \leq \frac{s^n}{1-s} d(x_0, x_1) < \epsilon$$

and hence $\{x_n\}$ is a Cauchy sequence. \square

Remark: In the proof above, we have

$$d(x_n, x_m) \leq \frac{s^n}{1-s} d(x_0, x_1)$$

for $m > n$. If we let $m \rightarrow \infty$, we get

$$(3.4.1) \quad d(x_n, a) \leq \frac{s^n}{1-s} d(x_0, x_1)$$

where a is the fixed point. This gives us complete control over the rate of convergence – if we know the contraction factor s and the length $d(x_0, x_1)$ of the first step in the iteration process, we have very precise information about how good an approximation x_n is of the fixed point. This is important in numerical applications of the method.

In Section 4.7 we shall use Banach's Fixed Point Theorem to prove the existence of solutions of differential equations and in Chapter 6 we shall use it to prove the Inverse and Implicit Function Theorems.

Remark: There is a methodological aspect of the proof above that is worth a comment: We have proved that the fixed point a exists without actually finding it – all we have worked with, are the terms $\{x_n\}$ of the given sequence. This is one of the great advantages of completeness; in a complete space you don't have to construct the limit object, you just have to check that the approximations form a Cauchy sequence. You may wonder what is the value of knowing that something exists when you don't know what it is, but it turns out that existence is of great value in itself, both mathematically and psychologically: Very few of us are willing to spend a lot of effort on studying the hypothetical properties of something that may not even exist!

Exercises to Section 3.4

1. Show that the discrete metric is always complete.
2. Assume that (X, d_X) and (Y, d_Y) are complete spaces, and give $X \times Y$ the metric d defined by

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

Show that $(X \times Y, d)$ is complete.

3. If A is a subset of a metric space (X, d) , the *diameter* $\text{diam}(A)$ of A is defined by

$$\text{diam}(A) = \sup\{d(x, y) \mid x, y \in A\}$$

Let $\{A_n\}$ be a collection of subsets of X such that $A_{n+1} \subseteq A_n$ and $\text{diam}(A_n) \rightarrow 0$, and assume that $\{a_n\}$ is a sequence such that $a_n \in A_n$ for each $n \in \mathbb{N}$. Show that if X is complete, the sequence $\{a_n\}$ converges.

4. Assume that d_1 and d_2 are two metrics on the same space X . We say that d_1 and d_2 are *equivalent* if there are constants K and M such that $d_1(x, y) \leq Kd_2(x, y)$ and $d_2(x, y) \leq Md_1(x, y)$ for all $x, y \in X$. Show that if d_1 and d_2 are equivalent, and one of the spaces (X, d_1) , (X, d_2) is complete, then so is the other.
5. Assume that $f: [0, 1] \rightarrow [0, 1]$ is a differentiable function and that there is a number $s < 1$ such that $|f'(x)| < s$ for all $x \in (0, 1)$. Show that there is exactly one point $a \in [0, 1]$ such that $f(a) = a$.
6. You are standing with a map in your hand inside the area depicted on the map. Explain that there is exactly one point on the map that is vertically above the point it depicts.
7. Assume that (X, d) is a complete metric space, and that $f: X \rightarrow X$ is a function such that $f^{\circ n}$ is a contraction for some $n \in \mathbb{N}$. Show that f has a unique fixed point.
8. A subset D of a metric space X is *dense* if for all $x \in X$ and all $\epsilon \in \mathbb{R}_+$ there is an element $y \in D$ such that $d(x, y) < \epsilon$. Show that if all Cauchy sequences $\{y_n\}$ from a dense set D converge in X , then X is complete.

3.5. Compact sets

We now turn to the study of compact sets. These sets are related both to closed sets and to the notion of completeness, and they are extremely useful in many applications.

Assume that $\{x_n\}$ is a sequence in a metric space X . If we have a strictly increasing sequence of natural numbers

$$n_1 < n_2 < n_3 < \dots < n_k < \dots,$$

we call the sequence $\{y_k\} = \{x_{n_k}\}$ a *subsequence* of $\{x_n\}$. A subsequence contains infinitely many of the terms in the original sequence, but usually not all.

I leave the first result as an exercise:

Proposition 3.5.1. *If the sequence $\{x_n\}$ converges to a , so do all subsequences.*

We are now ready to define compact sets:

Definition 3.5.2. *A subset K of a metric space (X, d) is called a compact set if every sequence in K has a subsequence converging to a point in K . The space (X, d) is compact if X is a compact set, i.e. if all sequences in X have a convergent subsequence.*

Remark: It is easy to overlook that the limit of the subsequence has to lie in the set K , but this is a crucial part of the definition.

Compactness is a rather complex notion that it takes a while to get used to. We shall start by relating it to other concepts we have already introduced. First a definition:

Definition 3.5.3. *A subset A of a metric space (X, d) is bounded if there is a number $M \in \mathbb{R}$ such that $d(a, b) \leq M$ for all $a, b \in A$.*

An equivalent definition is to say that there is a point $c \in X$ and a constant $K \in \mathbb{R}$ such that $d(a, c) \leq K$ for all $a \in A$ (it does not matter which point $c \in X$ we use in this definition). See Exercise 4.

Here is our first result on compact sets:

Proposition 3.5.4. *Every compact set K in a metric space (X, d) is closed and bounded.*

Proof. We argue contrapositively. First we show that if a set K is not closed, then it can not be compact, and then we show that if K is not bounded, it can not be compact.

Assume that K is not closed. Then there is a boundary point a that does not belong to K . For each $n \in \mathbb{N}$, there is an $x_n \in K$ such that $d(x_n, a) < \frac{1}{n}$. The sequence $\{x_n\}$ converges to $a \notin K$, and so do all its subsequences, and hence no subsequence can converge to a point in K .

Assume now that K is not bounded and pick a point $b \in K$. For every $n \in \mathbb{N}$ there is an element $x_n \in K$ such that $d(x_n, b) > n$. If $\{y_k\}$ is a subsequence of x_n , clearly $\lim_{k \rightarrow \infty} d(y_k, b) = \infty$. It is easy to see that $\{y_k\}$ can not converge to any element $y \in X$: According to the triangle inequality

$$d(y_k, b) \leq d(y_k, y) + d(y, b)$$

and since $d(y_k, b) \rightarrow \infty$, we must have $d(y_k, y) \rightarrow \infty$. Hence $\{x_n\}$ has no convergent subsequences, and K can not be compact. \square

In \mathbb{R}^n the converse of the result above holds:

Corollary 3.5.5. *A subset of \mathbb{R}^n is compact if and only if it is closed and bounded.*

Proof. We have to prove that a closed and bounded subset A of \mathbb{R}^n is compact. This is just a slight extension of the Bolzano-Weierstrass Theorem 2.3.3: A sequence $\{\mathbf{x}_n\}$ in A is bounded since A is bounded, and by the Bolzano-Weierstrass Theorem it has a subsequence converging to a point $\mathbf{a} \in \mathbb{R}^n$. Since A is closed, $\mathbf{a} \in A$. \square

Unfortunately, the corollary doesn't hold for metric spaces in general.

Example 1: Consider the metric space (\mathbb{N}, d) where d is the discrete metric. Then \mathbb{N} is complete, closed and bounded, but the sequence $\{n\}$ does not have a convergent subsequence.

We shall later see how we can strengthen the boundedness condition (to something called *total boundedness*) to get a characterization of compactness that holds in all complete metric spaces.

We next want to take a look at the relationship between completeness and compactness. Not all complete spaces are compact (\mathbb{R} is complete but not compact), but it turns out that all compact spaces are complete. To prove this, we need a lemma on subsequences of Cauchy sequences that is useful also in other contexts.

Lemma 3.5.6. *Assume that $\{x_n\}$ is a Cauchy sequence in a (not necessarily complete) metric space (X, d) . If there is a subsequence $\{x_{n_k}\}$ converging to a point a , then the original sequence $\{x_n\}$ also converges to a .*

Proof. We must show that for any given $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $d(x_n, a) < \epsilon$ for all $n \geq N$. Since $\{x_n\}$ is a Cauchy sequence, there is an $N \in \mathbb{N}$ such that $d(x_n, x_m) < \frac{\epsilon}{2}$ for all $n, m \geq N$. Since $\{x_{n_k}\}$ converges to a , there is a K such that $n_K \geq N$ and $d(x_{n_K}, a) \leq \frac{\epsilon}{2}$. For all $n \geq N$ we then have

$$d(x_n, a) \leq d(x_n, x_{n_K}) + d(x_{n_K}, a) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

by the triangle inequality. \square

Proposition 3.5.7. *Every compact metric space is complete.*

Proof. Let $\{x_n\}$ be a Cauchy sequence. Since X is compact, there is a subsequence $\{x_{n_k}\}$ converging to a point a . By the lemma above, $\{x_n\}$ also converges to a . Hence all Cauchy sequences converge, and X must be complete. \square

Here is another useful result:

Proposition 3.5.8. *A closed subset F of a compact set K is compact.*

Proof. Assume that $\{x_n\}$ is a sequence in F – we must show that $\{x_n\}$ has a subsequence converging to a point in F . Since $\{x_n\}$ is also a sequence in K , and K is compact, there is a subsequence $\{x_{n_k}\}$ converging to a point $a \in K$. Since F is closed, $a \in F$, and hence $\{x_n\}$ has a subsequence converging to a point in F . \square

We have previously seen that if f is a continuous function, the inverse images of open and closed sets are open and closed, respectively. The inverse image of a compact set need not be compact, but it turns out that the (direct) image of a compact set under a continuous function is always compact.

Proposition 3.5.9. *Assume that $f: X \rightarrow Y$ is a continuous function between two metric spaces. If $K \subseteq X$ is compact, then $f(K)$ is a compact subset of Y .*

Proof. Let $\{y_n\}$ be a sequence in $f(K)$; we shall show that $\{y_n\}$ has subsequence converging to a point in $f(K)$. Since $y_n \in f(K)$, we can for each n find an element $x_n \in K$ such that $f(x_n) = y_n$. Since K is compact, the sequence $\{x_n\}$ has a subsequence $\{x_{n_k}\}$ converging to a point $x \in K$. But then by Proposition 3.2.5, $\{y_{n_k}\} = \{f(x_{n_k})\}$ is a subsequence of $\{y_n\}$ converging to $y = f(x) \in f(K)$. \square

So far we have only proved technical results about the nature of compact sets. The next result gives the first indication why these sets are useful. It is a generalization of the Extreme Value Theorem of Calculus 2.3.4.

Theorem 3.5.10 (The Extreme Value Theorem). *Assume that K is a non-empty, compact subset of a metric space (X, d) and that $f: K \rightarrow \mathbb{R}$ is a continuous function. Then f has maximum and minimum points in K , i.e. there are points $c, d \in K$ such that*

$$f(d) \leq f(x) \leq f(c)$$

for all $x \in K$.

Proof. There is a quick way of proving this theorem by using the previous proposition (see the remark below), but I choose a slightly longer proof as I think it gives a better feeling for what is going on and how compactness arguments are used in practice. I only prove the maximum part and leave the minimum as an exercise.

Let

$$M = \sup\{f(x) \mid x \in K\}$$

(as we don't yet know that f is bounded, we must consider the possibility that $M = \infty$) and choose a sequence $\{x_n\}$ in K such that $\lim_{n \rightarrow \infty} f(x_n) = M$. Since K is compact, $\{x_n\}$ has a subsequence $\{x_{n_k}\}$ converging to a point $c \in K$. Then on the one hand $\lim_{k \rightarrow \infty} f(x_{n_k}) = M$, and on the other $\lim_{k \rightarrow \infty} f(x_{n_k}) = f(c)$ according to Proposition 3.2.9. Hence $f(c) = M$, and since $M = \sup\{f(x) \mid x \in K\}$, we see that c is a maximum point for f on K . \square

Remark: As already mentioned, it is possible to give a shorter proof of the Extreme Value Theorem by using Proposition 3.5.9. According to the proposition, the set $f(K)$ is compact and thus closed and bounded. This means that $\sup f(K)$ and $\inf f(K)$ belong to $f(K)$, and hence there are points $c, d \in K$ such that $f(c) = \sup f(K)$ and $f(d) = \inf f(K)$. Clearly, c is a maximum and d a minimum point for f .

Let us finally turn to the description of compactness in terms of total boundedness.

Definition 3.5.11. A subset A of a metric space X is called *totally bounded* if for each $\epsilon > 0$ there is a finite number $B(a_1, \epsilon), B(a_2, \epsilon), \dots, B(a_n, \epsilon)$ of balls with centers in A and radius ϵ that cover A (i.e. $A \subseteq B(a_1, \epsilon) \cup B(a_2, \epsilon) \cup \dots \cup B(a_n, \epsilon)$).

We first observe that a compact set is always totally bounded.

Proposition 3.5.12. Let K be a compact subset of a metric space X . Then K is totally bounded.

Proof. We argue contrapositively: Assume that A is *not* totally bounded, then there is an $\epsilon > 0$ such that no finite collection of ϵ -balls cover A . We shall construct a sequence $\{x_n\}$ in A that does not have a convergent subsequence. We begin by choosing an arbitrary element $x_1 \in A$. Since $B(x_1, \epsilon)$ does not cover A , we can choose $x_2 \in A \setminus B(x_1, \epsilon)$. Since $B(x_1, \epsilon)$ and $B(x_2, \epsilon)$ do not cover A , we can choose $x_3 \in A \setminus (B(x_1, \epsilon) \cup B(x_2, \epsilon))$. Continuing in this way, we get a sequence $\{x_n\}$ such that

$$x_n \in A \setminus (B(x_1, \epsilon) \cup B(x_2, \epsilon) \cup \dots \cup B(x_{n-1}, \epsilon))$$

This means that $d(x_n, x_m) \geq \epsilon$ for all $n, m \in \mathbb{N}$, $n \neq m$, and hence $\{x_n\}$ has no convergent subsequence. \square

We are now ready for the final theorem. Note that we have now added the assumption that X is complete – without this condition, the statement is false (see Exercise 10).

Theorem 3.5.13. A subset A of a complete metric space X is compact if and only if it is closed and totally bounded.

Proof. As we already know that a compact set is closed and totally bounded, it suffices to prove that a closed and totally bounded set A is compact. Let $\{x_n\}$ be a sequence in A . Our aim is to construct a convergent subsequence $\{x_{n_k}\}$. Choose balls $B_1^1, B_2^1, \dots, B_{k_1}^1$ of radius one that cover A . At least one of these balls must contain infinitely many terms from the sequence. Call this ball S_1 (if there are more than one such ball, just choose one). We now choose balls $B_1^2, B_2^2, \dots, B_{k_2}^2$ of radius $\frac{1}{2}$ that cover A . At least one of these ball must contain infinitely many of the terms from the sequence that lie in S_1 . If we call this ball S_2 , $S_1 \cap S_2$ contains infinitely many terms from the sequence. Continuing in this way, we find a sequence of balls S_k of radius $\frac{1}{k}$ such that

$$S_1 \cap S_2 \cap \dots \cap S_k$$

always contains infinitely many terms from the sequence.

We can now construct a convergent subsequence of $\{x_n\}$. Choose n_1 to be the first number such that x_{n_1} belongs to S_1 . Choose n_2 to be first number larger than n_1 such that x_{n_2} belongs to $S_1 \cap S_2$, then choose n_3 to be the first number larger than n_2 such that x_{n_3} belongs to $S_1 \cap S_2 \cap S_3$. Continuing in this way, we get a subsequence $\{x_{n_k}\}$ such that

$$x_{n_k} \in S_1 \cap S_2 \cap \dots \cap S_k$$

for all k . Since the S_k 's are shrinking, $\{x_{n_k}\}$ is a Cauchy sequence, and since X is complete, $\{x_{n_k}\}$ converges to a point a . Since A is closed, $a \in A$. Hence we have proved that any sequence in A has a subsequence converging to a point in A , and thus A is compact. \square

In the next section, we shall study yet another way to describe compact sets.

Problems to Section 3.5

1. Show that a space (X, d) with the discrete metric is compact if and only if X is a finite set.
2. Prove Proposition 3.5.1.
3. Prove the minimum part of Theorem 3.5.10.
4. Let A be a subset of a metric space X .
 - a) Show that if A is bounded, then for every point $c \in X$ there is a constant M_c such that $d(a, c) \leq M_c$ for all $a \in A$.
 - b) Assume that there is a point $c \in X$ and a number $M \in \mathbb{R}$ such that $d(a, c) \leq M$ for all $a \in A$. Show that A is bounded.
5. Let (X, d) be a metric space. For a subset A of X , let ∂A denote the set of all boundary points of A . Recall that the *closure* of A is the set $\bar{A} = A \cup \partial A$.
 - a) A subset A of X is called *precompact* if \bar{A} is compact. Show that A is precompact if and only if all sequences in A have convergent subsequences.
 - b) Show that a subset of \mathbb{R}^m is precompact if and only if it is bounded.
6. Assume that (X, d) is a metric space and that $f: X \rightarrow [0, \infty)$ is a continuous function. Assume that for each $\epsilon > 0$, there is a compact set $K_\epsilon \subseteq X$ such that $f(x) < \epsilon$ when $x \notin K_\epsilon$. Show that f has a maximum point.
7. Let (X, d) be a compact metric space, and assume that $f: X \rightarrow \mathbb{R}$ is continuous when we give \mathbb{R} the usual metric. Show that if $f(x) > 0$ for all $x \in X$, then there is a positive, real number a such that $f(x) > a$ for all $x \in X$.

8. Assume that $f: X \rightarrow Y$ is a continuous function between metric spaces, and let K be a compact subset of Y . Show that $f^{-1}(K)$ is closed. Find an example which shows that $f^{-1}(K)$ need not be compact.
9. Show that a totally bounded subset of a metric space is always bounded. Find an example of a bounded set in a metric space that is not totally bounded.
10. Let $d(x, y) = |x - y|$ be the usual metric on \mathbb{Q} . Let $A = \{q \in \mathbb{Q} \mid 0 \leq q \leq 2\}$. Show that A is a closed and totally bounded subset of \mathbb{Q} , but that A is not compact.
11. A metric space (X, d) is *locally compact* if there for each $a \in X$ is an $r > 0$ such that the closed ball $\overline{B}(a; r) = \{x \in X : d(a, x) \leq r\}$ is compact.
 - a) Show that \mathbb{R}^n is locally compact.
 - b) Show that if $X = \mathbb{R} \setminus \{0\}$, and $d: X \rightarrow \mathbb{R}$ is the metric defined by $d(x, y) = |x - y|$, then (X, d) is locally compact, but not complete.
 - c) As all compact spaces are complete, and convergence is a local property, it is easy to think that all locally compact spaces must also be complete, but the example in b) shows that this is not the case. What is wrong with the following argument for that all locally compact spaces X are complete?
We must show that all Cauchy sequences $\{x_n\}$ in X converge. Since we can get the distance $d(x_n, x_m)$ as small as we want by choosing n and m large enough, we can find an $r > 0$ such that $\overline{B}(x_n, r)$ is compact, and $x_m \in \overline{B}(x_n, r)$ for all $m \geq n$. Hence $\{x_m\}_{m \geq n}$ has a subsequence converging to a point $a \in \overline{B}(x_n, r)$, and this sequence is also a subsequence of $\{x_n\}$. Thus the Cauchy sequence $\{x_n\}$ has a subsequence converging to a , and hence $\{x_n\}$ also converges to a .
12. Let (X, d) be a metric space.
 - a) Assume that K_1, K_2, \dots, K_n is a finite collection of compact subsets of X . Show that the union $K_1 \cup K_2 \cup \dots \cup K_n$ is compact.
 - b) Assume that \mathcal{K} is a collection of compact subset of X . Show that the intersection $\bigcap_{K \in \mathcal{K}} K$ is compact.
13. Let (X, d) be a metric space. Assume that $\{K_n\}$ is a sequence of non-empty, compact subsets of X such that $K_1 \supseteq K_2 \supseteq \dots \supseteq K_n \supseteq \dots$. Prove that $\bigcap_{n \in \mathbb{N}} K_n$ is non-empty.
14. Let (X, d_X) and (Y, d_Y) be two metric spaces. Assume that (X, d_X) is compact, and that $f: X \rightarrow Y$ is bijective and continuous. Show that the inverse function $f^{-1}: Y \rightarrow X$ is continuous.
15. Assume that C and K are disjoint, compact subsets of a metric space (X, d) , and define

$$a = \inf\{d(x, y) \mid x \in C, y \in K\}$$

Show that a is strictly positive and that there are points $x_0 \in C$, $y_0 \in K$ such that $d(x_0, y_0) = a$. Show by an example that the result does not hold if we only assume that one of the sets C and K is compact and the other one closed.

16. Assume that (X, d) is compact and that $f: X \rightarrow X$ is continuous.
 - a) Show that the function $g(x) = d(x, f(x))$ is continuous and has a minimum point.
 - b) Assume in addition that $d(f(x), f(y)) < d(x, y)$ for all $x, y \in X$, $x \neq y$. Show that f has a unique fixed point. (*Hint: Use the minimum from a)*)

3.6. An alternative description of compactness

The descriptions of compactness that we studied in the previous section, suffice for most purposes in this book, but for some of the more advanced proofs there is another description that is more convenient. This alternative description is also the right one to use if one wants to extend the concept of compactness to even more general spaces, so-called *topological spaces*. In such spaces, sequences are not always an efficient tool, and it is better to have a description of compactness in terms of coverings by open sets.

To see what this means, assume that K is a subset of a metric space X . An *open covering* of K is simply a (finite or infinite) collection \mathcal{O} of open sets whose union contains K , i.e.

$$K \subseteq \bigcup \{O : O \in \mathcal{O}\}$$

The purpose of this section is to show that in metric spaces, the following property is equivalent to compactness.

Definition 3.6.1 (Open Covering Property). *Let K be a subset of a metric space X . Assume that for every open covering \mathcal{O} of K , there is a finite number of elements O_1, O_2, \dots, O_n in \mathcal{O} such that*

$$K \subseteq O_1 \cup O_2 \cup \dots \cup O_n$$

(we say that each open covering of K has a finite subcovering). Then the set K is said to have the open covering property.

The open covering property is quite abstract and may take some time to get used to, but it turns out to be a very efficient tool. Note that the term “open covering property” is not standard terminology, and that it will disappear once we have proved that it is equivalent to compactness.

Let us first prove that a set with the open covering property is necessarily compact. Before we begin, we need a simple observation: Assume that x is a point in our metric space X , and that no subsequence of the sequence $\{x_n\}$ converges to x . Then there must be an open ball $B(x; r)$ around x which only contains finitely many terms from $\{x_n\}$ (because if all balls around x contained infinitely many terms, we could use these terms to construct a subsequence converging to x).

Proposition 3.6.2. *If a subset K of a metric space X has the open covering property, then it is compact.*

Proof. We argue contrapositively, i.e., we assume that K is *not* compact and prove that it does not have the open covering property. Since K is not compact, there is a sequence $\{x_n\}$ that does not have any subsequence converging to points in K . By the observation above, this means that for each element $x \in K$, there is an open ball $B(x; r_x)$ around x which only contains finitely many terms of the sequence. The family $\{B(x, r_x) : x \in K\}$ is an open covering of K , but it cannot have a finite subcovering since any such subcovering $B(x_1, r_{x_1}), B(x_2, r_{x_2}), \dots, B(x_m, r_{x_m})$ only contains finitely many of the infinitely many terms in the sequence. \square

To prove the opposite implication, we shall use an elegant trick based on the Extreme Value Theorem, but first we need a lemma (the strange cut-off at 1 in the definition of $f(x)$ below is just to make sure that the function is finite):

Lemma 3.6.3. *Let \mathcal{O} be an open covering of a subset A of a metric space X . Define a function $f: A \rightarrow \mathbb{R}$ by*

$$f(x) = \sup\{r \in \mathbb{R} \mid r < 1 \text{ and } B(x; r) \subseteq O \text{ for some } O \in \mathcal{O}\}$$

Then f is continuous and strictly positive (i.e. $f(x) > 0$ for all $x \in A$).

Proof. The strict positivity is easy: Since \mathcal{O} is a covering of A , there is a set $O \in \mathcal{O}$ such that $x \in O$, and since O is open, there is an r , $0 < r < 1$, such that $B(x; r) \subseteq O$. Hence $f(x) \geq r > 0$.

To prove the continuity, it suffices to show that $|f(x) - f(y)| \leq d(x, y)$ as we can then choose $\delta = \epsilon$ in the definition of continuity. Observe first that if $f(x), f(y) \leq d(x, y)$, there is nothing to prove. Assume therefore that at least one of these values is larger than $d(x, y)$. Without out loss of generality, we may assume that $f(x)$ is the larger of the two. There must then be an $r > d(x, y)$ and an $O \in \mathcal{O}$ such that $B(x, r) \subseteq O$. For any such r , $B(y, r - d(x, y)) \subseteq O$ since $B(y, r - d(x, y)) \subset B(x, r)$. This means that $f(y) \geq f(x) - d(x, y)$. Since by assumption $f(x) \geq f(y)$, we have $|f(x) - f(y)| \leq d(x, y)$ which is what we set out to prove. \square

We are now ready for the main theorem (some authors refer to this as the *Heine-Borel Theorem*, while others only use this label when the metric space is \mathbb{R} or \mathbb{R}^n ; see Exercise 1):

Theorem 3.6.4. *A subset K of a metric space is compact if and only if it has the open covering property.*

Proof. It remains to prove that if K is compact and \mathcal{O} is an open covering of K , then \mathcal{O} has a finite subcovering. By the Extreme Value Theorem 3.5.10, the function f in the lemma attains a minimal value r on K , and since f is strictly positive, $r > 0$. This means that for all $x \in K$, the ball $B(x, \frac{r}{2})$ is contained in a set $O \in \mathcal{O}$. Since K is compact, it is totally bounded, and hence there is a finite collection of balls $B(x_1, \frac{r}{2}), B(x_2, \frac{r}{2}), \dots, B(x_n, \frac{r}{2})$ that covers K . Each ball $B(x_i, \frac{r}{2})$ is contained in a set $O_i \in \mathcal{O}$, and hence O_1, O_2, \dots, O_n is a finite subcovering of \mathcal{O} . \square

As usual, there is a reformulation of the theorem above in terms of closed sets. Let us first agree to say that a collection \mathcal{F} of sets has the *finite intersection property over K* if

$$K \cap F_1 \cap F_2 \cap \dots \cap F_n \neq \emptyset$$

for all finite collections F_1, F_2, \dots, F_n of sets from \mathcal{F} .

Corollary 3.6.5. *Assume that K is a subset of a metric space X . Then the following are equivalent:*

- (i) *K is compact.*

(ii) If a collection \mathcal{F} of closed sets has the finite intersection property over K , then

$$K \cap \left(\bigcap_{F \in \mathcal{F}} F \right) \neq \emptyset$$

Proof. Left to the reader (see Exercise 7). \square

Problems to Section 3.6

1. Assume that \mathcal{I} is a collection of open intervals in \mathbb{R} whose union contains $[0, 1]$. Show that there exists a finite collection I_1, I_2, \dots, I_n of sets from \mathcal{I} such that

$$[0, 1] \subseteq I_1 \cup I_2 \cup \dots \cup I_n$$

This is sometimes called the *Heine-Borel Theorem*.

2. Let $\{K_n\}$ be a decreasing sequence (i.e., $K_{n+1} \subseteq K_n$ for all $n \in \mathbb{N}$) of nonempty, compact sets. Show that $\bigcap_{n \in \mathbb{N}} K_n \neq \emptyset$. (This exactly the same problem as 3.5.13, but this time you should do it with the methods in this section).
3. Assume that $f: X \rightarrow Y$ is a continuous function between two metric spaces. Use the open covering property to show that if K is a compact subset of X , then $f(K)$ is a compact subset of Y .
4. Assume that K_1, K_2, \dots, K_n are compact subsets of a metric space X . Use the open covering property to show that $K_1 \cup K_2 \cup \dots \cup K_n$ is compact.
5. Use the open covering property to show that a closed subset of a compact set is compact.
6. Assume that $f: X \rightarrow Y$ is a continuous function between two metric spaces, and assume that K is a compact subset of X . We shall prove that f is *uniformly continuous* on K , i.e. that for each $\epsilon > 0$, there exists a $\delta > 0$ such that whenever $x, y \in K$ and $d_X(x, y) < \delta$, then $d_Y(f(x), f(y)) < \epsilon$ (this looks very much like ordinary continuity, but the point is that we can use the *same* δ at all points $x, y \in K$).
 - a) Given $\epsilon > 0$, explain that for each $x \in K$ there is a $\delta(x) > 0$ such that $d_Y(f(x), f(y)) < \frac{\epsilon}{2}$ for all y with $d(x, y) < \delta(x)$.
 - b) Explain that $\{B(x, \frac{\delta(x)}{2})\}_{x \in K}$ is an open covering of K , and that it has a finite subcovering $B(x_1, \frac{\delta(x_1)}{2}), B(x_2, \frac{\delta(x_2)}{2}), \dots, B(x_n, \frac{\delta(x_n)}{2})$.
 - c) Put $\delta = \min\{\frac{\delta(x_1)}{2}, \frac{\delta(x_2)}{2}, \dots, \frac{\delta(x_n)}{2}\}$, and show that if $x, y \in K$ with $d_X(x, y) < \delta$, then $d_Y(f(x), f(y)) < \epsilon$.
7. Prove Corollary 3.6.5. (*Hint:* Observe that $K \cap (\bigcap_{F \in \mathcal{F}} F) \neq \emptyset$ if and only if $\{F^c\}_{F \in \mathcal{F}}$ is an open covering of K .)

3.7. The completion of a metric space

Completeness is probably the most important notion in this book as most of the deep and interesting theorems about metric spaces only hold when the space is complete. In this section we shall see that it is always possible to make an incomplete space complete by adding new elements. Although this is definitely an interesting result from a philosophical perspective, it will not be needed later in the book, and you may skip this section if you want (it becomes quite technical after a while).

To describe completions, we need the following concept:

Definition 3.7.1. Let (X, d) be a metric space and assume that D is a subset of X . We say that D is dense in X if for each $x \in X$ there is a sequence $\{y_n\}$ from D converging to x .

We know that \mathbb{Q} is dense in \mathbb{R} – we may, e.g., approximate a real number by longer and longer parts of its decimal expansion. For $x = \sqrt{2}$ this would mean the approximating sequence

$$y_1 = 1.4 = \frac{14}{10}, \quad y_2 = 1.41 = \frac{141}{100}, \quad y_3 = 1.414 = \frac{1414}{1000}, \quad y_4 = 1.4142 = \frac{14142}{10000}, \dots$$

There is an alternative description of dense that we shall also need.

Proposition 3.7.2. A subset D of a metric space X is dense if and only if for each $x \in X$ and each $\delta > 0$, there is a $y \in D$ such that $d(x, y) \leq \delta$.

Proof. Left as an exercise. □

We can now return to our initial problem: How do we extend an incomplete metric space to a complete one? The following definition describes what we are looking for.

Definition 3.7.3. If (X, d_X) is a metric space, a completion of (X, d_X) is a metric space $(\bar{X}, d_{\bar{X}})$ such that:

- (i) (X, d_X) is a subspace of $(\bar{X}, d_{\bar{X}})$; i.e. $X \subseteq \bar{X}$ and $d_{\bar{X}}(x, y) = d_X(x, y)$ for all $x, y \in X$.
- (ii) X is dense $(\bar{X}, d_{\bar{X}})$.

The canonical example of a completion is that \mathbb{R} is the completion \mathbb{Q} . We also note that a complete metric space is its own (unique) completion.

An incomplete metric space will have more than one completion, but as they are all isometric³, they are the same for most practical purposes, and we usually talk about *the* completion of a metric space.

Proposition 3.7.4. Assume that (Y, d_Y) and (Z, d_Z) are completions of the metric space (X, d_X) . Then (Y, d_Y) and (Z, d_Z) are isometric.

Proof. We shall construct an isometry $i: Y \rightarrow Z$. Since X is dense in Y , there is for each $y \in Y$ a sequence $\{x_n\}$ from X converging to y . This sequence must be a Cauchy sequence in X and hence in Z . Since Z is complete, $\{x_n\}$ converges to an element $z \in Z$. The idea is to define i by letting $i(y) = z$. For the definition to work properly, we have to check that if $\{\hat{x}_n\}$ is another sequence in X converging to y , then $\{\hat{x}_n\}$ converges to z in Z . This is the case since $d_Z(x_n, \hat{x}_n) = d_X(x_n, \hat{x}_n) = d_Y(x_n, \hat{x}_n) \rightarrow 0$ as $n \rightarrow \infty$.

To prove that i preserves distances, assume that y, \hat{y} are two points in Y , and that $\{x_n\}, \{\hat{x}_n\}$ are sequences in X converging to y and \hat{y} , respectively. Then $\{x_n\}, \{\hat{x}_n\}$ converges to $i(y)$ and $i(\hat{y})$, respectively, in Z , and we have

$$d_Z(i(y), i(\hat{y})) = \lim_{n \rightarrow \infty} d_Z(x_n, \hat{x}_n) = \lim_{n \rightarrow \infty} d_X(x_n, \hat{x}_n)$$

³Recall from Section 3.1 that an *isometry* from (X, d_X) to (Y, d_Y) is a bijection $i: X \rightarrow Y$ such that $d_Y(i(x), i(y)) = d_X(x, y)$ for all $x, y \in X$. Two metric spaces are often considered “the same” when they are isometric; i.e. when there is an isometry between them.

$$= \lim_{n \rightarrow \infty} d_Y(x_n, \hat{x}_n) = d_Y(y, \hat{y})$$

(we are using repeatedly that if $\{u_n\}$ and $\{v_n\}$ are sequences in a metric space converging to u and v , respectively, then $d(u_n, v_n) \rightarrow d(u, v)$, see Exercise 3.2.8 b). It remains to prove that i is a bijection. Injectivity follows immediately from distance preservation: If $y \neq \hat{y}$, then $d_Z(i(y), i(\hat{y})) = d_Y(y, \hat{y}) \neq 0$, and hence $i(y) \neq i(\hat{y})$. To show that i is surjective, consider an arbitrary element $z \in Z$. Since X is dense in Z , there is a sequence $\{x_n\}$ from X converging to z . Since Y is complete, $\{x_n\}$ is also converging to an element y in Y . By construction, $i(y) = z$, and hence i is surjective. \square

We shall use the rest of the section to show that all metric spaces (X, d) have a completion. As the construction is longer and more complicated than most others in this book, I'll give you a brief preview first. We'll start with the set \mathcal{X} of all Cauchy sequences in X (this is only natural as what we want to do is add points to X such that all Cauchy sequences have something to converge to). Next we introduce an equivalence relation (recall Section 1.5) \sim on \mathcal{X} by defining

$$\{x_n\} \sim \{y_n\} \iff \lim_{n \rightarrow \infty} d(x_n, y_n) = 0$$

We let $[x_n]$ denote the equivalence class of the sequence $\{x_n\}$, and we let \bar{X} be the set of all equivalence classes. The next step is to introduce a metric \bar{d} on \bar{X} by defining

$$\bar{d}([x_n], [y_n]) = \lim_{n \rightarrow \infty} d(x_n, y_n)$$

The space (\bar{X}, \bar{d}) is our candidate for the completion of (X, d) . To prove that it works, we first observe that \bar{X} contains a copy D of the original space X : For each $x \in X$, let $\bar{x} = [x, x, x, \dots]$ be the equivalence class of the constant sequence $\{x, x, x, \dots\}$, and put

$$D = \{\bar{x} \mid x \in X\}$$

We then prove that D is dense in \bar{X} and that \bar{X} is complete. Finally, we can replace each element \bar{x} in D by the original element $x \in X$, and we have our completion.

So let us begin the work. The first lemma gives us the information we need to get started.

Lemma 3.7.5. *Assume that $\{x_n\}$ and $\{y_n\}$ are two Cauchy sequences in a metric space (X, d) . Then $\lim_{n \rightarrow \infty} d(x_n, y_n)$ exists.*

Proof. As \mathbb{R} is complete, it suffices to show that $\{d(x_n, y_n)\}$ is a Cauchy sequence. We have

$$\begin{aligned} |d(x_n, y_n) - d(x_m, y_m)| &= |d(x_n, y_n) - d(x_m, y_n) + d(x_m, y_n) - d(x_m, y_m)| \leq \\ &\leq |d(x_n, y_n) - d(x_m, y_n)| + |d(x_m, y_n) - d(x_m, y_m)| \leq d(x_n, x_m) + d(y_n, y_m) \end{aligned}$$

where we have used the inverse triangle inequality 3.1.4 in the final step. Since $\{x_n\}$ and $\{y_n\}$ are Cauchy sequences, we can get $d(x_n, x_m)$ and $d(y_n, y_m)$ as small as we wish by choosing n and m sufficiently large, and hence $\{d(x_n, y_n)\}$ is a Cauchy sequence. \square

As mentioned above, we let \mathcal{X} be the set of all Cauchy sequences in the metric space (X, d_X) , and we introduce a relation \sim on \mathcal{X} by

$$\{x_n\} \sim \{y_n\} \iff \lim_{n \rightarrow \infty} d(x_n, y_n) = 0$$

Lemma 3.7.6. *\sim is an equivalence relation.*

Proof. We have to check the three properties in Definition 1.5.2:

- (i) *Reflexivity:* Since $\lim_{n \rightarrow \infty} d(x_n, x_n) = 0$, the relation is reflexive.
- (ii) *Symmetry:* Since $\lim_{n \rightarrow \infty} d(x_n, y_n) = \lim_{n \rightarrow \infty} d(y_n, x_n)$, the relation is symmetric.
- (iii) *Transitivity:* Assume that $\{x_n\} \sim \{y_n\}$ and $\{y_n\} \sim \{z_n\}$. Then

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

and consequently

$$\begin{aligned} 0 &\leq \lim_{n \rightarrow \infty} d(x_n, z_n) \leq \lim_{n \rightarrow \infty} (d(x_n, y_n) + d(y_n, z_n)) = \\ &= \lim_{n \rightarrow \infty} d(x_n, y_n) + \lim_{n \rightarrow \infty} d(y_n, z_n) = 0 \end{aligned}$$

which shows that $\{x_n\} \sim \{z_n\}$, and hence the relation is transitive. \square

We denote the equivalence class of $\{x_n\}$ by $[x_n]$, and we let \bar{X} be the set of all equivalence classes. The next lemma will allow us to define a natural metric on \bar{X} .

Lemma 3.7.7. *If $\{x_n\} \sim \{\hat{x}_n\}$ and $\{y_n\} \sim \{\hat{y}_n\}$, then $\lim_{n \rightarrow \infty} d(x_n, y_n) = \lim_{n \rightarrow \infty} d(\hat{x}_n, \hat{y}_n)$.*

Proof. Since $d(x_n, y_n) \leq d(x_n, \hat{x}_n) + d(\hat{x}_n, \hat{y}_n) + d(\hat{y}_n, y_n)$ by the triangle inequality, and $\lim_{n \rightarrow \infty} d(x_n, \hat{x}_n) = \lim_{n \rightarrow \infty} d(\hat{y}_n, y_n) = 0$, we get

$$\lim_{n \rightarrow \infty} d(x_n, y_n) \leq \lim_{n \rightarrow \infty} d(\hat{x}_n, \hat{y}_n)$$

By reversing the rôles of elements with and without hats, we get the opposite inequality. \square

We may now define a function $\bar{d}: \bar{X} \times \bar{X} \rightarrow [0, \infty)$ by

$$\bar{d}([x_n], [y_n]) = \lim_{n \rightarrow \infty} d(x_n, y_n)$$

Note that by the previous lemma \bar{d} is *well-defined*; i.e. the value of $\bar{d}([x_n], [y_n])$ does not depend on which representatives $\{x_n\}$ and $\{y_n\}$ we choose from the equivalence classes $[x_n]$ and $[y_n]$.

We have reached our first goal:

Lemma 3.7.8. *(\bar{X}, \bar{d}) is a metric space.*

Proof. We need to check the three conditions in the definition of a metric space.

- (i) *Positivity:* Clearly $\bar{d}([x_n], [y_n]) = \lim_{n \rightarrow \infty} d(x_n, y_n) \geq 0$, and by definition of the equivalence relation, we have equality if and only if $[x_n] = [y_n]$.

(ii) *Symmetry*: Since the underlying metric d is symmetric, we have

$$\bar{d}([x_n], [y_n]) = \lim_{n \rightarrow \infty} d(x_n, y_n) = \lim_{n \rightarrow \infty} d(y_n, x_n) = \bar{d}([y_n], [x_n])$$

(iii) *Triangle inequality*: For all equivalence classes $[x_n], [y_n], [z_n]$, we have

$$\begin{aligned} \bar{d}([x_n], [z_n]) &= \lim_{n \rightarrow \infty} d(x_n, z_n) \leq \lim_{n \rightarrow \infty} d(x_n, y_n) + \lim_{n \rightarrow \infty} d(y_n, z_n) = \\ &= \bar{d}([x_n], [y_n]) + \bar{d}([y_n], [z_n]) \end{aligned} \quad \square$$

For each $x \in X$, let \bar{x} be the equivalence class of the constant sequence $\{x, x, x, \dots\}$. Since $\bar{d}(\bar{x}, \bar{y}) = \lim_{n \rightarrow \infty} d(x, y) = d(x, y)$, the mapping $x \rightarrow \bar{x}$ is an embedding (recall Definition 3.1.3) of X into \bar{X} . Hence \bar{X} contains a copy of X , and the next lemma shows that this copy is dense in \bar{X} .

Lemma 3.7.9. *The set*

$$D = \{\bar{x} : x \in X\}$$

is dense in \bar{X} .

Proof. Assume that $[x_n] \in \bar{X}$. By Proposition 3.7.2, it suffices to show that for each $\epsilon > 0$ there is an $\bar{x} \in D$ such that $\bar{d}(\bar{x}, [x_n]) < \epsilon$. Since $\{x_n\}$ is a Cauchy sequence, there is an $N \in \mathbb{N}$ such that $d(x_n, x_N) < \frac{\epsilon}{2}$ for all $n \geq N$. Put $x = x_N$. Then $\bar{d}([x_n], \bar{x}) = \lim_{n \rightarrow \infty} d(x_n, x_N) \leq \frac{\epsilon}{2} < \epsilon$. \square

It still remains to prove that (\bar{X}, \bar{d}) is complete. The next lemma is the first step in this direction.

Lemma 3.7.10. *Every Cauchy sequence in D converges to an element in \bar{X} .*

Proof. If $\{\bar{u}_k\}$ is a Cauchy sequence in D , then $\{u_k\}$ is a Cauchy sequence in X and gives rise to an element $[u_k]$ in \bar{X} . In order to avoid confusion later, we relabel this element $[u_n]$ – the name of the index doesn't matter. We need to check that $\{\bar{u}_k\}$ converges to $[u_n]$. Since $\{u_k\}$ is a Cauchy sequence, there is for every $\epsilon > 0$ an $N \in \mathbb{N}$ such that $d(u_k, u_n) < \frac{\epsilon}{2}$ whenever $k, n \geq N$. For $k \geq N$, we thus have $\bar{d}(\bar{u}_k, [u_n]) = \lim_{m \rightarrow \infty} d(u_k, u_m) \leq \frac{\epsilon}{2} < \epsilon$. As $\epsilon > 0$ is arbitrary, this means that $\{\bar{u}_k\} \rightarrow [u_n]$. \square

The lemma above isn't enough to conclude that \bar{X} is complete as \bar{X} will have “new” Cauchy sequences that don't correspond to Cauchy sequences in X . However, since D is dense, this is not a big problem as the following observation shows:

Lemma 3.7.11. *Assume that (Y, d) is a metric space and that D is a dense subset of Y . If all Cauchy sequences $\{z_n\}$ in D converges (to an element in Y), then (Y, d) is complete.*

Proof. Let $\{y_n\}$ be a Cauchy sequence in Y . Since D is dense in Y , there is for each n an element $z_n \in D$ such that $d(z_n, y_n) < \frac{1}{n}$. It is easy to check that since $\{y_n\}$ is a Cauchy sequence, so is $\{z_n\}$. By assumption, $\{z_n\}$ converges to an element in Y , and by construction $\{y_n\}$ must converge to the same element. Hence (Y, d) is complete. \square

As we can now conclude that (\bar{X}, \bar{d}) is complete, we have reached the main theorem.

Theorem 3.7.12. *Every metric space (X, d) has a completion.*

Proof. We have already proved that (\bar{X}, \bar{d}) is a complete metric space that contains $D = \{\bar{x} : x \in X\}$ as a dense subset. In addition, we know that D is a copy of X (more precisely, $x \rightarrow \bar{x}$ is an isometry from X to D). All we have to do, is to replace the elements \bar{x} in D by the original elements x in X , and we have found a completion of X . \square

Remark: The theorem above doesn't solve all problems with incomplete spaces as there may be additional structure we want the completion to reflect. If, e.g., the original space consists of functions, we may want the completion also to consist of functions, but there is nothing in the construction above that guarantees that this is possible. We shall return to this question in later chapters.

Problems to Section 3.7

1. Prove Proposition 3.7.2.
2. Let us write $(X, d_X) \sim (Y, d_Y)$ to indicate that the two spaces are isometric. Show that
 - (i) $(X, d_X) \sim (X, d_X)$
 - (ii) If $(X, d_X) \sim (Y, d_Y)$, then $(Y, d_Y) \sim (X, d_X)$
 - (iii) If $(X, d_X) \sim (Y, d_Y)$ and $(Y, d_Y) \sim (Z, d_Z)$, then $(X, d_X) \sim (Z, d_Z)$.
3. Show that the only completion of a complete metric space is the space itself.
4. Show that \mathbb{R} is the completion of \mathbb{Q} (in the usual metrics).
5. Assume that $i: X \rightarrow Y$ is an isometry between two metric spaces (X, d_X) and (Y, d_Y) .
 - (i) Show that a sequence $\{x_n\}$ converges in X if and only if $\{i(x_n)\}$ converges in Y .
 - (ii) Show that a set $A \subseteq X$ is open, closed, or compact if and only if $i(A)$ is open, closed, or compact.

Notes and references to Chapter 3

The notion of a metric space was introduced by Maurice Fréchet (1878-1973) in his doctoral thesis from 1906. It may be seen as part of a general movement at the time to try to extract and isolate the crucial ingredients of mathematical theories. In 1914, Felix Hausdorff (1868-1942) generalized the concept even further and introduced what today is known as a *Hausdorff space*, a special kind of topological space. This development toward abstraction reached its peak with an extremely influential group of (mainly French) mathematicians who wrote a long series of books *Éléments de Mathématique* under the pen name of Nicolas Bourbaki.

Fréchet also introduced compactness in terms of convergence of subsequences. The more abstract description in terms of open coverings was introduced by Pavel Sergeyevich Alexandrov (1896-1982) and Pavel Samuilovich Urysohn (1898-1924) in 1923.

Banach's Fixed Point Theorem was proved by Stefan Banach (1892-1945) in 1922. In addition to the applications you will meet later in the book, it plays a

central part in the theory of iterated function systems and fractals. See Barnsley's book [4] for a colorful presentation with lots of pictures.

Many books on real analysis have chapters on metric spaces. If you want to take a look at other presentations, you might try the books by Körner [21] and Tao [38] (in the latter case you'll get many interesting insights, but you'll have to make most of the proofs yourself!). If you want to take one step further in abstraction and look at topological spaces, Munkres' book [30] is a very readable introduction.

Most of the mathematics we shall discuss in this book was developed in the first half of the 20th century, but to a large extent as a reaction to problems and challenges raised by mathematicians of the 19th century. Gray's book [14] on analysis in the 19th century will give you an excellent introduction to these questions, and it will also provide you with a better understanding of why the theory looks the way it does. One of the lessons to be learned from the history of mathematics, is the importance of precise definitions – Gray's volume is (like any other serious treatise on the history of mathematics) full of discussions of what the mathematicians of the past really meant by the concepts they introduced and discussed. Modern definitions of the ϵ - δ kind may look unnecessarily complicated when you meet them in a calculus course, but they have the great advantage of being precise!

Spaces of Continuous Functions

In this chapter we shall apply the theory we developed in the previous chapter to spaces where the elements are functions. We shall study completeness and compactness of such spaces and take a look at some applications. But before we turn to these spaces, it will be useful to take a look at different notions of continuity and convergence and what they can be used for.

4.1. Modes of continuity

If (X, d_X) and (Y, d_Y) are two metric spaces, the function $f: X \rightarrow Y$ is continuous at a point a if for each $\epsilon > 0$ there is a $\delta > 0$ such that $d_Y(f(x), f(a)) < \epsilon$ whenever $d_X(x, a) < \delta$. If f is also continuous at another point b , we may need a different δ to match the same ϵ . A question that often comes up is when we can use the same δ for *all* points x in the space X . The function is then said to be *uniformly continuous* in X . Here is the precise definition:

Definition 4.1.1. *Let $f: X \rightarrow Y$ be a function between two metric spaces. We say that f is uniformly continuous if for each $\epsilon > 0$ there is a $\delta > 0$ such that for all points $x, y \in X$ with $d_X(x, y) < \delta$, we have $d_Y(f(x), f(y)) < \epsilon$.*

A function which is continuous at all points in X , but not uniformly continuous, is often called *pointwise continuous* when we want to emphasize the distinction.

Example 1: The function $f: \mathbb{R} \rightarrow \mathbb{R}$ defined by $f(x) = x^2$ is pointwise continuous, but not uniformly continuous. The reason is that the curve becomes steeper and steeper as $|x|$ goes to infinity, and that we hence need increasingly smaller δ 's to match the same ϵ (make a sketch!) See Exercise 1 for a more detailed discussion. ♣

If the underlying space X is compact, pointwise continuity and uniform continuity are the same. This means, e.g., that a continuous function defined on a closed and bounded subset of \mathbb{R}^n is always uniformly continuous.

Proposition 4.1.2. *Assume that X and Y are metric spaces. If X is compact, all continuous functions $f: X \rightarrow Y$ are uniformly continuous.*

Proof. We argue contrapositively: Assume that f is *not* uniformly continuous; we shall show that f is not continuous.

Since f fails to be uniformly continuous, there is an $\epsilon > 0$ we cannot match; i.e., for each $\delta > 0$ there are points $x, y \in X$ such that $d_X(x, y) < \delta$, but $d_Y(f(x), f(y)) \geq \epsilon$. Choosing $\delta = \frac{1}{n}$, there are thus points $x_n, y_n \in X$ such that $d_X(x_n, y_n) < \frac{1}{n}$ and $d_Y(f(x_n), f(y_n)) \geq \epsilon$. Since X is compact, the sequence $\{x_n\}$ has a subsequence $\{x_{n_k}\}$ converging to a point a . Since $d_X(x_{n_k}, y_{n_k}) < \frac{1}{n_k}$, the corresponding sequence $\{y_{n_k}\}$ of y 's must also converge to a . We are now ready to show that f is not continuous at a : Had it been, the two sequences $\{f(x_{n_k})\}$ and $\{f(y_{n_k})\}$ would both have converged to $f(a)$ according to Proposition 3.2.5, something they clearly cannot since $d_Y(f(x_n), f(y_n)) \geq \epsilon$ for all $n \in \mathbb{N}$. \square

There is an even more abstract form of continuity that will be important later. This time we are not considering a single function, but a whole collection of functions:

Definition 4.1.3. *Let (X, d_X) and (Y, d_Y) be metric spaces, and let \mathcal{F} be a collection of functions $f: X \rightarrow Y$. We say that \mathcal{F} is equicontinuous if for all $\epsilon > 0$, there is a $\delta > 0$ such that for all $f \in \mathcal{F}$ and all $x, y \in X$ with $d_X(x, y) < \delta$, we have $d_Y(f(x), f(y)) < \epsilon$.*

Note that in the case, the same δ should not only hold at all points $x, y \in X$, but also for all functions $f \in \mathcal{F}$.

Example 2: Let \mathcal{F} be the set of all contractions $f: X \rightarrow X$. Then \mathcal{F} is equicontinuous, since we can choose $\delta = \epsilon$. To see this, just note that if $d_X(x, y) < \delta = \epsilon$, then $d_X(f(x), f(y)) \leq d_X(x, y) < \epsilon$ for all $x, y \in X$ and all $f \in \mathcal{F}$. \clubsuit

Equicontinuous families will be important when we study compact sets of continuous functions in Section 4.8.

Exercises for Section 4.1

1. Show that the function $f(x) = x^2$ is not uniformly continuous on \mathbb{R} . (*Hint:* You may want to use the factorization $f(x) - f(y) = x^2 - y^2 = (x + y)(x - y)$).
2. Prove that the function $f: (0, 1) \rightarrow \mathbb{R}$ given by $f(x) = \frac{1}{x}$ is not uniformly continuous.
3. A function $f: X \rightarrow Y$ between metric spaces is said to be *Lipschitz-continuous with Lipschitz constant K* if $d_Y(f(x), f(y)) \leq K d_X(x, y)$ for all $x, y \in X$. Assume that \mathcal{F} is a collection of functions $f: X \rightarrow Y$ with Lipschitz constant K . Show that \mathcal{F} is equicontinuous.
4. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function and assume that the derivative f' is bounded. Show that f is uniformly continuous.

4.2. Modes of convergence

In this section we shall study two ways in which a sequence $\{f_n\}$ of functions can converge to a limit function f : *pointwise convergence* and *uniform convergence*. The distinction is rather similar to the distinction between pointwise and uniform continuity in the previous section – in the pointwise case, a condition can be satisfied in different ways for different x 's; in the uniform case, it must be satisfied in the same way for all x . We begin with pointwise convergence:

Definition 4.2.1. Let (X, d_X) and (Y, d_Y) be two metric spaces, and let $\{f_n\}$ be a sequence of functions $f_n: X \rightarrow Y$. We say that $\{f_n\}$ converges pointwise to a function $f: X \rightarrow Y$ if $f_n(x) \rightarrow f(x)$ for all $x \in X$. This means that for each x and each $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $d_Y(f_n(x), f(x)) < \epsilon$ when $n \geq N$.

Note that the N in the last sentence of the definition depends on x – we may need a much larger N for some x 's than for others. If we can use the *same* N for all $x \in X$, we have uniform convergence. Here is the precise definition:

Definition 4.2.2. Let (X, d_X) and (Y, d_Y) be two metric spaces, and let $\{f_n\}$ be a sequence of functions $f_n: X \rightarrow Y$. We say that $\{f_n\}$ converges uniformly to a function $f: X \rightarrow Y$ if for each $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that if $n \geq N$, then $d_Y(f_n(x), f(x)) < \epsilon$ for all $x \in X$.

At first glance, the two definitions may seem confusingly similar, but the difference is that in the last one, the *same* N should work simultaneously for all x , while in the first we can adapt N to each individual x . Hence uniform convergence implies pointwise convergence, but a sequence may converge pointwise but not uniformly. Before we look at an example, it will be useful to reformulate the definition of uniform convergence.

Proposition 4.2.3. Let (X, d_X) and (Y, d_Y) be two metric spaces, and let $\{f_n\}$ be a sequence of functions $f_n: X \rightarrow Y$. For any function $f: X \rightarrow Y$ the following are equivalent:

- (i) $\{f_n\}$ converges uniformly to f .
- (ii) $\sup\{d_Y(f_n(x), f(x)) \mid x \in X\} \rightarrow 0$ as $n \rightarrow \infty$.

Hence uniform convergence means that the “maximal” distance between f and f_n goes to zero.

Proof. (i) \implies (ii) Assume that $\{f_n\}$ converges uniformly to f . For any $\epsilon > 0$, we can find an $N \in \mathbb{N}$ such that $d_Y(f_n(x), f(x)) < \epsilon$ for all $x \in X$ and all $n \geq N$. This means that $\sup\{d_Y(f_n(x), f(x)) \mid x \in X\} \leq \epsilon$ for all $n \geq N$ (note that we may have unstrict inequality \leq for the supremum although we have strict inequality $<$ for each $x \in X$), and since ϵ is arbitrary, this implies that $\sup\{d_Y(f_n(x), f(x)) \mid x \in X\} \rightarrow 0$.

(ii) \implies (i) Assume that $\sup\{d_Y(f_n(x), f(x)) \mid x \in X\} \rightarrow 0$ as $n \rightarrow \infty$. Given an $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that $\sup\{d_Y(f_n(x), f(x)) \mid x \in X\} < \epsilon$ for all $n \geq N$. But then we have $d_Y(f_n(x), f(x)) < \epsilon$ for all $x \in X$ and all $n \geq N$, which means that $\{f_n\}$ converges uniformly to f . \square

Here is an example which shows clearly the distinction between pointwise and uniform convergence:

Example 1: Let $f_n: [0, 1] \rightarrow \mathbb{R}$ be the function in Figure 4.2.1. It is constant zero except on the interval $[0, \frac{1}{n}]$ where it looks like a tent of height 1.




Figure 4.2.1. The functions f_n in Example 1.

If you insist, the function is defined by

$$f_n(x) = \begin{cases} 2nx & \text{if } 0 \leq x < \frac{1}{2n} \\ -2nx + 2 & \text{if } \frac{1}{2n} \leq x < \frac{1}{n} \\ 0 & \text{if } \frac{1}{n} \leq x \leq 1 \end{cases}$$

but it is much easier just to work from the picture.


The sequence $\{f_n\}$ converges pointwise to 0, because at every point $x \in [0, 1]$ the value of $f_n(x)$ eventually becomes 0 (for $x = 0$, the value is always 0, and for $x > 0$ the “tent” will eventually pass to the left of x .) However, since the maximum value of all f_n is 1, $\sup\{d_Y(f_n(x), 0) \mid x \in [0, 1]\} = 1$ for all n , and hence $\{f_n\}$ does not converge uniformly to 0. ♣

When we are working with convergent sequences, we would often like the limit to inherit properties from the elements in the sequence. If, e.g., $\{f_n\}$ is a sequence of *continuous* functions converging to a limit f , we are often interested in showing that f is also continuous. The next example shows that this is not always the case when we are dealing with pointwise convergence.

Example 2: Let $f_n: \mathbb{R} \rightarrow \mathbb{R}$ be the function in Figure 4.2.2. It is defined by

$$f_n(x) = \begin{cases} -1 & \text{if } x \leq -\frac{1}{n} \\ nx & \text{if } -\frac{1}{n} < x < \frac{1}{n} \\ 1 & \text{if } \frac{1}{n} \leq x \end{cases}$$

The sequence $\{f_n\}$ converges pointwise to the function, f defined by

Figure 4.2.2. The functions f_n in Example 2.

$$f(x) = \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ 1 & \text{if } x > 0 \end{cases}$$

but although all the functions $\{f_n\}$ are continuous, the limit function f is not. ♣

If we strengthen the convergence from pointwise to uniform, the limit of a sequence of continuous functions is always continuous.

Proposition 4.2.4. *Let (X, d_X) and (Y, d_Y) be two metric spaces, and assume that $\{f_n\}$ is a sequence of continuous functions $f_n: X \rightarrow Y$ converging uniformly to a function f . Then f is continuous.*

Proof. Let $a \in X$. Given an $\epsilon > 0$, we must find a $\delta > 0$ such that $d_Y(f(x), f(a)) < \epsilon$ whenever $d_X(x, a) < \delta$. Since $\{f_n\}$ converges uniformly to f , there is an $N \in \mathbb{N}$ such that when $n \geq N$, $d_Y(f(x), f_n(x)) < \frac{\epsilon}{3}$ for all $x \in X$. Since f_N is continuous at a , there is a $\delta > 0$ such that $d_Y(f_N(x), f_N(a)) < \frac{\epsilon}{3}$ whenever $d_X(x, a) < \delta$. If $d_X(x, a) < \delta$, we then have

$$\begin{aligned} d_Y(f(x), f(a)) &\leq d_Y(f(x), f_N(x)) + d_Y(f_N(x), f_N(a)) + d_Y(f_N(a), f(a)) < \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon \end{aligned}$$

and hence f is continuous at a . \square

The technique in the proof above is quite common, and arguments of this kind are often referred to as $\frac{\epsilon}{3}$ -arguments. It's quite instructive to take a closer look at the proof to see where it fails for pointwise convergence.

Let us end this section by taking a look at a less artificial example than those we have studied so far.

Example 3: Consider the sequence $\{f_n\}$ where the functions are defined by

$$f_n(x) = nx(1-x)^n \quad \text{for } x \in [0, 1].$$

We want to study the behavior of the sequence as $n \rightarrow \infty$. To see if the sequence converges pointwise, we compute $\lim_{n \rightarrow \infty} nx(1-x)^n$. As $f_n(0) = f_n(1) = 0$ for all n , the limit is obviously 0 when $x = 0$ or $x = 1$. For $x \in (0, 1)$, the situation is more complicated as nx goes to infinity and $(1-x)^n$ goes to zero. As exponential growth is faster than linear growth, $(1-x)^n$ “wins” and the limit is zero also in this case. If you don’t trust this argument, you can use L’Hôpital’s rule instead (remember to differentiate with respect to n and not x):

$$\begin{aligned} \lim_{n \rightarrow \infty} nx(1-x)^n &= \lim_{n \rightarrow \infty} \frac{n}{(1-x)^{-n}} = \lim_{n \rightarrow \infty} \frac{n}{e^{-n \ln(1-x)}} \\ &\stackrel{L'H}{=} \lim_{n \rightarrow \infty} \frac{1}{e^{-n \ln(1-x)}(-\ln(1-x))} = - \lim_{n \rightarrow \infty} \frac{(1-x)^n}{\ln(1-x)} = 0 \end{aligned}$$

This means that $\{f_n\}$ converges pointwise to 0 as n goes to infinity, but is the convergence uniform? To answer this question, we have to check whether the *maximal* distance between $f_n(x)$ and 0 goes to zero. This distance is clearly equal to the maximal value of $f_n(x)$, and differentiating, we get

$$\begin{aligned} f'_n(x) &= n(1-x)^n + nx \cdot n(1-x)^{n-1}(-1) \\ &= n(1-x)^{n-1}(1 - (n+1)x) \end{aligned}$$

This expression is zero for $x = \frac{1}{n+1}$, and it is easy to check that $\frac{1}{n+1}$ is indeed the maximum point of f_n on $[0, 1]$. The maximum value is

$$f_n\left(\frac{1}{n+1}\right) = \frac{n}{n+1} \left(1 - \frac{1}{n+1}\right)^n \rightarrow e^{-1} \quad \text{as } n \rightarrow \infty$$

(use L’Hôpital’s rule on $(1 - \frac{1}{n+1})^n$ if you don’t recognize the limit). As the limit is not zero, the convergence is not uniform. Hence $\{f_n\}$ converges pointwise but not uniformly to 0 on $[0, 1]$. Figure 4.2.3 shows some of the functions in the sequence. Note the similarity in behavior to the functions in Example 1. ♣




Figure 4.2.3. Functions f_5 , f_{10} , and f_{20} in Example 3.

Exercises for Section 4.2

1. Let $f_n: \mathbb{R} \rightarrow \mathbb{R}$ be defined by $f_n(x) = \frac{x}{n}$. Show that $\{f_n\}$ converges pointwise, but not uniformly to 0.
2. Let $f_n: (0, 1) \rightarrow \mathbb{R}$ be defined by $f_n(x) = x^n$. Show that $\{f_n\}$ converges pointwise, but not uniformly to 0.
3. The function $f_n: [0, \infty) \rightarrow \mathbb{R}$ is defined by $f_n(x) = e^{-x} \left(\frac{x}{n}\right)^{ne}$.
 - a) Show that $\{f_n\}$ converges pointwise.
 - b) Find the maximum value of f_n . Does $\{f_n\}$ converge uniformly?
4. The function $f_n: (0, \infty) \rightarrow \mathbb{R}$ is defined by

$$f_n(x) = n(x^{1/n} - 1)$$

Show that $\{f_n\}$ converges pointwise to $f(x) = \ln x$. Show that the convergence is uniform on each interval $(\frac{1}{k}, k)$, $k \in \mathbb{N}$, but not on $(0, \infty)$.

5. Let $f_n: \mathbb{R} \rightarrow \mathbb{R}$ and assume that the sequence $\{f_n\}$ of continuous functions converges uniformly to $f: \mathbb{R} \rightarrow \mathbb{R}$ on all intervals $[-k, k]$, $k \in \mathbb{N}$. Show that f is continuous.
6. Assume that X is a metric space and that f_n, g_n are functions from X to \mathbb{R} . Show that if $\{f_n\}$ and $\{g_n\}$ converge uniformly to f and g , respectively, then $\{f_n + g_n\}$ converges uniformly to $f + g$.
7. Assume that $f_n: [a, b] \rightarrow \mathbb{R}$ are continuous functions converging uniformly to f . Show that

$$\int_a^b f_n(x) dx \rightarrow \int_a^b f(x) dx$$

Find an example which shows that this is not necessarily the case if $\{f_n\}$ only converges pointwise to f .

8. Let $f_n: \mathbb{R} \rightarrow \mathbb{R}$ be given by $f_n(x) = \frac{1}{n} \sin(nx)$. Show that $\{f_n\}$ converges uniformly to 0, but that the sequence $\{f'_n\}$ of derivatives does not converge. Sketch the graphs of f_n to see what is happening.
9. Let (X, d) be a metric space and assume that the sequence $\{f_n\}$ of continuous functions converges uniformly to f . Show that if $\{x_n\}$ is a sequence in X converging to x , then $f_n(x_n) \rightarrow f(x)$. Find an example which shows that this is not necessarily the case if $\{f_n\}$ only converges pointwise to f .
10. Assume that the functions $f_n: X \rightarrow Y$ converges uniformly to f , and that $g: Y \rightarrow Z$ is uniformly continuous. Show that the sequence $\{g \circ f_n\}$ converges uniformly. Find an example which shows that the conclusion does not necessarily hold if g is only pointwise continuous.
11. Assume that $\sum_{n=0}^{\infty} M_n$ is a convergent series of positive numbers. Assume that $f_n: X \rightarrow \mathbb{R}$ is a sequence of continuous functions defined on a metric space (X, d) . Show that if $|f_n(x)| \leq M_n$ for all $x \in X$ and all $n \in \mathbb{N}$, then the partial sums $s_N(x) = \sum_{n=0}^N f_n(x)$ converge uniformly to a continuous function $s: X \rightarrow \mathbb{R}$ as $N \rightarrow \infty$. (This is called *Weierstrass' M-test*).
12. In this exercise we shall prove:

Dini's Theorem. *If (X, d) is a compact space and $\{f_n\}$ is an increasing sequence of continuous functions $f_n: X \rightarrow \mathbb{R}$ converging pointwise to a continuous function f , then the convergence is uniform.*

- a) Let $g_n = f - f_n$. Show that it suffices to prove that $\{g_n\}$ decreases uniformly to 0.

Assume for contradiction that g_n does not converge uniformly to 0.

- b) Show that there is an $\epsilon > 0$ and a sequence $\{x_n\}$ such that $g_n(x_n) \geq \epsilon$ for all $n \in \mathbb{N}$.
- c) Explain that there is a subsequence $\{x_{n_k}\}$ that converges to a point $a \in X$.
- d) Show that there is an $N \in \mathbb{N}$ and an $r > 0$ such that $g_N(x) < \epsilon$ for all $x \in B(a; r)$.
- e) Derive the contradiction we have been aiming for.

4.3. Integrating and differentiating sequences

In this and the next section, we shall take a look at what different modes of convergence have to say for our ability to integrate and differentiate series. The fundamental question is simple: Assume that we have a sequence of functions $\{f_n\}$ converging to a limit function f . If we integrate the functions f_n , will the integrals converge to the integral of f ? And if we differentiate the f_n 's, will the derivatives converge to f' ?

We shall soon see that without any further restrictions, the answers to both questions are no, but that it is possible to put conditions on the sequences that turn the answers into yes.

Let us start with integration and the following example which is a slight variation of Example 1 in Section 4.2; the only difference is that the height of the “tent” is now n instead of 1.

Example 1: Let $f_n: [0, 1] \rightarrow \mathbb{R}$ be the function in Figure 4.3.1.




Figure 4.3.1. The function f_n

It is given by the formula

$$f_n(x) = \begin{cases} 2n^2x & \text{if } 0 \leq x < \frac{1}{2n} \\ -2n^2x + 2n & \text{if } \frac{1}{2n} \leq x < \frac{1}{n} \\ 0 & \text{if } \frac{1}{n} \leq x \leq 1 \end{cases}$$

but it is much easier just to work from the picture. The sequence $\{f_n\}$ converges pointwise to 0, but the integrals $\int_0^1 f_n(x) dx$ do not converge to 0. In fact,

$\int_0^1 f_n(x) dx = \frac{1}{2}$ since the value of the integral equals the area under the function graph, i.e. the area of a triangle with base $\frac{1}{n}$ and height n . ♣

The example above shows that if the functions f_n converge *pointwise* to a function f on an interval $[a, b]$, the integrals $\int_a^b f_n(x) dx$ need not converge to $\int_a^b f(x) dx$. The reason is that with pointwise convergence, the difference between f and f_n may be very large on small sets – so large that the integrals of f_n fail to converge to the integral of f . If the convergence is *uniform*, this can not happen:

Proposition 4.3.1. *Assume that $\{f_n\}$ is a sequence of continuous functions converging uniformly to f on the interval $[a, b]$. Then the functions*

$$F_n(x) = \int_a^x f_n(t) dt$$

converge uniformly to

$$F(x) = \int_a^x f(t) dt$$

on $[a, b]$.

Proof. We must show that for a given $\epsilon > 0$, we can always find an $N \in \mathbb{N}$ such that $|F(x) - F_n(x)| < \epsilon$ for all $n \geq N$ and all $x \in [a, b]$. Since $\{f_n\}$ converges uniformly to f , there is an $N \in \mathbb{N}$ such that $|f(t) - f_n(t)| < \frac{\epsilon}{b-a}$ for all $t \in [a, b]$. For $n \geq N$, we then have for all $x \in [a, b]$:

$$\begin{aligned} |F(x) - F_n(x)| &= \left| \int_a^x (f(t) - f_n(t)) dt \right| \leq \int_a^x |f(t) - f_n(t)| dt \leq \\ &\leq \int_a^x \frac{\epsilon}{b-a} dt \leq \int_a^b \frac{\epsilon}{b-a} dt = \epsilon \end{aligned}$$

This shows that $\{F_n\}$ converges uniformly to F on $[a, b]$. \square

In applications it is often useful to have the result above with a flexible lower limit.

Corollary 4.3.2. *Assume that $\{f_n\}$ is a sequence of continuous functions converging uniformly to f on the interval $[a, b]$. For any $x_0 \in [a, b]$, the functions*

$$F_n(x) = \int_{x_0}^x f_n(t) dt$$

converge uniformly to

$$F(x) = \int_{x_0}^x f(t) dt$$

on $[a, b]$.

Proof. Recall that

$$\int_a^x f_n(t) dt = \int_a^{x_0} f_n(t) dt + \int_{x_0}^x f_n(t) dt$$

regardless of the order of the numbers a, x_0, x , and hence

$$\int_{x_0}^x f_n(t) dt = \int_a^x f_n(t) dt - \int_a^{x_0} f_n(t) dt$$

The first integral on the right converges uniformly to $\int_a^x f(t) dt$ by the proposition, and the second integral converges (as a sequence of numbers) to $\int_a^{x_0} f(t) dt$. Hence $\int_{x_0}^x f_n(t) dt$ converges uniformly to

$$\int_a^x f(t) dt - \int_a^{x_0} f(t) dt = \int_{x_0}^x f(t) dt$$

as was to be proved. \square

Let us reformulate this result in terms of series. Recall that in calculus we say that a series of functions $\sum_{n=0}^{\infty} v_n(x)$ converges *pointwise* to a function f on a set I if the sequence $\{s_N(x)\}$ of partial sum $s_N(x) = \sum_{n=0}^N v_n(x)$ converges to $f(x)$ for every $x \in I$. Similarly, we say that the series converges *uniformly* to f on I if the sequence $\{s_N\}$ of partial sum $s_N(x) = \sum_{n=0}^N v_n(x)$ converges uniformly to f on I .

Corollary 4.3.3. *Assume that $\{v_n\}$ is a sequence of continuous functions such that the series $\sum_{n=0}^{\infty} v_n(x)$ converges uniformly on the interval $[a, b]$. Then for any $x_0 \in [a, b]$, the series $\sum_{n=0}^{\infty} \int_{x_0}^x v_n(t) dt$ converges uniformly and*

$$\int_{x_0}^x \sum_{n=0}^{\infty} v_n(t) dt = \sum_{n=0}^{\infty} \int_{x_0}^x v_n(t) dt$$

Proof. Assume that the series $\sum_{n=0}^{\infty} v_n(x)$ converges uniformly to the function f . This means that the partial sums $s_N(x) = \sum_{n=0}^N v_n(x)$ converge uniformly to f , and hence by Corollary 4.3.2,

$$\begin{aligned} \int_{x_0}^x \sum_{n=0}^{\infty} v_n(t) dt &= \int_{x_0}^x f(t) dt = \lim_{N \rightarrow \infty} \int_{x_0}^x s_N(t) dt \\ &= \lim_{N \rightarrow \infty} \int_{x_0}^x \sum_{n=0}^N v_n(t) dt = \lim_{N \rightarrow \infty} \sum_{n=0}^N \int_{x_0}^x v_n(t) dt = \sum_{n=0}^{\infty} \int_{x_0}^x v_n(t) dt \end{aligned}$$

by the definition of infinite sums. \square

The corollary tells us that if the series $\sum_{n=0}^{\infty} v_n(x)$ converges uniformly, we can integrate it term by term to get

$$\int_{x_0}^x \sum_{n=0}^{\infty} v_n(t) dt = \sum_{n=0}^{\infty} \int_{x_0}^x v_n(t) dt$$

This formula may look obvious, but it does not in general hold for series that only converge pointwise. As we shall see many times later in the book, interchanging integrals and infinite sums is quite a tricky business.

To use the corollary efficiently, we need to be able to determine when a series of functions converges uniformly. The following simple test is often helpful:

Proposition 4.3.4 (Weierstrass' M -test). *Let $\{v_n\}$ be a sequence of functions $v_n: A \rightarrow \mathbb{R}$ defined on a set A , and assume that there is a convergent series $\sum_{n=0}^{\infty} M_n$ of non-negative numbers such that $|v_n(x)| \leq M_n$ for all $n \in \mathbb{N}$ and all $x \in A$. Then the series $\sum_{n=0}^{\infty} v_n(x)$ converges uniformly on A .*

Proof. Let $s_n(x) = \sum_{k=0}^n v_k(x)$ be the partial sums of the original series. Since the series $\sum_{n=0}^{\infty} M_n$ converges, we know that its partial sums $S_n = \sum_{k=0}^n M_k$ form a Cauchy sequence. Since for all $x \in A$ and all $m > n$,

$$|s_m(x) - s_n(x)| = \left| \sum_{k=n+1}^m v_k(x) \right| \leq \sum_{k=n+1}^m |v_k(x)| \leq \sum_{k=n+1}^m M_k = |S_m - S_n|,$$

we see that $\{s_n(x)\}$ is a Cauchy sequence (in \mathbb{R}) for each $x \in A$ and hence converges to a limit $s(x)$. This defines a pointwise limit function $s: A \rightarrow \mathbb{R}$.

To prove that $\{s_n\}$ converges *uniformly* to s , note that for every $\epsilon > 0$, there is an $N \in \mathbb{N}$ such that if $S = \sum_{k=0}^{\infty} M_k$, then

$$\sum_{k=n+1}^{\infty} M_k = S - S_n < \epsilon$$

for all $n \geq N$. This means that for all $n \geq N$,

$$|s(x) - s_n(x)| = \left| \sum_{k=n+1}^{\infty} v_k(x) \right| \leq \sum_{k=n+1}^{\infty} |v_k(x)| \leq \sum_{k=n+1}^{\infty} M_k < \epsilon$$

for all $x \in A$, and hence $\{s_n\}$ converges uniformly to s on A . \square

Example 2: Consider the series $\sum_{n=1}^{\infty} \frac{\cos nx}{n^2}$. Since $|\frac{\cos nx}{n^2}| \leq \frac{1}{n^2}$, and $\sum_{n=0}^{\infty} \frac{1}{n^2}$ converges, the original series $\sum_{n=1}^{\infty} \frac{\cos nx}{n^2}$ converges uniformly to a function f on any closed and bounded interval $[a, b]$. Hence we may integrate termwise to get

$$\int_0^x f(t) dt = \sum_{n=1}^{\infty} \int_0^x \frac{\cos nt}{n^2} dt = \sum_{n=1}^{\infty} \frac{\sin nx}{n^3}$$



Let us now turn to differentiation of sequences. This is a much trickier business than integration as integration often helps to smoothen functions while differentiation tends to make them more irregular. Here is a simple example.

Example 3: The sequence (not series!) $\{\frac{\sin nx}{n}\}$ obviously converges uniformly to 0, but the sequence of derivatives $\{\cos nx\}$ does not converge at all. \clubsuit

The example shows that even if a sequence $\{f_n\}$ of differentiable functions converges uniformly to a differentiable function f , the derivatives f'_n need not converge to the derivative f' of the limit function. If you draw the graphs of the functions f_n , you will see why — although they live in an increasingly narrower strip around the x -axis, they are all equally steep at their steepest, and the derivatives do not converge to 0.

To get a theorem that works, we have to put the conditions on the derivatives. The following result may look ugly and unsatisfactory, but it gives us the information we shall need.

Proposition 4.3.5. *Let $\{f_n\}$ be a sequence of differentiable functions on the interval $[a, b]$. Assume that the derivatives f'_n are continuous and that they converge uniformly to a function g on $[a, b]$. Assume also that there is a point $x_0 \in [a, b]$ such that the sequence $\{f_n(x_0)\}$ converges. Then the sequence $\{f_n\}$ converges uniformly on $[a, b]$ to a differentiable function f such that $f' = g$.*

Proof. The proposition is just Corollary 4.3.2 in a convenient disguise. If we apply that proposition to the sequence $\{f'_n\}$, we see that the integrals $\int_{x_0}^x f'_n(t) dt$ converge uniformly to $\int_{x_0}^x g(t) dt$. By the Fundamental Theorem of Calculus, we get

$$f_n(x) - f_n(x_0) \rightarrow \int_{x_0}^x g(t) dt \quad \text{uniformly on } [a, b]$$

Since $f_n(x_0)$ converges to a limit b , this means that $f_n(x)$ converges uniformly to the function $f(x) = b + \int_{x_0}^x g(t) dt$. Using the Fundamental Theorem of Calculus again, we see that $f'(x) = g(x)$. \square

Also in this case it is useful to have a reformulation in terms of series:

Corollary 4.3.6. *Let $\sum_{n=0}^{\infty} u_n(x)$ be a series where the functions u_n are differentiable with continuous derivatives on the interval $[a, b]$. Assume that the series of derivatives $\sum_{n=0}^{\infty} u'_n(x)$ converges uniformly on $[a, b]$. Assume also that there is a point $x_0 \in [a, b]$ where we know that the series $\sum_{n=0}^{\infty} u_n(x_0)$ converges. Then the series $\sum_{n=0}^{\infty} u_n(x)$ converges uniformly on $[a, b]$, and*

$$\left(\sum_{n=0}^{\infty} u_n(x) \right)' = \sum_{n=0}^{\infty} u'_n(x)$$

Proof. Left to the reader. \square

The corollary tells us that under rather strong conditions, we can differentiate the series $\sum_{n=0}^{\infty} u_n(x)$ term by term.

It's time for an example that sums up most of what we have been looking at in this section.

Example 4: We shall study the series

$$\sum_{n=1}^{\infty} \log \left(1 + \frac{x}{n^2} \right) \quad \text{for } x \geq 0$$

It's easy to see that the series converges pointwise: For $x = 0$ all the terms are zero, and for $x > 0$ we get convergence by comparing the series to the convergent series $\sum_{n=1}^{\infty} \frac{1}{n^2}$ (recall the Limit Comparison Test from calculus):

$$\lim_{n \rightarrow \infty} \frac{\log \left(1 + \frac{x}{n^2} \right)}{\frac{1}{n^2}} \stackrel{L'H}{=} \lim_{n \rightarrow \infty} \frac{\frac{1}{1 + \frac{x}{n^2}} \cdot \left(-\frac{2x}{n^3} \right)}{-\frac{2}{n^3}} = x < \infty$$

This means that the sum

$$f(x) = \sum_{n=1}^{\infty} \log \left(1 + \frac{x}{n^2} \right)$$

exists for $x \geq 0$.

It turns out that the series doesn't converge uniformly on all of $[0, \infty)$ (the convergence gets slower and slower as x increases), but that it does converge uniformly on $[0, a]$ for all $a > 0$. To prove this, we shall use that $\log(1 + u) \leq u$ for all $u \geq 0$, and hence $\log \left(1 + \frac{x}{n^2} \right) \leq \frac{x}{n^2}$ (if you don't see that $\log(1 + u) \leq u$, just check that the function $g(u) = u - \log(1 + u)$ is increasing for $u \geq 0$ and that $g(0) = 0$). This means that for $x \in [0, a]$,

$$\log \left(1 + \frac{x}{n^2} \right) \leq \frac{x}{n^2} \leq \frac{a}{n^2}$$

and as $\sum_{n=1}^{\infty} \frac{a}{n^2}$ converges, our series $\sum_{n=1}^{\infty} \log \left(1 + \frac{x}{n^2} \right)$ converges uniformly on $[0, a]$ by Weierstrass' M -test 4.3.4. By Proposition 4.2.4, the limit function f is continuous on $[0, a]$. As any positive x is an element of $[0, a]$ for some a , f is continuous at all positive x .

Let us see if f is differentiable. If we differentiate the series term by term, we get

$$\sum_{n=1}^{\infty} \frac{1}{n^2 + x}$$

As $\frac{1}{n^2 + x} \leq \frac{1}{n^2}$ for all $x \geq 0$, this series converges uniformly on $[0, \infty)$ by Weierstrass' M -test 4.3.4, and hence

$$f'(x) = \sum_{n=1}^{\infty} \frac{1}{n^2 + x}$$

by Corollary 4.3.6 (note that to use this result, we need to know that the *differentiated* series $\sum_{n=1}^{\infty} \frac{1}{n^2 + x}$ converges uniformly). ♣

Exercises for Section 4.3

1. Show that $\sum_{n=0}^{\infty} \frac{\cos(nx)}{n^2 + 1}$ converges uniformly on \mathbb{R} .
2. Show that $\sum_{n=1}^{\infty} \frac{x}{n^2}$ converges uniformly on $[-1, 1]$.
3. Let $f_n: [0, 1] \rightarrow \mathbb{R}$ be defined by $f_n(x) = nx(1 - x^2)^n$. Show that $f_n(x) \rightarrow 0$ for all $x \in [0, 1]$, but that $\int_0^1 f_n(x) dx \rightarrow \frac{1}{2}$.
4. Explain in detail how Corollary 4.3.6 follows from Proposition 4.3.5.
5. a) Show, by summing a geometric series, that

$$\frac{1}{1 - e^{-x}} = \sum_{n=0}^{\infty} e^{-nx} \quad \text{for } x > 0$$

- b) Explain that we can differentiate the series above term by term to get

$$\frac{e^{-x}}{(1 - e^{-x})^2} = \sum_{n=1}^{\infty} n e^{-nx} \quad \text{for all } x > 0$$

- c) Does the series $\sum_{n=1}^{\infty} n e^{-nx}$ converge uniformly on $(0, \infty)$?
6. a) Show that the series $f(x) = \sum_{n=1}^{\infty} \frac{\sin \frac{x}{n}}{n}$ converges uniformly on \mathbb{R} .

b) Show that the limit function f is differentiable and find an expression for $f'(x)$.

7. One can show that

$$x = \sum_{n=1}^{\infty} \frac{2(-1)^{n+1}}{n} \sin(nx) \quad \text{for } x \in (-\pi, \pi)$$

If we differentiate term by term, we get

$$1 = \sum_{n=1}^{\infty} 2(-1)^{n+1} \cos(nx) \quad \text{for } x \in (-\pi, \pi)$$

Is this a correct formula?

8. a) Show that the sequence $\sum_{n=1}^{\infty} \frac{1}{n^x}$ converges uniformly on all intervals $[a, \infty)$ where $a > 1$.

b) Let $f(x) = \sum_{n=1}^{\infty} \frac{1}{n^x}$ for $x > 1$. Show that $f'(x) = -\sum_{n=1}^{\infty} \frac{\ln x}{n^x}$.

9. a) Show that the series

$$\sum_{n=1}^{\infty} \frac{1}{1+n^2x}$$

converges for $x > 0$ and diverges for $x = 0$.

b) Show that the series converges uniformly on $[a, \infty)$ for all $a > 0$.

c) Define $f: (0, \infty) \rightarrow \mathbb{R}$ by $f(x) = \sum_{n=1}^{\infty} \frac{1}{1+n^2x}$. Show that f is continuous.

d) Show that the series does not converge uniformly on $(0, \infty)$.

10. a) Show that the series

$$\sum_{n=1}^{\infty} \frac{\arctan(nx)}{n^2}$$

converges uniformly on \mathbb{R} , and explain that the function f defined by

$$f(x) = \sum_{n=1}^{\infty} \frac{\arctan(nx)}{n^2}$$

is continuous on \mathbb{R} .

b) Show that f is differentiable at all points $x \neq 0$. Express $f'(x)$ as a series.

c) Show that f is *not* differentiable at $x = 0$.

4.4. Applications to power series

In this section, we shall illustrate the theory in the previous section by applying it to the power series you know from calculus. If you are not familiar with \limsup and \liminf , you should read the discussion in Section 2.2 before you continue. The results in this section are not needed in the sequel.

Recall that a power series is a function of the form

$$f(x) = \sum_{n=0}^{\infty} c_n(x-a)^n$$

where a is a real number and $\{c_n\}$ is a sequence of real numbers. It is defined for the x -values that make the series converge. We define the *radius of convergence* of the series to be the number R such that

$$\frac{1}{R} = \limsup_{n \rightarrow \infty} \sqrt[n]{|c_n|}$$

with the interpretation that $R = 0$ if the limit is infinite, and $R = \infty$ if the limit is 0. To justify this terminology, we need the the following result.

Proposition 4.4.1. *If R is the radius of convergence of the power series*

$$\sum_{n=0}^{\infty} c_n(x-a)^n,$$

then the series converges for $|x-a| < R$ and diverges for $|x-a| > R$. If $0 < r < R$, the series converges uniformly on $[a-r, a+r]$.

Proof. Let us first assume that $|x-a| > R$. This means that $\frac{1}{|x-a|} < \frac{1}{R}$, and since $\limsup_{n \rightarrow \infty} \sqrt[n]{|c_n|} = \frac{1}{R}$, there must be arbitrarily large values of n such that $\sqrt[n]{|c_n|} > \frac{1}{|x-a|}$. Hence $|c_n(x-a)^n| > 1$, and consequently the series must diverge as the terms do not decrease to zero.

To prove the (uniform) convergence, assume that r is a number between 0 and R . Since $\frac{1}{r} > \frac{1}{R}$, we can pick a positive number $b < 1$ such that $\frac{b}{r} > \frac{1}{R}$. Since $\limsup_{n \rightarrow \infty} \sqrt[n]{|c_n|} = \frac{1}{R}$, there must be an $N \in \mathbb{N}$ such that $\sqrt[n]{|c_n|} < \frac{b}{r}$ when $n \geq N$. This means that $|c_n r^n| < b^n$ for $n \geq N$, and hence that $|c_n(x-a)^n| < b^n$ for all $x \in [a-r, a+r]$. Since $\sum_{n=N}^{\infty} b^n$ is a convergent geometric series, Weierstrass' M-test tells us that the series $\sum_{n=N}^{\infty} c_n(x-a)^n$ converges uniformly on $[a-r, a+r]$. Since only the tail of a sequence counts for convergence, the full series $\sum_{n=0}^{\infty} c_n(x-a)^n$ also converges uniformly on $[a-r, a+r]$. Since r is an arbitrary number less than R , we see that the series must converge on the open interval $(a-R, a+R)$, i.e. whenever $|x-a| < R$. \square

Remark: When we want to find the radius of convergence, it is occasionally convenient to compute a slightly different limit such as

$$\limsup_{n \rightarrow \infty} \sqrt[n+1]{c_n} \quad \text{or} \quad \limsup_{n \rightarrow \infty} \sqrt[n]{c_n}$$

instead of $\limsup_{n \rightarrow \infty} \sqrt[n]{c_n}$. This corresponds to finding the radius of convergence of the power series we get by either multiplying or dividing the original one by $(x-a)$, and gives the correct answer as multiplying or dividing a series by a non-zero number doesn't change its convergence properties.

The proposition above does not tell us what happens at the endpoints $a \pm R$ of the interval of convergence, but we know from calculus that a series may converge at both, one or neither of the endpoints. Although the convergence is uniform on all subintervals $[a-r, a+r]$, it is not in general uniform on $(a-R, a+R)$.

Let us now take a look at integration and differentiation of power series.

Corollary 4.4.2. *Assume that the power series $f(x) = \sum_{n=0}^{\infty} c_n(x-a)^n$ has radius of convergence R larger than 0. Then the function f is continuous and differentiable on the open interval $(a-R, a+R)$ with*

$$f'(x) = \sum_{n=1}^{\infty} n c_n(x-a)^{n-1} = \sum_{n=0}^{\infty} (n+1) c_{n+1}(x-a)^n \quad \text{for } x \in (a-R, a+R)$$

and

$$\int_a^x f(t) dt = \sum_{n=0}^{\infty} \frac{c_n}{n+1} (x-a)^{n+1} = \sum_{n=1}^{\infty} \frac{c_{n-1}}{n} (x-a)^n \quad \text{for } x \in (a-R, a+R)$$

Proof. Since the power series converges uniformly on each subinterval $[a-r, a+r]$, the sum is continuous on each such interval according to Proposition 4.2.4. Since each x in $(a-R, a+R)$ is contained in the interior of some of the subintervals $[a-r, a+r]$, we see that f must be continuous on the full interval $(a-R, a+R)$. The formula for the integral follows immediately by applying Corollary 4.3.3 on each subinterval $[a-r, a+r]$ in a similar way.

To get the formula for the derivative, we shall apply Corollary 4.3.6. To use this result, we need to know that the differentiated series $\sum_{n=1}^{\infty} (n+1)c_{n+1}(x-a)^n$ has the same radius of convergence as the original series; i.e. that

$$\limsup_{n \rightarrow \infty} \sqrt[n+1]{|(n+1)c_{n+1}|} = \limsup_{n \rightarrow \infty} \sqrt[n]{|c_n|} = \frac{1}{R}$$

(recall that by the remark above, we may use the $n+1$ -st root on the left hand side instead of the n -th root). Since $\lim_{n \rightarrow \infty} \sqrt[n+1]{n+1} = 1$, this is not hard to show (see Exercise 7). Applying Corollary 4.3.6 on each subinterval $[a-r, a+r]$, we now get the formula for the derivative at each point $x \in (a-r, a+r)$. Since each point in $(a-R, a+R)$ belongs to the interior of some of the subintervals, the formula for the derivative must hold at all points $x \in (a-R, a+R)$. \square

A function that is the sum of a power series, is called a *real analytic function*. Such functions have derivatives of all orders.

Corollary 4.4.3. Assume that $f(x) = \sum_{n=0}^{\infty} c_n(x-a)^n$ for $x \in (a-R, a+R)$. Then f is k times differentiable in $(a-R, a+R)$ for any $k \in \mathbb{N}$, and $f^{(k)}(a) = k!c_k$. Hence $\sum_{n=0}^{\infty} c_n(x-a)^n$ is the Taylor series

$$f(x) = \sum_{n=0}^{\infty} \frac{f^{(n)}(a)}{n!} (x-a)^n$$

Proof. Using the previous corollary, we get by induction that $f^{(k)}$ exists on $(a-R, a+R)$ and that

$$f^{(k)}(x) = \sum_{n=k}^{\infty} n(n-1) \cdots (n-k+1) c_n (x-a)^{n-k}$$

Putting $x = a$, we get $f^{(k)}(a) = k!c_k$, and the corollary follows. \square

Abel's Theorem

We have seen that the sum $f(x) = \sum_{n=0}^{\infty} c_n(x-a)^n$ of a power series is continuous in the interior $(a-R, a+R)$ of its interval of convergence. But what happens if the series converges at an endpoint $a \pm R$? This is a surprisingly intricate problem, but in 1826 Niels Henrik Abel (1802-1829) proved that the power series is necessarily continuous also at the endpoint.

Before we turn to the proof, we need a lemma that can be thought of as a discrete version of integration by parts.

Lemma 4.4.4 (Abel's Summation Formula). *Let $\{a_n\}_{n=0}^\infty$ and $\{b_n\}_{n=0}^\infty$ be two sequences of real numbers, and let $s_n = \sum_{k=0}^n a_k$. Then*

$$\sum_{n=0}^N a_n b_n = s_N b_N + \sum_{n=0}^{N-1} s_n (b_n - b_{n+1}).$$

If the series $\sum_{n=0}^\infty a_n$ converges, and $b_n \rightarrow 0$ as $n \rightarrow \infty$, then

$$\sum_{n=0}^\infty a_n b_n = \sum_{n=0}^\infty s_n (b_n - b_{n+1})$$

in the sense that either the two series both diverge or they converge to the same limit.

Proof. Note that $a_n = s_n - s_{n-1}$ for $n \geq 1$, and that this formula even holds for $n = 0$ if we define $s_{-1} = 0$. Hence

$$\sum_{n=0}^N a_n b_n = \sum_{n=0}^N (s_n - s_{n-1}) b_n = \sum_{n=0}^N s_n b_n - \sum_{n=0}^N s_{n-1} b_n$$

Changing the index of summation and using that $s_{-1} = 0$, we see that

$$\sum_{n=0}^N s_{n-1} b_n = \sum_{n=0}^{N-1} s_n b_{n+1}.$$

Putting this into the formula above, we get

$$\sum_{n=0}^N a_n b_n = \sum_{n=0}^N s_n b_n - \sum_{n=0}^{N-1} s_n b_{n+1} = s_N b_N + \sum_{n=0}^{N-1} s_n (b_n - b_{n+1})$$

and the first part of the lemma is proved. The second follows by letting $N \rightarrow \infty$. \square

We are now ready to prove:

Theorem 4.4.5 (Abel's Theorem). *The sum of a power series $f(x) = \sum_{n=0}^\infty c_n (x-a)^n$ is continuous in its entire interval of convergence. This means in particular that if R is the radius of convergence, and the power series converges at the right endpoint $a + R$, then $\lim_{x \uparrow a+R} f(x) = f(a + R)$, and if the power series converges at the left endpoint $a - R$, then $\lim_{x \downarrow a-R} f(x) = f(a - R)$.¹*

Proof. As we already know that f is continuous in the open interval $(a - R, a + R)$, we only need to check the endpoints. To keep the notation simple, we shall assume that $a = 0$ and concentrate on the right endpoint R . Thus we want to prove that $\lim_{x \uparrow R} f(x) = f(R)$.

¹I use $\lim_{x \uparrow b}$ and $\lim_{x \downarrow b}$ for one-sided limits, also denoted by $\lim_{x \rightarrow b-}$ and $\lim_{x \rightarrow b+}$.

Note that $f(x) = \sum_{n=0}^{\infty} c_n R^n \left(\frac{x}{R}\right)^n$. If we introduce $f_n(R) = \sum_{k=0}^n c_k R^k$ and assume that $|x| < R$, we may apply the second version of Abel's summation formula with $a_n = c_n R^n$ and $b_n = \left(\frac{x}{R}\right)^n$ to get

$$f(x) = \sum_{n=0}^{\infty} f_n(R) \left(\left(\frac{x}{R}\right)^n - \left(\frac{x}{R}\right)^{n+1} \right) = \left(1 - \frac{x}{R}\right) \sum_{n=0}^{\infty} f_n(R) \left(\frac{x}{R}\right)^n$$

We need a similar formula for $f(R)$. Summing a geometric series, we see that $\frac{1}{1 - \frac{x}{R}} = \sum_{n=0}^{\infty} \left(\frac{x}{R}\right)^n$. Multiplying by $f(R) \left(1 - \frac{x}{R}\right)$ on both sides, we get

$$f(R) = \left(1 - \frac{x}{R}\right) \sum_{n=0}^{\infty} f(R) \left(\frac{x}{R}\right)^n$$

Hence

$$|f(x) - f(R)| = \left| \left(1 - \frac{x}{R}\right) \sum_{n=0}^{\infty} (f_n(R) - f(R)) \left(\frac{x}{R}\right)^n \right|$$

Given an $\epsilon > 0$, we must find a $\delta > 0$ such that this quantity is less than ϵ when $R - \delta < x < R$. This may seem obvious due to the factor $(1 - x/R)$, but the problem is that the infinite sum may go to infinity when $x \rightarrow R$. Hence we need to control the tail of the series before we exploit the factor $(1 - x/R)$. Fortunately, this is not difficult: Since $f_n(R) \rightarrow f(R)$, we first pick an $N \in \mathbb{N}$ such that $|f_n(R) - f(R)| < \frac{\epsilon}{2}$ for $n \geq N$. Then

$$\begin{aligned} |f(x) - f(R)| &\leq \left(1 - \frac{x}{R}\right) \sum_{n=0}^{N-1} |f_n(R) - f(R)| \left(\frac{x}{R}\right)^n + \\ &\quad + \left(1 - \frac{x}{R}\right) \sum_{n=N}^{\infty} |f_n(R) - f(R)| \left(\frac{x}{R}\right)^n \leq \\ &\leq \left(1 - \frac{x}{R}\right) \sum_{n=0}^{N-1} |f_n(R) - f(R)| \left(\frac{x}{R}\right)^n + \left(1 - \frac{x}{R}\right) \sum_{n=0}^{\infty} \frac{\epsilon}{2} \left(\frac{x}{R}\right)^n = \\ &= \left(1 - \frac{x}{R}\right) \sum_{n=0}^{N-1} |f_n(R) - f(R)| \left(\frac{x}{R}\right)^n + \frac{\epsilon}{2} \end{aligned}$$

where we have summed a geometric series. As we now have a *finite* sum, the first term clearly converges to 0 when $x \uparrow R$. Hence there is a $\delta > 0$ such that this term is less than $\frac{\epsilon}{2}$ when $R - \delta < x < R$, and consequently $|f(x) - f(R)| < \epsilon$ for such values of x . \square

Let us take a look at a famous example.

Example 1: Summing a geometric series, we get

$$\frac{1}{1 + x^2} = \sum_{n=0}^{\infty} (-1)^n x^{2n} \quad \text{for } |x| < 1$$

Integrating, we obtain

$$\arctan x = \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{2n+1} \quad \text{for } |x| < 1$$

Using the Alternating Series Test, we see that the series converges even for $x = 1$. By Abel's Theorem we can take the limits on both sides to get

$$\arctan 1 = \sum_{n=0}^{\infty} (-1)^n \frac{1}{2n+1}$$

As $\arctan 1 = \frac{\pi}{4}$, we have proved

$$\frac{\pi}{4} = 1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \dots$$

This is often called Leibniz' or Gregory's formula for π , but it was actually first discovered by the Indian mathematician Madhava (ca. 1340 – ca. 1425). ♣

This example is rather typical; the most interesting information is often obtained at an endpoint, and we need Abel's Theorem to secure it.

It is natural to think that Abel's Theorem must have a converse saying that if $\lim_{x \uparrow a+R} \sum_{n=0}^{\infty} c_n x^n$ exists, then the sequence converges at the right endpoint $x = a + R$. This, however, is not true as the following simple example shows.

Example 2: Summing a geometric series, we have

$$\frac{1}{1+x} = \sum_{n=0}^{\infty} (-x)^n \quad \text{for } |x| < 1$$

Obviously, $\lim_{x \uparrow 1} \sum_{n=0}^{\infty} (-x)^n = \lim_{x \uparrow 1} \frac{1}{1+x} = \frac{1}{2}$, but the series does not converge for $x = 1$. ♣

It is possible to put extra conditions on the coefficients of the series to ensure convergence at the endpoint, see Exercise 8.

Exercises for Section 4.4

- Find power series with radius of convergence 0, 1, 2, and ∞ .
- Find power series with radius of convergence 1 that converge at both, one and neither of the endpoints.
- Show that for any polynomial P , $\lim_{n \rightarrow \infty} \sqrt[n]{|P(n)|} = 1$.
- Use the result in Exercise 3 to find the radius of convergence:
 - $\sum_{n=0}^{\infty} \frac{2^n x^n}{n^3 + 1}$
 - $\sum_{n=0}^{\infty} \frac{2n^2 + n - 1}{3n + 4} x^n$
 - $\sum_{n=0}^{\infty} n x^{2n}$
- Explain that $\frac{1}{1-x^2} = \sum_{n=0}^{\infty} x^{2n}$ for $|x| < 1$,
 - Show that $\frac{2x}{(1-x^2)^2} = \sum_{n=0}^{\infty} 2n x^{2n-1}$ for $|x| < 1$.
 - Show that $\frac{1}{2} \ln \left| \frac{1+x}{1-x} \right| = \sum_{n=0}^{\infty} \frac{x^{2n+1}}{2n+1}$ for $|x| < 1$.
- Explain why $\frac{1}{1+x} = \sum_{n=0}^{\infty} (-1)^n x^n$ for $|x| < 1$.
 - Show that $\ln(1+x) = \sum_{n=0}^{\infty} (-1)^n \frac{x^{n+1}}{n+1}$ for $|x| < 1$.
 - Show that $\ln 2 = \sum_{n=0}^{\infty} (-1)^n \frac{1}{n+1}$.

7. Let $\sum_{n=0}^{\infty} c_n(x-a)^n$ be a power series.

a) Show that the radius of convergence is given by

$$\frac{1}{R} = \limsup_{n \rightarrow \infty} \sqrt[n+k]{|c_n|}$$

for any integer k .

b) Show that $\lim_{n \rightarrow \infty} \sqrt[n+1]{n+1} = 1$ (write $\sqrt[n+1]{n+1} = (n+1)^{\frac{1}{n+1}}$).

c) Prove the formula

$$\limsup_{n \rightarrow \infty} \sqrt[n+1]{|(n+1)c_{n+1}|} = \limsup_{n \rightarrow \infty} \sqrt[n]{|c_n|} = \frac{1}{R}$$

in the proof of Corollary 4.4.2.

8. In this problem we shall prove the following partial converse of Abel's Theorem:

Tauber's Theorem. Assume that $s(x) = \sum_{n=0}^{\infty} c_n x^n$ is a power series with radius of convergence 1. Assume that $s = \lim_{x \uparrow 1} \sum_{n=0}^{\infty} c_n x^n$ is finite. If in addition $\lim_{n \rightarrow \infty} n c_n = 0$, then the power series converges for $x = 1$ and $s = s(1)$.

a) Explain that if we can prove that the power series converges for $x = 1$, then the rest of the theorem will follow from Abel's Theorem.

b) Show that $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^N n |c_n| = 0$.

c) Let $s_N = \sum_{n=0}^N c_n$. Explain that

$$s(x) - s_N = - \sum_{n=0}^N c_n (1 - x^n) + \sum_{n=N+1}^{\infty} c_n x^n$$

d) Show that $1 - x^n \leq n(1 - x)$ for $|x| < 1$.

e) Let N_x be the integer such that $N_x \leq \frac{1}{1-x} < N_x + 1$. Show that

$$\sum_{n=0}^{N_x} c_n (1 - x^n) \leq (1 - x) \sum_{n=0}^{N_x} n |c_n| \leq \frac{1}{N_x} \sum_{n=0}^{N_x} n |c_n| \rightarrow 0$$

as $x \uparrow 1$.

f) Show that

$$\left| \sum_{n=N_x+1}^{\infty} c_n x^n \right| \leq \sum_{n=N_x+1}^{\infty} n |c_n| \frac{x^n}{n} \leq \frac{d_x}{N_x} \sum_{n=0}^{\infty} x^n$$

where $d_x \rightarrow 0$ as $x \uparrow 1$. Show that $\sum_{n=N_x+1}^{\infty} c_n x^n \rightarrow 0$ as $x \uparrow 1$.

g) Prove Tauber's theorem.

4.5. Spaces of bounded functions

So far we have looked at functions individually or as part of a sequence. We shall now take a bold step and consider functions as elements in metric spaces. As we shall see later in the chapter, this will make it possible to use results from the theory of metric spaces to prove theorems about functions, e.g., to use Banach's Fixed Point Theorem 3.4.5 to prove the existence of solutions to differential equations. In

this section, we shall consider spaces of bounded functions and in the next section we shall look at the more important case of continuous functions.

If (X, d_X) and (Y, d_Y) are metric spaces, a function $f: X \rightarrow Y$ is *bounded* if the set of values $\{f(x) : x \in X\}$ is a bounded set, i.e. if there is a number $M \in \mathbb{R}$ such that $d_Y(f(u), f(v)) \leq M$ for all $u, v \in X$. An equivalent definition is to say that for any $a \in X$, there is a constant M_a such that $d_Y(f(a), f(x)) \leq M_a$ for all $x \in X$.

Note that if $f, g: X \rightarrow Y$ are two bounded functions, then there is a number K such that $d_Y(f(x), g(x)) \leq K$ for all $x \in X$. To see this, fix a point $a \in X$, and let M_a and N_a be numbers such that $d_Y(f(a), f(x)) \leq M_a$ and $d_Y(g(a), g(x)) \leq N_a$ for all $x \in X$. Since by the triangle inequality

$$\begin{aligned} d_Y(f(x), g(x)) &\leq d_Y(f(x), f(a)) + d_Y(f(a), g(a)) + d_Y(g(a), g(x)) \\ &\leq M_a + d_Y(f(a), g(a)) + N_a \end{aligned}$$

we can take $K = M_a + d_Y(f(a), g(a)) + N_a$.

We now let

$$B(X, Y) = \{f : X \rightarrow Y \mid f \text{ is bounded}\}$$

be the collection of all bounded functions from X to Y . We shall turn $B(X, Y)$ into a metric space by introducing a metric ρ . The idea is to measure the distance between two functions by looking at how far apart they can be at a point; i.e. by

$$\rho(f, g) = \sup\{d_Y(f(x), g(x)) \mid x \in X\}$$

Note that by our argument above, $\rho(f, g) < \infty$. Our first task is to show that ρ really is a metric on $B(X, Y)$.

Proposition 4.5.1. *If (X, d_X) and (Y, d_Y) are metric spaces,*

$$\rho(f, g) = \sup\{d_Y(f(x), g(x)) \mid x \in X\}$$

defines a metric ρ on $B(X, Y)$.

Proof. As we have already observed that $\rho(f, g)$ is always finite, we only have to prove that ρ satisfies the three properties of a metric: positivity, symmetry, and the triangle inequality. The first two are more or less obvious, and we concentrate on the triangle inequality: If f, g, h are three functions in $C(X, Y)$, we must show that

$$\rho(f, g) \leq \rho(f, h) + \rho(h, g)$$

For all $x \in X$,

$$d_Y(f(x), g(x)) \leq d_Y(f(x), h(x)) + d_Y(h(x), g(x)) \leq \rho(f, h) + \rho(h, g)$$

and taking supremum over all $x \in X$, we get

$$\rho(f, g) \leq \rho(f, h) + \rho(h, g)$$

and the proposition is proved. \square

Not surprisingly, convergence in $(B(X, Y), \rho)$ is just the same as uniform convergence.

Proposition 4.5.2. *A sequence $\{f_n\}$ converges to f in $(B(X, Y), \rho)$ if and only if it converges uniformly to f .*

Proof. According to Proposition 4.2.3, $\{f_n\}$ converges uniformly to f if and only if

$$\sup\{d_Y(f_n(x), f(x)) \mid x \in X\} \rightarrow 0$$

This just means that $\rho(f_n, f) \rightarrow 0$, which is to say that $\{f_n\}$ converges to f in $(B(X, Y), \rho)$. \square

The next result introduces an important idea that we shall see many examples of later: The space $B(X, Y)$ inherits completeness from Y .

Theorem 4.5.3. *Let (X, d_X) and (Y, d_Y) be metric spaces and assume that (Y, d_Y) is complete. Then $(B(X, Y), \rho)$ is also complete.*

Proof. Assume that $\{f_n\}$ is a Cauchy sequence in $B(X, Y)$. We must prove that f_n converges to a function $f \in B(X, Y)$.

Consider an element $x \in X$. Since $d_Y(f_n(x), f_m(x)) \leq \rho(f_n, f_m)$ and $\{f_n\}$ is a Cauchy sequence in $(B(X, Y), \rho)$, the function values $\{f_n(x)\}$ form a Cauchy sequence in Y . Since Y is complete, $\{f_n(x)\}$ converges to a point $f(x)$ in Y . This means that $\{f_n\}$ converges *pointwise* to a function $f: X \rightarrow Y$. We must prove that $f \in B(X, Y)$ and that $\{f_n\}$ converges to f in the ρ -metric.

Since $\{f_n\}$ is a Cauchy sequence, we can for any $\epsilon > 0$ find an $N \in \mathbb{N}$ such that $\rho(f_n, f_m) < \frac{\epsilon}{2}$ when $n, m \geq N$. This means that for all $x \in X$ and all $n, m \geq N$, $d_Y(f_n(x), f_m(x)) < \frac{\epsilon}{2}$. If we let $m \rightarrow \infty$, we see that for all $x \in X$ and all $n \geq N$

$$d_Y(f_n(x), f(x)) = \lim_{m \rightarrow \infty} d_Y(f_n(x), f_m(x)) \leq \frac{\epsilon}{2}$$

Hence $\rho(f_n, f) < \epsilon$ which implies that f is bounded (since f_n is) and that $\{f_n\}$ converges uniformly to f in $B(X, Y)$. \square

The metric ρ is mainly used for theoretical purpose, and we don't have to find the exact distance between two functions very often, but in some cases it's possible using techniques you know from calculus. If X is an interval $[a, b]$ and Y is the real line (both with the usual metric), the distance $\rho(f, g)$ is just the supremum of the function $h(t) = |f(t) - g(t)|$, something you can find by differentiation (at least if the functions f and g are reasonably nice).

Exercises to Section 4.5

1. Let $f, g: [0, 1] \rightarrow \mathbb{R}$ be given by $f(x) = x$, $g(x) = x^2$. Find $\rho(f, g)$.
2. Let $f, g: [0, 2\pi] \rightarrow \mathbb{R}$ be given by $f(x) = \sin x$, $g(x) = \cos x$. Find $\rho(f, g)$.
3. Show that the two ways of defining a bounded function are equivalent (one says that the set of values $\{f(x) : x \in X\}$ is a bounded set; the other one says that for any $a \in X$, there is a constant M_a such that $d_Y(f(a), f(x)) \leq M_a$ for all $x \in X$).
4. Complete the proof of Proposition 4.5.1 by showing that ρ satisfies the first two conditions of a metric (positivity and symmetry).
5. Check the claim at the end of the proof of Theorem 4.5.3: Why does $\rho(f_n, f) < \epsilon$ imply that f is bounded when f_n is?
6. Let c_0 be the set of all bounded sequences in \mathbb{R} . If $\{x_n\}, \{y_n\}$ are in c_0 , define

$$\rho(\{x_n\}, \{y_n\}) = \sup(|x_n - y_n| : n \in \mathbb{N})$$

Show that (c_0, ρ) is a complete metric space.

7. For $f \in B(\mathbb{R}, \mathbb{R})$ and $r \in \mathbb{R}$, we define a function f_r by $f_r(x) = f(x + r)$.
- a) Show that if f is uniformly continuous, then $\lim_{r \rightarrow 0} \rho(f_r, f) = 0$.
 - b) Show that the function g defined by $g(x) = \cos(\pi x^2)$ is not uniformly continuous on \mathbb{R} .
 - c) Is it true that $\lim_{r \rightarrow 0} \rho(f_r, f) = 0$ for all $f \in B(\mathbb{R}, \mathbb{R})$?

4.6. Spaces of bounded, continuous functions

The spaces of bounded functions that we worked with in the previous section are too large for many purposes. It may sound strange that a space can be too large, but the problem is that if a space is large, it contains very little information – just knowing that a function is bounded, gives us very little to work with. Knowing that a function is continuous contains a lot more information, and for that reason we shall now turn to spaces of continuous functions. It's a bit like geography; knowing that a person is in France contains much more information than knowing she's in Europe.

As before, we assume that (X, d_X) and (Y, d_Y) are metric spaces. We define

$$C_b(X, Y) = \{f: X \rightarrow Y \mid f \text{ is continuous and bounded}\}$$

to be the collection of all bounded and continuous functions from X to Y . As $C_b(X, Y)$ is a subset of $B(X, Y)$, the metric

$$\rho(f, g) = \sup\{d_Y(f(x), g(x)) \mid x \in X\}$$

that we introduced on $B(X, Y)$ is also a metric on $C_b(X, Y)$. We make a crucial observation:

Proposition 4.6.1. *$C_b(X, Y)$ is a closed subset of $B(X, Y)$.*

Proof. By Proposition 3.3.7, it suffices to show that if $\{f_n\}$ is a sequence in $C_b(X, Y)$ that converges to an element $f \in B(X, Y)$, then $f \in C_b(X, Y)$. Since by Proposition 4.5.2 $\{f_n\}$ converges uniformly to f , Proposition 4.2.4 tells us that f is continuous and hence in $C_b(X, Y)$. \square

The next result is a more useful version of Theorem 4.5.3.

Theorem 4.6.2. *Let (X, d_X) and (Y, d_Y) be metric spaces and assume that (Y, d_Y) is complete. Then $(C_b(X, Y), \rho)$ is also complete.*

Proof. Recall from Proposition 3.4.4 that a closed subspace of a complete space is itself complete. Since $B(X, Y)$ is complete by Theorem 4.5.3, and $C_b(X, Y)$ is a closed subset of $B(X, Y)$ by the proposition above, it follows that $C_b(X, Y)$ is complete. \square

The reason why we so far have restricted ourselves to the space $C_b(X, Y)$ of *bounded*, continuous functions and not worked with the space of *all* continuous functions, is that the supremum

$$\rho(f, g) = \sup\{d_Y(f(x), g(x)) \mid x \in X\}$$

can be infinite when f and g are just assumed to be continuous. As a metric is not allowed to take infinite values, this creates problems for the theory, and the simplest

solution is to restrict ourselves to *bounded*, continuous functions. Sometimes this is a small nuisance, and it is useful to know that the problem doesn't occur when X is compact:

Proposition 4.6.3. *Let (X, d_X) and (Y, d_Y) be metric spaces, and assume that X is compact. Then all continuous functions from X to Y are bounded.*

Proof. Assume that $f: X \rightarrow Y$ is continuous, and pick a point $a \in X$. It suffices to prove that the function

$$h(x) = d_Y(f(x), f(a))$$

is bounded, and this will follow from the Extreme Value Theorem 3.5.10 if we can show that it is continuous. By the Inverse Triangle Inequality 3.1.4

$$|h(x) - h(y)| = |d_Y(f(x), a) - d_Y(f(y), a)| \leq d_Y(f(x), f(y))$$

and since f is continuous, so is h (any δ that works for f will also work for h). \square

If we define

$$C(X, Y) = \{f: X \rightarrow Y \mid f \text{ is continuous}\},$$

the proposition above tells us that for compact X , the spaces $C(X, Y)$ and $C_b(X, Y)$ coincide. In most of our applications, the underlying space X will be compact (often a closed interval $[a, b]$), and we shall then just be working with the space $C(X, Y)$. The following theorem sums up the results above for X compact.

Theorem 4.6.4. *Let (X, d_X) and (Y, d_Y) be metric spaces, and assume that X is compact. Then*

$$\rho(f, g) = \sup\{d_Y(f(x), g(x)) \mid x \in X\}$$

defines a metric on $C(X, Y)$. If (Y, d_Y) is complete, so is $(C(X, Y), \rho)$.

Exercises to Section 4.6

1. Let $X, Y = \mathbb{R}$. Find functions $f, g \in C(X, Y)$ such that

$$\sup\{d_Y(f(x), g(x)) \mid x \in X\} = \infty$$

2. Assume that $X \subset \mathbb{R}^n$ is not compact. Show that there is an unbounded, continuous function $f: X \rightarrow \mathbb{R}$.
3. Assume that $f: \mathbb{R} \rightarrow \mathbb{R}$ is a bounded continuous function. If $u \in C([0, 1], \mathbb{R})$, we define $L(u): [0, 1] \rightarrow \mathbb{R}$ to be the function

$$L(u)(t) = \int_0^1 \frac{1}{1+t+s} f(u(s)) ds$$

- a) Show that L is a function from $C([0, 1], \mathbb{R})$ to $C([0, 1], \mathbb{R})$.
- b) Assume that

$$|f(u) - f(v)| \leq \frac{C}{\ln 2} |u - v| \quad \text{for all } u, v \in \mathbb{R}$$

for some number $C < 1$. Show that the equation $Lu = u$ has a unique solution in $C([0, 1], \mathbb{R})$.

4. When X is noncompact, we have defined our metric ρ on the space $C_b(X, Y)$ of *bounded* continuous function and not on the space $C(X, Y)$ of *all* continuous functions. As mentioned in the text, the reason is that for unbounded, continuous functions,

$$\rho(f, g) = \sup\{d_Y(f(x), g(x)) \mid x \in X\}$$

may be ∞ , and a metric can not take infinite values. Restricting ourselves to $C_b(X, Y)$ is one way of overcoming this problem. Another method is to change the metric on Y such that it never occurs. We shall now take a look at this alternative method.

If (Y, d) is a metric space, we define the *truncated metric* \bar{d} by:

$$\bar{d}(x, y) = \begin{cases} d(x, y) & \text{if } d(x, y) \leq 1 \\ 1 & \text{if } d(x, y) > 1 \end{cases}$$

- a) Show that the truncated metric is indeed a metric.
- b) Show that a set $G \subseteq Y$ is open in (Y, \bar{d}) if and only if it is open in (Y, d) . What about closed sets?
- c) Show that a sequence $\{z_n\}$ in Y converges to a in the truncated metric \bar{d} if and only if it converges in the original metric d .
- d) Show that the truncated metric \bar{d} is complete if and only if the original metric is complete.
- e) Show that a set $K \subseteq Y$ is compact in (Y, \bar{d}) if and only if it is compact in (Y, d) .
- f) Show that for a metric space (X, d_X) , a function $f: X \rightarrow Y$ is continuous with respect to \bar{d} if and only if it is continuous with respect to d . Show the same for functions $g: Y \rightarrow X$.
- g) For functions $f, g \in C(X, Y)$, define

$$\bar{\rho}(f, g) = \sup\{\bar{d}(f(x), g(x)) \mid x \in X\}$$

Show that $\bar{\rho}$ is a metric on $C(X, Y)$. Show that $\bar{\rho}$ is complete if d is.

4.7. Applications to differential equations

So far it may seem that we have been creating a theory of metric spaces for its own sake. It's an impressive construction where things fit together nicely, but is it of any use?

It's time to take a look at applications, and we start by showing how Banach's Fixed Point Theorem 3.4.5 and the completeness of the spaces $C([a, b], \mathbb{R}^n)$ can be used to prove existence and uniqueness of solutions of differential equations under quite general conditions.

Consider a system of differential equations

$$\begin{aligned} y_1'(t) &= f_1(t, y_1(t), y_2(t), \dots, y_n(t)) \\ y_2'(t) &= f_2(t, y_1(t), y_2(t), \dots, y_n(t)) \\ &\vdots \\ y_n'(t) &= f_n(t, y_1(t), y_2(t), \dots, y_n(t)) \end{aligned}$$

with initial conditions $y_1(0) = Y_1, y_2(0) = Y_2, \dots, y_n(0) = Y_n$. We begin by introducing vector notation to make the formulas easier to read:

$$\mathbf{y}(t) = \begin{pmatrix} y_1(t) \\ y_2(t) \\ \vdots \\ y_n(t) \end{pmatrix}$$

$$\mathbf{y}_0 = \begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{pmatrix}$$

and

$$\mathbf{f}(t, \mathbf{y}(t)) = \begin{pmatrix} f_1(t, y_1(t), y_2(t), \dots, y_n(t)) \\ f_2(t, y_1(t), y_2(t), \dots, y_n(t)) \\ \vdots \\ f_n(t, y_1(t), y_2(t), \dots, y_n(t)) \end{pmatrix}$$

In this notation, the system becomes

$$(4.7.1) \quad \mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t)), \quad \mathbf{y}(0) = \mathbf{y}_0$$

The next step is to rewrite the differential equation as an integral equation. If we integrate on both sides of (4.7.1), we get

$$\mathbf{y}(t) - \mathbf{y}(0) = \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

i.e.

$$(4.7.2) \quad \mathbf{y}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

On the other hand, if we start with a solution of (4.7.2) and differentiate, we arrive at (4.7.1). Hence solving (4.7.1) and (4.7.2) amounts to exactly the same thing, and for us it will be convenient to concentrate on (4.7.2).

Let us begin by putting an arbitrary, continuous function \mathbf{z} into the right hand side of (4.7.2). What we get out is another function \mathbf{u} defined by

$$\mathbf{u}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{z}(s)) ds$$

We can think of this as a function F mapping continuous functions \mathbf{z} to continuous functions $\mathbf{u} = F(\mathbf{z})$. From this point of view, a solution \mathbf{y} of the integral equation (4.7.2) is just a fixed point for the function F – we are looking for a \mathbf{y} such that $\mathbf{y} = F(\mathbf{y})$. (Don't worry if you feel a little dizzy; that's just normal at this stage! Note that F is a function acting on a function \mathbf{z} to produce a new function $\mathbf{u} = F(\mathbf{z})$ – it takes some time to get used to such creatures!)

Our plan is to use Banach's Fixed Point Theorem 3.4.5 to prove that F has a unique fixed point, but first we have to introduce a crucial condition. We say that

the function $\mathbf{f}: [a, b] \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is *uniformly Lipschitz with Lipschitz constant K on the interval $[a, b]$* if K is a real number such that

$$\|\mathbf{f}(t, \mathbf{y}) - \mathbf{f}(t, \mathbf{z})\| \leq K\|\mathbf{y} - \mathbf{z}\|$$

for all $t \in [a, b]$ and all $\mathbf{y}, \mathbf{z} \in \mathbb{R}^n$. Here is the key observation in our argument.

Lemma 4.7.1. *Assume that $\mathbf{y}_0 \in \mathbb{R}^n$ and that $\mathbf{f}: [0, \infty) \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is continuous and uniformly Lipschitz with Lipschitz constant K on $[0, \infty)$. If $a < \frac{1}{K}$, the map*

$$F: C([0, a], \mathbb{R}^n) \rightarrow C([0, a], \mathbb{R}^n)$$

defined by

$$F(\mathbf{z})(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{z}(s)) ds$$

is a contraction.

Remark: The notation here is rather messy. Remember that $F(\mathbf{z})$ is a function from $[0, a]$ to \mathbb{R}^n . The expression $F(\mathbf{z})(t)$ denotes the value of this function at the point $t \in [0, a]$.

Proof of Lemma 4.7.1. Let \mathbf{v}, \mathbf{w} be two elements in $C([0, a], \mathbb{R}^n)$, and note that for any $t \in [0, a]$

$$\begin{aligned} \|F(\mathbf{v})(t) - F(\mathbf{w})(t)\| &= \left\| \int_0^t (\mathbf{f}(s, \mathbf{v}(s)) - \mathbf{f}(s, \mathbf{w}(s))) ds \right\| \leq \\ &\leq \int_0^t \|\mathbf{f}(s, \mathbf{v}(s)) - \mathbf{f}(s, \mathbf{w}(s))\| ds \leq \int_0^t K\|\mathbf{v}(s) - \mathbf{w}(s)\| ds \leq \\ &\leq K \int_0^t \rho(\mathbf{v}, \mathbf{w}) ds \leq K \int_0^a \rho(\mathbf{v}, \mathbf{w}) ds = Ka \rho(\mathbf{v}, \mathbf{w}) \end{aligned}$$

Taking the supremum over all $t \in [0, a]$, we get

$$\rho(F(\mathbf{v}), F(\mathbf{w})) \leq Ka \rho(\mathbf{v}, \mathbf{w}).$$

Since $Ka < 1$, this means that F is a contraction. \square

We are now ready for the main theorem.

Theorem 4.7.2. *Assume that $\mathbf{y}_0 \in \mathbb{R}^n$ and that $\mathbf{f}: [0, \infty) \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is continuous and uniformly Lipschitz on $[0, \infty)$. Then the initial value problem*

$$(4.7.3) \quad \mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t)), \quad \mathbf{y}(0) = \mathbf{y}_0$$

has a unique solution \mathbf{y} on $[0, \infty)$.

Proof. Let K be the uniform Lipschitz constant, and choose a number $a < 1/K$. According to the lemma, the function

$$F: C([0, a], \mathbb{R}^n) \rightarrow C([0, a], \mathbb{R}^n)$$

defined by

$$F(\mathbf{z})(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{z}(s)) ds$$

is a contraction. Since $C([0, a], \mathbb{R}^n)$ is complete by Theorem 4.6.4, Banach's Fixed Point Theorem 3.4.5 tells us that F has a unique fixed point \mathbf{y} . This means that the integral equation

$$(4.7.4) \quad \mathbf{y}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

has a unique solution on the interval $[0, a]$. To extend the solution to a longer interval, we just repeat the argument on the interval $[a, 2a]$, using $\mathbf{y}(a)$ as initial value. The function we then get, is a solution of the integral equation (4.7.4) on the extended interval $[0, 2a]$ as we for $t \in [a, 2a]$ have

$$\begin{aligned} \mathbf{y}(t) &= \mathbf{y}(a) + \int_a^t \mathbf{f}(s, \mathbf{y}(s)) ds = \\ &= \mathbf{y}_0 + \int_0^a \mathbf{f}(s, \mathbf{y}(s)) ds + \int_a^t \mathbf{f}(s, \mathbf{y}(s)) ds = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds \end{aligned}$$

Continuing this procedure to new intervals $[2a, 3a]$, $[3a, 4a]$, we see that the integral equation (4.7.4) has a unique solution on all of $[0, \infty)$. As we have already observed that equation (4.7.3) has exactly the same solutions as equation (4.7.4), the theorem is proved. \square

In the exercises you will see that the conditions in the theorem are important. If they fail, the equation may have more than one solution, or a solution defined only on a bounded interval.

Remark: The proof of Theorem 4.7.2 is based on Banach's Fixed Point Theorem, and the fixed point in that theorem is obtained by iteration. This means that the solutions of our differential equation can be approximated by iterating the map F . In numerical analysis this way of obtaining an approximate solution is referred to as *Picard iteration* in honor of Émile Picard (1856-1941).

Exercises to Section 4.7

1. Solve the initial value problem

$$y' = 1 + y^2, \quad y(0) = 0$$

and show that the solution is only defined on the interval $[0, \pi/2)$.

2. Show that all the functions

$$y(t) = \begin{cases} 0 & \text{if } 0 \leq t \leq a \\ (t - a)^{\frac{3}{2}} & \text{if } t > a \end{cases}$$

where $a \geq 0$ are solutions of the initial value problem

$$y' = \frac{3}{2}y^{\frac{1}{3}}, \quad y(0) = 0$$

Remember to check that the differential equation is satisfied at $t = a$.

3. In this problem we shall sketch how the theorem in this section can be used to study higher order systems. Assume we have a second order initial value problem

$$u''(t) = g(t, u(t), u'(t)) \quad u(0) = a, u'(0) = b \quad (*)$$

where $g: [0, \infty) \times \mathbb{R}^2 \rightarrow \mathbb{R}$ is a given function. Define a function $\mathbf{f}: [0, \infty) \times \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by

$$\mathbf{f}(t, u, v) = \begin{pmatrix} v \\ g(t, u, v) \end{pmatrix}$$

Show that if

$$\mathbf{y}(t) = \begin{pmatrix} u(t) \\ v(t) \end{pmatrix}$$

is a solution of the initial value problem

$$\mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t)), \quad \mathbf{y}(0) = \begin{pmatrix} a \\ b \end{pmatrix},$$

then u is a solution of the original problem (*).

4.8. Compact sets of continuous functions

The compact subsets of \mathbb{R}^m are easy to describe – they are just the closed and bounded sets. This characterization is extremely useful as it is much easier to check that a set is closed and bounded than to check that it satisfies the definition of compactness. In the present section, we shall prove a similar kind of characterization of compact sets in $C(X, \mathbb{R}^m)$ – we shall show that a subset of $C(X, \mathbb{R}^m)$ is compact if and only if it is closed, bounded and equicontinuous. This is known as the Arzelà-Ascoli Theorem. But before we turn to it, we have a question of independent interest to deal with. We have already encountered the notion of a dense set in Section 3.7, but repeat it here:

Definition 4.8.1. Let (X, d) be a metric space and assume that A is a subset of X . We say that A is dense in X if for each $x \in X$ there is a sequence from A converging to x .

Recall (Proposition 3.7.2) that dense sets can also be described in a slightly different way: A subset D of a metric space X is dense if and only if for each $x \in X$ and each $\delta > 0$, there is a $y \in D$ such that $d(x, y) < \delta$.

We know that \mathbb{Q} is dense in \mathbb{R} – we may, e.g., approximate a real number by longer and longer parts of its decimal expansion. For $x = \sqrt{2}$ this would mean the approximating sequence

$$a_1 = 1.4 = \frac{14}{10}, \quad a_2 = 1.41 = \frac{141}{100}, \quad a_3 = 1.414 = \frac{1414}{1000}, \quad a_4 = 1.4142 = \frac{14142}{10000}, \dots$$

Recall from Section 1.6 that \mathbb{Q} is countable, but that \mathbb{R} is not. Still every element in the uncountable set \mathbb{R} can be approximated arbitrarily well by elements in the much smaller set \mathbb{Q} . This property turns out to be so useful that it deserves a name.

Definition 4.8.2. A metric set (X, d) is called separable if it has a countable, dense subset A .

Our first result is a simple, but rather surprising connection between separability and compactness.

Proposition 4.8.3. *All compact metric spaces (X, d) are separable. We can choose the countable dense set A in such a way that for any $\delta > 0$, there is a finite subset A_δ of A such that all elements of X are within distance less than δ of A_δ , i.e. for all $x \in X$ there is an $a \in A_\delta$ such that $d(x, a) < \delta$.*

Proof. We use that a compact space X is totally bounded (recall Proposition 3.5.12). This means that for all $n \in \mathbb{N}$, there is a finite number of balls of radius $\frac{1}{n}$ that cover X . The centers of all these balls (for all $n \in \mathbb{N}$) form a countable subset A of X (to get a listing of A , first list the centers of the balls of radius 1, then the centers of the balls of radius $\frac{1}{2}$ etc.). We shall prove that A is dense in X .

Let x be an element of X . To find a sequence $\{a_n\}$ from A converging to x , we first pick the center a_1 of one of the balls (there is at least one) of radius 1 that x belongs to, then we pick the center a_2 of one of the balls of radius $\frac{1}{2}$ that x belongs to, etc. Since $d(x, a_n) < \frac{1}{n}$, $\{a_n\}$ is a sequence from A converging to x .

To find the set A_δ , just choose $m \in \mathbb{N}$ so big that $\frac{1}{m} < \delta$, and let A_δ consist of the centers of the balls of radius $\frac{1}{m}$. \square

Remark: A compactness argument shows that the last part of the proposition (about A_δ) holds for *all* countable dense subsets A of a compact space, but we shall not be needing this (see Exercise 8).

We are now ready to turn to $C(X, \mathbb{R}^m)$. First we recall the definition of equicontinuous sets of functions from Section 4.1.

Definition 4.8.4. *Let (X, d_X) and (Y, d_Y) be metric spaces, and let \mathcal{F} be a collection of functions $f: X \rightarrow Y$. We say that \mathcal{F} is equicontinuous if for all $\epsilon > 0$, there is a $\delta > 0$ such that for all $f \in \mathcal{F}$ and all $x, y \in X$ with $d_X(x, y) < \delta$, we have $d_Y(f(x), f(y)) < \epsilon$.*

We begin with a lemma that shows that for equicontinuous sequences, it suffices to check convergence on a dense set.

Lemma 4.8.5. *Assume that (X, d_X) is a compact and (Y, d_Y) a complete metric space, and let $\{g_k\}$ be an equicontinuous sequence in $C(X, Y)$. Assume that $A \subseteq X$ is a dense set as described in Proposition 4.8.3 and that $\{g_k(a)\}$ converges for all $a \in A$. Then $\{g_k\}$ converges in $C(X, Y)$.*

Proof. Since $C(X, Y)$ is complete, it suffices to prove that $\{g_k\}$ is a Cauchy sequence. Given an $\epsilon > 0$, we must thus find an $N \in \mathbb{N}$ such that $\rho(g_n, g_m) < \epsilon$ when $n, m \geq N$. Since the sequence is equicontinuous, there exists a $\delta > 0$ such that if $d_X(x, y) < \delta$, then $d_Y(g_k(x), g_k(y)) < \frac{\epsilon}{4}$ for all k . Choose a finite subset A_δ of A such that any element in X is within less than δ of an element in A_δ . Since the sequences $\{g_k(a)\}$, $a \in A_\delta$, converge, they are all Cauchy sequences, and we can find an $N \in \mathbb{N}$ such that when $n, m \geq N$, $d_Y(g_n(a), g_m(a)) < \frac{\epsilon}{4}$ for all $a \in A_\delta$ (here we are using that A_δ is finite).

For any $x \in X$, we can find an $a \in A_\delta$ such that $d_X(x, a) < \delta$. But then for all $n, m \geq N$,

$$d_Y(g_n(x), g_m(x)) \leq$$

$$\begin{aligned}
&\leq d_Y(g_n(x), g_n(a)) + d_Y(g_n(a), g_m(a)) + d_Y(g_m(a), g_m(x)) < \\
&< \frac{\epsilon}{4} + \frac{\epsilon}{4} + \frac{\epsilon}{4} = \frac{3\epsilon}{4}
\end{aligned}$$

Since this holds for any $x \in X$, we must have $\rho(g_n, g_m) \leq \frac{3\epsilon}{4} < \epsilon$ for all $n, m \geq N$, and hence $\{g_k\}$ is a Cauchy sequence and converges in the complete space $C(X, Y)$. \square

We are now ready to prove the hard part of the Arzelà-Ascoli Theorem.

Proposition 4.8.6. *Assume that (X, d) is a compact metric space, and let $\{f_n\}$ be a bounded and equicontinuous sequence in $C(X, \mathbb{R}^m)$. Then $\{f_n\}$ has a subsequence converging in $C(X, \mathbb{R}^m)$.*

Proof. Since X is compact, there is a countable, dense subset

$$A = \{a_1, a_2, \dots, a_n, \dots\}$$

as in Proposition 4.8.3. According to the lemma, it suffices to find a subsequence $\{g_k\}$ of $\{f_n\}$ such that $\{g_k(a)\}$ converges for all $a \in A$.

We begin a little less ambitiously by showing that $\{f_n\}$ has a subsequence $\{f_n^{(1)}\}$ such that $\{f_n^{(1)}(a_1)\}$ converges (recall that a_1 is the first element in our listing of the countable set A). Next we show that $\{f_n^{(1)}\}$ has a subsequence $\{f_n^{(2)}\}$ such that both $\{f_n^{(2)}(a_1)\}$ and $\{f_n^{(2)}(a_2)\}$ converge. Continuing taking subsequences in this way, we shall for each $j \in \mathbb{N}$ find a sequence $\{f_n^{(j)}\}$ such that $\{f_n^{(j)}(a)\}$ converges for $a = a_1, a_2, \dots, a_j$. Finally, we shall construct the sequence $\{g_k\}$ by combining all the sequences $\{f_n^{(j)}\}$ in a clever way.

Let us start by constructing $\{f_n^{(1)}\}$. Since the sequence $\{f_n\}$ is bounded, $\{f_n(a_1)\}$ is a bounded sequence in \mathbb{R}^m , and by the Bolzano-Weierstrass Theorem 2.3.3, it has a convergent subsequence $\{f_{n_k}(a_1)\}$. We let $\{f_n^{(1)}\}$ consist of the functions appearing in this subsequence. If we now apply $\{f_n^{(1)}\}$ to a_2 , we get a new bounded sequence $\{f_n^{(1)}(a_2)\}$ in \mathbb{R}^m with a convergent subsequence. We let $\{f_n^{(2)}\}$ be the functions appearing in this subsequence. Note that $\{f_n^{(2)}(a_1)\}$ still converges as $\{f_n^{(2)}\}$ is a subsequence of $\{f_n^{(1)}\}$. Continuing in this way, we see that we for each $j \in \mathbb{N}$ have a sequence $\{f_n^{(j)}\}$ such that $\{f_n^{(j)}(a)\}$ converges for $a = a_1, a_2, \dots, a_j$. In addition, each sequence $\{f_n^{(j)}\}$ is a subsequence of the previous ones.

We are now ready to construct a sequence $\{g_k\}$ such that $\{g_k(a)\}$ converges for all $a \in A$. We do it by a diagonal argument, putting g_1 equal to the first element in the first sequence $\{f_n^{(1)}\}$, g_2 equal to the second element in the second sequence $\{f_n^{(2)}\}$ etc. In general, the k -th term in the g -sequence equals the k -th term in the k -th f -sequence $\{f_n^{(k)}\}$, i.e. $g_k = f_k^{(k)}$. Note that except for the first few elements, $\{g_k\}$ is a subsequence of *any* sequence $\{f_n^{(j)}\}$. This means that $\{g_k(a)\}$ converges for all $a \in A$, and the proof is complete. \square

As a simple consequence of this result we get:

Corollary 4.8.7. *If (X, d) is a compact metric space, all bounded, closed and equicontinuous sets \mathcal{K} in $C(X, \mathbb{R}^m)$ are compact.*

Proof. According to the proposition, any sequence in \mathcal{K} has a convergent subsequence. Since \mathcal{K} is closed, the limit must be in \mathcal{K} , and hence \mathcal{K} is compact. \square

As already mentioned, the converse of this result is also true, but before we prove it, we need a technical lemma that is quite useful also in other situations:

Lemma 4.8.8. *Assume that (X, d_X) and (Y, d_Y) are metric spaces and that $\{f_n\}$ is a sequence of continuous function from X to Y which converges uniformly to f . If $\{x_n\}$ is a sequence in X converging to a , then $\{f_n(x_n)\}$ converges to $f(a)$.*

Remark: This lemma is not as obvious as it may seem – it is not true if we replace uniform convergence by pointwise!

Proof of Lemma 4.8.8. Given $\epsilon > 0$, we must show how to find an $N \in \mathbb{N}$ such that $d_Y(f_n(x_n), f(a)) < \epsilon$ for all $n \geq N$. Since we know from Proposition 4.2.4 that f is continuous, there is a $\delta > 0$ such that $d_Y(f(x), f(a)) < \frac{\epsilon}{2}$ when $d_X(x, a) < \delta$. Since $\{x_n\}$ converges to a , there is an $N_1 \in \mathbb{N}$ such that $d_X(x_n, a) < \delta$ when $n \geq N_1$. Also, since $\{f_n\}$ converges uniformly to f , there is an $N_2 \in \mathbb{N}$ such that if $n \geq N_2$, then $d_Y(f_n(x), f(x)) < \frac{\epsilon}{2}$ for all $x \in X$. If we choose $N = \max\{N_1, N_2\}$, we see that if $n \geq N$,

$$d_Y(f_n(x_n), f(a)) \leq d_Y(f_n(x_n), f(x_n)) + d_Y(f(x_n), f(a)) < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon$$

and the lemma is proved. \square

We are finally ready to prove the main theorem:

Theorem 4.8.9 (The Arzelà-Ascoli Theorem). *Let (X, d_X) be a compact metric space. A subset \mathcal{K} of $C(X, \mathbb{R}^m)$ is compact if and only if it is closed, bounded and equicontinuous.*

Proof. It remains to prove that a compact set \mathcal{K} in $C(X, \mathbb{R}^m)$ is closed, bounded and equicontinuous. Since compact sets are always closed and bounded according to Proposition 3.5.4, it suffices to prove that \mathcal{K} is equicontinuous. We argue by contradiction: We assume that the compact set \mathcal{K} is *not* equicontinuous and show that this leads to a contradiction.

Since \mathcal{K} is not equicontinuous, there must be an $\epsilon > 0$ which can not be matched by any δ ; i.e. for any $\delta > 0$, there is a function $f \in \mathcal{K}$ and points $x, y \in X$ such that $d_X(x, y) < \delta$, but $d_{\mathbb{R}^m}(f(x), f(y)) \geq \epsilon$. If we put $\delta = \frac{1}{n}$, we get at function $f_n \in \mathcal{K}$ and points $x_n, y_n \in X$ such that $d_X(x_n, y_n) < \frac{1}{n}$, but $d_{\mathbb{R}^m}(f_n(x_n), f_n(y_n)) \geq \epsilon$. Since \mathcal{K} is compact, there is a subsequence $\{f_{n_k}\}$ of $\{f_n\}$ which converges (uniformly) to a function $f \in \mathcal{K}$. Since X is compact, the corresponding subsequence $\{x_{n_k}\}$ of $\{x_n\}$, has a subsequence $\{x_{n_{k_j}}\}$ converging to a point $a \in X$. Since $d_X(x_{n_{k_j}}, y_{n_{k_j}}) < \frac{1}{n_{k_j}}$, the corresponding sequence $\{y_{n_{k_j}}\}$ of y 's also converges to a .

Since $\{f_{n_{k_j}}\}$ converges uniformly to f , and $\{x_{n_{k_j}}\}, \{y_{n_{k_j}}\}$ both converge to a , the lemma tells us that

$$f_{n_{k_j}}(x_{n_{k_j}}) \rightarrow f(a) \quad \text{and} \quad f_{n_{k_j}}(y_{n_{k_j}}) \rightarrow f(a)$$

But this is impossible since $d_{\mathbb{R}^m}(f(x_{n_{k_j}}), f(y_{n_{k_j}})) \geq \epsilon$ for all j . Hence we have our contradiction, and the theorem is proved. \square

In the next section we shall see how we can use the Arzelà-Ascoli Theorem to prove the existence of solutions of differential equations.

Exercises for Section 4.8

1. Show that \mathbb{R}^n is separable for all n .
2. Show that a subset A of a metric space (X, d) is dense if and only if all open balls $B(a, r)$, $a \in X$, $r > 0$, contain elements from A .
3. Assume that (X, d) is a complete metric space, and that A is a dense subset of X . We let A have the subset metric d_A .
 - a) Assume that $f: A \rightarrow \mathbb{R}$ is uniformly continuous. Explain that if $\{a_n\}$ is a sequence from A converging to a point $x \in X$, then $\{f(a_n)\}$ converges. Show that the limit is the same for all such sequences $\{a_n\}$ converging to the same point x .
 - b) Define $\bar{f}: X \rightarrow \mathbb{R}$ by putting $\bar{f}(x) = \lim_{n \rightarrow \infty} f(a_n)$ where $\{a_n\}$ is a sequence from A converging to x . We call \bar{f} the *continuous extension of f to X* . Show that \bar{f} is uniformly continuous.
 - c) Let $f: \mathbb{Q} \rightarrow \mathbb{R}$ be defined by

$$f(q) = \begin{cases} 0 & \text{if } q < \sqrt{2} \\ 1 & \text{if } q > \sqrt{2} \end{cases}$$

Show that f is continuous on \mathbb{Q} (we are using the usual metric $d_{\mathbb{Q}}(q, r) = |q - r|$). Is f uniformly continuous?

- d) Show that f does not have a continuous extension to \mathbb{R} .
4. Let K be a compact subset of \mathbb{R}^n . Let $\{f_n\}$ be a sequence of contractions of K . Show that $\{f_n\}$ has a uniformly convergent subsequence.
5. A function $f: [-1, 1] \rightarrow \mathbb{R}$ is called *Lipschitz continuous with Lipschitz constant $K \in \mathbb{R}$* if

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

for all $x, y \in [-1, 1]$. Let \mathcal{K} be the set of all Lipschitz continuous functions with Lipschitz constant K such that $f(0) = 0$. Show that \mathcal{K} is a compact subset of $C([-1, 1], \mathbb{R})$.

6. Assume that (X, d_X) and (Y, d_Y) are two metric spaces, and let $\sigma: [0, \infty) \rightarrow [0, \infty)$ be a nondecreasing, continuous function such that $\sigma(0) = 0$. We say that σ is a *modulus of continuity* for a function $f: X \rightarrow Y$ if

$$d_Y(f(u), f(v)) \leq \sigma(d_X(u, v))$$

for all $u, v \in X$.

- a) Show that a family of functions with the same modulus of continuity is equicontinuous.
- b) Assume that (X, d_X) is compact, and let $x_0 \in X$. Show that if σ is a modulus of continuity, then the set

$$\mathcal{K} = \{f: X \rightarrow \mathbb{R}^n : f(x_0) = \mathbf{0} \text{ and } \sigma \text{ is modulus of continuity for } f\}$$
 is compact.
- c) Show that all functions in $C([a, b], \mathbb{R}^m)$ has a modulus of continuity.

7. A metric space (X, d) is called *locally compact* if for each point $a \in X$, there is a *closed* ball $\bar{B}(a; r)$ centered at a that is compact. (Recall that $\bar{B}(a; r) = \{x \in X : d(a, x) \leq r\}$). Show that \mathbb{R}^m is locally compact, but that $C([0, 1], \mathbb{R})$ is not.
8. Assume that A is a dense subset of a compact metric space (X, d) . Show that for each $\delta > 0$, there is a finite subset A_δ of A such that all elements of X are within distance less than δ of A_δ , i.e. for all $x \in X$ there is an $a \in A_\delta$ such that $d(x, a) < \delta$.

4.9. Differential equations revisited

In Section 4.7 we used Banach's Fixed Point Theorem to study initial value problems of the form

$$(4.9.1) \quad \mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t)), \quad \mathbf{y}(0) = \mathbf{y}_0$$

or equivalently

$$(4.9.2) \quad \mathbf{y}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

In this section we shall see how the Arzelà-Ascoli Theorem can be used to prove existence of solutions under weaker conditions than before. But in the new approach we shall also lose something – we can only prove that the solutions exist in small intervals, and we can no longer guarantee uniqueness.

The starting point is Euler's method for finding approximate solutions to differential equations. If we want to approximate the solution starting at \mathbf{y}_0 at time $t = 0$, we begin by partitioning time into discrete steps of length Δt ; hence we work with the time line

$$T = \{t_0, t_1, t_2, t_3, \dots\}$$

where $t_0 = 0$ and $t_{i+1} - t_i = \Delta t$. We start the approximate solution $\hat{\mathbf{y}}$ at \mathbf{y}_0 and move in the direction of the derivative $\mathbf{y}'(t_0) = \mathbf{f}(t_0, \mathbf{y}_0)$, i.e. we put

$$\hat{\mathbf{y}}(t) = \mathbf{y}_0 + \mathbf{f}(t_0, \mathbf{y}_0)(t - t_0)$$

for $t \in [t_0, t_1]$. Once we reach t_1 , we change directions and move in the direction of the new derivative $\mathbf{y}'(t_1) = \mathbf{f}(t_1, \hat{\mathbf{y}}(t_1))$ so that we have

$$\hat{\mathbf{y}}(t) = \hat{\mathbf{y}}(t_1) + \mathbf{f}(t_1, \hat{\mathbf{y}}(t_1))(t - t_1)$$

for $t \in [t_1, t_2]$. If we insert the expression for $\hat{\mathbf{y}}(t_1)$, we get:

$$\hat{\mathbf{y}}(t) = \mathbf{y}_0 + \mathbf{f}(t_0, \mathbf{y}_0)(t_1 - t_0) + \mathbf{f}(t_1, \hat{\mathbf{y}}(t_1))(t - t_1)$$

If we continue in this way, changing directions at each point in T , we get

$$\hat{\mathbf{y}}(t) = \mathbf{y}_0 + \sum_{i=0}^{k-1} \mathbf{f}(t_i, \hat{\mathbf{y}}(t_i))(t_{i+1} - t_i) + \mathbf{f}(t_k, \hat{\mathbf{y}}(t_k))(t - t_k)$$

for $t \in [t_k, t_{k+1}]$. If we observe that

$$\mathbf{f}(t_i, \hat{\mathbf{y}}(t_i))(t_{i+1} - t_i) = \int_{t_i}^{t_{i+1}} \mathbf{f}(t_i, \hat{\mathbf{y}}(t_i)) ds,$$

(note that the integrand is constant) we can rewrite this expression as

$$\hat{\mathbf{y}}(t) = \mathbf{y}_0 + \sum_{i=0}^{k-1} \int_{t_i}^{t_{i+1}} \mathbf{f}(t_i, \hat{\mathbf{y}}(t_i)) ds + \int_{t_k}^t \mathbf{f}(t_k, \hat{\mathbf{y}}(t_k)) ds$$

If we also introduce the notation

$$\underline{s} = \text{the largest } t_i \in T \text{ such that } t_i \leq s,$$

we may express this more compactly as

$$\hat{\mathbf{y}}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(\underline{s}, \hat{\mathbf{y}}(\underline{s})) ds$$

Note that we can also write this as

$$\hat{\mathbf{y}}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \hat{\mathbf{y}}(s)) ds + \int_0^t (\mathbf{f}(\underline{s}, \hat{\mathbf{y}}(\underline{s})) - \mathbf{f}(s, \hat{\mathbf{y}}(s))) ds$$

where the last term measures how much $\hat{\mathbf{y}}$ “deviates” from being a solution of equation (4.9.2) (observe that there is one s and one \underline{s} term in the last integral).

Intuitively, one would think that the approximate solution $\hat{\mathbf{y}}$ will converge to a real solution \mathbf{y} when the step size Δt goes to zero. To be more specific, if we let $\hat{\mathbf{y}}_n$ be the approximate solution we get when we choose $\Delta t = \frac{1}{n}$, we would expect the sequence $\{\hat{\mathbf{y}}_n\}$ to converge to a solution of (4.9.2). It turns out that in the most general case we can not quite prove this, but we can instead use the Arzelà-Ascoli Theorem 4.8.9 to find a *subsequence* converging to a solution.

Before we turn to the proof, it will be useful to see how integrals of the form

$$I_k(t) = \int_0^t \mathbf{f}(s, \hat{\mathbf{y}}_k(s)) ds$$

behave when the functions $\hat{\mathbf{y}}_k$ converge uniformly to a limit \mathbf{y} . The following lemma is a slightly more complicated version of Proposition 4.3.1.

Lemma 4.9.1. *Let $\mathbf{f}: [0, \infty) \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a continuous function, and assume that $\{\hat{\mathbf{y}}_k\}$ is a sequence of continuous functions $\hat{\mathbf{y}}_k: [0, a] \rightarrow \mathbb{R}^m$ converging uniformly to a function \mathbf{y} . Then the integral functions*

$$I_k(t) = \int_0^t \mathbf{f}(s, \hat{\mathbf{y}}_k(s)) ds$$

converge uniformly to

$$I(t) = \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

on $[0, a]$.

Proof. Since the sequence $\{\hat{\mathbf{y}}_k\}$ converges uniformly, it is bounded, and hence there is a constant K such that $|\hat{\mathbf{y}}_k(t)| \leq K$ for all $k \in \mathbb{N}$ and all $t \in [0, a]$ (prove this!). The continuous function \mathbf{f} is uniformly continuous on the compact set $[0, a] \times [-K, K]^m$, and hence for every $\epsilon > 0$, there is a $\delta > 0$ such that if $\|\mathbf{y} - \mathbf{y}'\| < \delta$, then $\|\mathbf{f}(s, \mathbf{y}) - \mathbf{f}(s, \mathbf{y}')\| < \frac{\epsilon}{a}$ for all $s \in [0, a]$. Since $\{\hat{\mathbf{y}}_k\}$ converges uniformly to \mathbf{y} , there is an $N \in \mathbb{N}$ such that if $n \geq N$, $|\hat{\mathbf{y}}_n(s) - \mathbf{y}(s)| < \delta$ for all $s \in [0, a]$. But then

$$\begin{aligned} \|I_n(t) - I(t)\| &= \left\| \int_0^t (\mathbf{f}(s, \hat{\mathbf{y}}_n(s)) - \mathbf{f}(s, \mathbf{y}(s))) ds \right\| \leq \\ &\leq \int_0^t \|\mathbf{f}(s, \hat{\mathbf{y}}_n(s)) - \mathbf{f}(s, \mathbf{y}(s))\| ds < \int_0^a \frac{\epsilon}{a} ds = \epsilon \end{aligned}$$

for all $t \in [0, a]$, and hence $\{I_k\}$ converges uniformly to I . \square

We are now ready for the main result.

Theorem 4.9.2. *Assume that $\mathbf{f}: [0, \infty) \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a continuous function and that $\mathbf{y}_0 \in \mathbb{R}^m$. Then there exist a positive real number a and a function $\mathbf{y}: [0, a] \rightarrow \mathbb{R}^m$ such that $\mathbf{y}(0) = \mathbf{y}_0$ and*

$$\mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t)) \quad \text{for all } t \in [0, a]$$

Remark: Note that there is no uniqueness statement (the problem may have more than one solution), and that the solution is only guaranteed to exist on a bounded interval (it may disappear to infinity after finite time).

Proof of Theorem 4.9.2. Choose a big, compact subset $C = [0, R] \times [-R, R]^m$ of $[0, \infty) \times \mathbb{R}^m$ containing $(0, \mathbf{y}_0)$ in its interior. By the Extreme Value Theorem, the components of \mathbf{f} have a maximum value on C , and hence there exists a number $M \in \mathbb{R}$ such that $|f_i(t, \mathbf{y})| \leq M$ for all $(t, \mathbf{y}) \in C$ and all $i = 1, 2, \dots, m$. If the initial value has components

$$\mathbf{y}_0 = \begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_m \end{pmatrix}$$

we choose $a \in \mathbb{R}$ so small that the set

$$A = [0, a] \times [Y_1 - Ma, Y_1 + Ma] \times [Y_2 - Ma, Y_2 + Ma] \times \cdots \times [Y_m - Ma, Y_m + Ma]$$

is contained in C . This may seem mysterious, but the point is that our approximate solutions $\hat{\mathbf{y}}$ of the differential equation can never leave the area

$$[Y_1 - Ma, Y_1 + Ma] \times [Y_2 - Ma, Y_2 + Ma] \times \cdots \times [Y_m - Ma, Y_m + Ma]$$

while $t \in [0, a]$ since all the derivatives are bounded by M .

Let $\hat{\mathbf{y}}_n$ be the approximate solution obtained by using Euler's method on the interval $[0, a]$ with time step $\frac{a}{n}$. The sequence $\{\hat{\mathbf{y}}_n\}$ is bounded since $(t, \hat{\mathbf{y}}_n(t)) \in A$, and it is equicontinuous since the components of \mathbf{f} are bounded by M . By Proposition 4.8.6, $\hat{\mathbf{y}}_n$ has a subsequence $\{\hat{\mathbf{y}}_{n_k}\}$ converging uniformly to a function \mathbf{y} . If we can prove that \mathbf{y} solves the integral equation

$$\mathbf{y}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

for all $t \in [0, a]$, we shall have proved the theorem.

From the calculations at the beginning of the section, we know that

$$(4.9.3) \quad \hat{\mathbf{y}}_{n_k}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \hat{\mathbf{y}}_{n_k}(s)) ds + \int_0^t (\mathbf{f}(s, \hat{\mathbf{y}}_{n_k}(s)) - \mathbf{f}(s, \mathbf{y}(s))) ds$$

and according to the lemma

$$\int_0^t \mathbf{f}(s, \hat{\mathbf{y}}_{n_k}(s)) ds \rightarrow \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds \quad \text{uniformly for } t \in [0, a]$$

If we can only prove that

$$(4.9.4) \quad \int_0^t (\mathbf{f}(\underline{s}, \hat{\mathbf{y}}_{n_k}(\underline{s})) - \mathbf{f}(s, \hat{\mathbf{y}}_{n_k}(s))) ds \rightarrow 0,$$

we will get

$$\mathbf{y}(t) = \mathbf{y}_0 + \int_0^t \mathbf{f}(s, \mathbf{y}(s)) ds$$

as $k \rightarrow \infty$ in (4.9.3), and the theorem will be proved

To prove (4.9.4), observe that since A is a compact set, \mathbf{f} is uniformly continuous on A . Given an $\epsilon > 0$, we thus find a $\delta > 0$ such that $\|\mathbf{f}(s, \mathbf{y}) - \mathbf{f}(s', \mathbf{y}')\| < \frac{\epsilon}{a}$ when $\|(s, \mathbf{y}) - (s', \mathbf{y}')\| < \delta$ (we are measuring the distance in the ordinary \mathbb{R}^{m+1} -metric). Since

$$\|(\underline{s}, \hat{\mathbf{y}}_{n_k}(\underline{s})) - (s, \hat{\mathbf{y}}_{n_k}(s))\| \leq \|(\Delta t, M\Delta t, \dots, M\Delta t)\| = \sqrt{1 + mM^2} \Delta t,$$

we can clearly get $\|(\underline{s}, \hat{\mathbf{y}}_{n_k}(\underline{s})) - (s, \hat{\mathbf{y}}_{n_k}(s))\| < \delta$ by choosing k large enough (and hence Δt small enough). For such k we then have

$$\left\| \int_0^t (\mathbf{f}(\underline{s}, \hat{\mathbf{y}}_{n_k}(\underline{s})) - \mathbf{f}(s, \hat{\mathbf{y}}_{n_k}(s))) ds \right\| < \int_0^a \frac{\epsilon}{a} ds = \epsilon$$

and hence

$$\int_0^t (\mathbf{f}(\underline{s}, \hat{\mathbf{y}}_{n_k}(\underline{s})) - \mathbf{f}(s, \hat{\mathbf{y}}_{n_k}(s))) ds \rightarrow 0$$

as $k \rightarrow \infty$. As already observed, this completes the proof. \square

Remark: An obvious question at this stage is why didn't we extend our solution beyond the interval $[0, a]$ as we did in the proof of Theorem 4.7.2? The reason is that in the present case we do not have control over the length of our intervals, and hence the second interval may be very small compared to the first one, the third one even smaller, and so on. Even if we add an infinite number of intervals, we may still only cover a finite part of the real line. There are good reasons for this: the differential equation may only have solutions that survive for a finite amount of time. A typical example is the equation

$$y' = (1 + y^2), \quad y(0) = 0$$

where the (unique) solution $y(t) = \tan t$ goes to infinity when $t \rightarrow \frac{\pi}{2}^-$. Note also that if we solve the equation with the more general initial condition $y(0) = y_0$, we get $y(t) = \tan(t + \arctan(y_0))$, which shows that the life span of the solution depends on the initial condition y_0 .

The proof above is a relatively simple(!), but typical example of a wide class of compactness arguments in the theory of differential equations. In such arguments one usually starts with a sequence of approximate solutions and then uses compactness to extract a subsequence converging to a solution. Compactness methods are strong in the sense that they can often prove local existence of solutions under very general conditions, but they are weak in the sense that they give very little information about the nature of the solution. But just knowing that a solution exists, is often a good starting point for further explorations.

Exercises for Section 4.9

1. Prove that if $\mathbf{f}_n: [a, b] \rightarrow \mathbb{R}^m$ are continuous functions converging uniformly to a function \mathbf{f} , then the sequence $\{\mathbf{f}_n\}$ is bounded in the sense that there is a constant $K \in \mathbb{R}$ such that $\|\mathbf{f}_n(t)\| \leq K$ for all $n \in \mathbb{N}$ and all $t \in [a, b]$ (this property is used in the proof of Lemma 4.9.1).
2. Go back to exercises 1 and 2 in Section 4.7. Show that the differential equations satisfy the conditions of Theorem 4.9.2. Comment.
3. It is occasionally useful to have a slightly more general version of Theorem 4.9.2 where the solution doesn't just start a given point, but passes through it:

Theorem Assume that $\mathbf{f}: \mathbb{R} \times \mathbb{R}^m \rightarrow \mathbb{R}^m$ is a continuous function. For any $t_0 \in \mathbb{R}$ and $\mathbf{y}_0 \in \mathbb{R}^m$, there exists a positive real number a and a function $\mathbf{y}: [t_0 - a, t_0 + a] \rightarrow \mathbb{R}^m$ such that $\mathbf{y}(t_0) = \mathbf{y}_0$ and

$$\mathbf{y}'(t) = \mathbf{f}(t, \mathbf{y}(t)) \quad \text{for all } t \in [t_0 - a, t_0 + a]$$

Prove this theorem by modifying the proof of Theorem 4.9.2 (run Euler's method "backwards" on the interval $[t_0 - a, t_0]$).

4.10. Polynomials are dense in the continuous function

From calculus we know that many continuous functions can be approximated by their Taylor polynomials, but to have Taylor polynomials of all orders, a function f has to be infinitely differentiable, i.e. the higher order derivatives $f^{(k)}$ have to exist for all k . Most continuous functions are not differentiable at all, and the question is whether they still can be approximated by polynomials. In this section we shall prove:

Theorem 4.10.1 (Weierstrass' Theorem). *The polynomials are dense in the space $C([a, b], \mathbb{R})$ for all $a, b \in \mathbb{R}$, $a < b$. In other words, for each continuous function $f: [a, b] \rightarrow \mathbb{R}$, there is a sequence of polynomials $\{p_n\}$ converging uniformly to f .*

I'll offer two proofs of this theorem, but you don't have to read them both. The first proof (due to the Russian mathematician Sergei Bernstein (1880-1968)) is quite surprising; it uses probability theory to establish the result for the interval $[0, 1]$, and then a straight forward scaling argument to extend it to all closed and bounded intervals. The second proof uses traditional, analytic methods, and should be more accessible to people who don't have a background in probability theory. Also in this case we first prove the theorem for the interval $[0, 1]$.

Proof 1: The probabilistic approach

The idea is simple: Assume that you are tossing a biased coin which has probability x of coming up "heads". If you toss it more and more times, you expect the proportion of times it comes up "heads" to stabilize around x . If somebody has promised you an award of $f(X)$ dollars, where X is the actual proportion of "heads" you have had during your (say) 1000 first tosses, you would expect your award to be close to $f(x)$ (assuming that f is continuous). If the number of tosses was increased to 10 000, you would feel even more certain.

Let us formalize this: Let Y_i be the outcome of the i -th toss in the sense that Y_i has the value 0 if the coin comes up “tails” and 1 if it comes up “heads”. The proportion of “heads” in the first N tosses is then given by

$$X_N = \frac{1}{N}(Y_1 + Y_2 + \cdots + Y_N)$$

Each Y_i is binomially distributed with mean $E(Y_i) = x$ and variance $\text{Var}(Y_i) = x(1-x)$. We thus have

$$E(X_N) = \frac{1}{N}(E(Y_1) + E(Y_2) + \cdots + E(Y_N)) = x$$

and (using that the Y_i 's are independent)

$$\text{Var}(X_N) = \frac{1}{N^2}(\text{Var}(Y_1) + \text{Var}(Y_2) + \cdots + \text{Var}(Y_N)) = \frac{1}{N}x(1-x)$$

(if you don't remember these formulas from probability theory, we shall derive them by analytic methods in Exercise 6). As N goes to infinity, we would expect X_N to converge to x with probability 1. If the “award function” f is continuous, we would also expect our average award $E(f(X_N))$ to converge to $f(x)$.

To see what this has to do with polynomials, let us compute the average award $E(f(X_N))$. Since the probability of getting exactly k heads in N tosses is $\binom{N}{k}x^k(1-x)^{N-k}$, we get

$$E(f(X_N)) = \sum_{k=0}^N f\left(\frac{k}{N}\right) \binom{N}{k} x^k (1-x)^{N-k}$$

Our expectation that $E(f(X_N)) \rightarrow f(x)$ as $N \rightarrow \infty$, can therefore be rephrased as

$$\sum_{k=0}^N f\left(\frac{k}{N}\right) \binom{N}{k} x^k (1-x)^{N-k} \rightarrow f(x) \quad N \rightarrow \infty$$

If we expand the parentheses $(1-x)^{N-k}$, we see that the expressions on the right hand side are just polynomials in x , and hence we have arrived at the hypothesis that the polynomials

$$p_N(x) = \sum_{k=0}^N f\left(\frac{k}{N}\right) \binom{N}{k} x^k (1-x)^{N-k}$$

converge to $f(x)$. We shall prove that this is indeed the case, and that the convergence is uniform.

Before we turn to the proof, we need some notation and a lemma. For any random variable X with expectation x and any $\delta > 0$, we shall write

$$\mathbf{1}_{\{|x-X| \geq \delta\}} = \begin{cases} 1 & \text{if } |x-X| \geq \delta \\ 0 & \text{otherwise} \end{cases}$$

and similarly for $\mathbf{1}_{\{|x-X| < \delta\}}$.

Lemma 4.10.2 (Chebyshev's Inequality). *For a bounded random variable X with mean x ,*

$$E(\mathbf{1}_{\{|x-X| \geq \delta\}}) \leq \frac{1}{\delta^2} \text{Var}(X)$$

Proof. Since $\delta^2 \mathbf{1}_{\{|x-X| \geq \delta\}} \leq (x-X)^2$, we have

$$\delta^2 \mathbf{E}(\mathbf{1}_{\{|x-X| \geq \delta\}}) \leq \mathbf{E}((x-X)^2) = \text{Var}(X)$$

Dividing by δ^2 , we get the lemma. \square

We are now ready to prove that the Bernstein polynomials converge.

Proposition 4.10.3. *If $f: [0, 1] \rightarrow \mathbb{R}$ is a continuous function, the Bernstein polynomials*

$$p_N(x) = \sum_{k=0}^N f\left(\frac{k}{N}\right) \binom{N}{k} x^k (1-x)^{N-k}$$

converge uniformly to f on $[0, 1]$.

Proof. Given $\epsilon > 0$, we must show how to find an N such that $|f(x) - p_n(x)| < \epsilon$ for all $n \geq N$ and all $x \in [0, 1]$. Since f is continuous on the compact set $[0, 1]$, it is uniformly continuous by Proposition 4.1.2, and hence we can find a $\delta > 0$ such that $|f(u) - f(v)| < \frac{\epsilon}{2}$ whenever $|u - v| < \delta$. Since $p_n(x) = \mathbf{E}(f(X_n))$, we have

$$|f(x) - p_n(x)| = |f(x) - \mathbf{E}(f(X_n))| = |\mathbf{E}(f(x) - f(X_n))| \leq \mathbf{E}(|f(x) - f(X_n)|)$$

We split the last expectation into two parts: the cases where $|x - X_n| < \delta$ and the rest:

$$\mathbf{E}(|f(x) - f(X_n)|) = \mathbf{E}(\mathbf{1}_{\{|x-X_n| < \delta\}} |f(x) - f(X_n)|) + \mathbf{E}(\mathbf{1}_{\{|x-X_n| \geq \delta\}} |f(x) - f(X_n)|)$$

The idea is that the first term is always small since f is continuous and that the second term will be small when N is large because X_N is then unlikely to deviate much from x . Here are the details:

By choice of δ , we have for the first term

$$\mathbf{E}(\mathbf{1}_{\{|x-X_n| < \delta\}} |f(x) - f(X_n)|) \leq \mathbf{E}\left(\mathbf{1}_{\{|x-X_n| < \delta\}} \frac{\epsilon}{2}\right) \leq \frac{\epsilon}{2}$$

For the second term, we first note that since f is a continuous function on a compact interval, it must be bounded by a constant M . Hence by Chebyshev's inequality

$$\begin{aligned} \mathbf{E}(\mathbf{1}_{\{|x-X_n| \geq \delta\}} |f(x) - f(X_n)|) &\leq 2M \mathbf{E}(\mathbf{1}_{\{|x-X_n| \geq \delta\}}) \\ &\leq \frac{2M}{\delta^2} \text{Var}(X_n) = \frac{2Mx(1-x)}{\delta^2 n} \leq \frac{M}{2\delta^2 n} \end{aligned}$$

where we in the last step used that $\frac{1}{4}$ is the maximal value of $x(1-x)$ on $[0, 1]$. If we now choose $N \geq \frac{M}{\delta^2 \epsilon}$, we see that we get

$$\mathbf{E}(\mathbf{1}_{\{|x-X_n| \geq \delta\}} |f(x) - f(X_n)|) < \frac{\epsilon}{2}$$

for all $n \geq N$. Combining all the inequalities above, we see that if $n \geq N$, we have for all $x \in [0, 1]$ that

$$\begin{aligned} |f(x) - p_n(x)| &\leq \mathbf{E}(|f(x) - f(X_n)|) \\ &= \mathbf{E}(\mathbf{1}_{\{|x-X_n| < \delta\}} |f(x) - f(X_n)|) + \mathbf{E}(\mathbf{1}_{\{|x-X_n| \geq \delta\}} |f(x) - f(X_n)|) \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \end{aligned}$$

and hence the Bernstein polynomials p_n converge uniformly to f . \square

We have now proved Weierstrass' result for the interval $[0, 1]$. As already mentioned, we shall use a simple change of variable to extend it to an arbitrary interval $[a, b]$, but as we also need this change of variable argument in the analytic approach, we postpone it till after we have finished the analytic proof.

Proof 2: The analytic approach

Also in this case, we shall first prove the theorem for the interval $[0, 1]$. We first observe that it is enough to prove that all continuous functions $f: [0, 1] \rightarrow \mathbb{R}$ with $f(0) = f(1) = 0$ can be approximated by a sequence of polynomials. The reason is that if $g: [0, 1] \rightarrow \mathbb{R}$ is an arbitrary continuous function, then $f(x) = g(x) - g(1)x - g(0)(1 - x)$ satisfies the condition, and if f can be approximated uniformly by polynomials $p_n(x)$, then g can be approximated uniformly by the polynomials $q_n(x) = p_n(x) + g(1)x + g(0)(1 - x)$.

In the rest of the proof, $f: [0, 1] \rightarrow \mathbb{R}$ is a continuous function with $f(0) = f(1) = 0$, and it will be convenient to extend f to all of \mathbb{R} by letting $f(x) = 0$ for $x \notin [0, 1]$. Note that f is uniformly continuous on \mathbb{R} – it is uniformly continuous on $[0, 1]$ by Proposition 4.1.2 and the extension doesn't destroy this property.




Figure 4.10.1. The kernels P_n

We shall make use of the polynomials

$$P_n(x) = c_n(1 - x^2)^n$$

where c_n is the constant such that $\int_{-1}^1 P_n(x) dx = 1$. Note that P_n is positive on $[-1, 1]$. Figure 4.10.1 shows P_n for $n = 5$ (the most spread out curve), $n = 50$, and $n = 200$ (the spikiest curve). This behavior is essential for the proof: The integral $\int_{-1}^1 P_n(x) dx$ always equals one, but the essential contributions to the integral come from a narrower and narrower interval around 0.

Before we turn to the proof of the theorem, we need some information about the size of c_n :

Lemma 4.10.4. $c_n < \sqrt{n}$.

Proof. We shall first show that $(1 - x^2)^n \geq 1 - nx^2$ for all $x \in [-1, 1]$ (note that the right hand side is what we get if we expand the left hand side by the Binomial Theorem and only keep the first two terms). If we put

$$h(x) = (1 - x^2)^n - (1 - nx^2)$$

we have $h'(x) = n(1 - x^2)^{n-1}(-2x) + 2nx = 2nx(1 - (1 - x^2)^{n-1})$ which is positive for $x \in (0, 1]$ and negative for $x \in [-1, 0)$. Since $h(0) = 0$, this means that $h(x) \geq 0$ for all $x \in [-1, 1]$, and hence $(1 - x^2)^n \geq 1 - nx^2$. Using this, we get

$$\begin{aligned} 1 &= c_n \int_{-1}^1 (1 - x^2)^n dx \geq c_n \int_{-\frac{1}{\sqrt{n}}}^{\frac{1}{\sqrt{n}}} (1 - nx^2) dx \\ &= c_n \left[x - n \frac{x^3}{3} \right]_{-\frac{1}{\sqrt{n}}}^{\frac{1}{\sqrt{n}}} = \frac{4}{3} \frac{c_n}{\sqrt{n}} > \frac{c_n}{\sqrt{n}} \end{aligned} \quad \square$$

We now define

$$p_n(x) = \int_{-1}^1 f(x+t)P_n(t) dt$$

Note that since $\int_{-1}^1 P_n(x) dx = 1$, we can think of $p_n(x)$ as an average of values of f . As n increases, the values close to x get more and more weight (remember how the polynomials P_n look), and we would expect $p_n(x)$ to converge to $f(x)$.

We need to prove that the p_n 's are polynomials and that they converge *uniformly* to f . To see that they are polynomials, introduce a new variable $y = x + t$ and note that

$$p_n(x) = \int_{x-1}^{x+1} f(y)P_n(y-x) dy = \int_0^1 f(y)P_n(y-x) dy$$

where we in the last step have used that f is 0 outside $[0, 1]$. If you think of what happens when you expand the expression $P_n(y-x) = c_n(1 - (y-x)^2)^n$ and then integrate with respect to y , you will see that $p_n(x)$ is indeed a polynomial in x (if you don't see it, try to carry out the calculation for $n = 1$ and $n = 2$).

Proposition 4.10.5. *The polynomials p_n converge uniformly to f on $[0, 1]$.*

Proof. Combining

$$p_n(x) = \int_{-1}^1 f(x+t)P_n(t) dt$$

and

$$f(x) = f(x) \cdot 1 = f(x) \int_{-1}^1 P_n(t) dt = \int_{-1}^1 f(x)P_n(t) dt$$

we get

$$|p_n(x) - f(x)| = \left| \int_{-1}^1 (f(x+t) - f(x))P_n(t) dt \right| \leq \int_{-1}^1 |f(x+t) - f(x)|P_n(t) dt$$

We need to show that given an $\epsilon > 0$, we can get $|p_n(x) - f(x)| < \epsilon$ for all $x \in [0, 1]$ by choosing n large enough. Since f is uniformly continuous, there is a $\delta > 0$ such

that $|f(y) - f(x)| < \frac{\epsilon}{4}$ when $|y - x| < \delta$. Hence

$$\begin{aligned} |p_n(x) - f(x)| &= \int_{-1}^1 |f(x+t) - f(x)| P_n(t) dt = \int_{-1}^{-\delta} |f(x+t) - f(x)| P_n(t) dt \\ &\quad + \int_{-\delta}^{\delta} |f(x+t) - f(x)| P_n(t) dt + \int_{\delta}^1 |f(x+t) - f(x)| P_n(t) dt \\ &\leq \int_{-1}^{-\delta} |f(x+t) - f(x)| P_n(t) dt + \frac{\epsilon}{2} + \int_{\delta}^1 |f(x+t) - f(x)| P_n(t) dt \end{aligned}$$

Let us take a closer look at the term $\int_{\delta}^1 |f(x+t) - f(x)| P_n(t) dt$. If M is the supremum of $f(x)$, $x \in [0, 1]$, we have

$$\begin{aligned} \int_{\delta}^1 |f(x+t) - f(x)| P_n(t) dt &\leq 2M \int_{\delta}^1 P_n(t) dt = \\ &= 2M \int_{\delta}^1 c_n(1-t^2)^n dt \leq 2M \int_{\delta}^1 \sqrt{n}(1-\delta^2)^n dt \leq 2M\sqrt{n}(1-\delta^2)^n \end{aligned}$$

By a totally similar argument,

$$\int_{-1}^{-\delta} |f(x+t) - f(x)| P_n(t) dt \leq 2M\sqrt{n}(1-\delta^2)^n,$$

and hence

$$|p_n(x) - f(x)| \leq 4M\sqrt{n}(1-\delta^2)^n + \frac{\epsilon}{2}$$

As $\sqrt{n}(1-\delta^2)^n \rightarrow 0$ when $n \rightarrow \infty$ (you can check this by L'Hôpital's rule if you don't see it immediately), we can get $4M\sqrt{n}(1-\delta^2)^n$ less than $\frac{\epsilon}{2}$ by choosing n large enough. Hence p_n converges uniformly to f and the proposition is proved. \square

Extension to arbitrary intervals

To get Weierstrass' result on a general interval, we just have to move functions from the interval $[a, b]$ to $[0, 1]$ and back. The function

$$T(x) = \frac{x-a}{b-a}$$

maps $[a, b]$ bijectively to $[0, 1]$, and the inverse function

$$T^{-1}(y) = a + (b-a)y$$

maps $[0, 1]$ back to $[a, b]$. If f is a continuous function on $[a, b]$, the function $\hat{f} = f \circ T^{-1}$ is a continuous function on $[0, 1]$ taking exactly the same values in the same order. If $\{q_n\}$ is a sequence of polynomials converging uniformly to \hat{f} on $[0, 1]$, then the functions $p_n = q_n \circ T$ converge uniformly to f on $[a, b]$. Since

$$p_n(x) = q_n\left(\frac{x-a}{b-a}\right)$$

the p_n 's are polynomials, and hence Weierstrass' theorem is proved.

Remark: Weierstrass' Theorem is important because many mathematical arguments are easier to perform on polynomials than on continuous functions in general. If the property we study is preserved under uniform limits (i.e. if the limit f of a

uniformly convergent sequence of functions $\{f_n\}$ always inherits the property from the f_n 's), we can use Weierstrass' Theorem to extend the argument from polynomials to all continuous functions. In the next section, we shall study an extension of the result called the Stone-Weierstrass Theorem which generalizes the result to many more settings.

Exercises for Section 4.10

1. Show that there is no sequence of polynomials that converges uniformly to the continuous function $f(x) = \frac{1}{x}$ on $(0, 1)$.
2. Show that there is no sequence of polynomials that converges uniformly to the function $f(x) = e^x$ on \mathbb{R} .
3. In this problem

$$f(x) = \begin{cases} e^{-1/x^2} & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$

- a) Show that if $x \neq 0$, then the n -th derivative has the form

$$f^{(n)}(x) = e^{-1/x^2} \frac{P_n(x)}{x^{N_n}}$$

where P_n is a polynomial and $N_n \in \mathbb{N}$.

- b) Show that $f^{(n)}(0) = 0$ for all n .
- c) Show that the Taylor polynomials of f at 0 do not converge to f except at the point 0.
4. Assume that $f: [a, b] \rightarrow \mathbb{R}$ is a continuous function such that $\int_a^b f(x)x^n dx = 0$ for all $n = 0, 1, 2, 3, \dots$
 - a) Show that $\int_a^b f(x)p(x) dx = 0$ for all polynomials p .
 - b) Use Weierstrass' theorem to show that $\int_a^b f(x)^2 dx = 0$. Conclude that $f(x) = 0$ for all $x \in [a, b]$.
5. In this exercise we shall show that $C([a, b], \mathbb{R})$ is a separable metric space, i.e. that it has a countable, dense subset.
 - a) Assume that (X, d) is a metric space, and that $S \subseteq T$ are subsets of X . Show that if S is dense in (T, d_T) and T is dense in (X, d) , then S is dense in (X, d) .
 - b) Show that for any polynomial p , there is a sequence $\{q_n\}$ of polynomials with rational coefficients that converges uniformly to p on $[a, b]$.
 - c) Show that the polynomials with rational coefficients are dense in $C([a, b], \mathbb{R})$.
 - d) Show that $C([a, b], \mathbb{R})$ is separable.
6. In this problem we shall reformulate Bernstein's proof in purely analytic terms, avoiding concepts and notation from probability theory. You should keep the Binomial Formula

$$(a + b)^N = \sum_{k=0}^N \binom{N}{k} a^k b^{N-k}$$

and the definition $\binom{N}{k} = \frac{N(N-1)(N-2)\dots(N-k+1)}{1 \cdot 2 \cdot 3 \cdot \dots \cdot k}$ in mind.

- a) Show that $\sum_{k=0}^N \binom{N}{k} x^k (1-x)^{N-k} = 1$.
- b) Show that $\sum_{k=0}^N \frac{k}{N} \binom{N}{k} x^k (1-x)^{N-k} = x$ (this is the analytic version of the probabilistic formula $E(X_N) = \frac{1}{N}(E(Y_1) + E(Y_2) + \dots + E(Y_N)) = x$)

- c) Show that $\sum_{k=0}^N \left(\frac{k}{N} - x\right)^2 \binom{N}{k} x^k (1-x)^{N-k} = \frac{1}{N} x(1-x)$ (this is the analytic version of $\text{Var}(X_N) = \frac{1}{N} x(1-x)$). *Hint:* Write

$$\left(\frac{k}{N} - x\right)^2 = \frac{1}{N^2} (k(k-1) + (1-2xN)k + N^2 x^2)$$

and use points b) and a) on the second and third term in the sum.

- d) Show that if p_n is the n -th Bernstein polynomial, then

$$|f(x) - p_n(x)| \leq \sum_{k=0}^n |f(x) - f(k/n)| \binom{n}{k} x^n (1-x)^{n-k}$$

- e) Given $\epsilon > 0$, explain why there is a $\delta > 0$ such that $|f(u) - f(v)| < \epsilon/2$ for all $u, v \in [0, 1]$ such that $|u - v| < \delta$. Explain why

$$\begin{aligned} |f(x) - p_n(x)| &\leq \sum_{\{k: |\frac{k}{n} - x| < \delta\}} |f(x) - f(k/n)| \binom{n}{k} x^n (1-x)^{n-k} + \\ &+ \sum_{\{k: |\frac{k}{n} - x| \geq \delta\}} |f(x) - f(k/n)| \binom{n}{k} x^n (1-x)^{n-k} \leq \\ &< \frac{\epsilon}{2} + \sum_{\{k: |\frac{k}{n} - x| \geq \delta\}} |f(x) - f(k/n)| \binom{n}{k} x^n (1-x)^{n-k} \end{aligned}$$

- f) Show that there is a constant M such that $|f(x)| \leq M$ for all $x \in [0, 1]$. Explain all the steps in the calculation:

$$\begin{aligned} &\sum_{\{k: |\frac{k}{n} - x| \geq \delta\}} |f(x) - f(k/n)| \binom{n}{k} x^n (1-x)^{n-k} \leq \\ &\leq 2M \sum_{\{k: |\frac{k}{n} - x| \geq \delta\}} \binom{n}{k} x^n (1-x)^{n-k} \leq \\ &\leq 2M \sum_{k=0}^n \left(\frac{\frac{k}{n} - x}{\delta}\right)^2 \binom{n}{k} x^n (1-x)^{n-k} \leq \frac{2M}{n\delta^2} x(1-x) \leq \frac{M}{2n\delta^2} \end{aligned}$$

- g) Explain why we can get $|f(x) - p_n(x)| < \epsilon$ by choosing n large enough, and explain why this proves Proposition 4.10.3.

4.11. The Stone-Weierstrass Theorem

In this section, we shall generalize Weierstrass' Theorem from an interval $[a, b]$ to a general, compact metric space X . In the new setting, we don't have any polynomials, and the key problem is to figure out what properties of the polynomials make Weierstrass' Theorem work. It turns out that the central notion is that of an algebra of functions.

Definition 4.11.1. Let (X, d) be a compact metric space. A nonempty subset \mathcal{A} of $C(X, \mathbb{R})$ is called an algebra (of functions) if the following conditions are satisfied:

- (i) If $f \in \mathcal{A}$ and $c \in \mathbb{R}$, then $cf \in \mathcal{A}$.
- (ii) If $f, g \in \mathcal{A}$, then $f + g \in \mathcal{A}$ and $fg \in \mathcal{A}$.

We say that \mathcal{A} is closed if it is closed as a subset of $C(X, \mathbb{R})$.

It's useful to take a look at some examples:

Example 1: The set \mathcal{P} of all polynomials form an algebra in $C([a, b], \mathbb{R})$. This is obvious since if we multiply a polynomial by a number, we still have a polynomial, and the same is the case if we add or multiply two polynomials. By Weierstrass' Theorem, \mathcal{P} is dense in $C([a, b], \mathbb{R})$, and hence the closure of \mathcal{P} is all of $C([a, b], \mathbb{R})$. Since there are continuous functions that are not polynomials, this means that \mathcal{P} is not closed. ♣

The next two examples will be important motivations for the conditions we shall introduce later in the section.

Example 2: Let X be a compact metric space, and choose a point $a \in X$. Then

$$\mathcal{A} = \{f \in C(X, \mathbb{R}) : f(a) = 0\}$$

is a closed algebra (you will prove this in Exercise 1). Note that \mathcal{A} is not all of $C(X, \mathbb{R})$. ♣

Example 3: Let X be a compact metric space, and choose two different points $a, b \in X$. Then

$$\mathcal{A} = \{f \in C(X, \mathbb{R}) : f(a) = f(b)\}$$

is a closed algebra (you will prove this in Exercise 2). Note that \mathcal{A} is not all of $C(X, \mathbb{R})$. ♣

We are now ready to begin. Our first task is to sort out some important properties of algebras.

Lemma 4.11.2. *Assume that (X, d) is a compact metric space and that \mathcal{A} is an algebra of functions in $C(X, \mathbb{R})$. If*

$$P(t) = a_n t^n + a_{n-1} t^{n-1} + \cdots + a_1 t$$

is a polynomial with constant term 0, and f is an element in \mathcal{A} , then the function

$$(P \circ f)(x) = a_n f(x)^n + a_{n-1} f(x)^{n-1} + \cdots + a_1 f(x)$$

is in \mathcal{A} .

Proof. First note that since constant functions are not necessarily in \mathcal{A} , we have to assume that the polynomials have constant term 0. Once this has been observed, the lemma follows immediately from the definition of an algebra. If you want to, you can make a formal proof by induction on the degree of the polynomial. □

For the next lemmas, we need the algebra to be closed.

Lemma 4.11.3. *Assume that (X, d) is a compact metric space and that \mathcal{A} is a closed algebra of functions in $C(X, \mathbb{R})$. If $f \in \mathcal{A}$, then $|f| \in \mathcal{A}$.*

Proof. Note that since \mathcal{A} is closed, it suffices to show that for any $\epsilon > 0$, there is a $g \in \mathcal{A}$ such that $\rho(|f|, g) < \epsilon$. Note also that since X is compact, the Extreme Value

Theorem 3.5.10 tells us that f is bounded, and hence there is an $N \in \mathbb{R}$ such that $-N \leq f(x) \leq N$ for all $x \in X$. Applying Weierstrass' theorem for $C([-N, N], \mathbb{R})$ to the absolute value function $t \mapsto |t|$, we find a polynomial p such that $||t| - p(t)| < \frac{\epsilon}{2}$ for all $t \in [-N, N]$. Note that $|p(0)| < \frac{\epsilon}{2}$, and hence $P(x) = p(x) - p(0)$ is a polynomial with constant term 0 such that $||t| - P(t)| < \epsilon$ for all $t \in [-N, N]$ (the point of this maneuver is that we need a polynomial with constant term 0 to apply the lemma above).

By construction, $||f(x)| - (P \circ f)(x)| < \epsilon$ for all $x \in X$ and hence $\rho(|f|, P \circ f) < \epsilon$. Since $P \circ f \in \mathcal{A}$ by the previous lemma, we have found what we were looking for. \square

The next lemma explains why it was important to get hold of the absolute values.

Lemma 4.11.4. *Assume that (X, d) is a compact metric space and that \mathcal{A} is a closed algebra of functions in $C(X, \mathbb{R})$. If the functions f_1, f_2, \dots, f_n are in \mathcal{A} , then the functions f_{\max} and f_{\min} defined by*

$$f_{\max}(x) = \max\{f_1(x), f_2(x), \dots, f_n(x)\}$$

and

$$f_{\min}(x) = \min\{f_1(x), f_2(x), \dots, f_n(x)\}$$

are also in \mathcal{A} .

Proof. If we only have two functions, f_1 and f_2 , this follows immediately from the formulas

$$f_{\max}(x) = \frac{1}{2}(f_1(x) + f_2(x)) + \frac{1}{2}|f_1(x) - f_2(x)|$$

and

$$f_{\min}(x) = \frac{1}{2}(f_1(x) + f_2(x)) - \frac{1}{2}|f_1(x) - f_2(x)|$$

By what we have already proved, it is easy to see that f_{\max} and f_{\min} are in \mathcal{A} . We can now extend to more than two functions by induction (see Exercise 4 for a little help). \square

It's time to take a closer look at what we are aiming for: We want to find conditions that guarantee that our closed algebra \mathcal{A} is all of $C(X, \mathbb{R})$. Examples 2 and 3 are obvious stumbling blocks as they describe closed algebras that are *not* all of $C(X, \mathbb{R})$, but it turns out that they are the only obstacles in our way. Here is the terminology we need:

Definition 4.11.5. *Assume that (X, d) is a compact metric space and that \mathcal{A} is an algebra of functions in $C(X, \mathbb{R})$. We say that:*

- a) \mathcal{A} separates points if for all distinct points $a, b \in X$, there is a function $f \in \mathcal{A}$ such that $f(a) \neq f(b)$.
- b) \mathcal{A} does not vanish anywhere if for all $a \in X$, there is a function $g \in \mathcal{A}$ such that $g(a) \neq 0$.

We shall prove two versions of the main theorem. Here is the first one:

Theorem 4.11.6 (The Stone-Weierstrass Theorem, version 1). *Assume that (X, d) is a compact metric space and that \mathcal{A} is a closed algebra of functions in $C(X, \mathbb{R})$ that separates points and does not vanish anywhere. Then $\mathcal{A} = C(X, \mathbb{R})$.*

Before we turn to the proof of the theorem, we need two more technical results.

Lemma 4.11.7. *Assume that (X, d) is a compact metric space and that \mathcal{A} is an algebra of functions in $C(X, \mathbb{R})$ that separates points and does not vanish anywhere. If a, b are two distinct points in X , there is a function $v \in \mathcal{A}$ such that $v(a) = 1$ and $v(b) = 0$.*

Proof. Assume we can find a function $u \in \mathcal{A}$ such that $u(a) \neq u(b)$ and $u(a) \neq 0$, then

$$v(x) = \frac{u(x)^2 - u(b)u(x)}{u(a)^2 - u(b)u(a)}$$

will do the job.

To construct u , observe that since \mathcal{A} separates points and does not vanish anywhere, there are function $f, g \in \mathcal{A}$ such that $f(a) \neq f(b)$ and $g(a) \neq 0$. If $f(a) \neq 0$, we can just put $u = f$, and hence we can concentrate on the case where $f(a) = 0$. The plan is to put

$$u(x) = f(x) + \lambda g(x)$$

for a suitable nonzero constant λ . Since $\lambda \neq 0$, we automatically have $u(a) = f(a) + \lambda g(a) = \lambda g(a) \neq 0$, and we only have to choose λ such that $u(a) \neq u(b)$. In order to have $u(a) = u(b)$, we need $\lambda g(a) = f(b) + \lambda g(b)$, i.e. $\lambda(g(a) - g(b)) = f(b)$. As $f(b) \neq f(a) = 0$, it is clearly possible to choose $\lambda \neq 0$ such that this does not happen (indeed, every λ except $\frac{f(b)}{g(a) - g(b)}$ will do). \square

Corollary 4.11.8. *Assume that (X, d) is a compact metric space and that \mathcal{A} is an algebra of functions in $C(X, \mathbb{R})$ that separates points and does not vanish anywhere. If a, b are two distinct points in X and α, β are two real numbers, there is a function $f \in \mathcal{A}$ such that $f(a) = \alpha$ and $f(b) = \beta$.*

Proof. By the lemma above, there are functions $v_1, v_2 \in \mathcal{A}$ such that $v_1(a) = 1$, $v_1(b) = 0$ and $v_2(b) = 1$, $v_2(a) = 0$. We now just put $f(x) = \alpha v_1(x) + \beta v_2(x)$. \square

We are now ready for the proof of the theorem. It's a quite instructive, double compactness argument using the open covering description of compactness (see Theorem 3.6.4).

Proof of Theorem 4.11.6. Assume that $f \in C(X, \mathbb{R})$. Since \mathcal{A} is closed, it suffices to show that given an $\epsilon > 0$, there is a function $h \in \mathcal{A}$ such that $|f(y) - h(y)| < \epsilon$ for all $y \in X$.

In the first part of the proof, we shall prove that for each $x \in X$ there is a function $g_x \in \mathcal{A}$ such that $g_x(x) = f(x)$ and $g_x(y) > f(y) - \epsilon$ for all $y \in X$. To this end, note that by the corollary above, we can for each $z \in X$ find a function $h_z \in \mathcal{A}$ such that $h_z(x) = f(x)$ and $h_z(z) = f(z)$. Since h_z and f are continuous, there is an open neighborhood O_z of z where $h_z(y) > f(y) - \epsilon$ for all $y \in O_z$. The family $\{O_z\}_{z \in X}$ is an open covering of the compact set X ,

and by Theorem 3.6.4, there is a finite subcovering $O_{z_1}, O_{z_2}, \dots, O_{z_n}$. If we put $g_x(y) = \max\{h_{z_1}(y), h_{z_2}(y), \dots, h_{z_n}(y)\}$, we have $g_x \in \mathcal{A}$ by Lemma 4.11.4, and by construction, $g_x(x) = f(x)$ and $g_x(y) > f(y) - \epsilon$ for all $y \in X$.

We are now ready to construct a function $h \in \mathcal{A}$ such that $|f(y) - h(y)| < \epsilon$ for all $y \in X$. Since $g_x(x) = f(x)$ and g_x and f are continuous, there is an open neighborhood G_x of x such that $g_x(y) < f(y) + \epsilon$ for all $y \in G_x$. The family $\{G_x\}_{x \in X}$ is an open covering of the compact set X , and by Theorem 3.6.4, there is a finite subcovering $G_{x_1}, G_{x_2}, \dots, G_{x_m}$. If we put $h(y) = \min\{g_{x_1}(y), g_{x_2}(y), \dots, g_{x_m}(y)\}$, h is in \mathcal{A} by Lemma 4.11.4. By construction $f(y) - \epsilon < h(y) < f(y) + \epsilon$ for all $y \in X$, and hence we have found our function h . \square

There is a problem with Theorem 4.11.6: in practice, our algebras are seldom closed. The following variation of the theorem fixes this problem, and it also makes the result look more like Weierstrass' Theorem.

Theorem 4.11.9 (The Stone-Weierstrass Theorem, version 2). *Assume that (X, d) is a compact metric space and that \mathcal{A} is an algebra of functions in $C(X, \mathbb{R})$ that separates points and does not vanish anywhere. Then \mathcal{A} is dense in $C(X, \mathbb{R})$.*

Proof. Let $\bar{\mathcal{A}}$ be the closure of \mathcal{A} as a subset of $C(X, \mathbb{R})$. If $\bar{\mathcal{A}}$ is an algebra, the first version of the Stone-Weierstrass Theorem applies and tells us that $\bar{\mathcal{A}} = C(X, \mathbb{R})$, which means that \mathcal{A} is dense in $C(X, \mathbb{R})$. Hence it suffices to prove that $\bar{\mathcal{A}}$ is an algebra.

I only sketch the argument and leave the details to the reader (see Exercise 8 for help). That $f \in \bar{\mathcal{A}}$ means that there is a sequence $\{f_n\}$ of functions in \mathcal{A} that converges uniformly to f . To show that the conditions in Definition 4.11.1 are satisfied, it suffices to show that if $\{f_n\}$ and $\{g_n\}$ are sequences in $C(X, \mathbb{R})$ that converge uniformly to f and g , respectively, then $\{cf_n\}$ converges uniformly to cf for all constants c , and $\{f_n + g_n\}$ and $\{f_n g_n\}$ converge uniformly to $f + g$ and fg , respectively. \square

As an example of how the Stone-Weierstrass Theorem is used, let us see how we can apply it to extend Weierstrass' Theorem to two dimensions.

Example 4: A polynomial in two variables is a function $p: \mathbb{R}^2 \rightarrow \mathbb{R}$ of the form

$$p(x, y) = \sum_{\substack{0 \leq n \leq N \\ 0 \leq m \leq M}} c_{nm} x^n y^m$$

If $a < b$ and $c < d$, we want to show that the set \mathcal{P} of all such polynomials is dense in $C([a, b] \times [c, d], \mathbb{R})$. This follows from the Stone-Weierstrass Theorem as \mathcal{P} is an algebra that separates points and does not vanish anywhere (check this!). \clubsuit

There is also a complex-valued version of the Stone-Weierstrass Theorem that is sometimes useful. We are now interested in subalgebras \mathcal{A} of the space $C(X, \mathbb{C})$ of all continuous, *complex-valued* functions on our compact space X . An algebra is defined just as before (see Definition 4.11.1) except that the functions and the

numbers are allowed to take complex values. It turns out that in the complex case, we need one more condition.

Definition 4.11.10. *A subset \mathcal{A} of $C(X, \mathbb{C})$ is closed under conjugation if $\bar{f} \in \mathcal{A}$ whenever $f \in \mathcal{A}$ (here \bar{f} is the function defined from f by complex conjugation: $\bar{f}(x) = \overline{f(x)}$).*

We are now ready to state and prove the complex version of our theorem. Fortunately, we don't have to start anew, but can reduce the problem to the real case.

Theorem 4.11.11 (The Stone-Weierstrass Theorem, complex version). *Assume that (X, d) is a compact metric space and that \mathcal{A} is an algebra of functions in $C(X, \mathbb{C})$ that is closed under conjugation, separates points, and does not vanish anywhere. Then \mathcal{A} is dense in $C(X, \mathbb{C})$.*

Proof. If $\mathcal{A}_{\mathbb{R}}$ denotes the set of all real-valued functions in \mathcal{A} , then $\mathcal{A}_{\mathbb{R}}$ is clearly a subalgebra of $C(X, \mathbb{R})$. Observe also that if $f = u + iv$ is a function in \mathcal{A} , its real part u and imaginary part v belong to $\mathcal{A}_{\mathbb{R}}$. This is because the adjoint $\bar{f} = u - iv$ belongs to \mathcal{A} by assumption, and

$$u = \frac{1}{2}(f + \bar{f}) \quad \text{and} \quad v = \frac{1}{2i}(f - \bar{f})$$

The next step is to show that $\mathcal{A}_{\mathbb{R}}$ separates points and doesn't vanish anywhere. Since \mathcal{A} separates points, the complex version of Lemma 4.11.7 (which can be proved exactly like the real case) tells us that if a, b are two distinct points in X , there is a function f in \mathcal{A} such that $f(a) = 1$ and $f(b) = 0$. As the real part u of f is in $\mathcal{A}_{\mathbb{R}}$ and satisfies $u(a) = 1$ and $u(b) = 0$, $\mathcal{A}_{\mathbb{R}}$ separates points. Since \mathcal{A} doesn't vanish anywhere, there is for each $a \in X$ a function $f \in \mathcal{A}$ such that $f(a) = \gamma \neq 0$. Then $g = \bar{\gamma}f$ is a function in \mathcal{A} such that $g(a) = |\gamma|^2 \in \mathbb{R} \setminus \{0\}$, and hence the real part of g is a function in $\mathcal{A}_{\mathbb{R}}$ that doesn't vanish at a . As a was an arbitrary point in X , this shows that $\mathcal{A}_{\mathbb{R}}$ doesn't vanish anywhere.

Theorem 4.11.9 now tells us that $\mathcal{A}_{\mathbb{R}}$ is dense in $C(X, \mathbb{R})$. This means that for any function $w \in C(X, \mathbb{R})$, there is a sequence $\{w_n\}$ of functions in $\mathcal{A}_{\mathbb{R}}$ that converges uniformly to w . If $f = u + iv$ is a function in $C(X, \mathbb{C})$, we can thus find sequences $\{u_n\}$ and $\{v_n\}$ in $\mathcal{A}_{\mathbb{R}}$ that converge uniformly to u and v , respectively, and hence the functions $f_n = u_n + iv_n$ form a sequence of functions in \mathcal{A} that converges uniformly to f . This shows that \mathcal{A} is dense in $C(X, \mathbb{C})$, and the theorem is proved. \square

In the proof above, we used the assumption that \mathcal{A} is closed under conjugation to show that the real and complex part of a function in \mathcal{A} belongs to $\mathcal{A}_{\mathbb{R}}$. In Exercise 12 you will find an example that shows that the theorem is false if remove the conjugation assumption.

We shall end with a look at an application of the complex Stone-Weierstrass Theorem that is useful in Fourier analysis (we shall return to it in Chapter ??). Let

$$X = \{\mathbf{x} \in \mathbb{R}^2 : \|\mathbf{x}\| = 1\}$$

be the unit circle in \mathbb{R}^2 . Using polar coordinates, we can use a single number $\theta \in [-\pi, \pi]$ to describe a point on the circle.² Note that $\theta = -\pi$ and $\theta = \pi$ describes the same point on the circle, and hence any function f in $C(X, \mathbb{C})$ corresponds in a natural way to a function f in $C([-\pi, \pi], \mathbb{C})$ with $f(-\pi) = f(\pi)$. We shall move back and forth between these two representations without further ado, sometimes thinking of the functions as defined on X and sometimes as defined on $[-\pi, \pi]$.

A *trigonometric polynomial* is a function defined on X by an expression of the form

$$p(\theta) = \sum_{n=-N}^N c_n e^{in\theta}$$

where $N \in \mathbb{N}$ and the c_n 's are complex numbers. This may seem a strange name, but note that if we use the identities

$$e^{in\theta} = (e^{i\theta})^n = (\cos \theta + i \sin \theta)^n$$

and multiply out all parentheses, p becomes a complex polynomial in $\cos \theta$ and $\sin \theta$. Note also that $p(-\pi) = p(\pi)$, and that p hence can be thought of as a continuous function on X .

If we add or multiply two trigonometric polynomials, we get a new trigonometric polynomial, and hence the set \mathcal{A} of all trigonometric polynomials is an algebra in $C(X, \mathbb{C})$. As the conjugate

$$\overline{p(\theta)} = \overline{\sum_{n=-N}^N c_n e^{in\theta}} = \sum_{n=-N}^N \overline{c_n} e^{-in\theta}$$

is also a trigonometric polynomial, \mathcal{A} is closed under conjugation, and since \mathcal{A} contains all constant functions, it doesn't vanish anywhere. To show that \mathcal{A} separates points, note that $q(\theta) = e^{i\theta} = \cos \theta + i \sin \theta$ is a trigonometric polynomial, and that $q(a) \neq q(b)$ if a and b are two different points in $(-\pi, \pi]$. Hence all the conditions of the complex version of the Stone-Weierstrass theorem are satisfied, and we have proved:

Proposition 4.11.12. *If X is the unit circle in \mathbb{R}^2 , the trigonometric polynomials are dense in $C(X, \mathbb{C})$.*

We can reformulate this result as a statement about periodic functions on the interval $[-\pi, \pi]$ that will be handy when we get to Fourier analysis.

Corollary 4.11.13. *If C_P is the set of all continuous functions $f: [-\pi, \pi] \rightarrow \mathbb{C}$ such that $f(-\pi) = f(\pi)$, then the trigonometric polynomials are dense in C_P .*

Exercises for Section 4.11

1. Show that \mathcal{A} in Example 2 is a closed algebra.
2. Show that \mathcal{A} in Example 3 is a closed algebra.
3. Use an induction argument to give a detailed proof of Lemma 4.11.2.

²The reason why I am using $[-\pi, \pi]$ instead of the more natural interval $[0, 2\pi]$, is that it fits in better with what we shall later do in Fourier analysis.

4. Carry out the induction arguments in the proof of Lemma 4.11.4. It may be useful to observe that

$$\max\{f_1, f_2, \dots, f_k, f_{k+1}\} = \max\{\max\{f_1, f_2, \dots, f_k\}, f_{k+1}\}$$

and similarly for min.

5. Let \mathcal{A} be the algebra of all polynomials. Show that \mathcal{A} separates points and doesn't vanish anywhere in $C([0, 1], \mathbb{R})$.
6. Show that the set \mathcal{P} in Example 4 is an algebra that separates points and does not vanish anywhere.
7. Explain carefully why the function h in the proof of Theorem 4.11.6 satisfies $f(y) - \epsilon < h(y) < f(y) + \epsilon$ for all $y \in X$.
8. Assume that (X, d) is a compact space, and assume that $\{f_n\}$ and $\{g_n\}$ are sequences in $C(X, \mathbb{R})$ that converge to f and g , respectively.
- Show that for any real number c , the sequence $\{cf_n\}$ converges to cf in $C(X, \mathbb{R})$.
 - Show that $\{f_n + g_n\}$ converges to $f + g$ in $C(X, \mathbb{R})$.
 - Show that $\{f_n g_n\}$ converges to fg in $C(X, \mathbb{R})$ (note that since X is compact, all the functions are bounded).
 - Write out the proof of Theorem 4.11.9 in full detail.
9. A complex polynomial is a function of the form $p(x) = c_n x^n + c_{n-1} x^{n-1} + \dots + c_1 x + c_0$ where $c_n, c_{n-1}, \dots, c_1, c_0 \in \mathbb{C}$. Show that the complex polynomials are dense in $C([0, 1], \mathbb{C})$. (*Warning:* It's important here that we are working over a *real* interval $[0, 1]$. It's tempting to assume that the result will continue to hold if we replace $[0, 1]$ by any compact subset X of \mathbb{C} , but that's not the case – see Exercise 12f) below.)
10. Assume that (X, d_X) and (Y, d_Y) are two compact metric spaces, and let d be the metric on $X \times Y$ defined by

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

- Show that $(X \times Y, d)$ is compact.
- Let \mathcal{A} consist of all function $h: X \times Y \rightarrow \mathbb{R}$ of the form

$$h(x, y) = \sum_{i=1}^N f_i(x)g_i(y)$$

where $f_i: X \rightarrow \mathbb{R}$, $g_i: Y \rightarrow \mathbb{R}$ are continuous functions. Show that \mathcal{A} is an algebra of continuous functions.

- Show that \mathcal{A} separates points in $X \times Y$.
 - Show that \mathcal{A} doesn't vanish anywhere in $X \times Y$.
 - Assume that $k: X \times Y \rightarrow \mathbb{R}$ is continuous and that $\epsilon > 0$. Show that there are continuous functions $f_1, f_2, \dots, f_N: X \rightarrow \mathbb{R}$ and $g_1, g_2, \dots, g_N: Y \rightarrow \mathbb{R}$ such that $|k(x, y) - \sum_{i=1}^N f_i(x)g_i(y)| < \epsilon$ for all $x \in X, y \in Y$.
11. Let \mathcal{A} be the collection of all functions $f: \mathbb{C} \rightarrow \mathbb{C}$ of the form

$$f(z) = \sum_{\substack{0 \leq n \leq N \\ 0 \leq m \leq M}} c_{nm} z^n \bar{z}^m$$

where N, M are nonnegative integers, and c_{nm} are complex numbers. Show that if K is a nonempty, compact subset of \mathbb{C} , then \mathcal{A} is dense in $C(K, \mathbb{C})$.

12. In this problem, we shall look more deeply into the trigonometric polynomials in Proposition 4.11.12 and also into ordinary, complex polynomials. Along the way, we shall need to integrate complex valued functions, and we shall do it componentwise:

If $f(\theta) = u(\theta) + iv(\theta)$ where $u, v: [-\pi, \pi] \rightarrow \mathbb{R}$ are continuous functions, we define the indefinite integral by

$$\int f(\theta) d\theta = \int u(\theta) d\theta + i \int v(\theta) d\theta$$

and correspondingly for definite integrals.

a) Show that

$$\int e^{ia\theta} d\theta = \frac{1}{ia} e^{ia\theta} + C$$

for all real $a \neq 0$.

b) Show that if $n \in \mathbb{Z}$, then

$$\int_{-\pi}^{\pi} e^{ik\theta} d\theta = \begin{cases} 0 & \text{if } k \neq 0 \\ 2\pi & \text{if } k = 0 \end{cases}$$

c) Show that if $p(\theta) = \sum_{n=-N}^N c_n e^{in\theta}$ is a trigonometric polynomial, then

$$\int_{-\pi}^{\pi} |p(\theta)|^2 d\theta = 2\pi \sum_{n=-N}^N |c_n|^2$$

(recall that $|p(\theta)|^2 = p(\theta)\overline{p(\theta)}$).

d) Let \mathcal{A} consist of all trigonometric polynomials $f: [\pi, \pi] \rightarrow \mathbb{C}$ of the form $p(\theta) = \sum_{n=0}^N c_n e^{in\theta}$ (note that we are only using nonnegative indices). Show that \mathcal{A} is an algebra in $C(X, \mathbb{C})$ that separates points and does not vanish anywhere (as in the text, X is the unit circle in \mathbb{R}^2).

e) Show that $\int_{-\pi}^{\pi} |e^{-i\theta} - p(\theta)|^2 d\theta \geq 2\pi$ for all $p \in \mathcal{A}$.

f) Explain that $\overline{\mathcal{A}} \neq C(X, \mathbb{C})$. This shows that the complex Stone-Weierstrass Theorem 4.11.11 doesn't hold if we remove the condition that \mathcal{A} is closed under conjugation.

g) Let $D = \{z \in \mathbb{C} : |z| \leq 1\}$ be the unit disk in \mathbb{C} , and let \mathcal{P} be the algebra of complex polynomials. Show that the conjugate function $z \mapsto \bar{z}$ is *not* in the closure of \mathcal{P} , and hence that the complex polynomials are not dense in $C(D, \mathbb{C})$. (Hint: If $p(z) = c_n z^n + c_{n-1} z^{n-1} + \dots + c_1 z + c_0$ is a complex polynomial, $p(e^{i\theta}) = c_n e^{in\theta} + c_{n-1} e^{i(n-1)\theta} + \dots + c_1 e^{i\theta} + c_0$ when $z = e^{i\theta}$ is a point on the unit circle.)

Notes and references to Chapter 4

Although uniform convergence is a crucial tool in the study of series, the concept wasn't discovered till the middle of the 19th century. Weierstrass used it in a study of power series written in 1841, but as the paper was first published many years later, the concept was mainly disseminated through his lectures. Philipp Ludwig von Seidel (1821-1896) and George Gabriel Stokes (1819-1903) seem to have discovered the notion independently. Uniform continuity is also a concept associated with Weierstrass and his school, but it seems to have been discussed by Bolzano already in the 1830's (there is some disagreement on how Bolzano's writings should be interpreted). Equicontinuity was introduced by Giulio Ascoli (1843-1896) who used it to prove the first half of the Ascoli-Arzelà Theorem in 1884. The second half was proved by Cesare Arzelà (1847-1912) in 1895. The books by Gray [14] and Bressoud [7] will show you how difficult the different notions of convergence

and continuity were to grasp for even the greatest mathematicians of the 18th and 19th century.

The idea of function spaces – spaces where the elements are functions – may first have occurred in Vito Volterra’s (1860-1940) work on the calculus of variations, a part of mathematics where it’s quite natural to think of functions as variables as one often wants to find the function that maximizes or minimizes a certain expression. Volterra’s ideas were taken up by Jacques Hadamard (1865-1963) who coined the word “functional” for functions that take other functions as variables.

The proof of the existence of solutions to differential equations that we studied in Section 4.7 goes back to Émile Picard (1856-1941) and Ernst Lindelöf (1870-1946), and is often referred to as *Picard iteration*. The alternative approach in Section 4.9 originates with another member of the Italian school, Giuseppe Peano (1858-1932).

Weierstrass proved his approximation theorem in 1885. It was generalized by Marshall H. Stone (1903-1989) in 1937 to what is now known as the Stone-Weierstrass Theorem. Bernstein’s probabilistic proof of Weierstrass’ theorem has given rise to an important subfield of numerical analysis known as *constructive function theory*.

If you want to go deeper into the theory of differential equations and dynamical systems, Meiss’ book [27] is an excellent place to start. The classical text by Arnold [3] relies more on geometric intuition, but is full of insights. If you want to know more about the calculus of variations, the books by van Brunt [41] and Kot [23] are excellent introductions on a suitable level.

Bibliography

- [1] R. Abraham, J.E. Marsden, T. Raitu, *Manifolds, Tensor Analysis, and Applications*, 2nd Edition, Springer-Verlag, 1988.
- [2] Tom M. Apostol, *Calculus, Volume 1*, 2nd Edition, Wiley, 1967.
- [3] Vladimir I. Arnold, *Ordinary Differential Equations*, Springer-Verlag, 1992.
- [4] Michael F. Barnsley, *Fractals Everywhere*, 2nd Edition, Academic Press, 1993.
- [5] Patrick Billingsley, *Probability and Measure*, 3rd Edition, John Wiley & Sons, 1999. (Steer clear of the Anniversary Edition of 2012 – it is full of misprints.)
- [6] David M. Bressoud, *A Radical Approach to Lebesgue's Theory of Integration*, Cambridge University Press, 2008.
- [7] ———, *A Radical Approach to Real Analysis*, 2nd Edition, Mathematical Association of America, 2007.
- [8] ———, *Second Year Calculus*, Springer-Verlag, 1991.
- [9] Henri P. Cartan, *Differential Calculus*, Kershaw, 1971 (if you read French, you may prefer the original version: *Calcul Différentiel*, Hermann, 1967).
- [10] Donald L. Cohn, *Measure Theory*, 2nd Edition, Birkhäuser, 2013.
- [11] Rodney Coleman, *Calculus on Normed Vector Spaces*, Springer-Verlag, 2012.
- [12] Kenneth R. Davidson, Allan P. Donsig, *Real Analysis and Applications*, Springer-Verlag, 2009.
- [13] Gerald B. Folland, *Real Analysis: Modern Techniques and Applications*, 2nd Edition, John Wiley & Sons, 1999.
- [14] Jeremy Gray, *The Real and the Complex: A History of Analysis in the 19th Century*, Springer-Verlag, 2015.
- [15] D.H. Griffel, *Applied Functional Analysis*, 2nd Edition, Dover Books, 2002 (first published by Ellis Horwood 1985).
- [16] Paul R. Halmos, *Naïve Set Theory*, Springer-Verlag, 1974 (first published by Van Nostrand 1960).
- [17] Richard Hammack, *Book of Proof*, 2nd Edition, 2013. Privately published, downloadable from <http://www.people.vcu.edu/~rhammack/BookOfProof/>
- [18] Thomas Hawkins, *Lebesgue's Theory of Integration: Its Origin and Development*, Chelsea Publishing, 1975.
- [19] John Hubbard, Barbara Burke Hubbard, *Vector Calculus, Linear Algebra, and Differential Forms: A Unified Approach*, 5th Edition, Matrix Editions, 2015
- [20] Victor J. Katz, *A History of Mathematics*, 3rd Edition, Pearson, 2014.
- [21] T.W. Körner, *A Companion to Analysis. A Second First or a First Second Course in Analysis*, American Mathematical Society, 2004.
- [22] ———, *Fourier Analysis*, Cambridge University Press, 1988.
- [23] Mark Kot, *A First Course in the Calculus of Variations*, American Mathematical Society, 2014.

-
- [24] Steven G. Krantz, Harold R. Parks, *The Implicit Function Theorem: History, Theory, and Applications*, Birkhäuser, 2003.
 - [25] Tamara J. Lakins, *The Tools of Mathematical Reasoning*, American Mathematical Society, 2016.
 - [26] John N. McDonald, Neil A. Weiss, *A Course in Real Analysis*, 2nd Edition, Academic Press, 2013.
 - [27] James D. Meiss, *Differential Dynamical Systems*, Society for Industrial and Applied Mathematics (SIAM), 2007.
 - [28] Frank Morgan, *Real Analysis*, American Mathematical Society, 2005.
 - [29] James R. Munkres, *Analysis on Manifolds*, Addison-Wesley, 1991.
 - [30] ———, *Topology*, 2nd Edition, Prentice-Hall, 2000.
 - [31] Joel W. Robbin, On the existence theorem for differential equations, *Proceedings of the American Mathematical Society*, **19**, 1968, pp. 1005-1006.
 - [32] Bryan P. Rynne, Martin A. Youngson, *Linear Functional Analysis*, Springer-Verlag, 2008.
 - [33] René L. Schilling, *Measures, Integrals, and Martingales*, Cambridge University Press, 2005.
 - [34] Michael Spivak, *Calculus*, 3rd Edition, Cambridge University Press, 1994.
 - [35] ———, *Calculus on Manifolds*, Benjamin, 1965.
 - [36] Elias M. Stein, Rami Shakarchi, *Fourier Analysis*, Princeton University Press, 2013.
 - [37] Terence Tao, *Analysis 1*, 3rd Edition, Springer-Verlag, 2016 (first published by Hindustan Book Agency, 2006).
 - [38] ———, *Analysis 2*, 3rd Edition, Springer-Verlag, 2016 (first published by Hindustan Book Agency, 2006).
 - [39] ———, *Measure Theory*, American Mathematical Society, 2011. Can be downloaded from:
<https://terrytao.files.wordpress.com/2012/12/gsm-126-tao5-measure-book.pdf>
 - [40] Georgi P. Tolstov, *Fourier Series*, Dover Books, 1976.
 - [41] Bruce van Brunt, *The Calculus of Variations*, Springer-Verlag, 2006.
 - [42] John B. Walsh: *Knowing the Odds: An introduction to Probability*, American Mathematical Society, 2012.

Index

A° , 54
 $B(X, Y)$, 99
 $C(X, Y)$, 102
 $C_b(X, Y)$, 101
 X/\sim , 19
 \Longleftrightarrow , 6
 \implies , 6
 \circ , 13
 \inf , 30
 \liminf , 33
 \limsup , 33
 \mapsto , 13
 \mathbb{C} , 9
 \mathbb{Q} , 8
 \mathbb{Z} , 8
 $B(a; r)$, 49, 52
 \mathbb{N} , 8
 \overline{A} , 54
 $\overline{B}(a; r)$, 53
 \mathbb{R} , 8
 \mathbb{R}^n , 9
 ρ , 99, 101
 \setminus , 10
 \sim , 18
 \sup , 30
 \times , 10
 c , 10
 d_A , 46
 $f(A)$, 14
 f^{-1} (inverse function), 16
 $f^{-1}(B)$ (inverse image), 14

Abel's Summation Formula, 95

Abel's Theorem, 95
Abel, Niels Henrik (1802-1829), 94
Alexandrov, Pavel Sergeyevich
(1896-1982), 76
algebra
 of functions, 123
 of sets, 12
Arzelà, Cesare (1847-1912), 131
Arzelà-Ascoli Theorem, 110, 113
Ascoli, Giulio (1843-1896), 131

ball
 closed, 53
 open, 49, 52
Banach's Fixed Point Theorem, 61
Banach, Stefan (1892-1945), 76
Bernoulli, Jakob (1654-1705), 42
Bernoulli, Johann (1667-1748), 42
Bernstein polynomials, 118
Bernstein, Sergei (1880-1968), 116
bijection, 16
bijective, 16
Bolzano, Bernhard (1781-1848), 42, 131
Bolzano-Weierstrass Theorem, 38
Boole, George (1815-1864), 9
boolean operations, 9
boundary point, 53
bounded
 function, 99
 sequence, 32
 set, 63
bounded above, 30
bounded below, 30

- Bourbaki, Nicolas, 76
- Cantor's diagonal argument, 21
- Cantor, Georg (1845-1918), 22
- cartesian product, 10
- Cauchy sequence, 59
 - in \mathbb{R}^m , 34
- Cauchy, Augustin Louis (1789-1857), 42
- Chebyshev's inequality, 117
- closed
 - ball, 53
 - in terms of sequences, 55
 - set, 53
- closed under conjugation, 128
- closure, 54
- compact set, 63
 - in \mathbb{R}^m , 64
 - in $C(X, \mathbb{R}^m)$, 110
- compactness, 63
 - and total boundedness, 66
 - in terms of open coverings, 70
 - in terms of the finite intersection property, 70
- complement, 10
- complete, 60
- completeness
 - in metric space, 60
 - of \mathbb{R}^n , 34
 - of $B(X, Y)$, 100
 - of $C(X, Y)$, 102
 - of $C_b(X, Y)$, 101
- Completeness Principle, 30, 36
- completion
 - of a metric space, 72, 76
- composite function, 13
- continuous, 51
 - at a point in \mathbb{R} , 26
 - at a point in \mathbb{R}^m , 28
 - at a point in a metric space, 49
 - equi-, 80, 108
 - in terms of closed sets, 56
 - in terms of open sets, 56
 - in terms of sequences, 50
 - pointwise, 79
 - uniformly, 79
- contraction, 61
- contraction factor, 61
- contrapositive proof, 6
- converge
 - in \mathbb{R}^m , 25
 - in metric space, 48
 - of series, 88
 - pointwise, 81, 88
 - uniformly, 81, 83, 87–89, 99
- convergence
 - radius of, 92
- countability of \mathbb{Q} , 21
- countable, 20
- covering, 69
- De Morgan's laws, 10, 11
- decreasing sequence, 32
- Dedekind, Richard (1831-1916), 42
- dense, 72, 107
- differential equation, 103, 112
- Dini's Theorem, 85
- discrete metric, 46
- disjoint, 9
- distributive laws, 9, 11
- embedding, 47
- equicontinuous, 80, 108
- equivalence class, 18
- equivalence relation, 18
- equivalent, 6
- Euler's method, 112
- Euler, Leonhard (1707-1783), 42
- exterior point, 53
- Extreme Value Theorem
 - in \mathbb{R} , 39
 - in metric spaces, 65
- family, 11
- finite intersection property, 70
- finite subcovering, 69
- fixed point, 60
- Fréchet, Maurice (1878-1973), 76
- function, 13
- Gregory's formula for π , 97
- Hadamard, Jacques (1865-1963), 132
- Hamming-metric, 45
- Hausdorff, Felix (1868-1942), 76
- Heine-Borel Theorem, 42, 70, 71
- if and only if, 6
- if...then, 5
- image, 14
 - forward, 14
 - inverse, 14
 - of compact set, 65
- imply, 6
- increasing sequence, 32
- indexed set, 12

- Induction Principle, 7
- infimum, 30
- initial conditions, 104
- injection, 16
- injective, 15, 16
- integral equation, 104
- interior, 54
- interior point, 53
- Intermediate Value Theorem, 37
- intersection, 9, 11
- inverse function, 16
- inverse image, 14
- inverse triangle inequality, 47
- isometry, 46
- iterate, 60

- Lagrange, Joseph Louis (1736-1813), 42
- Laplace, Pierre-Simon (1749-1827), 42
- Leibniz' formula for π , 97
- Leibniz, Gottfried Wilhelm (1646-1716), 42
- limit inferior, 33
- limit superior, 33
- Lindelöf, Ernst (1870-1946), 132
- locally compact, 68
- lower bound, 30

- Madhava's formula for π , 97
- Manhattan metric, 45
- map, 13
- mapping, 13
- maximum point, 39
- Mean Value Theorem
 - in \mathbb{R} , 40
- metric, 44
 - discrete, 46
 - truncated, 103
- metric space, 44
- minimum point, 39
- monotone sequence, 32

- neighborhood, 55
- Newton, Isaac (1642-1727), 42

- one-to-one correspondence, 16
- one-to-one function, 16
- onto, 16
- open
 - ball, 49, 52
 - covering, 69
 - set, 53
- Open Covering Property, 69, 126

- partition
 - of a set, 18
- partition classes, 18
- Peano, Giuseppe (1858-1932), 132
- Picard iteration, 106
- Picard, Émile (1856-1941), 106, 132
- pointwise continuous, 79
- pointwise convergence, 81
- power series, 92
 - differentiation of, 93
 - integration of, 93
- proof, 5
 - by contradiction, 6
 - by induction, 7
 - contrapositive, 6

- quotient construction, 19

- radius of convergence, 92
- real analytic function, 94
- reflexive relation, 18
- relation, 17
 - equivalence, 18
 - reflexive, 18
 - symmetric, 18
 - transitive, 18
- Rolle's Theorem, 40

- separable, 107
- separates points, 125
- sequence, 48
- sequences
 - differentiating, 90
 - integrating, 87
- series
 - differentiating, 90
 - integrating, 88
 - power, 92
- set, 8
- set theoretic difference, 10
- Stokes, George Gabriel (1819-1903), 131
- Stone, Marshall H. (1903-1989), 132
- Stone-Weierstrass Theorem, 126, 127
 - complex case, 128
- subsequence, 38, 63
- subspace metric, 46
- supremum, 30
- surjection, 16
- surjective, 16
- symmetric relation, 18

- Tauber's Theorem, 98
- Taylor series

- in \mathbb{R} , 94
- totally bounded, 66
- transitive relation, 18
- triangle inequality
 - in \mathbb{R}^n , 24
 - in metric spaces, 44
- inverse, 47
- trigonometric polynomial, 129
- truncated metric, 103

- uncountability of \mathbb{R} , 21
- uncountable, 20
- uniform convergence, 81, 83, 87–89, 99
- uniformly
 - continuous, 79
 - Lipschitz, 105
- union, 9, 11
- universe, 10
- upper bound, 30
- Urysohn, Pavel Samuilovich
(1898-1924), 76

- vanish anywhere
 - does not, 125
- Volterra, Vito (1860-1940), 132
- von Seidel, Philipp Ludwig (1821-1896),
131

- Weierstrass' M -test, 89
- Weierstrass' Theorem, 116
- Weierstrass, Karl Theodor Wilhelm
(1815-1897), 42, 131